Public control of rational and unpredictable epidemics

Charles Sims (cbsims@utk.edu), Faculty Fellow at the Howard H. Baker Jr. Center for Public Policy and an assistant professor in the Department of Economics at the University of Tennessee, 1640 Cumberland Ave, Knoxville, TN 37920.

David Finnoff (finnoff@uwyo.edu), Associate professor in the Department of Economics and Finance, 1000 E. University Avenue, University of Wyoming, Laramie, WY 82071.

Suzanne M. O'Regan (s.m.oregan@gmail.com), Postdoctoral Research Fellow, National Institute for Mathematical and Biological Synthesis, University of Tennessee

Abstract: Efforts to improve disease forecasting are often justified on the grounds that accurate and reliable forecasts offer opportunities to improve response to epidemics. Yet little is known about how disease forecasts influence the responsiveness to public health interventions. This paper couples an optimal stopping model and rational disease prevention with a stochastic compartmental epidemiological model of disease spread to develop a framework for evaluating the timing of public health interventions in response to rational but unpredictable epidemics. Optimal behavior by a public health agency is characterized by a pair of critical prevalence thresholds that trigger implementation and suspension of a public prevention program. Unlike existing economic thresholds for disease prevention, our decision thresholds account for time-varying probabilities of different levels of infection, prevention by private individuals, and the likelihood of crossing epidemiological thresholds that determine disease persistence.

Keywords: infectious disease, public health interventions, option value, bifurcation, hysteresis

JEL Codes: D81, Q57, Q58

Introduction

The inability of societies to eliminate many treatable disease can be traced, in some part, to a disconnect between individual and societal incentives to prevent the spread of disease. Individual efforts to prevent disease are costly only to the individual yet convey benefits to others (Gersovitz and Hammer 2004, Gersovitz and Hammer 2005). These uncaptured external benefits have aspects of a public good and result in under-provision of prevention by private individuals which places public health interventions in a central role. Existing work on the economic incentives associated with public health interventions typically adopt a static framework (Brito, Sheshinski et al. 1991, Xu 1999, Kureishi 2009), assume individuals have perfect information about disease outcomes (Wickwire 1977, Sethi and Staats 1978, Geoffard and Philipson 1996, Francis 1997, Geoffard and Philipson 1997, Gersovitz and Hammer 2005), or focus exclusively on individual preventative behavior (Auld 2003, Chen 2009, Aadland, Finnoff et al. 2013). We relax these assumptions to investigate how the inability to perfectly predict disease outcomes influences the responsiveness of public sector prevention programs.

The economic efficiency of public health interventions requires individuals to forecast the future evolution of the disease in order to weigh the present value of program costs and infections avoided through intervention (i.e., program benefits). Forecasting efforts typically utilize compartmental epidemiological models (Shaman and Karspeck 2012), agent-based models (Nsoesie, Marathe et al. 2013), or parametric statistical models (Soebiyanto, Adimi et al. 2010) to predict future disease metrics such timing, peak, and intensity. However, diseases are notoriously difficult to predict (Laporte 1993, Brooks, Farrow et al. 2014).¹ Efforts to improve

¹Many groups are now utilizing techniques from meteorology to develop epidemiological models that could predict impending disease activity. Google Flu Trends was originally developed for this purpose in 2008 but was decommissioned on 2015. In 2013, the Centers for Disease Control and Prevention (CDC) launched the “Predict the Influenza Season Challenge” with the goal of forecasting the timing, peak, and intensity of the 2013-2014 flu season.

disease forecasting are often justified on the grounds accurate and reliable forecasts offer opportunities to improve preparedness and response to epidemics (Laporte 1993, Brooks, Farrow et al. 2014). Yet little is known about how disease forecasts influence prevention decisions. For individual prevention decisions such as washing hands, wearing a mask or using a condom, pessimistic disease forecasts (higher expected incidence of disease) induce more current risky behavior (Auld 2003). However, public prevention programs differ from individual prevention decisions in two key ways.

First, disease prevention programs often impose sunk costs on society in the form of new facilities, staff training, or the development of websites and computer applications.² When future disease outcomes are unpredictable, there is an incentive to delay these sunk costs in order to respond to new information about the disease – a disease prevention option value (Pindyck 2007). The option value creates a path dependence, in which the optimality of public health interventions depends not only on the current state of the epidemic but also on the history of past interventions. The option value also implies that both the first (trend) and second (variance) moment of the disease forecast influence public prevention decisions. Thus, random fluctuations in key epidemiological parameters (Gualtieri and Hecht 2011, Allen and Lahodny Jr 2012) induced by changes in the number of individuals (demographic stochasticity) or due to random parameter variation resulting from external forces such as climate (environmental stochasticity) must be accounted for when evaluating the responsiveness of public health interventions.

Second, from the perspective of public health agencies, the complexity of individual prevention decisions adds an additional behavioral source of uncertainty to disease forecasts. Existing economic literature focuses on two types of preventative actions an individual may take

² For example, the federal government's Obamacare enrollment website has cost about \$2.1 billion (Wayne 2014).

in response to disease: vaccination (Geoffard and Philipson 1996, Francis 1997, Gersovitz and Hammer 2004, Nævdal 2012, Chen and Toxvaerd 2014) and reductions in potentially infectious contacts (Philipson and Posner 1993, Kremer 1996, Fenichel, Castillo-Chavez et al. 2011, Aadland, Finnoff et al. 2013). Both types of prevention decisions may be difficult for public health agencies to predict owing to information (Klein, Laxminarayan et al. 2007, Chen 2009), differences in the formation of disease expectations (Auld 2003, Fenichel and Wang 2013) and a divergence between subjective and objective risk of infection (Carman and Kooreman 2014). Since stochastic thresholds for disease elimination may differ from their deterministic counterparts (Allen and Lahodny Jr 2012), different sources of stochasticity may lead to different conclusions about the persistence of disease and thus the need for public health interventions. Thus efforts to account for individual prevention decisions when forecasting disease outcomes may improve forecasts on average but may also increase the program's option value leading to additional delays.

We couple optimal stopping (Chow, Robbins et al. 1971) and the rational epidemics literature (Geoffard and Philipson 1996) with a stochastic compartmental epidemiological model of disease spread (Gray, Greenhalgh et al. 2011) to develop a framework for evaluating the timing of public health interventions in response to rational but unpredictable epidemics. Private prevention is modeled as a reduction in the rate that susceptible individuals contact others (Gersovitz and Hammer 2004). Because contact rate reductions are a public good, a costly public prevention program can be implemented that further reduces the contact rate. However, three different sources of stochasticity make future forecasts of disease inaccurate, potentially compromising the cost-effectiveness of public investments in disease prevention. Using numerical examples based on gonorrhea and pneumococcus infections, computational techniques

are used to identify for a pair of critical prevalence thresholds that trigger implementation and suspension of a public prevention program. Implementing and suspending the prevention program when prevalence reaches these thresholds maximizes the expected discounted social welfare.

