P.G. Ballard N.G. Bean, J.V. Ross
Intervention to maximise the probability of epidemic fade-out
Mathematical Biosciences, 2017; 293:1-10

Final publication at http://dx.doi.org/10.1016/j.mbs.2017.08.003

Intervention to maximise the probability of epidemic fade-out

P. G. Ballard[1,*], N. G. Bean[1], J. V. Ross[1]

Abstract

The emergence of a new strain of a disease, or the introduction of an existing strain to a naive population, can give rise to an epidemic. We consider how to maximise the probability of epidemic fade-out – that is, disease elimination in the trough between the first and second waves of infection – in the Markovian SIR-with-demography epidemic model. We assume we have an intervention at our disposal that results in a lowering of the transmission rate parameter, β, and that an epidemic has commenced. We determine the optimal stage during the epidemic in which to implement this intervention. This may be determined using Markov decision theory, but this is not always practical, in particular if the population size is large. Hence, we also derive a formula that gives an almost optimal solution, based upon the approximate deterministic behaviour of the model. This formula is explicit, simple, and, perhaps surprisingly, independent of β and the effectiveness of the intervention. We demonstrate that this policy can give a substantial increase in the probability of epidemic fade-out, and we also show that it is relatively robust to a less than ideal implementation.

Keywords: SIR infection model, stochastic model, Markov decision theory, epidemic control

1. Introduction

One of the key goals of epidemiology is to take action to minimise the impact of epidemic outbreaks. With this in mind, many studies have investigated ways to optimise the control of an outbreak. Good overviews can be found in Kar and Batabyal [12] and Yaesoubi and Cohen [28].

Studies tend to take one of two approaches: investigating either the use of vaccination [25] (reducing the susceptible population), or prophylactic measures to reduce the spread of the infection [22]. Prophylactic measures include antivirals [4, 9, 17, 18], or non-pharmaceutical interventions such as school closures

*Corresponding author

Email addresses: peter.ballard@adelaide.edu.au (P. G. Ballard), nigel.bean@adelaide.edu.au (N. G. Bean), joshua.ross@adelaide.edu.au (J. V. Ross)

[1]The University of Adelaide, School of Mathematical Sciences and ARC Centre of Excellence for Mathematical and Statistical Frontiers, Adelaide SA 5005, AUSTRALIA.

[10] or the teaching of basic personal health habits [21]. Vaccination is usually the ideal, but it is often not available in the early stages of a novel strain/disease. Antivirals, or other measures to reduce the spread of an infection, are therefore an important tool in attempting to control an outbreak.

Previous studies have concentrated on either the initial stages of an infection – and measures to prevent the infection becoming an outbreak – or an established infection, endemic to a population. In this paper, we instead examine *epidemic fade-out*, which has been nominated as an area requiring more research [5, 6], and has not previously been investigated in terms of control.

Epidemic fade-out refers to the case in which an infection has a large initial outbreak, and it is eliminated from the population in the first trough after that initial outbreak [19]. Therefore, techniques to maximise the probability of epidemic fade-out offer the opportunity to prevent an infection from becoming established - that is, endemic - in a population.

We use the Markovian SIR-with-demography infection model [20]. Important previous work was by van Herwaarden [26] and Meerson and Sasorov [19], who both provided methods for approximating the probability of epidemic fade-out for this model. van Herwaarden used the Fokker-Plank approximation, while Meerson and Sasorov used the WKB appoximation. Both of these papers gave explicit formulae for the probability of epidemic fade-out, to a good degree of accuracy. In a previous paper [3] we outlined a more accurate numerical approximation method, and also presented a range of results from our calculations. These results showed that the probability of epidemic fade-out is non-monotonic in the transmission rate parameter β. Typically, a lower value of β increases the probability of epidemic fade-out, which is the intuitive result (less transmission $\rightarrow$ higher probability of fade-out). But in some situations, perhaps counter-intuitively, a reduction in the value of β causes the probability of epidemic fade-out to decrease.

Similar examples of non-monotonicity in epidemics – of a reduction in transmission or an increase in treatment actually increasing the total epidemic size, or making the epidemic more likely to persist – have been reported by others, but in different contexts. Feng et al. [7], Rozhnova et al. [24], and Lee and Chowell [15] all reported non-monotonicity in the context of seasonal forcing. Xiao et al. [27] saw it in the case of multiple strains of an infection. Grigorieva and Khailov [8] is perhaps the closest analogue to this paper. In a deterministic SIR model, they showed that not reducing β early in the infection cycle can minimise the total epidemic size.

For epidemic fade-out, the non-monotonicity in the transmission rate parameter β suggests that there are two or more competing effects, and that in some states a higher β will maximise the probability of epidemic fade-out, and in other states a lower β will maximise it. So it should be possible to find the optimal policy for choosing higher or lower β. Finding this optimal policy is the topic of this paper.

We will show that this optimal policy entails delaying the implementation of measures to reduce β, resulting in more individuals being infected in the short term. In lethal epidemics, even if the long term result is a more likely fade-out

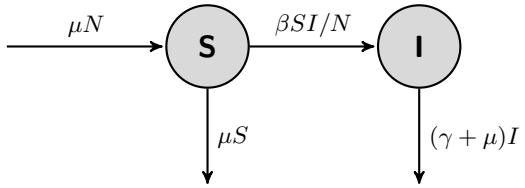

Figure 1: The SIR-with-demography epidemic model. S is the number of susceptibles and I is the number of infectious individuals. N is a fixed parameter, but the population size is not fixed. β is the transmission rate parameter, μ is the per-capita birth/death rate, and γ is the recovery rate.

and hence the elimination of the disease from the population, this is likely to be impossible to do ethically. Therefore the applicability of this method will probably be limited to situations of non-lethal infections, or diseases among animals.

We examine two different control scenarios: an idealised scenario in Section 3 and a more realistic scenario in Section 4. Section 3.3 is the most significant contribution of the paper, where we derive a simple control policy that is a close approximation to the optimal control policy in the idealised scenario. Section 4.3 supplements Section 3.3, by showing that the same simple control policy is also a close approximation to the optimal control policy in the realistic scenario. Effectively, we provide an explicit, simple rule for when to implement an intervention. Perhaps surprisingly, this rule is independent of the transmission rate parameter β and the effectiveness of the intervention. The results, which show a significant increase in the probability of epidemic fade-out when using any of these methods, are presented in Section 5.

2. Model and definitions

2.1. The SIR-with-demography model

We use the Markovian SIR-with-demography model, as described in Fig. 1 and Table 1. S and I represent the number of "susceptible" and "infectious" individuals respectively. The parameters β, γ and μ are all strictly positive. The number of "recovered" individuals (R) is usually included in the model, but is redundant and can be removed from the analysis by considering "death of infectious" (at rate μI) and "recovery of infectious" (at rate γI) as equivalent [11]. We use a common death rate μ, corresponding to a non-lethal infection, as this is the original and most common model [1]. If a different death rate μ_I is used for infectious individuals, then the analysis in the rest of this paper follows similarly, if one replaces all references to $(\gamma + \mu)$ with $(\gamma + \mu_I)$.

Description	Transition	Rate
Infection	$(S, I) \to (S-1, I+1)$	$\beta SI/N$
Birth of susceptible	$(S, I) \to (S+1, I)$	μN
Death of susceptible	$(S, I) \to (S-1, I)$	μS
Removal of infectious	$(S, I) \to (S, I-1)$	$(\gamma+\mu)I$

Table 1: Transition rates for the Markovian SIR-with-demography epidemic model displayed in Fig. 1.

