



# Optimised prophylactic vaccination in metapopulations

Mingmei Teo<sup>a</sup>, Nigel Bean<sup>b,c</sup>, Joshua V. Ross<sup>b,\*</sup>

<sup>a</sup> Defence Science and Technology Group, Eveleigh, Australia

<sup>b</sup> School of Mathematical Sciences, The University of Adelaide, Adelaide, SA 5005, Australia

<sup>c</sup> ARC Centre of Excellence for Mathematical and Statistical Frontiers, School of Mathematical Sciences, The University of Adelaide, SA 5005, Australia

## ARTICLE INFO

### Keywords:

Final epidemic size  
Importation  
Stochastic model

## ABSTRACT

A highly effective method for controlling the spread of an infectious disease is vaccination. However, there are many situations where vaccines are in limited supply. The ability to determine, under this constraint, a vaccination strategy which minimises the number of people that become infected over the course of a potential epidemic is essential. Two questions naturally arise: when is it best to allocate vaccines, and to whom should they be allocated? We address these questions in the context of metapopulation models of disease spread. We discover that in practice it is generally optimal to distribute all vaccines prophylactically, rather than withholding until infection is introduced. For small metapopulations, we provide a method for determining the optimal prophylactic allocation. As the optimal strategy becomes computationally intensive to obtain when the population size increases, we detail an approximation method to determine an approximately optimal vaccination scheme. We find that our approximate strategy is consistently at least as good as three strategies reported in the literature across a wide range of parameter values.

## 1. Introduction

Infectious diseases have devastating impacts on society. Examples include the 1918 ‘flu pandemic and the Ebola epidemic in West Africa in 2014. Vaccination is an effective method for controlling the spread of an infectious disease, and in fact eradicated smallpox. However, vaccines are expensive to produce and for some diseases, vaccines may not yet exist and may only be developed during an epidemic. Both of these scenarios result in a limited supply of vaccines. Here we provide new insights and a new algorithm which can be used by public health officials to help find a solution to the challenging problem of how best to allocate a limited supply of vaccines.

Our focus for this paper is on determining an optimal allocation of a finite number of vaccines to a stochastic SIR (*susceptible-infectious-recovered*) metapopulation model, where optimal means minimising the mean final epidemic size (the expected number of people that become infected following an import attempt). The optimal allocation of a limited supply of vaccines to a population has been investigated in a range of different contexts, such as metapopulation models (Ball et al., 1997, 2004; Ball and Lyne, 2002; Keeling and Shattock, 2012; Duijzer et al., 2015, 2016; Keeling and Ross, 2015; Yuan et al., 2015) and age-structured models (Medlock and Galvani, 2009; Wallinga et al.,

2010). Stochastic models similar to the one we consider have been studied previously, but with a focus on reducing the household basic reproduction number,  $R_*$  (Ball et al., 1997, 2004; Ball and Lyne, 2002; Keeling and Ross, 2015). The objective of minimising the final epidemic size has been considered, but in the context of deterministic models (Keeling and Shattock, 2012; Duijzer et al., 2015, 2016), or for stochastic models in only specialised cases (Yuan et al., 2015; Keeling and Ross, 2015).

Three main strategies arise from these studies: the *equalising* strategy (Ball et al., 1997; Ball and Lyne, 2002); the *deterministic* strategy (Keeling and Shattock, 2012); and, a *fair (pro-rata)* strategy (Azman and Lessler, 2014), used by many as a basis for comparison. We will present these strategies in more detail in Section 4.3.

This paper contributes to the literature by considering a Markovian metapopulation model and the objective of minimising the mean final epidemic size subject to a constraint on the number of vaccines available. An important and novel feature of our work concerns the import of infection. Here, we assume that there is a single *attempted* import of infection into the population. This means there is a chance the infection may be blocked if infectious contact is with a vaccinated individual, especially in the case of prophylactic vaccination. Hence, this type of seeding (a single attempted import) is an important consideration.

\* Corresponding author.

E-mail address: [joshua.ross@adelaide.edu.au](mailto:joshua.ross@adelaide.edu.au) (J.V. Ross).

<https://doi.org/10.1016/j.epidem.2020.100420>

Received 12 February 2019; Received in revised form 19 August 2020; Accepted 15 November 2020

Available online 30 November 2020

1755-4365/© 2020 The Author(s).

Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license

(<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

The first question we answer is whether it is better to vaccinate some of the population before the attempted import, or to wait until the location of first import is known. Both of these approaches may have benefits. Vaccinating before infection is present in the population allows for a reduction in the number of susceptible individuals in the population which might reduce the chance of successful import; it also limits the spread of the disease if import is successful. On the other hand, vaccination post-import might allow for vaccine distribution based on the location of the infection(s). We found that for the question of *when* to allocate, it is in practice generally better to allocate vaccines before the presence of infection in the population.