Two key results emerge from this approach. First, consistent with (Geoffard and Philipson 1996), public prevention programs may occur too late in the epidemic. However, the consideration of the prevention program's option value also suggests that programs may be implemented too soon. Public health programs must be implemented within a disease prevalence window. Acting too late can make the programs cost-prohibitive but acting too soon risks committing funds to prevent a disease that is less problematic than anticipated. Even modest degrees of disease forecast error can lead to substantial delays in program implementation which offers explanation for the common criticism that public sector responses to disease are too slow. Second, improvements in disease forecasting may not expedite public health interventions. Accounting for human behavior in response to disease will improve disease forecasts by lowering the expected path of the disease – a first moment forecast improvement. This improvement in forecast accuracy lowers the returns from immediate investment in public prevention and delays public health intervention. However, accounting for human behavior will also improve disease forecasts by lowering the standard deviation or error of the forecast - a second moment forecast improvement. This second moment improvement lowers the expected returns of *delaying* prevention programs and expedites public health interventions. The overall impact of incorporating human behavior into epidemiological models depends on whether the first moment or second moment effects dominate.

Public and private prevention of disease spread

A population N is divided into two disjoint classes: susceptible $S(t)$ and infectious $I(t)$. Individuals receive income M in the absence of the disease. Susceptible individuals contact others at rate $C(t)$ and become infected and infectious after successful contact with an infectious individual. Infection incurs a current money loss p_I such as lost wages or the monetary value of pain and suffering due to infection.

The average individual's contact rate decision is depicted in Figure 1. The individual receives benefits from making contacts (income earned through physical interactions at work, utility gained through social interactions) and in the absence of disease selects C_0 contacts to maximize benefits gained through contacts. The presence of disease makes contacts costly due to the risk of infection. As the prevalence of the disease rises, the benefits to the individual of contacts that perpetuate the spread of the disease decrease. This triggers a decline in the rate that susceptibles contact others: $dC/dt = -a_0(C, I) < 0$. For instance, individuals may choose to avoid public places, wash hands more frequently, or forgo school or work to avoid exposure to infected individuals.³ The private provision of preventative behaviors lowers the overall risk of infection, conveying benefits to others. As there is this public good aspect of the private disease prevention, there are welfare gains from public health interventions that encourage additional preventative behavior. For instance, government agencies may offer subsidies to encourage medical treatments, close schools and businesses, or impose quarantines in extreme cases. When public sector prevention programs are active, contact rates decline at rate $a_1(C, I) > a_0(C, I)$.

³ Extending the model to individual decisions to become immunized and public sector immunization programs requires an additional compartment in the epidemiological model that accounts for the vaccinated population.

Designing cost-effective public health interventions requires forecasting the trajectory of the disease and the human response to disease (Auld 2003). Based on expectation of future disease prevalence and human behavior in response to that prevalence, officials can implement a public health program when disease prevalence reaches a threshold level I^* , which instantly encourages additional preventative behavior. This program imposes two costs on society. First, the program requires an initial one-time sunk cost K . Implementing the program also initiates a stream of costs which may reflect total subsidy payments, expenditures on information campaigns, or operating costs for facilities and medical staff. Since prevention programs typically target susceptible individuals, the flow of program costs is $p_C(C)S(t)$ where $p_C(C)$ is the cost per susceptible individual. This per unit cost is decreasing in the contact rate ($\partial p_C / \partial C < 0$) as it will be more costly to incentivize additional preventative behavior when private prevention has already been supplied (see Figure 1). If p_C is interpreted as a subsidy, this is consistent with the majority of public-sector programs that provide subsidies to susceptible individuals to engage in preventative behavior (Geoffard and Philipson 1997). If the disease prevalence declines to I' , the program can be suspended which terminates the stream of costs $p_C(C)S$ and allows disease spread to be dictated by private incentives for prevention. If the disease re-emerges again at some future time $I(t) \geq I^*$, the program may be reinstated which once again encourages preventative behavior and starts the stream of costs $p_C(C)S$.

A risk-neutral social planner chooses if and when to implement and suspend a disease prevention program to maximize the expected present value of social welfare $W(I, C)$ given by the total income earned in the population net of the costs of the disease and the cost of the public health interventions. The problem is one of optimally switching between a regime where the

program is active and a regime where the program is suspended.⁴ Noting that $S = N - I$, the flow of returns to the social planner is

$$W(I, C) = \begin{cases} MN - p_I I & \text{if the program is inactive} \\ MN - p_I I - p_C(C)(N - I) & \text{if the program is active.} \end{cases} \quad (1)$$

The policy maker must determine when to move from the active regime to the inactive regime by choosing I^* . The policy maker must also determine when to move from the active regime to the inactive regime by choosing I' .

Based on traditional benefit-cost analysis, this program would be implemented when the expected present value of the avoided infection costs equals or exceeds the sunk cost and expected present value of the flow costs associated with the disease prevention program. However, when future prevalence of the disease and individual decision to engage in preventative behavior are uncertain, there is an incentive (an option value) to delay public health interventions longer than suggested by benefit-cost analysis. The delay allows public health officials to gain economically valuable information on the infectiveness of the disease and how individuals will respond to the disease. The benefit of this additional information must be balanced by the likelihood of additional infections.

The social planner must evaluate, at each instant in time, whether or not the program should be implemented given all future decisions are made optimally. Given the discount rate ρ , I^* satisfies

$$V(I_0, C_0) = \max_{I^*} E_0 \left[\int_0^{t^*(I^*)} W(I, C) e^{-\rho t} dt + \{[V(I, C) - K] e^{-\rho t^*}\} \right] \quad (2)$$

⁴ The model is similar to the entry and exit model of Dixit (1989) where market entry and exit is akin to entering the private market for disease prevention .

subject to $\frac{dI}{dt}, \frac{dC}{dt}$, $I(0) = I_0$, $C(0) = C_0$, where $t^*(I^*)$ is the expected time disease prevalence reaches I^* . The evaluation at each instant in time maximizes expected social welfare from that point forward by choosing to delay program implementation (whose payoff is defined as V_D) or to implement the program at cost K (whose payoff is defined as V_P).

Disease Dynamics

The disease is assumed to follow a susceptible-infected-susceptible (SIS) model which is the basis for most other epidemiological models and characterizes many STD's and bacterial infections where there is no protective immunity (Van den Driessche and Watmough 2000).⁵ Assuming the population is closed $\frac{dS}{dt} + \frac{dI}{dt} = \frac{dN}{dt} = 0$ implies per-capita birth rates and natural mortality rates are equal. Susceptibles are born at rate μN and die at rate μS . Susceptible individuals become infectious at rate βCSI where β is the probability that contact with an infectious individual leads to infection. Infected individuals recover from the disease at rate γI and die from the infection at rate μI . The deterministic SIS model is characterized by two differential equations:⁶

$$dS = \left[\mu N - \beta CS \frac{I}{N} + \gamma I - \mu S \right] dt \quad (3)$$

$$dI = \left[\beta CS \frac{I}{N} - \gamma I - \mu I \right] dt \quad (4)$$

⁵ In general, diseases caused by bacteria are of the SIS type, while diseases caused by virus are of the susceptible-infected-recovered (SIR) type (Brauer and Castillo-Chavez 2011). Applications of the SIS model include pneumococcus (Lamb, Greenhalgh et al. 2011), gonorrhea (Yorke, Hethcote et al. 1978), tuberculosis, measles, and the spread of computer viruses.

⁶ The closed system in (3)-(4) can be captured by a single differential equation equivalent to a logistic growth function where $\beta C - (\mu + \gamma)$ is the growth rate and $N \left(1 - \frac{\mu + \gamma}{\beta C} \right)$ is the carrying capacity. Similar epidemiological models have been used to model the diffusion of technological innovation (Stoneman 1983) and the spread of information through financial markets (Shiller and Pound 1989).

The average lifetime for the population is $1/\mu$ and $1/\gamma$ is the average infectious period.