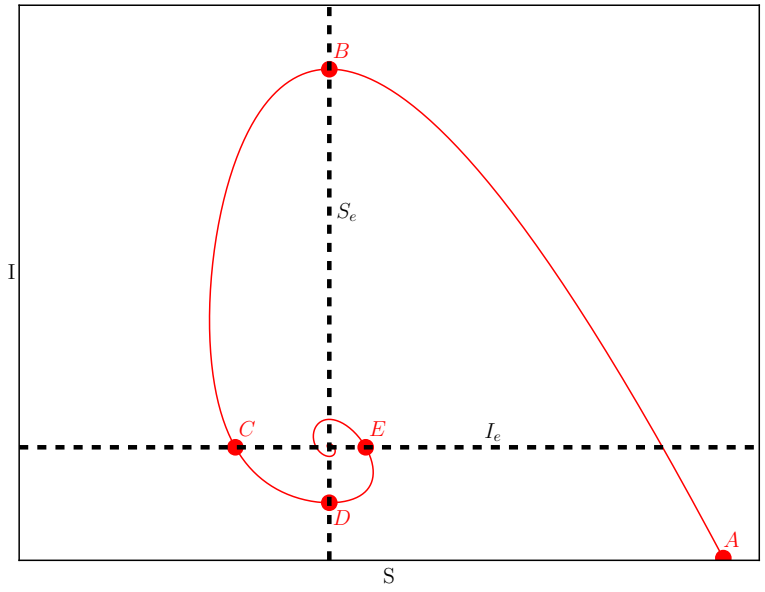

Figure 2: I versus S plot of the deterministic approximation of a typical SIR-with-demography model. It starts at point A, rises to B, then falls through C to D before rising to E and eventually converging on the endemic point (S_e, I_e). An actual stochastic realisation may fade out to $I = 0$ near point D, an effect known as epidemic fade-out. It follows from (7) that points B and D are both at $S = S_e$.

In the limit as N becomes large, a suitably scaled version of the stochastic process converges (uniformly in probability over finite time intervals) to a deterministic process [14]; this provides an approximation to the expected dynamics, for finite N, governed by the differential equations:

$$\begin{aligned} \frac{dS}{dt} &= \mu(N-S) - \beta SI/N, \\ \frac{dI}{dt} &= \beta SI/N - (\gamma+\mu)I. \end{aligned} \tag{1}$$

We refer to this as the *deterministic approximation*.

In a naive population, $S \approx N$. So R_0, the basic reproduction number, is given by:

$$R_0 = \frac{\beta}{\gamma + \mu}. \tag{2}$$

We are only concerned with cases in which $R_0 > 1$, when a major outbreak may occur. In these cases, the endemic point is where both derivatives in (1) are equal to zero, and is given by:

$$(S_e, I_e) = N\left(\frac{\gamma + \mu}{\beta}, \frac{\mu(\beta - \gamma - \mu)}{\beta(\gamma + \mu)}\right). \tag{3}$$

The stability analysis of (1), linearised around the endemic point (S_e, I_e), determines that the eigenvalues λ obey,

$$\lambda^2 + R_0\mu\lambda + \mu(\gamma + \mu)(R_0 - 1) = 0. \tag{4}$$

The trajectory of the deterministic approximation is oscillatory if and only if these eigenvalues are complex [13], which in turn requires,

$$\frac{\gamma + \mu}{\mu} > \frac{{R_0}^2}{4(R_0 - 1)}. \tag{5}$$

In any realistic system, $\gamma \gg \mu$ and inequality (5) is comfortably met. In that case, the trajectory of the deterministic approximation of a typical outbreak is shown in Fig. 2. It starts at point A, rises to a peak (B), falls through C to a first local minimum (D), and converges in a spiral towards the endemic point.

Since the stochastic model has discrete states, it is sometimes convenient to round the endemic state values up to the next highest integer pair:

$$(S_d, I_d) = (\lceil S_e \rceil, \lceil I_e \rceil). \tag{6}$$

It can also be shown from (1) and (3) that dI/dt is positive for $S > S_e$ and negative for $S < S_e$, that is:

$$\operatorname{sgn}\left(\frac{dI}{dt}\right) = \operatorname{sgn}\left(S - S_e\right). \tag{7}$$

As we mentioned above, the expected behaviour of the CTMC (continuous-time Markov chain) tracks the deterministic approximation as N becomes large. However, a stochastic realisation may fade out at the start (near point A), or in the first trough after the initial outbreak (near point D). The latter situation, known as epidemic fade-out, is the topic of this paper.

The initial state (point A) is (S_0, I_0). In all our calculations, I_0 is small and $S_0 = N - I_0$; this represents the beginning of an outbreak in a naive population.

2.2. Different transmission rate parameters

During the initial outbreak, a higher value of β causes S to fall to a lower value, which in turn can cause I to fall to a lower value during the first trough, increasing the probability of epidemic fade-out. On the other hand, when the first wave of infection subsides and the number of infectious individuals becomes very low, a lower value of β causes an increase in the probability of epidemic fade-out.

Therefore we conjecture that the optimal strategy is to allow a higher β early in the outbreak, and to implement the lower β later in the outbreak, as the CTMC approaches the first trough. (This is confirmed in the results in Section 5.1).

The initial transmission rate parameter, corresponding to intervention measures not being in place, is denoted $\beta^{(1)}$, and the value whilst the intervention is implemented is denoted $\beta^{(2)}$. We specify that $\beta^{(1)} > \beta^{(2)}$.

We also use a superscript in parentheses to represent variables corresponding to the use of $\beta^{(1)}$ or $\beta^{(2)}$; so for instance $P^{(k)}$ is the transition probability matrix when using $\beta = \beta^{(k)}$, for $k = 1, 2$.

2.3. Definition of epidemic fade-out

Informally, epidemic fade-out refers to fade-out during the first trough after the initial substantial wave of infection, roughly between points C and E in Fig. 2. But for calculations, it is important to have a precise definition. (In general the exact definition is not overly critical, as long as it is used consistently). To do so, we need to define the pre-condition (that a substantial outbreak has commenced), and then need to define what constitutes fade-out in the first trough (or conversely, what constitutes an escape from the first trough).

We define $\mathbb{S}$ to be the state space of all possible (S, I) values, and we define T to be:

$$T = \{(S, I) \in \mathbb{S} | S = S_d^{(1)} - 1\}. \tag{8}$$

T is illustrated by the green dotted line in Fig. 3. We define that the initial wave of infection has occurred if the CTMC reaches a state in T. We use this definition because if the CTMC satisfies this condition, it can be called a substantial outbreak, so a fade-out in the subsequent trough can reasonably be called epidemic fade-out.

Given that the CTMC reaches T, it will almost surely eventually fall to a state for which $I < I_d^{(1)}$. Therefore we define a two boundary hitting problem: epidemic fade-out occurs if the CTMC reaches a lower absorbing boundary L, before it reaches an artificial upper absorbing boundary U. The lower absorbing boundary is

$$L = \{(S, I) \in \mathbb{S} | I = 0\}.$$

In a CTMC in which β is constant, the deterministic approximation of the CTMC converges to a point near (S_d, I_d), as given by (6), taking an anticlockwise path in Fig. 2. So in that case, the line $U = \{(S, I) \in \mathbb{S} | S \geq S_d, I = I_d\}$ would be a natural definition of the upper absorbing boundary [3, 26].

