A decision-theoretic framework for adaptive management of epidemiological intervention

Daniel Merl
Applied Mathematics & Statistics
University of California, Santa Cruz

Marc Mangel
Applied Mathematics & Statistics
University of California, Santa Cruz

Corresponding Author:

Daniel Merl

Baskin School of Engineering

University of California, Santa Cruz

1156 High Street

Santa Cruz CA 95064

email: dmerl@ams.ucsc.edu

Running head: Adaptive management of epidemiological interventions

Keywords: decision theory, SIR models, optimal vaccination, stochastic dynamic programming

Abstract

Increased national concern over the possibility of sudden and wide-spread outbreak of infectious disease motivates the modeling of pathogen transmission for various characterizations of the underlying population structure. A critical, but often overlooked, component to preparation of this sort is the development of a methodology for defining optimal treatment policies to be deployed in the event of such an outbreak. In realistic scenarios, such a policy must weigh the costs of available treatments (e.g. vaccination or quarantine) against the current state of the epidemic, in terms of numbers of susceptible and infected individuals, and current estimates of the parameters of the transmission model. We present a decision-theoretic framework for adaptive management of epidemiological interventions. We derive an expression for the optimal adaptive policy conditional upon current information about the epidemic, evaluation of which is computed using stochastic dynamic programming. We present simulation studies to demonstrate the advantages, in terms of reduced expected loss and more efficient resource management, of the optimal adaptive policy over the optimal non-adaptive policy.

1 Introduction

The use of mathematical modeling in the study of emerging infectious disease has yielded tremendous insight into the governing dynamics of disease transmission [1, 16, 22]. However, relatively little research has been devoted to incorporating a conceptual understanding of disease transmission into a pragmatic policy tool. Due to the nonlinear and stochastic nature of modern epidemiological models, optimal control theory can only be applied in relatively simple cases [11, 23]. The increased concern over widespread outbreak of diseases such as bioterrorist smallpox, pandemic influenza, and foot-and-mouth disease, has recently led to several new approaches for deducing optimal vaccination strategies, based on explicit network-type models of the underlying population structure [2, 8, 9, 21].

Methods such as these are concerned with rapid, strategically distributed vaccination. Many scenarios exist, however, where this may not be possible, such as those in which a vaccine is not immediately available, or when some aspect of the transmission model is unknown (e.g. population structure or transmission rate). In these situations, an adaptive policy that can incorporate new information as it becomes available would be desirable. Additionally, fixed-proportion strategies such as those previously described may not take into account the natural tendency of an epidemic to die out once the number of infected individuals per susceptible individual falls below a critical threshold (i.e. $R_0 < 1$) [1]. In other words, having treated an epidemic for some time, reduced intervention efforts may reduce the total cost and conserve resources, while still arriving at an identical end-state.

A general strategy for epidemiological intervention must also consider multiple possible forms of treatment, of which vaccination is but one. Meltzer et al describe a smallpox-prevention model that allows for both quarantine and vaccination, though their approach is neither adaptive nor

optimal [15]. The availability of alternative treatments necessitates a policy framework capable of weighing the tradeoffs between treatments. Bauch and Earn describe a game theoretic strategy in the context of diseases for which vaccination is optional, and has an associated mortality risk [3]. In a more general intervention strategy, with multiple treatment types, the decision to vaccinate will depend on a complex relationship between the cost of vaccination, the rate of recovery, the cost of maintaining an infected individual, the cost of quarantining an infected, and the rate of transmission, viz., all parameters characterizing the dynamics of the disease.

In this paper, we present a decision-theoretic framework for minimizing the total cost associated with an epidemiological intervention while allowing for complex tradeoffs between costs of treatment and disease dynamics. Assuming the parameters of the disease model to be known, we describe an efficient Monte Carlo algorithm for exploring the expected cost surface, minimization over which produces the optimal intervention strategy. Simulations demonstrate the advantages of an adaptive strategy over a non-adaptive (fixed) strategy, in terms of reduced expected cost, and more efficient use of vaccine resources. The latter is noteworthy, as it has been suggested that mismanagement of vaccines (e.g. over-vaccination) can lead to increased virulence [5, 6, 7, 19].

2 Methods

We begin by describing the underlying disease dynamics using a standard S–I–R transmission model for a closed population of initial size N, where $S(t)$, $I(t)$, $R(t)$ represent the numbers of susceptible, infected, and recovered individuals at time t respectively (with $S(t) + I(t) + R(t) = N$ for all t). We assume a negative binomial transmission function with rate of transmission b, overdispersion parameter k, and rate of recovery ν, where ν accounts for both recovery from the infected state and disease-induced mortality [14]. The continuous time formulation of this model

is described by the set of differential equations

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -kS\log\left(1+\frac{bI}{k}\right) \quad (1)$$

$$\frac{\mathrm{d}I}{\mathrm{d}t} = kS\log\left(1+\frac{bI}{k}\right) - \nu I \quad (2)$$

$$\frac{\mathrm{d}R}{\mathrm{d}t} = \nu I. \quad (3)$$

We assume recovered individuals attain permanent immunity against subsequent infection, though this assumption may be relaxed within the framework we are describing The negative binomial transmission function accounts for social/network effects in the population structure via the overdispersion parameter k. Small values of k correspond to more localized interactions, whereas for large k the nature of the interactions approaches a mass action model. Keeling has suggested that generalized mass action models such as this achieve a good trade off between model complexity and realistic modeling of social interactions [10].

By holding the number of infected individuals ($I(t)$) constant, we can integrate (1) over a unit time interval to obtain a convenient discretization of the disease transmission process,

$$S(t+1) = \left(\frac{k}{k+bI(t)}\right)^k S(t). \quad (4)$$

Thus, the complete discrete time stochastic disease model is defined by two binomial random variables: $\tilde{i}(t)$ and $\tilde{r}(t)$, representing the numbers of newly infected and newly recovered individuals at time t, where

$$\tilde{i}(t) \sim \mathcal{B}\left(S(t), 1-\left(\frac{k}{k+bI(t)}\right)^k\right) \quad (5)$$

$$\tilde{r}(t) \sim \mathcal{B}\left(I(t), 1-e^{-\nu}\right). \quad (6)$$

The values $\tilde{i}(t)$ and $\tilde{r}(t)$ will correspond to the numbers of newly infected and newly recovered individuals arising between time steps t and $t+1$. Our simulation studies indicate that discrete

time simulations of this form provide reasonable approximations to the exact numerical solution to the system of differential equations. The reason underlying our preference for the stochastic formulation will be made evident below.