Deterministic disease dynamics presume public health officials are able to perfectly forecast the future evolution of the disease. In fact, multiple sources of stochasticity will make disease spread difficult to predict. Chance independent events of individual mortality and reproduction cause random fluctuations in the population growth rate – demographic stochasticity. Following Allen (2003), we use a diffusion approximation to a continuous time Markov chain process to account for individual-level random events that alter birth and death processes.⁷ The infectiveness of the disease, β , is likely a random variable due to stochastic factors that simultaneously impact each individual in the population. For example many infections transmit more readily between individuals in cold conditions – environmental stochasticity. Following Gray, Greenhalgh et al. (2011), the transmission rate over dt is normally distributed with mean $\tilde{\beta}dt$ and variance $\sigma_\beta^2 dt$: $\beta dt = \tilde{\beta}dt + \sigma_\beta dz_\beta$ where dz_β is the increment of a standard Wiener process.

When accounting for these two key sources of stochasticity, the epidemiological system (3)-(4) becomes:

$$dS = \left[-\tilde{\beta}CS \frac{I}{N} + (\mu + \gamma)I \right] dt + \sqrt{\mu N} dz_1 - \sqrt{\mu S} dz_2 - \sigma_\beta CS \frac{I}{N} dz_\beta \quad (5)$$

$$dI = \left[\tilde{\beta}CS \frac{I}{N} - (\mu + \gamma)I \right] dt - \sqrt{\mu I} dz_3 + \sigma_\beta CS \frac{I}{N} dz_\beta \quad (6)$$

where dz_1 , dz_2 , and dz_3 are Wiener processes capturing stochastic shocks to births, susceptible deaths, and infected deaths. Equations (5) and (6) imply that the future number of susceptible and

⁷ See appendix for details. For simplicity, we assume that random events do not alter the recovery process. In our epidemiological model, we assume that environmental stochasticity dominates demographic stochasticity due to the random nature of the contact process.

infected individuals is uncertain. The closed system in (5) and (6) can be captured by a single differential equation

$$dI = \left[\tilde{\beta} C \left(1 - \frac{I}{N} \right) - (\mu + \gamma) \right] I dt - \sqrt{\mu I} dz_3 + \sigma_\beta C (N - I) \frac{I}{N} dz_\beta \quad (7)$$

The inability to predict the infectiveness of the disease causes future prevalence to be most unpredictable at intermediate levels. When the disease is close to dying out or completely infecting the entire population, random variation in infectiveness makes little difference in the prevalence at the next instant in time.⁸

A well-known summary statistic for disease dynamics is the basic reproduction number R_0 - the expected number of secondary cases produced by a single newly infected individual entering a disease free population at equilibrium (Jacquez and O'Neill 1991). In epidemiological models, $R_0 = 1$ serves as a transcritical bifurcation where the endemic and disease-free steady states meet and exchange stability (O'Regan and Drake 2013). If $R_0 < 1$, $I(t)$ tends to 0 exponentially (i.e., the disease dies out within the population). However, if $R_0 > 1$, the disease persists as $I(t)$ rises to an endemic level. For the deterministic model described by equations (3) and (4) $R_0^D = \frac{\beta C}{(\mu + \gamma)}$. If the transmission rate, βC , exceeds the rate of mortality in the infected population, the disease persists and eventually reaches an endemic population $I_e^D = N \left(1 - \frac{\gamma + \mu}{\beta C} \right)$. When environmental stochasticity is added to this dynamic epidemiological system, $R_0^S = R_0^D - \frac{\sigma_\beta^2 N^2}{2(\mu + \gamma)}$. Compared to the deterministic version of the model, the disease is more likely to die out if future disease

⁸ Due to the popularity of arithmetic and geometric Brownian motion processes, it is common in economic applications to assume an additive ($\sigma_\beta dz_\beta$) or multiplicative ($\sigma_\beta I dz_\beta$) noise term. The former case implies the standard deviation of I is constant while the latter assumes the standard deviation increases over time. Here we see that accounting for the specific source of the uncertainty in I (random variations in the parameter β) leads to very different conclusions.

infectiveness is very uncertain or the population is large (Gray, Greenhalgh et al. 2011). If the disease persists, the counterpart of the endemic equilibrium for the stochastic system is $I_e^S < I_e^D$.

Adaptive human behavior in response to increasing disease prevalence

In the epidemiological literature, C is often treated as a fixed parameter that incorporates both the rate of contact and the inherent infectiousness of an infected individual. This is consistent with the assumption that susceptible and infected individuals randomly match (homogenous mixing). However, as the prevalence of the disease rises, the benefits of preventative behavior increase causing susceptible individuals to make efforts to avoid contact with infectious individuals (Gersovitz and Hammer 2004). The change in the contact rate over time influences emergence and elimination of disease (O'Regan and Drake 2013) and the timing of public health interventions (Geoffard and Philipson 1996).⁹ Unfortunately, the data needed to specify the relationship between contact rates and disease prevalence is often lacking. For exposition we assume the following simple relationship: $C(t) = 1/I(t)$. We also make two additional assumptions that reflect information constraints likely facing individuals within the population. First, stochastic shocks to $I(t)$ are unobservable to private individuals. This implies the population level contact rate follows the expected infected population: $C(t) = 1/E[I(t)]$. Second, individuals view the total population at risk as much larger than N which is selected by public health officials. This implies that individuals will ignore natural limits to infection within the population and approximate $E[I(t)]$ by $I_0 e^{a_0 t}$. These simplifying assumptions imply that the percent decrease in the contact rate over time is constant: $dC/dt = -a_0 C$.

⁹ If $|\dot{C}|$ is very small in comparison to the disease dynamics, the model with adaptive human behavior may be characterized as a fast-slow system (O'Regan and Drake 2013). In this case, there is a short-run endemic equilibrium and a long-run disease-free equilibrium.

However, day-to-day changes in the contact rate may be difficult for public health officials to predict. While individual preferences for health states are routinely estimated (Evans and Viscusi 1991), data requirements make such approaches infeasible for informing population responses to a disease. Even if such data were available, an individual's subjective risk of infection is likely to differ from objective measures (Carman and Kooreman 2014). Also contact rates can be influenced by unexpected medical innovations which make predicting the evolution of disease even more difficult.¹⁰ To capture the unpredictable nature of adaptive human behavior we assume C evolves according to:

$$dC = -a_0 C dt + \sigma_C C dz_C \quad (8)$$

where $a_0 \geq 0$ captures responsiveness of the population to increased infection risk, $\sigma_C \geq 0$ captures the degree of uncertainty in individual preferences, dz_C is the increment of a standard Wiener process and $E[dz_C dz_\beta] = \delta dt$ is the covariance between shocks to the infectiveness of the disease and shocks to private prevention efforts with $\delta \in [-1, 1]$. If $\delta > 0$, a positive shock to the infectiveness of the disease is likely to be accompanied by a positive shock to contact rates and $\sigma = 0$ reflects completely independent sources of uncertainty. For example, lower temperatures increase the infectiveness of many viruses and may also encourage individuals to remain in doors in close contact to infected individuals. Equation (8) implies the average contact rate in the population is a log-normal random variable with mean $C_0 e^{-a_0 t}$ and variance $C_0^2 e^{-2a_0 t} (e^{\sigma_C^2 t} - 1)$.

¹⁰ For example, LaRiviere and Wolff (2015) find that the advent of Viagra led to an increase in gonorrhea rates among elderly men.