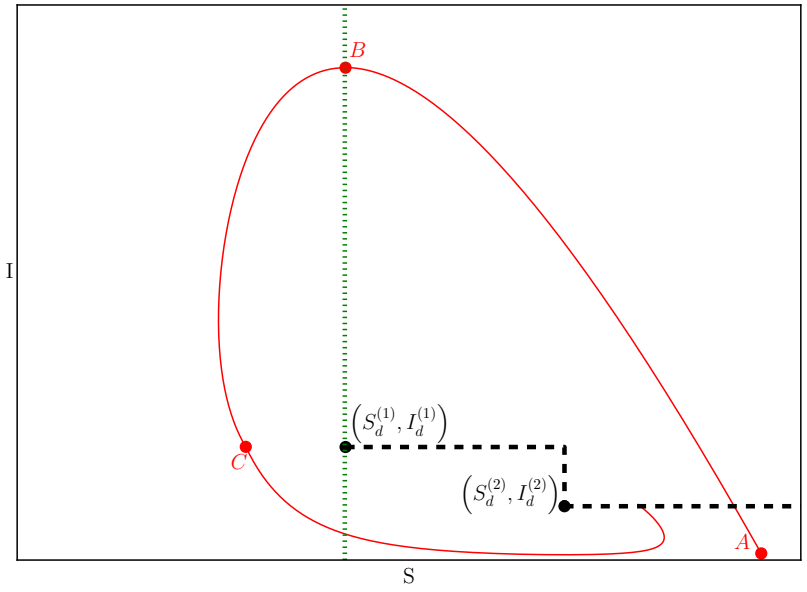

Figure 3: The probability of epidemic fade-out is the probability of the CTMC being absorbed at L ($I = 0$) before next reaching U (the black (dashed) line), given that it reaches T (the green (dotted) line). The red (solid) line shows the behaviour of the deterministic approximation to the CTMC, starting at point A; with the transmission rate parameter changing from $\beta^{(1)}$ to $\beta^{(2)}$ at point C.

However in a CTMC in which β can take two different values, there are two possible solutions to (6), depending on the value of β: $\left(S_d^{(1)}, I_d^{(1)}\right)$ (given by setting $\beta = \beta^{(1)}$ in (3)); and $\left(S_d^{(2)}, I_d^{(2)}\right)$ (given by setting $\beta = \beta^{(2)}$ in (3)); where $S_d^{(1)} \leq S_d^{(2)}$ and $I_d^{(1)} \geq I_d^{(2)}$.

To account for the possibility that β may be either $\beta^{(1)}$ or $\beta^{(2)}$, we end the first trough at $I = I_d^{(1)}$ for $S_d^{(1)} \leq S < S_d^{(2)}$, and at $I = I_d^{(2)}$ for $S_d^{(2)} \leq S$. We join the two boundaries with a vertical boundary at $S = S_d^{(2)}$ for $I_d^{(2)} \leq I \leq I_d^{(1)}$, giving:

$$
\begin{aligned}
&U = U_1 \cup U_2 \cup U_3, \text{ where} \\
&U_1 = \{(S, I) \in \mathbb{S} | I = I_d^{(1)}, S_d^{(1)} \leq S \leq S_d^{(2)}\}, \\
&U_2 = \{(S, I) \in \mathbb{S} | S = S_d^{(2)}, I_d^{(2)} \leq I \leq I_d^{(1)}\}, \text{ and} \\
&U_3 = \{(S, I) \in \mathbb{S} | I = I_d^{(2)}, S_d^{(2)} \leq S\}.
\end{aligned}
$$

This is illustrated in Fig. 3. We therefore define, p_0, the probability of epidemic fade-out, as the probability that the process is absorbed at L before reaching a state in U, given that it reaches T.

3. Idealised scenario: activation and de-activation of $\beta^{(2)}$

We first consider an idealised scenario, in which it is possible to switch an unlimited number of times between using $\beta^{(1)}$ and $\beta^{(2)}$, corresponding to activation and de-activation of the intervention measure, in a state dependent manner.

3.1. Definition of the problem

To give a practical solution, we limit the state space to be finite. So for the (infinite) state space $\mathbb{S}$, we define $\mathbb{S}'$ as the finite set of states,

$$
\mathbb{S}' = \{(S, I) | 0 \leq S \leq (1.1)N, 0 \leq I \leq (1.1)N\}. \tag{9}
$$

We enforce this by modifying the CTMC so that in Table 1, "Infection" events may not occur if $I \geq (1.1)N$ and "Birth of susceptible" events may not occur if $S \geq (1.1)N$. In all but the smallest systems, reaching these states is extremely improbable, so this modification has a negligible impact on the results.

We also define M as the set of transient states in $\mathbb{S}'$, that is all states in neither absorbing boundary:

$$
M = \mathbb{S}' \setminus (U \cup L). \tag{10}
$$

Then for some set $V \subseteq M$, the policy is:

$$
\beta = \begin{cases} \beta^{(2)} & \text{if } (S, I) \in V \text{ and the CTMC has previously reached } T; \\ \beta^{(1)} & \text{otherwise.} \end{cases}
$$

Problem. *Find the set V which maximises p_0.*

3.2. Optimal solution

This idealised scenario can be regarded as a Markov decision process, with an infinite horizon and no discounting, in which the "reward" is gained by reaching absorption at L. It can be solved by using a policy iteration algorithm [23, Section 7.2.5].

We define two transition probability matrices, $P^{(1)}$ and $P^{(2)}$ on M, corresponding to the jump chain of the CTMC [2]. The transition probabilities are calculated from the rates specified in Table 1; where $P_{ij}^{(k)}$ is the transition probability from state i to state j when using $\beta^{(k)}$ as the transmission rate parameter, for $k = 1, 2$. Transitions to absorbing states are not included, so some rows will sum to less than 1. These matrices are sparse, with at most four non-zero entries in each row.

Since we do not include the absorbing states in $P^{(1)}$ and $P^{(2)}$, the "reward" is only earned from states adjacent to the $I = 0$ boundary. We set up the respective reward vectors, $R^{(1)}$ and $R^{(2)}$, in which the reward is the probability of being absorbed at $I = 0$ on the next step. (Hence values will be non-zero only for states with $I = 1$.)

We create a decision vector $D(n)$ for each policy iteration step n. It has one entry per state in M, and every entry must be either 1 or 2. For each step n we also create the matrix $P(n)$ and the vector $R(n)$ according to the rule,

$$\begin{aligned} P_{ij}(n) &= P_{ij}^{(D_i(n))}, \\ R_i(n) &= R_i^{(D_i(n))}, \end{aligned} \tag{11}$$

for all $i, j \in M$.

The goal is to find the vector $D(n)$ which maximises p_0. Then V is the set of states $i \in M$ for which $D_i(n) = 2$.

The policy iteration algorithm solves this problem as follows:

1. Set $n = 1$ and initialise $D(n)$ to any permissible vector.
2. Build $P(n)$ and $R(n)$, according to (11).
3. Determine the column vector $v(n)$ by solving:

$$(\mathbb{I} - P(n))v(n) = R(n), \tag{12}$$

where $\mathbb{I}$ is the identity matrix. Then for each state $i \in M$, let

$$v_i^{(1)}(n) = R_i^{(1)} + \sum_{j \in M} P_{ij}^{(1)} v_j(n), \tag{13}$$

$$v_i^{(2)}(n) = R_i^{(2)} + \sum_{j \in M} P_{ij}^{(2)} v_j(n), \tag{14}$$

$$z_i(n) = \operatorname{sgn}\left(v_i^{(2)}(n) - v_i^{(1)}(n)\right). \tag{15}$$

(We use $v_i^{(1)}(n)$ and $v_i^{(2)}(n)$ on the left hand side of (13) and (14) because these equations are re-calculating $v_i(n)$ assuming $D_i(n) = 1$ and $D_i(n) = 2$ respectively, with the rest of the elements in $D(n)$ unchanged).

4. Update the policy: for each state i in D,

$$D_i(n+1) = \begin{cases} 1 & \text{if } z_i(n) = -1, \\ 2 & \text{if } z_i(n) = 1, \\ D_i(n) & \text{if } z_i(n) = 0. \end{cases} \tag{16}$$

5. If $D(n+1) = D(n)$, then the algorithm terminates, and V is the set of states i for which $D_i(n) = 2$. Otherwise, increment n and repeat from Step 2.