We define $\boldsymbol{\alpha}_t = (\alpha_{v,t}, \alpha_{r,t})$ to be the epidemiological intervention at time t consisting of two decisions: $\alpha_{v,t}$, the fraction of susceptible individuals to be vaccinated; and $\alpha_{r,t}$, the fraction infected individuals to be removed from the population of infected individuals (as by curing, culling or quarantine). The sequence of decisions $\{\boldsymbol{\alpha}_t\}_{t=0}^{\infty}$ collectively defines the intervention strategy. If $\boldsymbol{\alpha}_t$ is constrained to a single value $\boldsymbol{\alpha}_f$ for all t, we will refer to $\boldsymbol{\alpha}_f$ as a fixed intervention strategy. Alternatively, if $\boldsymbol{\alpha}_t$ is allowed to vary with the current state of the epidemic, we refer to $\{\boldsymbol{\alpha}_t\}$ as an adaptive intervention strategy (implicitly depending upon $S(t) = s$ and $I(t) = i$). We assume that vaccination guarantees permanent immunity to infection, and that removed individuals are no longer infectious. We define the cost per vaccination to be c_v, the cost per removal to be c_r, and the cost per remaining infected for one time interval to be c_i.

For any chosen valuation of these costs, we can recursively express the total cost of an intervention strategy, conditional on the current state of the epidemic $S(t) = s$ and $I(t) = i$ (and, implicitly, the parameters of the disease model), as

$$\begin{aligned} C(\boldsymbol{\alpha}_t, s, i) &= c_v \alpha_{v,t} s + c_r \alpha_{r,t} i + c_i (1 - \alpha_{r,t}) i \\ &\quad + E\{C(\boldsymbol{\alpha}_{t+1}, s', i')\}, \end{aligned} \tag{7}$$

where the state of the disease at time $t + 1$ is

$$s' = (1 - \alpha_{v,t}) s - \tilde{i}, \text{ and} \tag{8}$$

$$i' = (1 - \alpha_{r,t}) i + \tilde{i} - \tilde{r}. \tag{9}$$

The expectation term in (7) accounts for the stochastic nature of $\tilde{i}(t)$ and $\tilde{r}(t)$. The combination of

this random component and the nonlinear dynamics of the epidemic prevent a closed form solution for the total cost. Note that $\boldsymbol{\alpha}_t$ is actually a function of the current state of the epidemic at time t, but for notational convenience we index the decision by t. Thus (7) should be regarded as time-independent when conditioned on the current values of $S(t)$ and $I(t)$.

The optimal intervention (adaptive or fixed) is therefore the sequence $\{\boldsymbol{\alpha}_t\}$ that minimizes (7). Along the boundary case of $I(t) = 0$, the minimum cost of intervention is 0 and the solution is trivially $\boldsymbol{\alpha}_t = (0, 0)$ for any such t. For all other values, we employ a stochastic dynamic programming algorithm, augmented by a policy iteration step [4, 12]. Policy iteration is a method for converting a time-independent consistency equation such as (7) into a convergent sequence of consistency equations whose limit is the solution to the original equation. The theory states that by defining

$$\begin{aligned} C_0(s,i) &= c_v\alpha_v s + c_r\alpha_r i + c_i(1-\alpha_r)i \\ &\quad + E\{C_0((1-\alpha_v)s - \tilde{i}, (1-\alpha_r)i + \tilde{i} - \tilde{r})\}, \end{aligned} \tag{10}$$

for a fixed initial value of $\boldsymbol{\alpha}$ (e.g. $\boldsymbol{\alpha} = (0, 0)$), then by successively minimizing the sequence of equations given by

$$\begin{aligned} C_n(s,i) &= \min_{\{\boldsymbol{\alpha}_t\}} c_v\alpha_v s + c_r\alpha_r i + c_i(1-\alpha_r)i \\ &\quad + E\{C_{n-1}((1-\alpha_v)s - \tilde{i}, (1-\alpha_r)i + \tilde{i} - \tilde{r})\}, \end{aligned} \tag{11}$$

the solutions will converge to the desired solution under very general conditions.

Each minimization step is accomplished via stochastic dynamic programming. Since the expectation term is evaluated through Monte Carlo simulation from the appropriate distributions on $\tilde{i}$ and $\tilde{r}$, the solutions to the sequence of equations do not converge to a single minimum cost matrix $C_{\min}(s,i)$, but rather to a stationary distribution on minimum cost matrices. Note that

for the case of a fixed policy, this procedure amounts repeated forward simulation of the epidemic given initial conditions $(S(0) = s, I(0) = i)$ under each intervention policy $\boldsymbol{\alpha}_f$. In this case, we can increase efficiency by evaluating the expected cost once for each point on a fixed grid of $\boldsymbol{\alpha}_f$ values, and employing a loess smoother to fit a smooth surface through these points. The optimal fixed policy can be found through minimization of this surface. This procedure, known as curve-fitting Monte Carlo (CFMC), has been shown to improve the stability of optimizations while reducing the computational burden [17, 18, 20]. For the case of adaptive interventions, the sequential nature of the algorithm bears similarities to a Gibbs sampler for generating draws from a stationary distribution of optimal policies. After discarding a suitable number of pre-convergence iterations, we compute the minimum expected costs and associated policy decisions for each initial condition as the average over post-convergence iterations. We measured convergence using mean squared error between successive solution matrices. A variety of simulation studies (not shown) indicated that convergence was achieved rapidly, often after fewer than 10 iterations.