Computational solution to optimal switching problem

The problem is one of optimal switching with a two-dimensional state-space. According to Brekke and Øksendal (1994), the optimal switching problem can be rewritten as a set of variational inequalities. Prior to implementing the program, the value function when the program is inactive and the implementation curve, $I^*(C)$, satisfy the following Bellman equation

$$\begin{aligned} \rho V_D \geq MN - p_I I + \left[\tilde{\beta} C \left(1 - \frac{I}{N} \right) - (\mu + \gamma) \right] \frac{\partial V_D}{\partial I} - a_0 \frac{\partial V_D}{\partial C} + \frac{\sigma_C^2 C^2}{2} \frac{\partial^2 V_D}{\partial C^2} \\ + \frac{\left[\sigma_\beta C (N - I) \frac{I}{N} - \sqrt{\mu I} \right]^2}{2} \frac{\partial^2 V_D}{\partial I^2} + \left[\sigma_\beta C (N - I) \frac{I}{N} - \sqrt{\mu I} \right] \sigma_C C \delta \frac{\partial^2 V_D}{\partial I \partial C} \end{aligned} \quad (9)$$

and value matching condition

$$V_D[I(t), C(t)] \geq V_P[I(t), C(t)] - K \quad (10)$$

with one of the conditions satisfied at each point in the state space of $I(t)$ and $C(t)$. The left-hand side of (9) is the return a public health official would require to delay program implementation over the time interval dt . The right-hand side is the expected return from delaying program implementation over the interval dt . If (9) holds as an equality, it is optimal to delay program implementation. Equation (10) compares the total payoff when the program is inactive and active and acts as a boundary condition for the inactive regime. If (10) holds as an equality, it is optimal to immediately implement the program. The program implementation decision curve is the set of points where both conditions are met.

With the disease prevention program currently active, the value function and the curve that defines optimal program suspension, $I'(C)$, satisfy

$$\begin{aligned}
\rho V_P \geq & MN - p_I I - p_C(C)(N-1) + \left[\tilde{\beta} C \left(1 - \frac{I}{N} \right) - (\mu + \gamma) \right] \frac{\partial V_P}{\partial I} - a_1 \frac{\partial V_P}{\partial C} + \frac{\sigma_C^2 C^2}{2} \frac{\partial^2 V_P}{\partial C^2} \\
& + \frac{\left[\sigma_\beta C(N-I) \frac{I}{N} - \sqrt{\mu I} \right]^2}{2} \frac{\partial^2 V_P}{\partial I^2} + \left[\sigma_\beta C(N-I) \frac{I}{N} - \sqrt{\mu I} \right] \sigma_C C \delta \frac{\partial^2 V_P}{\partial I \partial C}
\end{aligned} \tag{11}$$

and

$$V_P[I(t), C(t)] \geq V_D[I(t), C(t)]. \tag{12}$$

Equation (11) compares the required and expected return from delaying program suspension. If (11) holds as an equality, it is optimal to keep the program active. Much like equation (10), equation (12) acts a boundary condition for the active regime. If (12) holds as an equality, it is optimal to suspend the program.

Mathematically, V_D and V_P are the solutions to the partial differential equations in (9) and (11). The two stochastic variables I and C and the dual program regimes require numerical methods to approximate these unknown value functions (Judd 1998, Miranda and Fackler 2002). We approximate $V_D[I(t), C(t)]$ and $V_P[I(t), C(t)]$ over a subset of the state space using piecewise linear basis functions and a linear spline constructed using upwind finite difference approximations (Balikcioglu, Fackler et al. 2011, Marten and Moore 2011). The approximation procedure solves for the $2 \times n^2$ basis function coefficients which satisfy (9) – (12) at a set of n nodal points spread evenly over the two-dimensional state-space. For more information see Miranda and Fackler (p. 129, 2002). The resulting complementarity problem is solved in Matlab using the smoothing-Newton root finding method (Qi and Liao 1999). Instead of an explicit solution, the public health program implementation, $I^*(C)$, and suspension, $I'(C)$, curves are sets

of n points where these conditions are met.

Illustrative Examples

Given the necessity of numeric solution techniques, we apply the model and approximation procedure to a realistic but hypothetical private and public response to disease (see Table 1). Specifically, the contact rate in the population is assumed to decrease by over 6% each year ($a_0 = 0.00017$) in response to increasing prevalence of the disease with a 1% annual volatility around this trend ($\sigma_C = 0.01$). The disease prevention program is assumed to double the rate that contacts decline in response to disease. Income in absence of the disease is treated as the numeraire ($M = 1$) and infection is assumed to completely eliminate one's ability to work and earn an income ($p_I = 1$). The sunk cost associated with implementing the public health program is equal to 200 times an individual's income ($K = 200$). Lastly we assume that the flow cost per susceptible individual triggered through the implementation of the program increases linearly as private prevention is undertaken: $p_C(C_0 - C)$. This is consistent with a linear demand for contacts (with slope p_C and intercept $p_C C_0$) and a subsidy program that must pay susceptible individuals their willingness to pay for contacts to encourage contacts reductions (i.e., prevention). We then investigate the sensitivity of model results to changes in economic parameters.

To highlight the role of epidemiology in determining the timing of public health interventions, this public and private response is considered in the context of two disease outbreaks: gonorrhea and pneumococcus (Gray, Greenhalgh et al. 2011). Gonorrhea incidence peaked in the mid-1970s and fell for many years thanks to a national gonorrhea control program

implemented in the late-1970s. However, cases have risen in the past few years. Between 2011 and 2012, gonorrhea incidence in the United States increased by 4.1% between 2011 and 2012 to a total of 334,826 cases. Pneumococcal disease is responsible for 4 million illness episodes, 445,000 hospitalizations and 22,000 deaths annually for a total direct medical cost of \$3.5 billion (Huang, Johnson et al. 2011). Pneumococcus is a bacterium that may be passed from person to person through direct contact and respiratory droplets (e.g., coughs and sneezes) of an infected person. While pneumococcus is commonly carried in the upper throat of children causing no infection, it can cause various infections such as ear and sinus infections, pneumonia, meningitis, and septicaemia.

Parameter values and reproductive numbers for both diseases are presented in Table 1. Epidemiological parameters for the model in (5) and (6) have been estimated or inferred using data on gonorrhea incidence in the United States (Hethcote and Yorke 1984) and data on a 2003 pneumococcus outbreak among children under two years in Scotland (McChlery, Scott et al. 2005). For the gonorrhea example, N is the sexually active population, C_0 is interpreted as the initial rate of unprotected sexual encounters, and $\tilde{\beta}$ represents the likelihood that unprotected sex results in gonorrhea transmission. For the pneumococcus example, N is the total population, C_0 is interpreted as more general contact (e.g., shaking hands), and the lower level of mean infectiveness reflects the lower likelihood of catching the disease through this more general contact.