Notice that $v_i(n)$ is the probability of hitting L before U (for transition matrix $P(n)$ and reward vector $R(n)$), given that the CTMC is in state i. The policy iteration algorithm finds the $P(n)$ and $R(n)$ which give the maximum $v_i(n)$ for all $i \in M$ [23, Proposition 7.2.14], and hence finds the policy which gives the maximum p_0, regardless of the initial state (S_0, I_0).

3.3. Simplifying the policy iteration algorithm

The technique in Section 3.2 is useful for small populations, but becomes impractical for even moderate population sizes. For instance, a population with $N = 1000$ would have approximately 10^6 states in M. The corresponding two-dimensional matrix $(\mathbb{I} - P(n))$ in (12) is then approximately $10^6 \times 10^6$, and even though it is sparse, solving (12) takes significant computing resources. Therefore it would be beneficial to find a solution method which avoids the need to solve (12).

Let q_{ij} be the transition rate between any two states i and j, $i \neq j$, if $\beta^{(1)}$ is the transmission rate parameter (where $i \in M$ and $j \in \mathbb{S}'$). For notational convenience we let $q_{ii} = 0$. Also let $q_i = \sum_{j \in \mathbb{S}'} q_{ij}$ be the sum of all transition rates out of state i when $\beta^{(1)}$ is the transmission rate parameter. So $P_{ij}^{(1)} = \dfrac{q_{ij}}{q_i}$ and $R_i^{(1)} = \dfrac{\sum_{j \in L} q_{ij}}{q_i}$. In that case we may rewrite (13) as,

$$v_i^{(1)}(n) = \sum_{j \in L} \left(\frac{q_{ij}}{q_i} \right) + \sum_{j \in M} \left(\frac{q_{ij}}{q_i} \right) v_j(n).$$

In order to unify these two sums, we also define $v_j(n) = 1$ for $j \in L$, and $v_j(n) = 0$ for $j \in U$. Then,

$$v_i^{(1)}(n) = \frac{\sum_{j \in \mathbb{S}'} q_{ij} v_j(n)}{q_i}. \tag{17}$$

For every state $i \in M$, except those for which $S_i = 0$, let h be the state that is reached from i by an infection event, and let $\delta_i = (\beta^{(1)} - \beta^{(2)}) S_i I_i / N$, where S_i and I_i are the S and I values corresponding to state i. That is, $q_{ih} - \delta_i$ is the transition rate from state i to state h when using $\beta^{(2)}$ as the transmission rate parameter. (The $S_i = 0$ case is excluded because in that case no infection event

is possible, so $v_i^{(1)}(n) = v_i^{(2)}(n)$, so (15) always evaluates to zero and it never matters whether or not i is in V.) Thus, $S_i > 0$ ensures that $\delta_i > 0$. Then (14) becomes:

$$v_i^{(2)}(n) = \frac{\sum_{\{j \in \mathbb{S}', j \neq h\}} q_{ij} v_j(n)}{q_i - \delta_i} + \frac{(q_{ih} - \delta_i)\, v_h(n)}{q_i - \delta_i}$$

$$\Rightarrow v_i^{(2)}(n) = \frac{\sum_{j \in \mathbb{S}'} q_{ij} v_j(n)}{q_i - \delta_i} - \frac{\delta_i v_h(n)}{q_i - \delta_i}. \tag{18}$$

Substituting in (17) gives,

$$v_i^{(2)}(n) = \left(\frac{q_i}{q_i - \delta_i}\right) v_i^{(1)}(n) - \frac{\delta_i v_h(n)}{q_i - \delta_i}$$

$$\Rightarrow v_i^{(2)}(n) - v_i^{(1)}(n) = \left(\frac{q_i}{q_i - \delta_i} - 1\right) v_i^{(1)}(n) - \frac{\delta_i v_h(n)}{q_i - \delta_i}$$

$$\Rightarrow v_i^{(2)}(n) - v_i^{(1)}(n) = \frac{\delta_i \left(v_i^{(1)}(n) - v_h(n)\right)}{q_i - \delta_i}. \tag{19}$$

Then noting that $\delta_i > 0$ and $q_i - \delta_i > 0$,

$$\operatorname{sgn}\left(v_i^{(2)}(n) - v_i^{(1)}(n)\right) = \operatorname{sgn}\left(v_i^{(1)}(n) - v_h(n)\right). \tag{20}$$

We can also substitute (17) and then (18) into the right hand side of (19), to give:

$$v_i^{(2)}(n) - v_i^{(1)}(n) = \frac{\delta_i}{q_i - \delta_i}\left(\frac{\sum_{j \in \mathbb{S}'} q_{ij} v_j(n)}{q_i} - v_h(n)\right)$$

$$= \frac{\delta_i}{q_i - \delta_i}\left(\frac{\sum_{j \in \mathbb{S}'} q_{ij} v_j(n)}{q_i} - \frac{\delta_i v_h(n)}{q_i} - \frac{(q_i - \delta_i) v_h(n)}{q_i}\right)$$

$$= \frac{\delta_i}{q_i - \delta_i}\left(\frac{(q - \delta_i) v_i^{(2)}(n)}{q_i} - \frac{(q_i - \delta_i) v_h(n)}{q_i}\right)$$

$$\Rightarrow v_i^{(2)}(n) - v_i^{(1)}(n) = \frac{\delta_i \left(v_i^{(2)}(n) - v_h(n)\right)}{q_i}$$

$$\Rightarrow \operatorname{sgn}\left(v_i^{(2)}(n) - v_i^{(1)}(n)\right) = \operatorname{sgn}\left(v_i^{(2)}(n) - v_h(n)\right). \tag{21}$$

Comparing (12) to (13) and (14) tells us that $v_i(n)$ is equal to either $v_i^{(1)}(n)$ or $v_i^{(2)}(n)$, so (20) and (21) combine to give,

$$\operatorname{sgn}\left(v_i^{(2)}(n) - v_i^{(1)}(n)\right) = \operatorname{sgn}\left(v_i(n) - v_h(n)\right),$$

which can then be substituted into (15). So (12), (13) and (14) can be removed from Step 3 of the policy iteration algorithm, which simplifies to:

3. For each state i,

$$z_i(n) = \operatorname{sgn}\left(v_i(n) - v_h(n)\right). \tag{22}$$

The meaning of (22) and (16) is that we should reduce the transition rate from i to h only if $v_i(n) > v_h(n)$, which is a reasonably intuitive result.

An important feature of (22) is that it does not necessarily include a full matrix calculation. This opens the possibility of simpler ways to calculate an optimal, or close to optimal, policy. For instance, (22) could be evaluated for a small number of states, using an approximate method as in [3] to calculate $v_i(n)$ and $v_h(n)$.

3.3.1. A simplified policy based on the deterministic local minimum

A further advantage of (22) is that we do not need to calculate $v_i(n)$ and $v_h(n)$ at all. We only need to calculate which is greater.

We can get a very quick approximation of $\operatorname{sgn}\left(v_i(n) - v_h(n)\right)$ in (22) by taking advantage of a property which we reported previously [3]: for a state x in the region where dI/dt of the deterministic approximation is negative or zero (which means $S \leq S_e^{(1)}$, by (7)), $v_x(n)$ is generally negatively correlated to the minimum I value of the deterministic curve beginning at x.

That is, to a good approximation, the closer the curve of the deterministic approximation comes to an absorbing boundary, the more likely the process is to be absorbed at that boundary. A deterministic curve with a lower minimum passes closer to the absorbing boundary, and is closer to that boundary for a longer time; both of these effects contribute to making epidemic fade-out more probable.