3 Results

We first simulate an emerging infectious disease for a closed population of $N = 500$ individuals. Disease model parameters were defined to be $(b, k, \nu) = (0.01, 0.01, 0.05)$. This choice of parameter values corresponds to a population in which encounters between individuals are highly aggregated. The pathogen in this case is such that the rate of transmission and rate of recovery are fairly low. Assuming the initial number of infected individuals to be 20, the probability of infection at $t = 0$ for the remaining 480 susceptible individuals is only 0.03. The probability of an infected individual recovering is 0.05. Thus, at the onset of the epidemic, the initial risk of infection is still low, but climbs rapidly with the number of infected individuals, each of which remains infectious

for an expected 20 days. By assuming the costs per untreated infected individual, vaccination, and removal to be $(c_i, c_v, c_r) = (1, 16, 20)$, we infer the optimal fixed and adaptive intervention strategies using the Monte Carlo algorithm described above. For the fixed strategy, the expected cost surface is minimized at $\boldsymbol{\alpha}_f = (0, 1)$ (figure 1), indicating a policy of immediate removal of infected individuals. For the adaptive case, the strategy is mixed, involving a period of vaccination, followed by a period of removals (figure 2). For this simulation, the separation between vaccination and removal, occurs almost exactly along the R_0 boundary, beneath which the size of the epidemic will continue to decrease naturally. The simulated disease dynamics under nonintervention, optimal fixed, and optimal adaptive policies, and comparison of total accrued costs appear in figure 3 and figure 4. Here the adaptive policy leads to a substantial decrease in total expected cost.

Next, we consider another epidemic in the same a closed population of $N = 500$ with disease model parameters $(b, k, \nu) = (0.5, 0.05, 0.1)$. The natural dynamics of this epidemic are such that contact between individuals are less aggregated than for the first simulation, and the emerging disease more highly infectious and with a faster rate of recovery. We now assume the costs per untreated infected individual, vaccination, and removal to be $(c_i, c_v, c_r) = (1, 11, 15)$. Assuming 20 initial infected individuals, we proceed as above. Minimization over the expected cost surface indicated the optimal fixed strategy to be $\boldsymbol{\alpha}_f = (0.53, 0)$ (figure 5). The adaptive strategy for the same epidemiological scenario appears in figure figure 6. In this case, the adaptive strategy involves little or no removal of infected individuals, but a decreasing schedule of vaccination as the risk of infection decreases. Simulated disease dynamics under nonintervention, optimal fixed, and optimal adaptive strategies, and comparison of total accrued costs appear in figure 7 and figure 8. In this scenario, the decrease in total cost associated with the adaptive strategy is not nearly as substantial as for the first simulation, however, we note the substantial decrease in the total

number of vaccinations required to halt the spread of the epidemic.

4 Discussion

Our simulation studies demonstrate quantitative differences between fixed and adaptive intervention strategies. For cases such as the first simulation, in which the adaptive policy results in a strict reduction in accrued costs compared to the fixed policy, the advantages of the adaptive policy are obvious. The savings incurred under the adaptive policy are due the ability of the adaptive strategy to "learn" the point at which the epidemic has crossed a critical threshold, beyond which additional interventive effort does not contribute significantly to the diminishing size of the epidemic. This threshold, referred to as R_0 in the disease literature, is defined as the expected number of new infections per infected individual. For $R_0 < 1$, $\frac{\mathrm{d}I}{\mathrm{d}t} < 0$, and the epidemic will have a natural tendency to die out, though the rate at which this occurs will still depend on the model parameters. An adaptive policy can account for this through reduced vaccinations in favor of removals in the region $R_0 < 1$ (figure 2). A pragmatic alternative to the fixed strategies we have described here would be a semi-adaptive fixed strategy, in which the fixed policy $\boldsymbol{\alpha}_f$ is discontinued at some non-terminal (i.e. $I(t) \neq 0$) state. This would require calculation of the optimal stopping point for each $\boldsymbol{\alpha}_f$, and could be incorporated into the framework we have described.

In cases such as the second simulation, in which the expected costs under the fixed and adaptive policy are approximately equal, there remain qualitative differences between results, which are best interpreted through contextualization of the epidemic. There are many situations in which it may be desirable to conserve the amount of vaccine available. There are other situations in which unnecessary removal of infected individuals should be avoided (e.g. if removal corresponds to culling as opposed to quarantine). In the second simulation, the fixed policy continues to

vaccinate remaining susceptible individuals, even after risk of infection is very low, unnecessarily depleting resources. If we liken this simulation to an emerging flu epidemic, for which natural recovery corresponds to an infected individual's return to a healthy state, the adaptive strategy may be preferable to the fixed. Adaptive intervention policies, through a strategic combination of vaccination and removal, tend to avoid the scenarios of over-vaccination or over-removal.

There is an important choice to be made in assigning values to c_v, c_r, and c_i. A monetary valuation scheme is the most straightforward, but it may be difficult to construct such a scheme that represents all aspects of the decision. One alternative would be a valuation in which each cost is chosen to represent a probability of mortality. In this way, the cost to be minimized is the total loss of life for the epidemic under a given intervention strategy. By assuming $\nu = \mu + \rho$ where μ is the rate of disease-induced mortality and ρ is the rate of natural recovery from the infected state, we can set $c_i = (1 - e^{-\nu})\frac{\mu}{\mu+\rho}$, so that $c_i(1 - \alpha_r)i$ from (7) is the number of infected individuals that die in a unit of time. Similarly, situations exist where it is reasonable to assign a probability of mortality to vaccination, as in the case of smallpox, and to removal, as in the case of culling livestock for foot-and-mouth disease.

An interesting extension to our adaptive policy model would involve applying a monetary constraint to a loss-of-life cost function. Assume p_i, p_v, and p_r are the probabilities of mortality associated with untreated infected individuals, vaccination, and removal as just described, and that d is the monetary resources available for the intervention. Now the optimal intervention will be the one that minimizes the total loss-of-life subject to the total spending constraint d:

$$C(s, i, d) = \min_{\{\boldsymbol{\alpha}_t\}} p_i(1 - \alpha_{r,t})i + p_v\alpha_{v,t}s + p_r\alpha_{r,t}i \tag{12}$$

$$+E\{C((s', i', d - c_i(1 - \alpha_{r,t})i - c_v\alpha_{v,t}s - c_r\alpha_{r,t}i)\}. \tag{13}$$

Modifying the existing stochastic dynamic programming algorithm to allow for this type of constraint is straightforward, however, the resulting increase in dimensionality of the solution space reduces the speed of the algorithm.