Disease and private prevention forecasts

Expected and stochastic disease outcomes are presented in Figure 2 for gonorrhea and Figure 3 for pneumococcus. The expected contact rate (Figures 2A and 3A) and infected

population (Figures 2C and 3C) along with three Monte Carlo simulations are shown when private demand for prevention is inelastic (contact rate is constant) and elastic (contact rate declines as diseases prevalence increases). The volatility in the infected population in Figures 2C and 3C reflect uncertainty in disease infectiveness and random events in births and mortality. The volatility in contact rates in Figures 2A and 3A reflects public health officials' inability to accurately predict human behavior. The expected and stochastic disease outcomes can be mapped into the 2-dimensional state space to illustrate the likelihood of observing specific combinations of $I(t)$ and $C(t)$. With private prevention lowering the contact rate over 6% per year, the expected infected population $E[I]$ for each contact rate follows the dashed lines in Figures 2D and 3D. When gonorrhea is first detected in the population ($I_0 = 100$), the average number of unprotected sexual encounters per year is 18 ($C_0 = 0.05$). As the infection spreads, the contact rate declines. Private prevention eventually causes the infected population to peak in a little over 2 years at 1700 infected individuals. When pneumococcus is first detected in the population ($I_0 = 100$), the average number of contacts per year is 1,569 ($C_0 = 4.30$). Private prevention eventually causes the expected infected population to peak in a little over 18 months at 67,405 infected individuals.

When private demand for prevention is inelastic, both diseases are expected to persist: $R_0^D > 1$. With private demand for prevention that is elastic, private prevention lowers the contact rate causing the economic-epidemiological system to eventually reach the disease-free equilibrium. By hastening this decline in contact rates, the role of public health programs is to hasten the expected time to eradication within the population. However, the economic return from investing to hasten eradication depends on the likelihood of different disease and prevention outcomes. Points in Figures 2D and 3D show combinations of contact rate and prevalence

generated by the three Monte Carlo simulations in Figures 2A, 3A, 2C, and 3C. Returns from the public prevention program will be less than expected when the infected population is lower than expected. An outbreak that is less severe than expected means fewer infections avoided due to the public health program and larger than expected program costs. Points above the dashed lines in Figures 2D and 3D indicate outcomes where the infected population is higher than expected while points below the dashed line indicate outcomes where the infected population is lower than expected. The points for pneumococcus are more dispersed along the y-axis indicating a relatively greater amount of demographic stochasticity and uncertainty in disease infectiveness relative to uncertainty in contact behavior. Thus a large part of the risk associated with investing in a pneumococcus prevention program is associated with observing fewer avoided damages than expected.

Returns from the public prevention program will also be less than expected when the contact rate is lower than expected. When private prevention is more aggressive than expected, program costs are higher than expected and there is a greater likelihood that a negative shock to the contact rate may unexpectedly lead to eradication by crossing the bifurcation threshold. Points above and to the right of the dashed lines in Figures 2D and 3D indicate outcomes where the contact rate is higher than expected while points above and to the left of the dashed line indicate outcomes where the contact rate is lower than expected. Points below the dashed line may represent contact rates that are higher or lower than expected depending on whether prevalence is rising or falling. The points for gonorrhea are more dispersed left to right indicating a relatively greater amount of uncertainty in contact behavior relative to environmental or demographic sources of stochasticity. Thus, for gonorrhea, the riskiness of a prevention program

is associated with the likelihood a negative shock will move the contact rate across the bifurcation threshold, stabilize the disease-free equilibrium, and render the prevention program obsolete.

Optimal timing of gonorrhea and pneumococcus prevention programs

The optimal implementation and suspension curves for gonorrhea and pneumococcus prevention programs are presented in Figures 4A and 4B respectively. The optimality of public health intervention depends on both the incidence of disease $I(t)$ and the private provision of disease prevention $C(t)$. When the current state of the epidemic, defined by the pair (I, C) , is in the hysteresis region, no change should be made to the prevention program. A sufficiently large increase in the contact rate (rightward movement along the x-axis) or a sufficiently large increase in the infected population (upward movement along the y-axis) will justify a public prevention program. Once the program has been implemented, a sufficiently large decrease in the contact rate (leftward movement along the x-axis) or a sufficiently large decrease in the infected population (downward movement along the y-axis) will justify suspending the prevention program.

Both implementation thresholds are downward sloping in the contact rate indicating a greater propensity to intervene in an epidemic when the supply of prevention from the private sector is low. Our results suggest that a gonorrhea prevention program should never be implemented if private prevention lowers the contact rate below 0.033 (12 unprotected sexual contacts per year). Likewise, a pneumococcus prevention program should never be implemented if private prevention lowers the contact rate below 2. The costs of the program outweigh the benefits from avoided infections when contact rates are already very low. Along this curve, the

cost savings from terminating the program are larger than the increase in expected damages that arise from a relatively slower decline in contact rates.

The implementation thresholds in Figure 4A and 4B account for the probability of each disease and private prevention outcome illustrated in Figures 2 and 3. They also account for the change in system stability that arises due to changes in the contact rate. There are two equilibria associated with the epidemiological model in equation (7). Figure 4C and 4D are bifurcation diagrams that show the stability of the endemic and disease-free equilibria implied from equation (7). When the contact rate is relatively high, the endemic equilibrium is stable and the disease-free equilibrium is unstable. The endemic equilibrium declines as private prevention lowers the contact rate. When the contact rate reaches 0.036 for gonorrhea and 2.1 for pneumococcus, the endemic equilibrium becomes unstable and the disease-free equilibrium becomes stable. As might be expected, program implementation is more likely when the endemic equilibrium is stable. Furthermore, a negative shock to the contact rate makes it more likely the disease is eradicated and less likely the public prevention program is implemented. However, a public prevention program may still be justified even if the disease is expected to die out due to the unpredictability of the epidemic.

Figure 5 overlays the expected and stochastic paths in Figure 3 over the controlled state space in Figure 4 to illustrate the likelihood of different public health intervention strategies. Dashed lines represent the expected path of prevalence and contact rate. If the disease and private prevention follow their expected trajectories, the gonorrhea prevention program should be delayed nearly 1 year until the infected population reaches 731. The pneumococcus prevention program should be delayed over 7 months until the infected population reaches 9,500. The more expedient response to pneumococcus is due to its more rapid accumulation of infection losses (due to its

more rapid spread) and the lower program costs associated with encouraging hand washing and wearing masks compared to the costs of encouraging protected sex.

Geoffard and Philipson (1996) note that public health interventions may not be Pareto-optimal if they are undertaken too late. Our numerical results suggest that public health officials have a window of opportunity in which to intervene in the management of disease. Similar to Geoffard and Philipson, delaying gonorrhea and pneumococcus prevention programs too long raises the program's costs which may eventually overwhelm potential benefits from preventing infection in a closed population. However, implementing programs too soon risks committing scarce funds to prevent diseases that are less damaging than anticipated either because environmental or demographic variability significantly lower the prevalence of disease or private individuals adopt more preventative behavior than expected. For example, while gonorrhea and pneumococcus prevention programs are justified based on expected values of $I(t)$ and $C(t)$, the green simulation in the gonorrhea example and blue simulation in the pneumococcus example never cross the implementation threshold. This implication of forecast error reflects a reality for many federal agencies forced to sell or donate excess H1N1 vaccines after dire predictions of a swine flu outbreak did not materialize.

Implications for program design and disease forecasting

While we account for two different diseases, our numerical results are based on a single prevention program and a fixed degree of private prevention. Figure 6, investigates the sensitivity of the model to changes in private prevention behavior and characteristics of the public prevention program. The star in the center of the figure indicates the threshold for program implementation based on expected trajectories of the two stochastic variables. Each symbol indicates how a

change in a parameter shifts these implementation thresholds with filled symbols indicating an increase in the benchmark parameter value and open symbols indicating a decrease. Because programs are implemented when disease prevalence is on the rise, optimally implementing the program at a higher prevalence also means delaying until more private prevention has undertaken.