Furthermore, it is possible to reduce this comparison to a formula. It is preferable to reduce β when $v_h < v_i$, which means (by our approximation) that h must be on a "higher" deterministic curve than i. A line from i to h (that is, from (S, I) to $(S-1, I+1)$) has a slope of -1. Since $dI/dt < 0$ in this region, a step of slope -1 goes to a "higher" deterministic curve when $dS/dt > -dI/dt$; that is, if:

$$(\gamma + \mu)I < \mu(N - S). \tag{23}$$

For states in the region where dI/dt of the deterministic approximation is positive (which means at least for $S > S_e^{(2)}$), we cannot use this approximation because the deterministic minimum has already been passed. However h is always on a higher (further from $I = 0$) curve than i, as well as having a higher I value; so $\beta = \beta^{(2)}$ is always preferred.

In the region $S_e^{(1)} < S \leq S_e^{(2)}$: if at any point we "try" $\beta = \beta^{(1)}$, this means that $S_e = S_e^{(1)}$, so $S > S_e$. Then we find (by the analysis in the previous paragraph) that $\beta = \beta^{(2)}$ is preferred. So this means that $\beta = \beta^{(2)}$ is preferred for all $S > S_e^{(1)}$.

Putting this together gives the following set V, the set of states in which to

use $\beta^{(2)}$:

$$\begin{aligned} &V = V_1 \cup V_2 \cup V_3, \text{ where} \\ &V_1 = \{(S,I) \in \mathbb{S} | S \leq S_e^{(1)}, (\gamma+\mu)I < \mu(N-S)\}, \\ &V_2 = \{(S,I) \in \mathbb{S} | S_e^{(1)} < S \leq S_e^{(2)}, I < I_e^{(1)}\}, \text{ and} \\ &V_3 = \{(S,I) \in \mathbb{S} | S_e^{(2)} < S, I < I_e^{(2)}\}. \end{aligned} \tag{24}$$

This formula is explicit, and it is quick to calculate regardless of the population size. It specifies V with a simple line, as illustrated by the green (top) line in Fig. 4.

If the CTMC roughly follows the deterministic approximation (Figs. 2 and 3) then V is first entered when $S < S_e^{(1)}$. So the most important component of (24) is V_1, as specified in (23). Note, importantly and possibly surprisingly, that (23) is independent of the values of $\beta^{(1)}$ and $\beta^{(2)}$. So the condition for initially using $\beta^{(2)}$ does not depend on the values of $\beta^{(1)}$ and $\beta^{(2)}$.

The formula also tells us, at least assuming the approximation used here, that the optimal policy cannot be improved by allowing three or more values of β. The optimum is always to use the highest available β for states not in V (corresponding to no intervention), and the lowest available β (corresponding to the most effective set of interventions) for states in V.

4. Realistic scenario: activation only of $\beta^{(2)}$

4.1. Definition of the problem

The idealised scenario in Section 3 corresponds to the most effective intervention possible, but it is not realistic. It allows the CTMC to repeatedly switch between using $\beta^{(1)}$ and $\beta^{(2)}$ as the state changes. In most real-world situations it would not be practical to start and stop infection-reducing measures as the process changes state near the boundary of V.

A more realistic situation, which we refer to as the *realistic scenario*, allows activation of $\beta^{(2)}$ only once. In this scenario, once $\beta = \beta^{(2)}$ is used, $\beta = \beta^{(2)}$ is *always* used (until the boundary L or U is reached), even if the CTMC subsequently leaves the region V. This is more practical because, in a typical application, measures to reduce the infection rate would be kept in place for a reasonable length of time once they are implemented.

As in Section 3.2, T is defined in (8), and M is defined in (10). Then for some set $V \subseteq M$, the policy is:

Initially, $\beta = \beta^{(1)}$. When the CTMC reaches a state in V, having previously been in a state in T, it permanently uses $\beta = \beta^{(2)}$.

Problem. *Find the set V which maximises p_0.*

4.2. Optimal solution

Since it is a relatively simple task to calculate the absorption probability once we are permanently using $\beta = \beta^{(2)}$, the realistic scenario can be regarded as an "optimal stopping" problem. The optimal stopping algorithm is as follows [23, Section 7.2.8]:

1. Build $P^{(1)}$, $R^{(1)}$, $P^{(2)}$ and $R^{(2)}$ as in Section 3.2.
2. Find $v^{(2)}$, the solution to

$$\left(\mathbb{I} - P^{(2)}\right) v^{(2)} = R^{(2)}. \tag{25}$$

   Now, $v^{(2)}$ is the vector of absorption probabilities assuming the transmission rate parameter is fixed at $\beta^{(2)}$, which form the "stopping rewards". So in state i, the CTMC can "stop" (switch to $\beta = \beta^{(2)}$) and take the "stopping reward" $v_i^{(2)}$.
3. Find, by linear programming, the vector v with the minimum $\sum_i v_i$, subject to the constraints:

$$v_i \geq \sum_{j \in M} P_{ij}^{(1)} v_j + R_i^{(1)} \quad \text{and} \quad v_i \geq v_i^{(2)}, \quad \forall i \in M. \tag{26}$$

4. Create V to represent the optimal policy. For each state i, if $v_i = v_i^{(2)}$, then the optimal policy in state i is to switch to using $\beta^{(2)}$, so i is added to V. If $v_i > v_i^{(2)}$, then the optimal policy in state i is to continue using $\beta^{(1)}$, so i is not added to V.

As in Section 3.2, v_i is the probability of hitting L before U, given that the CTMC is in state i. The vector $v = (v_i, i \in M)$, can also be calculated from a given V without running the algorithm: row i of P is zero if $i \in V$, and is otherwise equal to row i of $P^{(1)}$; element i of R is equal to $v_i^{(2)}$ if $i \in V$, and is otherwise equal to $R_i^{(1)}$; and v is the solution to,

$$(\mathbb{I} - P)\, v = R. \tag{27}$$

4.3. Simplifying the optimal stopping algorithm

If we use the same definitions for q_{ij}, q_i, δ_i and h as in Section 3.3, then the solution to (25) and (26) satisfies:

$$v_i^{(1)} = \frac{\sum_{j \in \mathbb{S}'} q_{ij} v_j}{q_i}, \tag{28}$$

$$v_i^{(2)} = \frac{\sum_{j \in \mathbb{S}'} q_{ij} v_j^{(2)}}{q_i - \delta_i} - \frac{\delta_i v_h^{(2)}}{q_i - \delta_i}, \quad \text{and} \tag{29}$$

$$v_i = \max\left(v_i^{(1)}, v_i^{(2)}\right), \tag{30}$$

for all $i \in M$.

Another way of expressing (30) is to say,

$$z_i = \text{sgn}\left(v_i^{(2)} - v_i^{(1)}\right); \tag{31}$$

where $v_i = v_i^{(2)}$ if $z_i = 1$, $v_i = v_i^{(1)}$ if $z_i = -1$, and v_i may be either if $z_i = 0$. In that case, we see that the (28), (29) and (31) are identical to (17), (18) and (15) respectively with the "(n)" postscripts removed, with the exception of the use of $v_i^{(2)}$ instead of v_i on the right hand side of (29). Then,