The simulations presented in this paper were designed to represent interesting epidemics, because our goal was to explore the implications of different types of epidemiological interventions for different epidemiological scenarios. Indeed, as May put it, "[t]he virtue of mathematics in such a context is that it forces clarity and precision upon the conjecture, thus enabling meaningful comparison between the consequences of basic assumptions and the empirical facts" [13]. In operational situations, a great deal of uncertainty may surround estimates of disease model parameters, and that uncertainty should be reflected in the policy decision. A simulation-based framework for deriving optimal intervention strategies can be easily integrated into a Bayesian MCMC estimation procedure, for which the results are draws from posterior distributions on the relevant parameters. The unification of policy optimization and inference procedures for emerging infectious disease models remains an important goal of mathematical biology.

Acknowledgments

This research was partially supported by DARPA Contract DAAD19-03-0162 and NSF Grant DMS 03-10542. We also thank Ottar Bjørnstad and Leah Johnson for many helpful conversations.

References

[1] R. Anderson and R. May. *Infectious diseases of humans: dynamics and control*. Oxford University Press, 1991.

[2] F. Ball and O. Lyne. Optimal vaccination policies for stochastic epidemics among a population of households. *Mathematical Biosciences*, 177&178:333–354, 2002.

[3] C. Bauch and D. Earn. Vaccination and the theory of games. *Proceedings of the National Academy of Sciences*, 101(36):13391–13394, 2004.

[4] D. Bertsekas. *Dynamic programming and optimal control*. Athena Scientific, 2001.

[5] D. Ebert and J. Bull. Challenging the trade-off model for the evolution of virulence: is virulence management feasible? *Trends in Microbiology*, 11(1):15–20, 2003.

[6] S. Gandon, M. Mackinnon, S. Nee, and A.Read. Imperfect vaccines and the evolution of pathogen virulence. *Nature*, 414:751–756, 2001.

[7] S. Gandon, M. Mackinnon, S. Nee, and A. Read. Imperfect vaccination: some epidemiological and evolutionary consequences. *Proceedings of the Royal Society of London B*, 270:1129–1136, 2003.

[8] M. E. Halloran, I. M. Longini Jr., A. Nizam, and Y. Yang. Containing bioterrorist smallpox. *Science*, 298:1428–1432, 2002.

[9] L. Hufnagel, D. Brockmann, and T. Geisel. Forecast and control of epidemics in a globalized world. *Proceedings of the National Academy of Sciences*, 101(42):15124–15129, 2004.

[10] M. Keeling. The implications of network structure for epidemic dynamics. *Theoretical Population Biology*, 67:1–8, 2005.

[11] A. Korobeinikov and P. K. Maini. Non-linear incidence and stability of infectious disease models. *Mathematical Medicine and Biology*, 2005. Advance Access published March 18 2005.

[12] M. Mangel and C. Clark. *Dynamic modeling in behavioral ecology*. Princeton University Press, 1988.

[13] R. May. Uses and abuses of mathematics in biology. *Science*, 303:790–793, 2004.

[14] H. McCallum, N. Barlow, and J. Hone. How should pathogen transmission be modelled? *Trends in Ecology and Evolution*, 16(6):295–300, 2001.

[15] M. Meltzer, I. Damon, J. LeDuc, and J. D. Millar. Modeling potential responses to smallpox as a bioterrorist weapon. *Emerging infectious diseases*, 7(6):959–969, 2001.

[16] D. Morens, G. Folkers, and A. Fauci. The challenge of emerging and re-emerging infectious diseases. *Nature*, 430:242–249, 2004.

[17] P. Müller. Simulation-based optimal design. In J. Bernardo, J. Berger, A. Dawid, and A. Smith, editors, *Bayesian Statistics 6*, pages 459–474. Oxford University Press, 1999.

[18] P. Müller and G. Parmigiani. Optimal design via curve fitting of monte carlo experiments. *Journal of the American Statistical Association*, 90(432):1322–1330, 1995.

[19] S. R. Palumbi. Humans as the world's greatest evolutionary force. *Science*, 293:1786–1790, 2001.

[20] G. Parmigiani. *Modeling in Medical Decision Making: A Bayesian Approach*. John Wiley & Sons, Ltd, 2002.

[21] R. Patel, I. M. Longini Jr., and M. E. Halloran. Finding optimal vaccination strategies for pandemic influenza using genetic algorithms. *Journal of Theoretical Biology*, 234:201–212, 2005.

[22] S. Schrag and P. Wiener. Emerging infectious disease: what are the relative roles of ecology and evolution. *Trends in Ecology and Evolution*, 10(8):319–324, 1995.

[23] S. M. White and K. A. J. White. Applications of biological control in resistant host-pathogen systems. *Mathematical Medicine and Biology*, 2005. Advance Access published March 22, 2005.

A Figures

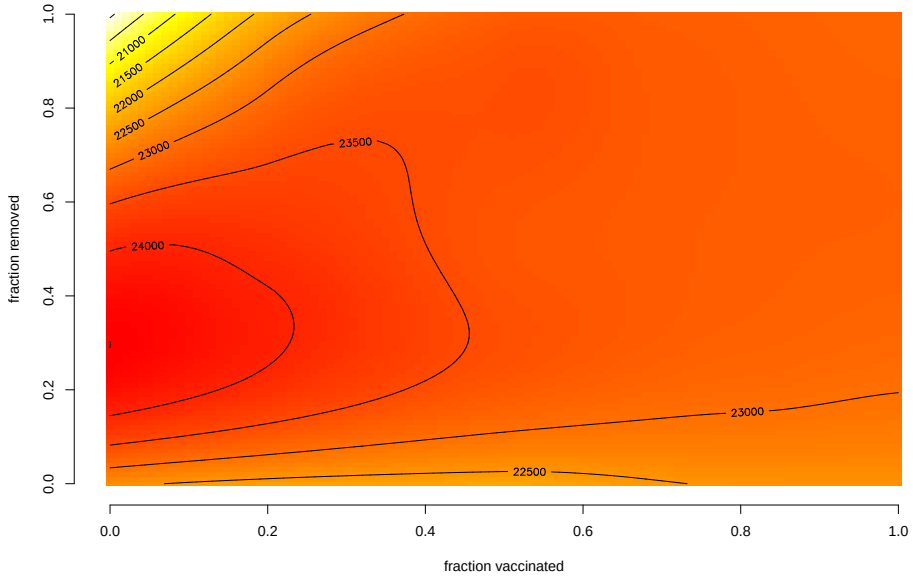

Figure 1: Expected cost surface for fixed intervention policies for simulation 1. Minimum cost is achieved at $\boldsymbol{\alpha}_f = (0, 1)$.