As expected increases in program costs (K, p_c), decreases in losses from infection (p_I) or decreases in program effectiveness (a_1) make it optimal to delay program implementation until a larger number of people have become infected relative to the benchmark prevalence. These parameters shift the implementation and suspension curves in figure 5 but do not alter the expected path of the disease (dashed lines in figure 5), leading to predictable impacts on responsiveness of public intervention. If a_1 is selected to internalize the public good provided by individual prevention decisions, this result suggests that private prevention behavior that generates a larger public good will also trigger a more expedient public health intervention.

The model also allows us to investigate how forecasts of $I(t)$ and $C(t)$ influence the responsiveness of public sector prevention programs (Figure 6). For our purposes, disease forecasts involve the mean ($\tilde{\beta}$ and a_0) and standard deviation (σ_c and σ_β) of the two stochastic variables. Changes in these forecast parameters shift the implementation and suspension curves in figure 5 and alter the expected path of the disease (dashed lines in figure 5). Compared to the benchmark, more optimistic forecasts of gonorrhea infectiveness (10% decrease in $\tilde{\beta}$) suggests the public prevention program should not be implemented in expectation. In contrast, more optimistic forecasts of pneumococcus infectiveness delays program implementation in expectation.

For both diseases, more optimistic forecasts of private prevention behavior (10% increase in a_0) delay the implementation of public prevention programs. In fact, if private demand for prevention is sufficiently elastic, the expected path of disease prevalence will never cross the implementation threshold – the public prevention program should not be implemented in expectation. This result is consistent with the idea that public prevention programs and increasing disease prevalence compete to induce protective activity (Geoffard and Philipson 1996). If forecasts assume the private demand for prevention is inelastic $a_0 = 0$, the gonorrhea prevention program would be implemented when the infected population reaches 125 instead of 751 – nearly eleven months too soon. Likewise, the pneumococcus prevention program would be implemented when the infected population reaches 170 instead of 9,500 – nearly seven months too soon. Failure to account for adaptive human behavior in epidemiological models may overestimate the expected return from public health interventions and encourage premature implementation of public health programs. This effect will be more pronounced the more elastic the actual private demand for prevention.

Changes in the forecast standard deviation or error are reflected by varying σ_C and σ_β . Here we find that forecast error tends to delay program implementation due to the program's option value. For the gonorrhea example, ignoring the error in private prevention forecasts suggests the program should be implemented when the infected population reaches 452 instead of 731. In the pneumococcus example, ignoring forecast error suggests the program should be implemented when the infected population reaches 1,025 instead of 9,500. In both cases, this is a difference of nearly 4 months in expectation. However, the impact of reducing forecast error depends on the disease context. For gonorrhea, reducing the forecast error associated with

infectiveness has a minimal impact whereas reducing forecast error associated with private prevention decisions has a more substantial impact on the timing of program implementation.

Conclusions

Evaluating the cost-effectiveness of public health interventions requires forecasts of disease spread in order to predict the costs and benefits of prevention programs. However, inaccuracies in disease forecasts can lead to misleading conclusions about the persistence of disease, the tendency for private individuals to engage in prevention, and subsequently the returns provided from investing public funds in disease prevention programs. Considerable effort has been devoted to improve the accuracy of disease forecasts under the assumption public sector prevention programs will respond more quickly to disease outbreaks. We combine an optimal stopping model of public investment in disease prevention with stochastic models of disease and private prevention that provides a framework for quantifying the effect of improvements in disease forecasting on the timeliness of public responses to disease.

We find two key results. First, public health programs must be implemented within a prevalence-contact window to ensure these programs efficiently utilize public funds. Program implementation that lags substantial reductions in the contact rate may render the program cost-prohibitive as it is forced to compete with private incentives for prevention. Implementing programs too soon risks committing funds to prevent a disease that is less problematic than anticipated. This serves as an opportunity cost of prevention since sunk costs associated with specific prevention programs cannot be reinvested toward other public health challenges.

Delaying the implementation of the program maximizes the expected discounted social welfare

by lowering the possibility that significant sunk costs are incurred to prevent a disease that is less damaging than initially expected. Thus pessimistic disease forecasts are bad news for public health but good news for the return from public health programs – a public health analog to Bernanke’s “bad news principle” (Bernanke 1983). This is in sharp contrast to previous findings about the implications of disease forecasts on private prevention decisions (Auld 2003).

Second, improvements in disease forecasting may not make public health agencies more responsive to disease. If improvements in disease forecasting arise from a more complete account of individual incentives to engage in preventative behavior, public health interventions are likely to be less responsive to disease outbreaks. Elastic private demand for prevention implies lower expected returns from investing in prevention. In contrast, efforts to reduce the standard error of disease forecasts will tend to make the public sector more responsive to disease outbreaks. These second-moment improvements in disease forecasting tend to lower the option value associated with public health programs which lowers the burden of proof agency officials require to invest public funds in prevention.

While the paper provides an initial indication of the implications of uncertainty and disease forecasting on the responsiveness of public health interventions, there is considerable work in this area that remains. First, the model utilizes a reduced form representation of individual human behavior that abstracts from individual risk preferences and heterogeneity among susceptible individuals. Second, the model considers various sources of stochasticity but there are likely other important sources of uncertainty influencing individual decisions to prevent the spread of disease. Imperfect monitoring and detection leads to questions about the current extent of the disease. Key parameters governing disease dynamics and associated economic damages may not be known early in the outbreak but will be better understood as the epidemic

progresses. Third, the model considers a very specific set of diseases that are described by an SIS epidemiological model. The dynamics of many infections do not conform well to predictions of the SIS model. This argues for future work directed at other types of compartmental epidemiological models. Finally the model focuses on preventative public health programs which reduce the contact rate. Many public health interventions involve immunization programs. Extending our model to account for this difference is left for future work.

Appendix

Demographic stochasticity following a Markov Chain process

We derive a system of stochastic differential equations that accounts for demographic stochasticity using a diffusion approximation to a continuous time Markov chain process. We consider the changes in both S and I populations and the probabilities associated with each change in population size ("event") in a small time interval Δt . Individual-level events are discrete and can occur randomly. In a sufficiently small time interval, at most one event can occur. For example, individuals are born with probability $(\mu N)\Delta t$ and the population size increases by one, i.e., $S(t + \Delta t) \rightarrow S(t) + 1$, or equivalently, $\Delta N = (\Delta S, \Delta I) = (1, 0)^T$. The population changes and their probabilities obtained from the above assumptions are listed in Table A1.