$$\begin{aligned}
v_i^{(2)} - v_i^{(1)} &= \frac{\sum_{j\in\mathbb{S}'} q_{ij} v_j^{(2)}}{q_i - \delta_i} - \frac{\delta_i v_h^{(2)}}{q_i - \delta_i} - v_i^{(1)} \\
&= \frac{\sum_{j\in\mathbb{S}'} q_{ij}\left(v_j^{(2)} - v_j\right)}{q_i - \delta_i} + \frac{\sum_{j\in\mathbb{S}'} q_{ij} v_j}{q_i - \delta_i} - \frac{\delta_i v_h^{(2)}}{q_i - \delta_i} - v_i^{(1)} \\
&= \frac{\sum_{j\in\mathbb{S}'} q_{ij}\left(v_j^{(2)} - v_j\right)}{q_i - \delta_i} + \frac{q_i v_i^{(1)}}{q_i - \delta_i} - \frac{\delta_i v_h^{(2)}}{q_i - \delta_i} - v_i^{(1)} \\
&= \frac{\sum_{j\in\mathbb{S}'} q_{ij}\left(v_j^{(2)} - v_j\right)}{q_i - \delta_i} + \frac{\delta_i\left(v_i^{(1)} - v_h^{(2)}\right)}{q_i - \delta_i} \\
&= \frac{\delta_i\left(v_i^{(1)} - v_h\right)}{q_i - \delta_i} + \frac{\left(\sum_{j\in\mathbb{S}'} q_{ij}\left(v_j^{(2)} - v_j\right)\right) - \delta_i\left(v_h^{(2)} - v_h\right)}{q_i - \delta_i} \\
\Rightarrow z_i &= \text{sgn}\left(\delta_i\left(v_i^{(1)} - v_h\right) + \left[\left(\sum_{j\in\mathbb{S}'} q_{ij}\left(v_j^{(2)} - v_j\right)\right) - \delta_i\left(v_h^{(2)} - v_h\right)\right]\right).
\end{aligned} \tag{32}$$

Although it may not be immediately obvious, (32) is quite similar to (22). The only differences are: the expression in square brackets; the presence of $v_i^{(1)}$ instead of v_i; and the inclusion of δ_i. Note that since $v_i \geq v_i^{(1)}$ and $v_i \geq v_i^{(2)}$ for all i, the right hand side of (32) cannot be greater than the right hand side of (22), so the criterion for including a state i in V is always more stringent in the realistic scenario than in the idealised scenario.

The presence of $v_i^{(1)}$ is due to the fact that the change is one-way, so for the initial application of the policy (that is, for the first iterative change from $v_i = v_i^{(1)}$ to $v_i = v_i^{(2)}$ in Section 3.2) the two equations are identical if the expression in square brackets is zero. The expression in square brackets accounts for whether the states surrounding state i are in V, and the q_{ij} and δ_i terms act as weighting factors.

However, when we look at the actual optimal policies generated by (22) (Fig. (4)), we see that most states in V are surrounded by other states in V, making the expression in square brackets equal to zero. So that suggests that

(32) will produce a V very similar to the V corresponding to the optimal policy for the realistic scenario – that is, that the optimal policies for the two scenarios will have very similar sets V.

This in turn suggests that using the set V defined in (24) will also be a good approximation of the optimal policy for the realistic scenario, as we see in the following section.

5. Results

We refer to the policy calculated in Section 3.2 as the *idealised scenario – optimal policy*, the policy calculated in Section 4.2 as the *realistic scenario – optimal policy*, and the policy calculated using (24) in Section 3.3.1 as the *simplified policy*.

5.1. Comparison of optimal policies to the simplified policy

Due to the computational requirements mentioned in Section 3.3, it is only feasible to calculate the optimal policies for small N. Policies were calculated using the methods in Sections 3.2, 4.2 and 3.3.1 for a range of parameters for $N \leq 300$. The corresponding sets V were calculated, and a typical result is shown in Fig. 4. The use of a small N necessitates choosing an unrealistically small value of γ/μ to illustrate the policies. However, the pre-condition (5) is still met.

We see that the idealised and realistic scenarios give very similar policies, and that the simplified policy is a close approximation of both. A similar result was seen with other sets of parameters. These results confirmed the prediction of Section 4.3, that the simplified policy in Section 3.3.1, is a good approximation for either the idealised scenario or the realistic scenario.

The result for the idealised scenario also confirmed the conjecture made in Section 2.2: that the higher β is preferable early in the outbreak, and the lower β is preferable as I falls to a low value. This confirms that in the realistic scenario we should switch from $\beta^{(1)}$ to $\beta^{(2)}$, not the other way around.

For small N it is also possible to calculate p_0 exactly: for the idealised scenario, (12) is solved and then $p_0 = v_0(n)$; for the realistic scenario, (27) is solved and then $p_0 = v_0$. Some typical results are shown in Fig. 5. It compares the optimal policies for the two different scenarios, as well as the simplified policy under the realistic scenario. Also shown are the outcomes with no change to β.

We see that the results from the three optimisation scenarios are extremely similar. This was a result we observed consistently over a wide range of parameters. This confirms another result of Sections 3.3 and 4.3: that the simplified policy is very nearly as good as the optimal policy, in either scenario.

Therefore we conclude that the simplified policy is a good practical choice, because it is easy to calculate, and so we use it in the further tests in the following sections.

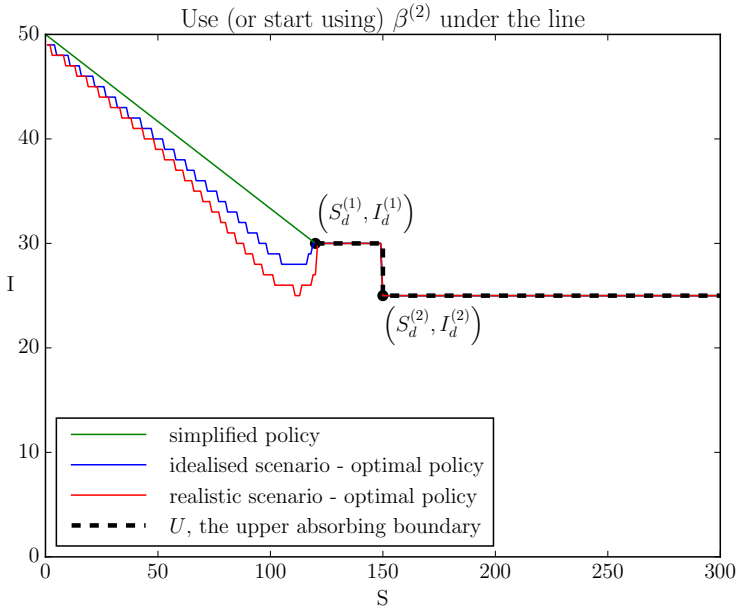

Figure 4: Comparison of the policies for exact methods of Sections 3.2 and 4.2, and the simplified method of Section 3.3.1. The parameters are $N = 300$, $\beta^{(1)} = 3$, $\beta^{(2)} = 2.4$, $\gamma = 1$ and $\mu = 0.2$. In each scenario, the policy (V) is the set of states below the respective line. Lines are in the same vertical order as the legend box. All lines are coincident along U. The endemic points corresponding to $\beta^{(1)}$ and $\beta^{(2)}$, $\left(S_d^{(1)}, I_d^{(1)}\right)$ and $\left(S_d^{(2)}, I_d^{(2)}\right)$ respectively, are marked.

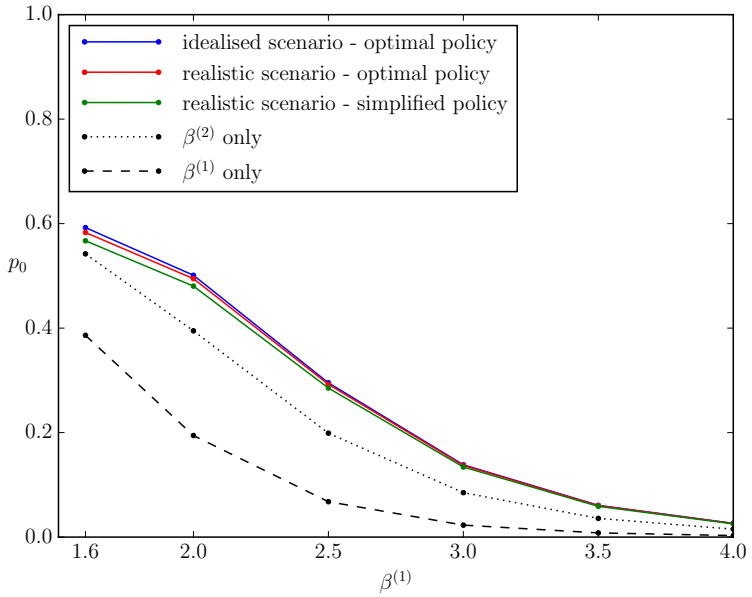

Figure 5: Comparison of p_0 versus $\beta^{(1)}$ for the three different policies, as well as for no change in β. The parameters are $N = 300$, $\beta^{(2)} = (0.8)\beta^{(1)}$, $\gamma = 1$, $\mu = 0.2$ and $(S_0, I_0) = (N-1, 1)$. $\beta^{(1)} = 3$ corresponds to the policies in Fig. 4. (Lines are in the same vertical order as the legend box.)