(a)

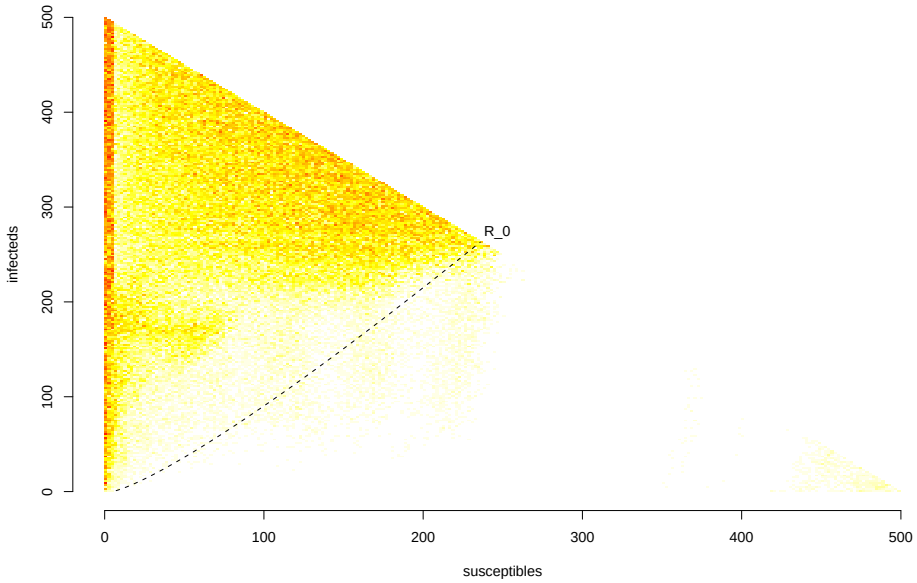

(b)

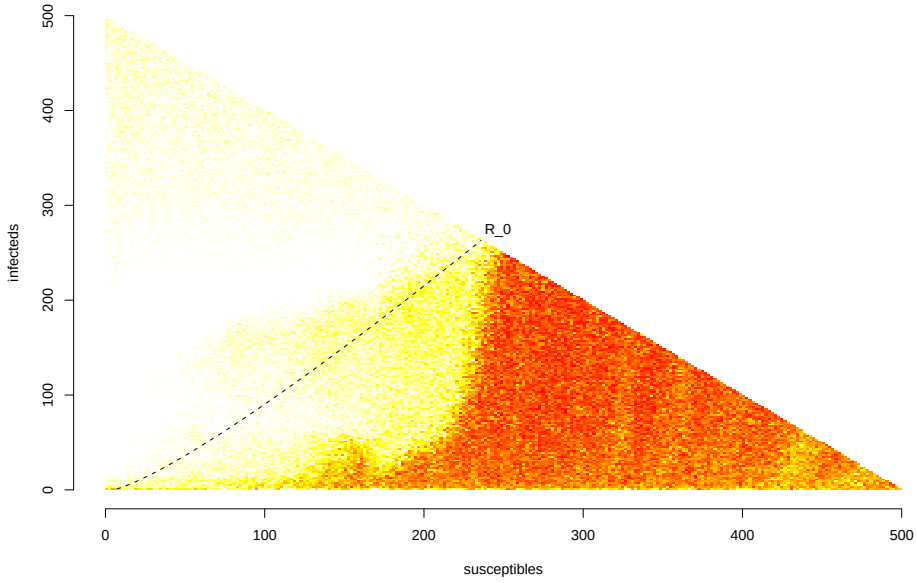

Figure 2: (a) Fraction of susceptible individuals to vaccinate and (b) fraction of infected individuals to remove under the adaptive intervention policy for simulation 1. The color spectrum ranges from white (0%) to red (100%). The region beneath the curve ($s < \frac{\nu i}{k \log(\frac{k+bi}{k})}$) corresponds to $R_0 < 1$.

(a)

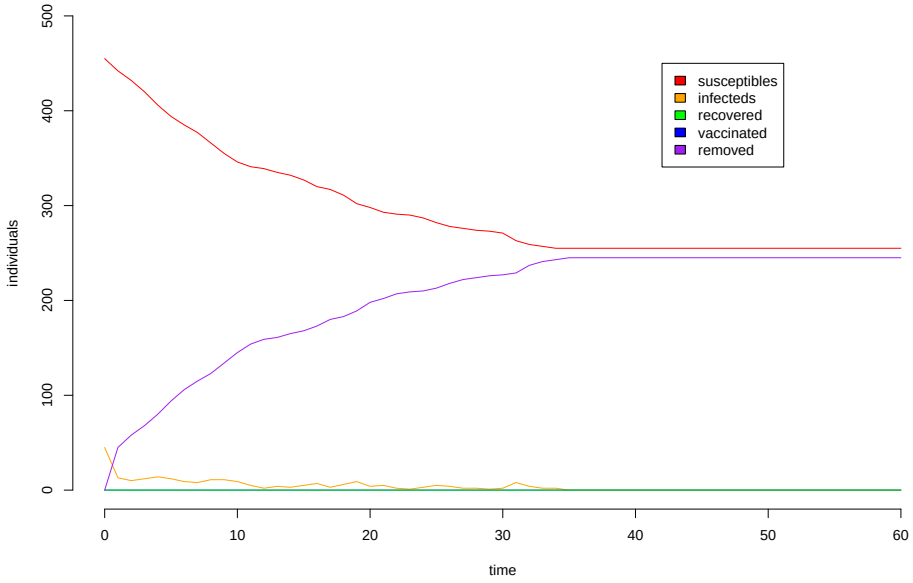

(b)