We calculate the expectation and covariance for the changes in the two populations N_i from the probabilities in Table A1. The expectation is $E(\Delta N) = \sum_{i=1}^4 p_i (\Delta N)_i$, given by the 2×1 vector

$$\mu \Delta t = \Delta t \begin{pmatrix} \mu I - \beta CS \frac{I}{N} \\ \beta CS \frac{I}{N} - \mu I \end{pmatrix}.$$

If N is sufficiently large, then ΔN has an approximate normal distribution with mean vector $\mu\Delta t$ and covariance matrix Σ (Allen 2003). To derive the 2×2 covariance matrix, we assume the number of events that occur in a time increment Δt is Poisson distributed. For example, the number of births of susceptible individuals is Poisson distributed with mean $(\mu N)\Delta t$. Provided the transition probabilities (the rates at which events occur) are sufficiently large, we can approximate Poisson random variables by normal random variables. For example, we can write the recruitment rate as a normal random variable,

$$(\Delta N)_1 = (\mu N)\Delta t + \sqrt{(\mu N)\Delta t}\eta_1,$$

where η_1 is a unit normal random variable. Then the number of births is normally distributed with mean $(\mu N)\Delta t$ and variance $(\mu N)\Delta t$. Assuming the number of events in a time interval Δt are Poisson distributed, we write down the following set of equations,

$$S(t + \Delta t) = S(t) + \sum_{j=1}^3 \text{Poisson}(p_j \Delta t) (\Delta S)_j$$

$$I(t + \Delta t) = I(t) + \sum_{j=1}^3 \text{Poisson}(p_j \Delta t) (\Delta I)_j$$

Using the normal distribution approximation to the Poisson distribution and denoting unit normal random variables by η_j , we obtain,

$$S(t + \Delta t) = S(t) + \sum_{j=1}^3 \left(p_j \Delta t + \sqrt{p_j \Delta t} \right) (\Delta S)_j \eta_j$$

$$I(t + \Delta t) = I(t) + \sum_{j=1}^3 \left(p_j \Delta t + \sqrt{p_j \Delta t} \right) (\Delta I)_j \eta_j$$

yielding

$$S(t + \Delta t) = S(t) + \mu_1 \Delta t + \sum_{j=1}^3 G_{1j} \sqrt{\Delta t} \eta_j$$

$$I(t + \Delta t) = I(t) + \mu_2 \Delta t + \sum_{j=1}^3 G_{2j} \sqrt{\Delta t} \eta_j$$

where $G_{ij} = (\Delta N_i)_j \sqrt{p_j} / \sqrt{\Delta t}$. The resulting matrix

$$G = \begin{pmatrix} \sqrt{\mu N} & -\sqrt{\mu S} & 0 \\ 0 & 0 & -\sqrt{\mu I} \end{pmatrix}$$

arising from the normal approximation to a Poisson distribution has the property that $GG^T = V$

(Allen et al. 2008). If $\Delta t \rightarrow 0$ and assuming the stochastic integral exists and is unique, then

$\sqrt{\Delta t} \eta_j \rightarrow dz_j$, where dz_j is a Wiener process. The equations converge to the system of stochastic differential equations,

$$dS = \left(\mu I - \beta CS \frac{I}{N} \right) dt + \sqrt{\mu N} dz_1 - \sqrt{\mu S} dz_2$$

$$dI = \left(\beta CS \frac{I}{N} - \mu I \right) dt - \sqrt{\mu I} dz_3.$$

Table A1. Change in S and I populations and the probabilities associated with each change in the population size (event) in a small time interval Δt

Event i	Change in population size $(\Delta N)_i$	Transition probability p_i
1	$(1,0)^T$	$\mu N \Delta t$
2	$(-1,0)^T$	$\mu S \Delta t$
3	$(0,-1)^T$	$\mu I \Delta t$

4	$(0,0)^T$	$1 - \sum_{i=1}^3 p_i$
---	-----------	------------------------

Table 1. Epidemiological and economic parameters

		<i>Gonorrhea</i>		<i>Pneumococcus</i>	
Description		Value	Source	Value	Source
Epidemiology					
N	Population	10,000	Gray et al. 2011	150,000	Lamb et al. 2011
I_0	Initial infected population	100	-	100	-
C_0	Initial contact rate	0.05	Yorke et al. 1978	4.30	Zhang et al. 2004
$\tilde{\beta}$	Infectiveness mean	0.5	Thin 1970	0.01	-
σ_{β}	Infectiveness standard deviation	1e-6/ C_0	-	1e-6/ C_0	-
γ	Rate of recovery	0.018	Yorke et al. 1978	0.02	Lamb et al. 2011
μ	Mortality/birth rate	6.8e-5	Yorke et al. 1978	1.4e-3	Zhang et al. 2004
R_0^D	Deterministic reproductive number	1.40	Yorke et al. 1978	2.00	-
R_0^S	Stochastic reproductive number	0.35	-	1.97	-
Private response to disease					
a_0	Private % reduction in contact rate	1.7e-4/day (6.2% annually)			
σ_C	% volatility in contact rates	0.01			
Public health program					
a_1	Public % reduction in contact rate	3.4e-4/day (12.4% annually)			
r	Social planner discount rate	0.02 annually			
K	Sunk cost of prevention program	100			
M	Income per person	1			
p_I	Cost per infected individual	1			
p_C	Cost of prevention per susceptible	0.06($C_0 - C$)		0.005($C_0 - C$)

Figure 1. Private provision of disease prevention in the face of a growing epidemic

Figure 2. Stability and expected outcomes of stochastic epidemiological model for gonorrhea. (A) Expected (solid) and stochastic (dashed) paths for contacts rates when private demand for prevention is inelastic and elastic. (B) 45 degree line. (C) Expected (solid) and stochastic (dashed) paths for infected population when private demand for prevention is inelastic and elastic. (D) Disease outcomes in the two dimensional state space (I, C).

Figure 3. Stability and expected outcomes of stochastic epidemiological model for pneumococcus. (A) Expected (solid) and stochastic (dashed) paths for contacts rates when private demand for prevention is inelastic and elastic. (B) 45 degree line. (C) Expected (solid) and stochastic (dashed) paths for infected population when private demand for prevention is inelastic and elastic. (D) Disease outcomes in the two dimensional state space (I, C).

Figure 4. Thresholds for disease persistence and public health intervention when demand for prevention is elastic. (A-B) Public prevention program implementation and suspension thresholds. (C-D) Bifurcation diagram for the stability of the endemic and disease-free equilibria.

Figure 5. Likelihood of implementing disease prevention program. Dashed lines represent the expected path of prevalence and contact rate. Dots show combinations of contact rate and prevalence generated by three Monte Carlo simulations of the stochastic model.

Figure 6. Impact of economic and uncertainty parameters on the expected program implementation thresholds for (A) gonorrhea and (B) pneumococcus. Legend indicates percent change in each parameter. Open symbols indicate a percent decrease in the benchmark parameter value while filled symbols indicate a percent increase in the benchmark parameter value

References

- Aadland, D., D. C. Finnoff and K. X. Huang (2013). "Syphilis Cycles." The BE Journal of Economic Analysis & Policy **13**(1): 297-348.
- Allen, L. J. and G. E. Lahodny Jr (2012). "Extinction thresholds in deterministic and stochastic epidemic models." Journal of Biological Dynamics **6**(2): 590-611.
- Auld, M. C. (2003). "Choices, beliefs, and infectious disease dynamics." Journal of Health Economics **22**(3): 361-377.
- Balikcioglu, M., P. L. Fackler and R. S. Pindyck (2011). "Solving optimal timing problems in environmental economics." Resource and Energy Economics **33**(3): 761-768.
- Bernanke, B. S. (1983). "Irreversibility, Uncertainty, and Cyclical Investment." The Quarterly Journal of Economics **98**(1): 85-106.
- Brauer, F. and C. Castillo-Chavez (2011). Mathematical Models in Population Biology and Epidemiology, Springer.
- Brekke, K. and B. Øksendal (1994). "Optimal Switching in an Economic Activity under Uncertainty." SIAM Journal on Control and Optimization **32**(4): 1021-1036.
- Brito, D. L., E. Sheshinski and M. D. Intriligator (1991). "Externalities and compulsory vaccinations." Journal of Public Economics **45**(1): 69-90.
- Brooks, L. C., D. C. Farrow, S. Hyun, R. J. Tibshirani and R. Rosenfeld (2014). "Flexible Modeling of Epidemics with an Empirical Bayes Framework." PLOS Computational Biology **11**(8).
- Carman, K. and P. Kooreman (2014). "Probability perceptions and preventive health care." Journal of Risk and Uncertainty **49**(1): 43-71.
- Chen, F. and F. Toxvaerd (2014). "The economics of vaccination." Journal of Theoretical Biology **363**: 105-117.
- Chen, F. H. (2009). "Modeling the effect of information quality on risk behavior change and the transmission of infectious diseases." Mathematical Biosciences **217**(2): 125-133.
- Chow, Y. S., H. Robbins and D. Siegmund (1971). Great Expectations: The Theory of Optimal Stopping, Houghton Mifflin Boston.
- Dixit, A. (1989). "Entry and Exit Decisions under Uncertainty." The Journal of Political Economy **97**(3): 620-638.
- Evans, W. N. and W. K. Viscusi (1991). "Estimation of state-dependent utility functions using survey data." The Review of Economics and Statistics: 94-104.