5.2. Effectiveness of the simplified policy

We examine the effectiveness of the simplified policy in the realistic scenario. Note that for large N, exact calculation of p_0 is impractical, so we calculate p_0 using the approximate solution method we previously reported, which has an average error of less than 1% [3].

Fig. 6 shows a typical result, varying $\beta^{(1)}$ and $\beta^{(2)}/\beta^{(1)}$ for a given N, γ and μ. (Although we only show the realistic scenario, the results for the idealised scenario are extremely close, to the point that the plots look identical.) We see that dramatic improvements in p_0 can be achieved with a relatively small reduction in β. Again, we tested a wide range of parameters, and the simplified policy consistently gave significant improvement.

In passing, note that the "$\beta^{(2)}/\beta^{(1)} = 1.0$" curve in Fig. 6 shows non-monotonicity in β, with a local maximum near $\beta/(\gamma+\mu) = 2$, as previously reported [3].

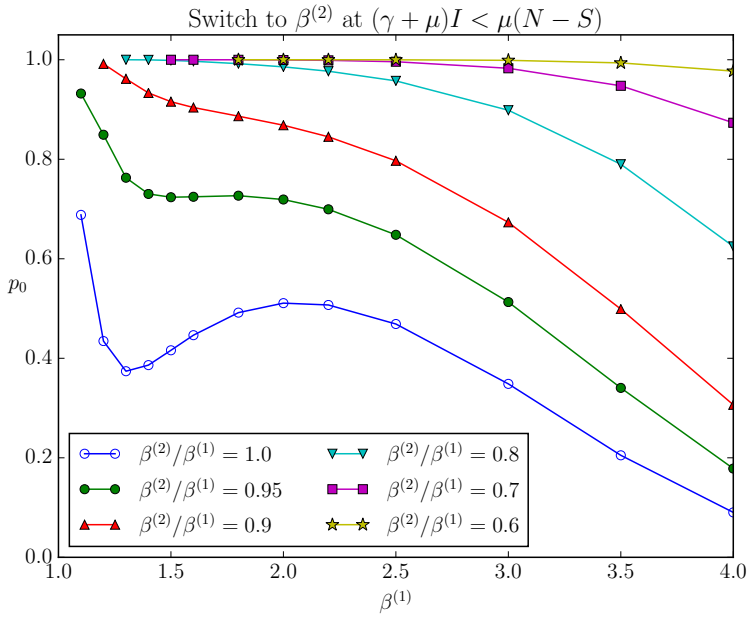

Figure 6: Plot of p_0 versus $\beta^{(1)}$, for different values of $\beta^{(2)}/\beta^{(1)}$ (realistic scenario, simplified policy). The parameters are $N = 30000$, $\gamma = 1$, $\mu = 0.025$ and $(S_0, I_0) = (N-1, 1)$.

5.3. Comparison of the simplified policy to other simple policies

The simplified policy of (24) is easy to calculate, but a practical problem is that in an outbreak scenario where we might wish to implement our policy, the precise values of the epidemiological parameters are only estimates, and the precise epidemiological status of the population (in terms of the numbers of susceptible and infectious individuals) can once again only be estimated.

Therefore we investigated the robustness of the policy to implementing the intervention at other stages of the epidemic. These results are shown in Fig. 7. We considered the $\beta^{(1)} = 3$ case of Fig. 6, keeping the same colour and marker scheme, but switched from $\beta = \beta^{(1)}$ to $\beta = \beta^{(2)}$ at other points in the cycle: points A through to E in Fig. 2. A refers to the case of using $\beta^{(2)}$ exclusively; B is the deterministic maximum I point ($S = S_d^{(1)}$ while $I > I_d^{(1)}$); C is when I falls to I_e; D is when I reaches its deterministic local minimum ($S = S_d^{(1)}$ after point C) and E refers to the case where β is always $\beta^{(1)}$. The simplified policy, denoted by $*$, is between points B and C.

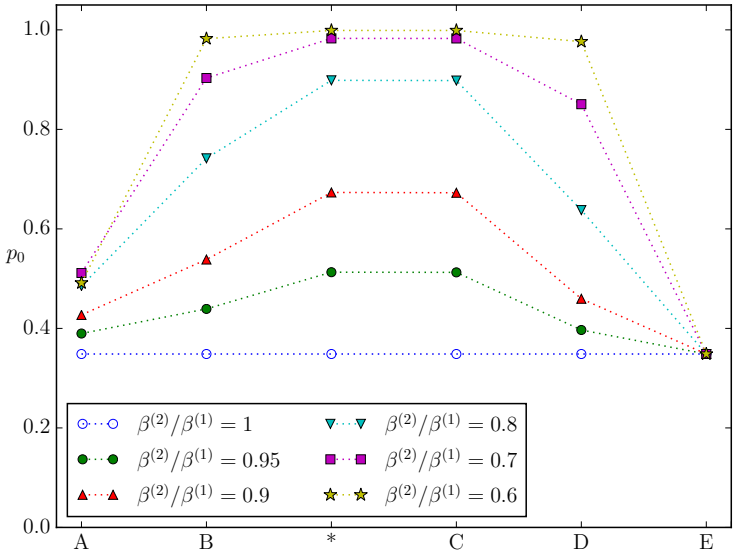

Figure 7: Comparison of p_0 for different switch points. $*$ is the simplified policy of (24). Points A to E are shown in Fig. 2. The $*$ results correspond to the $\beta^{(1)} = 3$ results in Fig. 6. The parameters are $N = 30000$, $\beta^{(1)} = 3$, $\gamma = 1$, $\mu = 0.025$ and $(S_0, I_0) = (N - 1, 1)$.

As predicted, the simplified policy always gives some improvement over using either $\beta^{(1)}$ or $\beta^{(2)}$ exclusively. However we found that although the simplified policy gave the best results, points B and C also gave significant improvement. On the other hand, intervening to reduce the transmission rate parameter, β, too late (point D) may or may not be preferable to using $\beta = \beta^{(2)}$ always (point A). (In the example in Fig. 7, it is always preferable, but tests with other parameter values have indicated that this is not always the case).

So this shows that there is a wide range of switch points which give some improvement over using either $\beta = \beta^{(1)}$ or $\beta = \beta^{(2)}$ exclusively. So long as a switch point is chosen after the peak I (point B) and some time before the deterministic local minimum of I (point D), a significant increase in p_0 will be

achievable.

6. Conclusion

In the SIR-with-demography model, reducing the transmission rate parameter from $\beta^{(1)}$ to $\beta^{(2)}$ at an appropriate point can give a substantial increase in the probability of epidemic fade-out, over that when using $\beta^{(1)}$ or $\beta^{(2)}$ exclusively. We believe that this has applications for timing the implementation of epidemic control measures, making it more likely for an epidemic to fade out before it becomes endemic. For instance, if there is a large outbreak, control measures (such as the allocation of antivirals) might be delayed until the epidemic is waning, approximately meeting the condition given by (23).