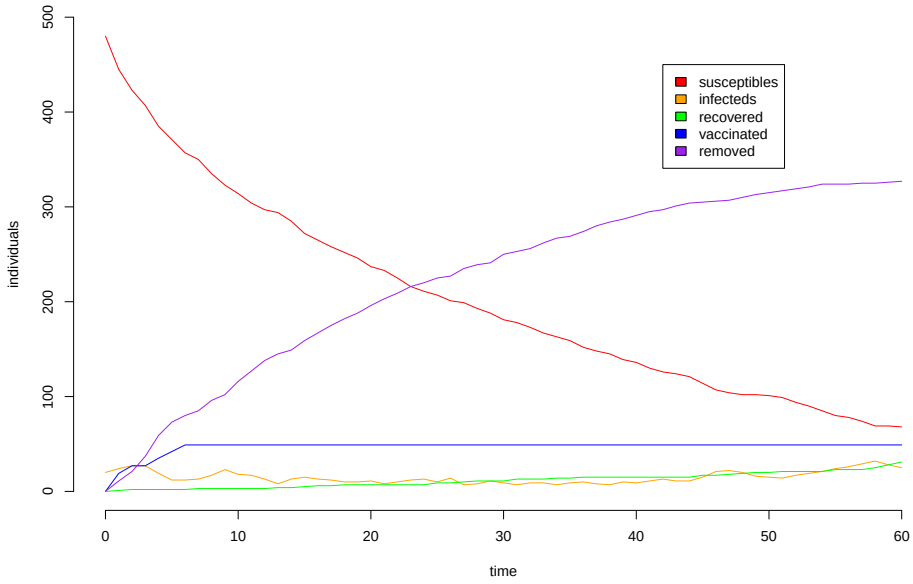

Figure 3: Comparison of disease dynamics under (a) fixed and (b) adaptive intervention policies for simulation 1.

(a)

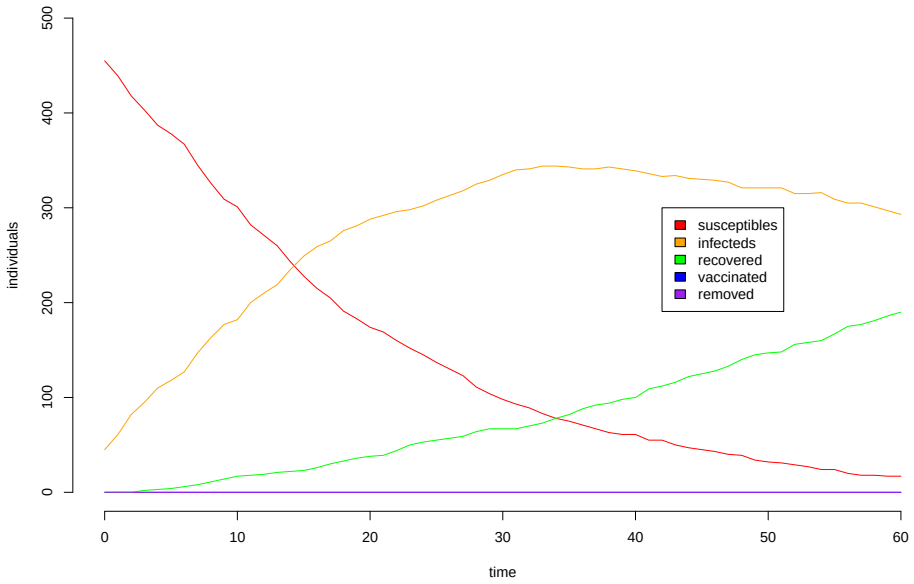

(b)

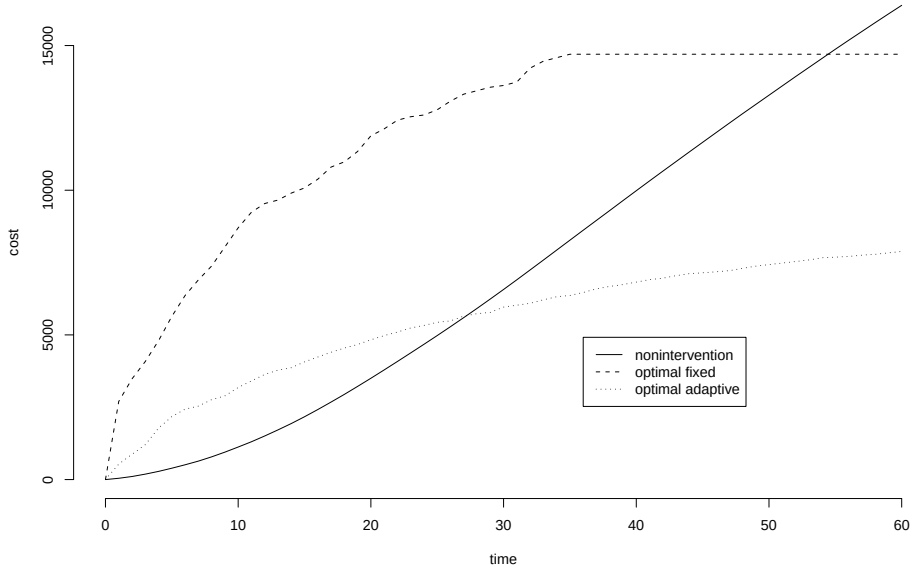

Figure 4: (a) Natural disease dynamics under a nonintervention policy, and (b) comparison of total accrued costs under nonintervention, fixed, and adaptive policies for simulation 1.