Next, we explore the problem of *where* to allocate. That is, we consider the optimal pre-allocation of vaccines to the population. As it is computationally intensive to use exact methods to determine the optimal allocation of vaccines for large populations, we develop an approximately optimal allocation. We compare our approximate strategy, as well as three existing strategies in the literature, with the optimal strategy. We found that our approximate strategy compared well with the optimal strategy and also compared favourably with other strategies from the literature, especially in the case of *medium-sized* populations (for example, overall population of 1800 individuals). Overall, we found that when the optimal strategy cannot be calculated, as is the case when the population size is large, our approximate strategy is preferred and for very large population sizes, a deterministic strategy performs well.

## 2. Model

A metapopulation model, where individuals live in *patches* and have different rates of transmission *between* patches and *within* patches, is a useful extension to the general homogeneous population model. Such two-level mixing models can be used to capture people living in countries, cities and/or households, and animals on farms; the grouping/aggregation of individuals incorporates more realism in the model. Hence, we use a metapopulation structure for any similar grouping of individuals. For this model, we have also chosen to consider SIR disease dynamics, where each individual in the population is classified as susceptible, infectious or recovered. Under these dynamics, an individual is initially susceptible, and has some chance of becoming infected if contact is made with an infectious individual. After some period of time, the infectious individual recovers and becomes immune to the disease. Note that we have not considered an exposed class, where individuals are infected but cannot transmit the disease, as we are interested in the mean final epidemic size which is independent of the exposed class (Andersson and Britton, 2000).

Consider a metapopulation model with  $m$  patches and SIR disease dynamics. This can be modelled by a continuous-time Markov chain with state space,

$$\mathbb{S} = \{(s_1, i_1, \dots, s_k, i_k, \dots, s_m, i_m) : 0 \leq s_k, i_k, s_k + i_k \leq N_k, \forall k = 1, \dots, m\},$$

where  $N_k$  is the population size of patch  $k$  and  $s_k$  and  $i_k$  represent the number of susceptible and infectious individuals in patch  $k$ , respectively, for  $k = 1, \dots, m$ . The transition rates for this model, which describe the infection and recovery rates in patch  $k$ , for  $k = 1, \dots, m$ , are given by,

$$q_{x,x+\text{inf}_k} = \frac{\beta s_k i_k}{N_k - 1} + \sum_{j \neq k}^m \left( \frac{\alpha_{jk} s_k i_j}{N_k} + \frac{\alpha_{kj} s_k i_j}{N_j} \right),$$

which corresponds to a new infection in patch  $k$ , and

$$q_{x,x+\text{rec}_k} = \gamma i_k,$$

which corresponds to a recovery in patch  $k$ , where,

- $\mathbf{x} = (s_1, i_1, \dots, s_k, i_k, \dots, s_m, i_m)$ ,

- $\text{inf}_k = \mathbf{e}_{2k} - \mathbf{e}_{2k-1}$ , where  $\mathbf{e}_k$  is a vector of 0s with a 1 in the  $k$ th position,
- $\text{rec}_k = -\mathbf{e}_{2k}$ ,
- $\gamma$  is the recovery rate for each infectious individual,
- $\beta$  is the effective transmission rate parameter for transmission within a patch, and
- $\alpha_{jk}$  is the effective transmission rate parameter for contact between an individual from patch  $j$  with an individual in patch  $k$ .

These transition rates are used to form the generator matrix  $Q$  for this continuous-time Markov chain. Note, throughout we assume the transmission rates between patches, e.g.  $\alpha_{jk}$ , to be considerably smaller than the transmission rate within a patch,  $\beta$ , corresponding to transmission within a patch being much more likely than transmission between patches — e.g. transmission to an individual within the same city compared to an individual in a different city. Further, we focus on the case where the basic reproduction number,  $R_0$ , which we approximate as  $\beta/\gamma$ , is greater than 1, ensuring a chance of a major outbreak.

Recall that we have chosen to use the mean final epidemic size, which is the expected number of people that become infected over the course of an epidemic, as our measure to minimise. For the SIR model, the mean final epidemic size can also be thought of as the expected number of recovered individuals at the end of the epidemic and can be calculated using path integrals of continuous-time Markov chains (Pollett and Stefanov, 2002; Pollett, 2003). Here, we state the equation used to calculate the mean final epidemic size and refer the reader to (Teo, 2017) for further details. The mean final epidemic size,  $\zeta(\mathbf{x})$ , for every state  $\mathbf{x} \in \mathbb{S}$ , can be obtained by solving,

$$Q_* \zeta = -\mathbf{f},$$

where the vector  $\mathbf{f}$  contains the number of recovered individuals for elements corresponding to absorbing states and 0's elsewhere. To construct the matrix  $Q_*$ , we let  $B = \{(s_1, 0, \dots, s_j, 0, \dots, s_m, 0) : 0 \leq s_j \leq N_j, \forall j = 1, \dots, m\}$  be the set of states in  $\mathbb{S}$  with no infectious individuals and introduce a new absorbing state, call it 0. For this state, we set  $q_{0,x} = 0, \forall \mathbf{x} \in \mathbb{S}, q_{x,0} = 1, \forall \mathbf{x} \in B$  and  $q_{x,0} = 0, \forall \mathbf{x} \in \mathbb{S} \setminus B$ . Then,  $Q_*$  consists of the transition rates in  $Q$  and the additional rates described above, but restricted to the original state space  $\mathbb{S}$ , which are now the transient states. Note that this is the overall mean final epidemic size not the mean final epidemic size in each patch.