This method is effective because it allows the epidemic to progress longer without intervention and infect more individuals in the initial outbreak, but this is balanced against the long term gain of epidemic fade-out. As we previously noted, this is not likely to be ethically possible in lethal epidemics. However it could have applications in situations of non-lethal infections, or diseases among animals, where losses can be economically measured [16]. Possible future research is to investigate the tradeoffs between such policies in terms of total infections. The result of Grigorieva and Khailov with a deterministic SIR model [8], that not reducing β early in the infection cycle can minimise the total infection size, suggests that there might be a similar result when using a stochastic model.

Optimal policies may be calculated using Markov decision process theory, but these are impractical for all but the smallest systems. We have presented a simplified policy (24) which gives a very close to optimal solution. The key factor in determining this policy, the inequality $(\gamma + \mu)I < \mu(N - S)$, is easy to test and is independent of the transmission rate parameter, or the effectiveness of the control measures.

We also observed that even a sub-optimal switch point can give a substantial increase in the probability of epidemic fade-out. This should be useful for practical applications where the exact state of the system is not easily observed.

The method of calculating the simplified policy is based on using the deterministic local minimum to estimate the relative probability of epidemic fade-out. This technique should be amenable to many Markov process problems which concern optimising the probability of hitting one boundary before another. A possibility for future work is to apply this technique to related problems, such as more complicated models, or to the evaluation of the probability of fade-out at other points in the epidemic cycle.

7. Acknowledgments

This work is supported by an APA Scholarship (PB), an Australian Research Council Future Fellowship (JVR; FT130100254), and the NHMRC (JVR; CRE PRISM^2).

8. References

[1] R. M. Anderson and R. M. May. *Infectious Diseases of Humans: dynamics and control*. Oxford University Press, Oxford, 1991.

[2] W. J. Anderson. *Continuous-time Markov chains : an applications-oriented approach*. Springer-Verlag, New York, 1991.

[3] P. G. Ballard, N. G. Bean, and J. V. Ross. The probability of epidemic fade-out is non-monotonic in transmission rate for the Markovian SIR model with demography. *Journal of Theoretical Biology*, 393:170–178, 2016. doi: 10.1016/j.jtbi.2016.01.012.

[4] A. J. Black, T. House, M. J. Keeling, and J. V. Ross. Epidemiological consequences of household-based antiviral prophylaxis for pandemic influenza. *Journal of the Royal Society Interface*, 10:20121019, 2013. doi: 10.1098/rsif.2012.1019.

[5] T. Britton, T. House, A. L. Lloyd, D. Mollison, S. Riley, and P. Trapman. Five challenges for stochastic epidemic models involving global transmission. *Epidemics*, 10:54–57, 2015. doi: 10.1016/j.epidem.2014.05.002.

[6] O. Diekmann and J. A. P. Heesterbeek. *Mathematical Epidemiology of Infectious Diseases: Model Building, Analysis and Interpretation*. John Wiley & Sons, Chichester, 2000.

[7] Z. Feng, S. Towers, and Y. Yang. Modeling the effects of vaccination and treatment on pandemic influenza. *The AAPS Journal*, 13(3):427–437, September 2011. doi: 10.1208/s12248-011-9284-7.

[8] E. V. Grigorieva and E. N. Khailov. Optimal vaccination, treatment, and preventive campaigns in regard to the SIR epidemic model. *Mathematical Modelling of Natural Phenomena*, 9(4):105–121, 2014. doi: 10.1051/mmnp/20149407.

[9] A. Handel, I. M. Longini, and R. Antia. Antiviral resistance and the control of pandemic influenza: The roles of stochasticity, evolution and model details. *Journal of Theoretical Biology*, 256:117–125, 2009. doi: 10.1016/j.jtbi.2008.09.021.

[10] C. Jackson, P. Mangtani, J. Hawker, B. Olowokure, and E. Vynnycky. The effects of school closures on influenza outbreaks and pandemics: Systematic review of simulation studies. *PLoS ONE*, 9(5):e97297, 2014. doi: 10.1371/journal.pone.0097297.

[11] A. Kamenev and B. Meerson. Extinction of an infectious disease: A large fluctuation in a nonequilibrium system. *Physical Review E*, 77:061107, 2008. doi: 10.1103/PhysRevE.77.061107.

[12] T .K. Kar and A. Batabyal. Stability analysis and optimal control of an SIR epidemic model with vaccination. *BioSystems*, 104:127–135, 2011. doi: 10.1016/j.biosystems.2011.02.001.

[13] E. Kreyszig. *Advanced Engineering Mathematics*. John Wiley & Sons, New York, 8th edition, 1999.

[14] T. G. Kurtz. Solutions of ordinary differential equations as limits of pure jump Markov processes. *Journal of Applied Probability*, 7(1):49–58, 1970. doi: 10.2307/3212147.

[15] S. Lee and G. Chowell. Exploring optimal control strategies in seasonally varying flu-like epidemics. *Journal of Theoretical Biology*, 2016. doi: http://dx.doi.org/10.1016/j.jtbi.2016.09.023.

[16] C. Lefèvre. Optimal control of a birth and death epidemic process. *Operations Research*, 29(5):971–982, 1981.

[17] I. M. Longini, M. E. Halloran, A. Nizam, and Y. Yang. Containing pandemic influenza with antiviral agents. *American Journal of Epidemiology*, 159(7):623–633, 2004. doi: 10.1093/aje/kwh092.

[18] M. Lydeamore, N. G. Bean, A. J. Black, and J. V. Ross. Choice of antiviral allocation scheme for pandemic influenza depends on strain transmissibility, delivery delay and stockpile size. *Bulletin on Mathematical Biology*, 78(2): 293–321, 2016. doi: 10.1007/s11538-016-0144-6.

[19] B. Meerson and P. V. Sasorov. WKB theory of epidemic fade-out in stochastic populations. *Physical Review E*, 80:041130, 2009. doi: 10.1103/PhysRevE.80.041130.

[20] I. Nåsell. Stochastic models of some endemic infections. *Mathematical Biosciences*, 179:1–19, 2002. doi: 10.1016/S0025-5564(02)00098-6.

[21] R. L. Miller Neilan, E. Schaefer, H. Gaff, K. R. Fister, and S. Lenhart. Modeling optimal intervention strategies for cholera. *Bulletin of Mathematical Biology*, 72:2004–2018, 2010. doi: 10.1007/s11538-010-9521-8.

[22] A. B. Piunovskiy and D. Clancy. An explicit optimal intervention policy for a deterministic epidemic model. *Optimal Control Applications and Methods*, 29:413–428, 2008. doi: 10.1002/oca.834.

[23] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, New York, 1994.

[24] G. Rozhnova, C. J. E. Metcalf, and B. T. Grenfell. Characterizing the dynamics of rubella relative to measles: the role of stochasticity. *Journal of the Royal Society Interface*, 10:20130643, 2013.

[25] M. W. Tanner, L. Sattenspiel, and L. Ntaimo. Finding optimal vaccination strategies under parameter uncertainty using stochastic programming. *Mathematical Biosciences*, 215:144–151, 2008. doi: 10.1016/j.mbs.2008.07.006.

[26] O. A. van Herwaarden. Stochastic epidemics: the probability of extinction of an infectious disease at the end of a major outbreak. *Journal of Mathematical Biology*, 35:793–813, 1997.

[27] Y. Xiao, F. Brauer, and S. Moghadas. Can treatment increase the epidemic size? *Journal of Mathematical Biology*, 72:343–361, 2016. doi: 10.1007/s00285-015-0887-y.

[28] R. Yaesoubi and T. Cohen. Generalized Markov models of infectious disease spread: A novel framework for developing dynamic health policies. *European Journal of Operational Research*, 215:679–687, 2011. doi: 10.1016/j.ejor.2011.07.016.