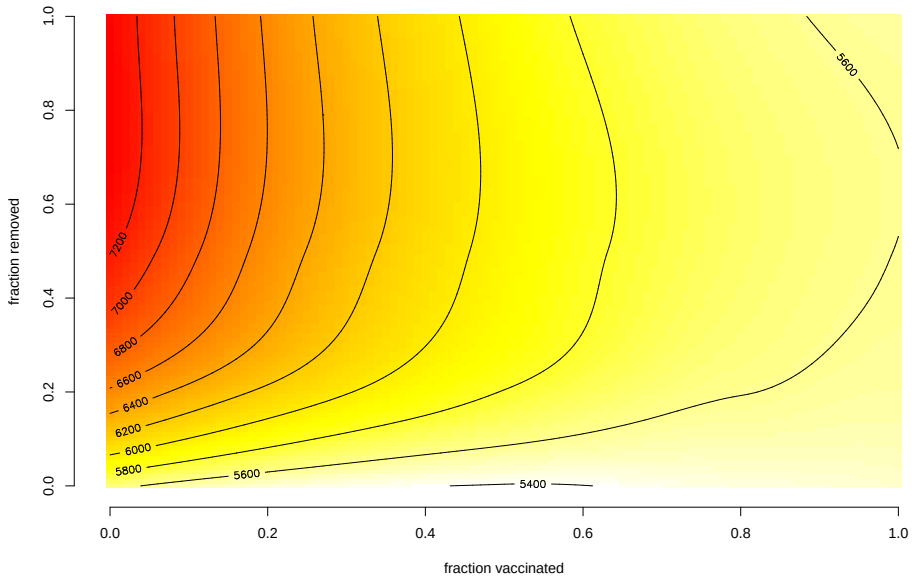

Figure 5: Expected cost surface for fixed intervention policies for a simulation 2. Minimum cost is achieved at $\boldsymbol{\alpha}_f = (0.53, 0)$.

(a)

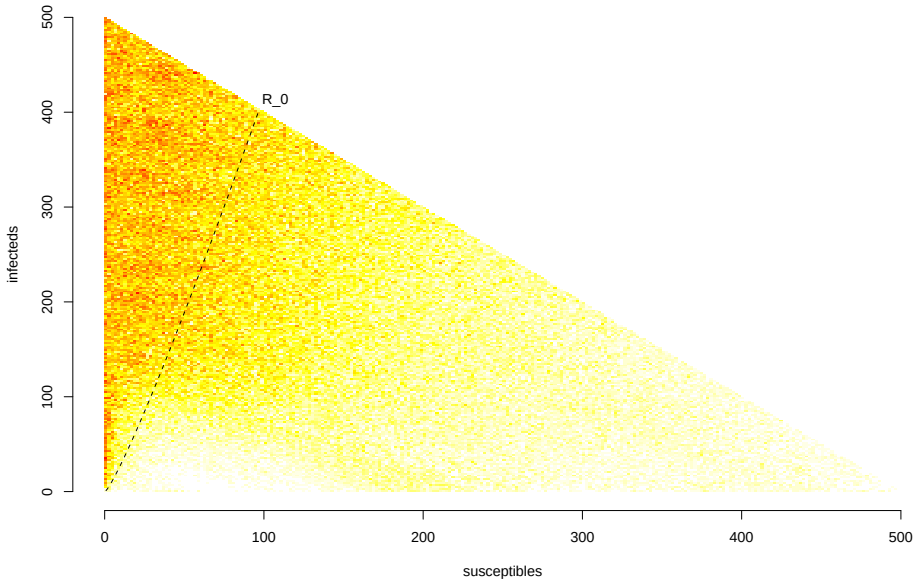

(b)

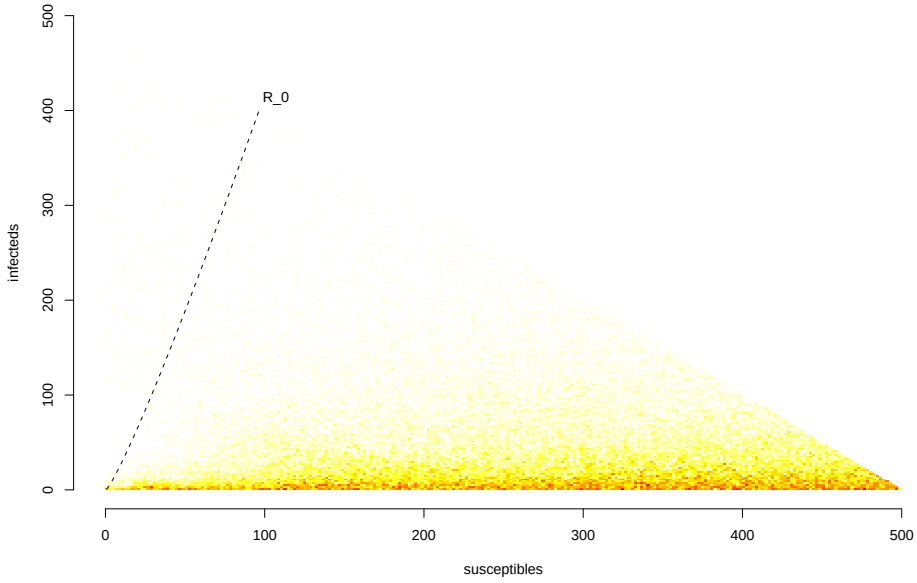

Figure 6: (a) Fraction of susceptible individuals to vaccinate and (b) fraction of infected individuals to remove under the adaptive intervention policy for simulation 2. The color spectrum ranges from white (0%) to red (100%). The region beneath the curve ($s < \frac{\nu i}{k \log(\frac{k+bi}{k})}$) corresponds to $R_0 < 1$.

(a)

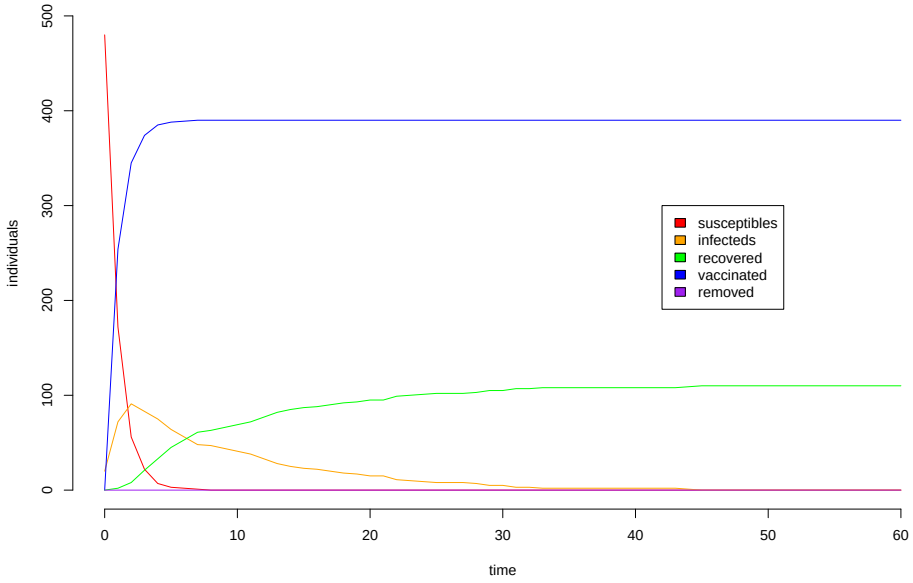

(b)