Fenichel, E. P., C. Castillo-Chavez, M. G. Ceddia, G. Chowell, P. A. G. Parra, G. J. Hickling, G. Holloway, R. Horan, B. Morin, C. Perrings, M. Springborn, L. Velazquez and C. Villalobos (2011). "Adaptive human behavior in epidemiological models." Proceedings of the National Academy of Sciences **108**(15): 6306-6311.

Fenichel, E. P. and X. Wang (2013). The mechanism and phenomena of adaptive human behavior during an epidemic and the role of information. Modeling the Interplay Between Human Behavior and the Spread of Infectious Diseases, Springer: 153-168.

Francis, P. J. (1997). "Dynamic epidemiology and the market for vaccinations." Journal of Public Economics **63**(3): 383-406.

Geoffard, P.-Y. and T. Philipson (1996). "Rational epidemics and their public control." International Economic Review: 603-624.

Geoffard, P.-Y. and T. Philipson (1997). "Disease eradication: private versus public vaccination." The American Economic Review: 222-230.

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

Gersovitz, M. and J. S. Hammer (2005). "Tax/subsidy policies toward vector-borne infectious diseases." Journal of Public Economics **89**(4): 647-674.

Gray, A., D. Greenhalgh, L. Hu, X. Mao and J. Pan (2011). "A stochastic differential equation SIS epidemic model." SIAM Journal on Applied Mathematics **71**(3): 876-902.

Gualtieri, A. F. and J. P. Hecht (2011). "Dynamics and Control of Infectious Diseases in Stochastic Metapopulation Models." Journal of Life Sciences **5**(7): 503-508.

Hethcote, H. W. and J. A. Yorke (1984). Gonorrhea Transmission Dynamics and Control. Berlin, Springer-Verlag.

Huang, S. S., K. M. Johnson, G. T. Ray, P. Wroe, T. A. Lieu, M. R. Moore, E. R. Zell, J. A. Linder, C. G. Grijalva and J. P. Metlay (2011). "Healthcare utilization and cost of pneumococcal disease in the United States." Vaccine **29**(18): 3398-3412.

Jacquez, J. A. and P. O'Neill (1991). "Reproduction numbers and thresholds in stochastic epidemic models I. Homogeneous populations." Mathematical Biosciences **107**(2): 161-186.

Judd, K. (1998). Numerical Methods in Economics, MIT Press.

Klein, E., R. Laxminarayan, D. L. Smith and C. A. Gilligan (2007). "Economic incentives and mathematical models of disease." Environment and Development Economics **12**(05): 707-732.

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

Kureishi, W. (2009). "Partial vaccination programs and the eradication of infectious diseases." Economics Bulletin **29**(4): 2758-2769.

Lamb, K. E., D. Greenhalgh and C. Robertson (2011). "A simple mathematical model for genetic effects in pneumococcal carriage and transmission." Journal of Computational and Applied Mathematics **235**(7): 1812-1818.

Laporte, R. E. (1993). "How to improve monitoring and forecasting of disease patterns." BMJ **307**(6919): 1573-1574.

LaRiviere, J. and H. Wolff (2015). "The Power of the Little Blue Pill: Innovations and Implications of Lifestyle Drugs in an Aging Population." Economic Inquiry **53**(1): 540-556.

Marten, A. L. and C. C. Moore (2011). "An options based bioeconomic model for biological and chemical control of invasive species." Ecological Economics **70**(11): 2050-2061.

McChlery, S. M., K. J. Scott and S. C. Clarke (2005). "Clonal analysis of invasive pneumococcal isolates in Scotland and coverage of serotypes by the licensed conjugate polysaccharide pneumococcal vaccine: possible implications for UK vaccine policy." European Journal of Clinical Microbiology and Infectious Diseases **24**(4): 262-267.

Miranda, M. and P. L. Fackler (2002). Applied Computational Economics and Finance, The MIT Press.

Nævdal, E. (2012). "Fighting Transient Epidemics—Optimal Vaccination Schedules Before And After An Outbreak." Health Economics **21**(12): 1456-1476.

Nsoesie, E., M. Marathe and J. Brownstein (2013). "Forecasting peaks of seasonal influenza epidemics." PLoS Currents **5**.

O'Regan, S. M. and J. M. Drake (2013). "Theory of early warning signals of disease emergence and leading indicators of elimination." Theoretical Ecology **6**(3): 333-357.

Philipson, T. and R. Posner (1993). Private choices and public health: The AIDS epidemic in an economic perspective. Cambridge, MA, Harvard University Press.

Pindyck, R. S. (2007). "Uncertainty in Environmental Economics." Rev Environ Econ Policy **1**(1): 45-65.

Qi, H. and L. Liao (1999). "A Smoothing Newton Method for Extended Vertical Linear Complementarity Problems." SIAM Journal on Matrix Analysis and Applications **21**(1): 45-66.

Sethi, S. P. and P. W. Staats (1978). "Optimal control of some simple deterministic epidemic models." Journal of the Operational Research Society: 129-136.

Shaman, J. and A. Karspeck (2012). "Forecasting seasonal outbreaks of influenza." Proceedings of the National Academy of Sciences **109**(50): 20425-20430.

- Shiller, R. J. and J. Pound (1989). "Survey evidence on diffusion of interest and information among investors." Journal of Economic Behavior & Organization **12**(1): 47-66.
- Soebiyanto, R. P., F. Adimi and R. K. Kiang (2010). "Modeling and predicting seasonal influenza transmission in warm regions using climatological parameters." PloS one **5**(3): e9450.
- Stoneman, P. (1983). The economic analysis of technological change, Oxford University Press.
- Van den Driessche, P. and J. Watmough (2000). "A simple SIS epidemic model with a backward bifurcation." Journal of Mathematical Biology **40**(6): 525-540.
- Wayne, A. (2014). "Obamacare Website Costs Exceed \$2 Billion, Study Finds." Bloomberg Business, from <http://www.bloomberg.com/news/articles/2014-09-24/obamacare-website-costs-exceed-2-billion-study-finds>.
- Wickwire, K. (1977). "Mathematical models for the control of pests and infectious diseases: a survey." Theoretical Population Biology **11**(2): 182-238.
- Xu, X. (1999). "Technological improvements in vaccine efficacy and individual incentive to vaccinate." Economics Letters **65**(3): 359-364.
- Yorke, J. A., H. W. Hethcote and A. Nold (1978). "Dynamics and control of the transmission of gonorrhea." Sexually Transmitted Diseases **5**(2): 51-56.