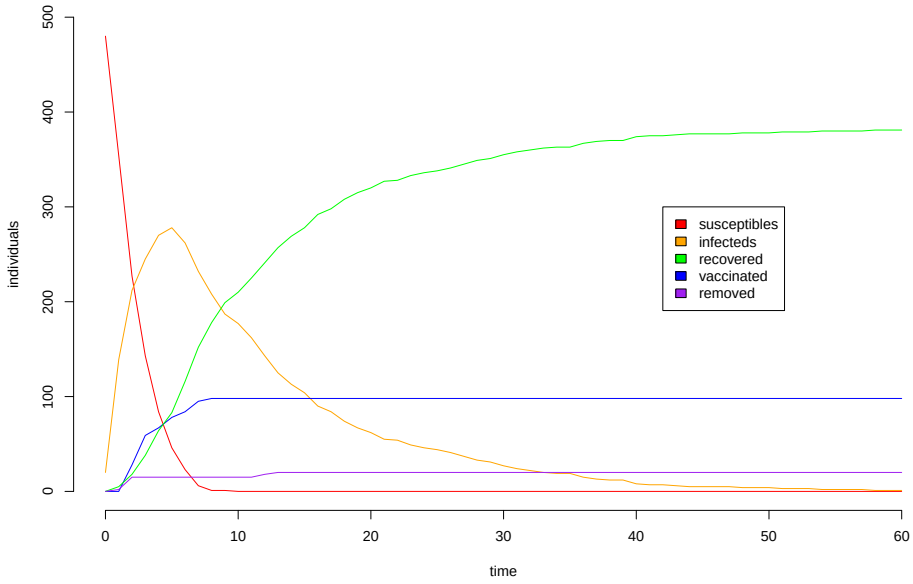

Figure 7: Comparison of disease dynamics under (a) fixed and (b) adaptive intervention policies for simulation 2.

(a)

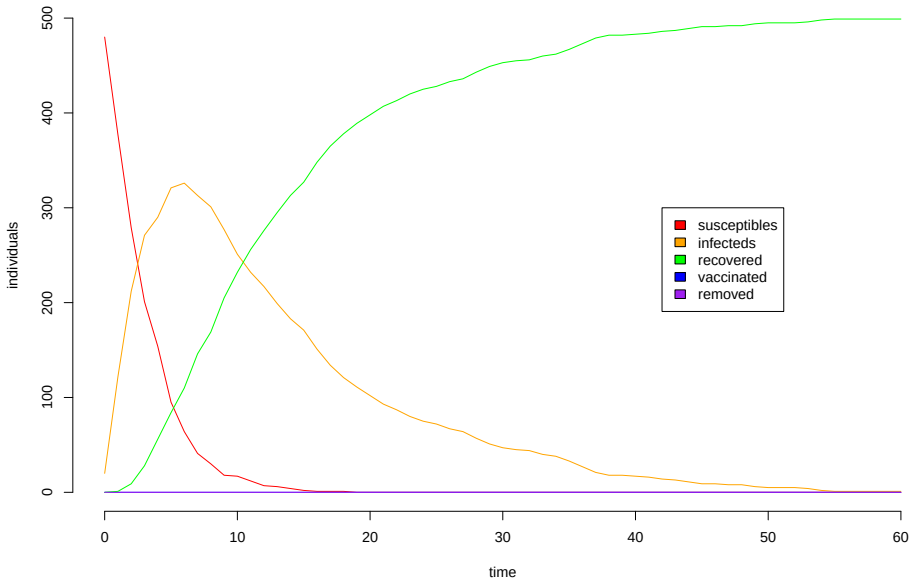

(b)

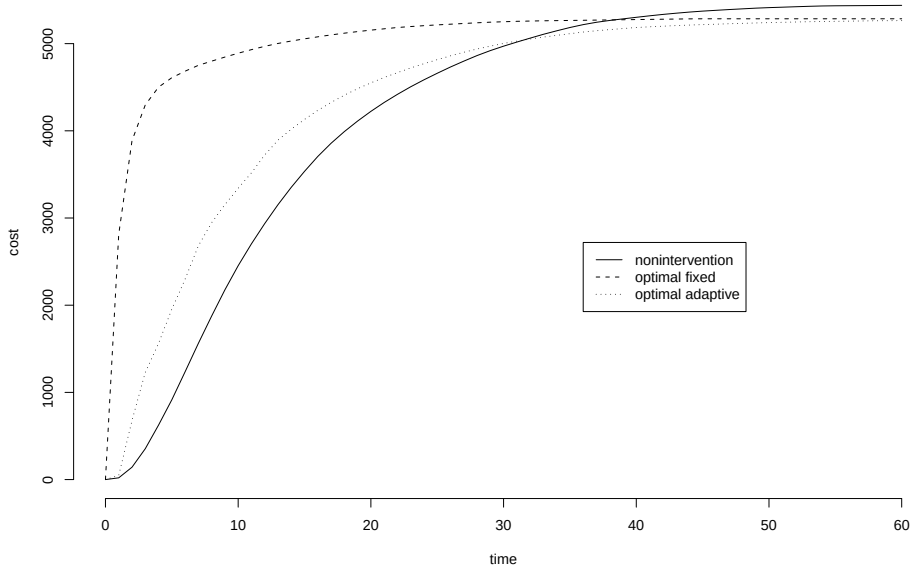

Figure 8: (a) Natural disease dynamics under a nonintervention policy, and (b) comparison of total accrued costs under nonintervention, fixed, and adaptive policies for simulation 2.