Next, we consider the impact of vaccination on the mean final epidemic size. The effect of vaccination on susceptible individuals can be modelled in a variety of ways. For example, vaccinating a susceptible individual can render the individual immune to the disease (Ball et al., 1997; Ball and Lyne, 2002; Keeling and Shattock, 2012; Duijzer et al., 2015, 2016; Keeling and Ross, 2015; Yuan et al., 2015) or vaccinating can reduce an individual's susceptibility to the disease (Ball and Lyne, 2006; Ball and Sirl, 2018). Here, we consider *perfect* vaccination, that is, when a susceptible individual is vaccinated, they become immune to the disease. For the SIR model, this can be represented by moving a susceptible individual to the recovered class. This means that in the calculation of the mean final epidemic size, the number of vaccinated individuals must be deducted from the number of recovered individuals at the end of the epidemic.

We assume that the metapopulation is disease-free before attempted import and infection may occur through a single attempted import of infection into the population. To model this, we let  $\rho_k$  be the probability that the single attempted import of infection is in patch  $k$ . This can be represented by  $\rho_k = \frac{\epsilon_k}{\sum_{j=1}^m \epsilon_j}$ , where  $\epsilon_k$  is the rate of import into patch  $k$ .

Note that we consider the rates  $\epsilon_k$  to be very small and we are interested in a single attempted import. This form for  $\rho_k$  allows for the consideration of different rates of import into different patches, corresponding to different levels of connectivity to the source of infection, say corresponding to air traffic volume with infected countries. If import is attempted at patch  $k$ , an individual is selected uniformly at random from

the patch; hence, import is successful in patch  $k$ , if a susceptible individual is chosen, with probability  $\frac{s_k}{N_k}$ . Hence, the probability of an individual in patch  $k$  becoming infected is  $\rho_k \times \left(\frac{s_k}{N_k}\right)$ , otherwise the infection fails to invade and the final epidemic size is 0. Note that if the entire population is initially susceptible, this probability reduces to  $\rho_k$  as  $s_k = N_k$ ; in this scenario, unsurprisingly, the optimal allocation is identical to the scenario of forced infection, where infection is always assumed to invade (rather than potentially blocked) as in this case a susceptible individual is chosen uniformly at random. This also remains true in the case  $\varepsilon_k \propto N_k$  because the mean final epidemic size for the latter case is a scaled version of the mean final epidemic size for the former. However, if the rate of import  $\varepsilon_k$  is not proportional to the patch size  $N_k$ , then there is a fundamental difference in the optimal strategy. Fig. 1 illustrates this phenomenon.

Consider a simple example with two patches of size  $N_1 = 9$  and  $N_2 = 15$ , where there are  $V = 6$  vaccines. We have chosen to fix the recovery rate  $\gamma = \frac{1}{2}$  and assume that the cross-patch infection rates are identical, that is,  $\alpha_{12} = \alpha_{21} = \alpha$ . Further, we have also chosen to vary the within-patch infection rate  $\beta \in [0.5, 5]$  and the cross-patch infection rate  $\alpha \in [0.01, 0.1] \times \beta$ . Fig. 1 consists of the optimal strategies (obtained using the method detailed in Section 4.1) for each  $(\alpha, \beta)$  pair and the following probabilities of attempted import into patch  $k$ ,  $\rho_k$ , for  $k = 1, 2$ ,

- $\rho_k = N_k / (N_1 + N_2)$  for  $k = 1, 2$  (Fig. 1b),
- $\rho_1 = 0.35$  and  $\rho_2 = 0.65$  (Fig. 1c),

- $\rho_1 = 0.4$  and  $\rho_2 = 0.6$  (Fig. 1d),
- $\rho_1 = 0.5$  and  $\rho_2 = 0.5$  (Fig. 1e), and
- $\rho_1 = 0.75$  and  $\rho_2 = 0.25$  (Fig. 1f).

Fig. 1a consists of the optimal strategies for the case of forced infection. This illustrates the importance of considering attempted import when determining an optimal prophylactic vaccination.

### 3. Prophylactic is better than reactive vaccination

First, we explore the question of *when* it is best to allocate vaccines. To do so, we consider three cases:

- allocating all vaccines before infection is present in the population (strict-prophylactic vaccination);
- withholding at least one vaccine until after infection is present in the population (predominantly-prophylactic vaccination); and,
- withholding all vaccines until after infection is present in the population (reactive vaccination).

To compare between these cases and determine the optimal vaccination strategy, we develop a method, the *modified BDP* algorithm (Appendix A.1 Algorithm 1), which utilises Backward Dynamic Programming (Powell, 2011; Puterman, 1994). The modified BDP algorithm is the standard backward dynamic programming algorithm, with a post-processing step that accounts for prophylactic vaccination and

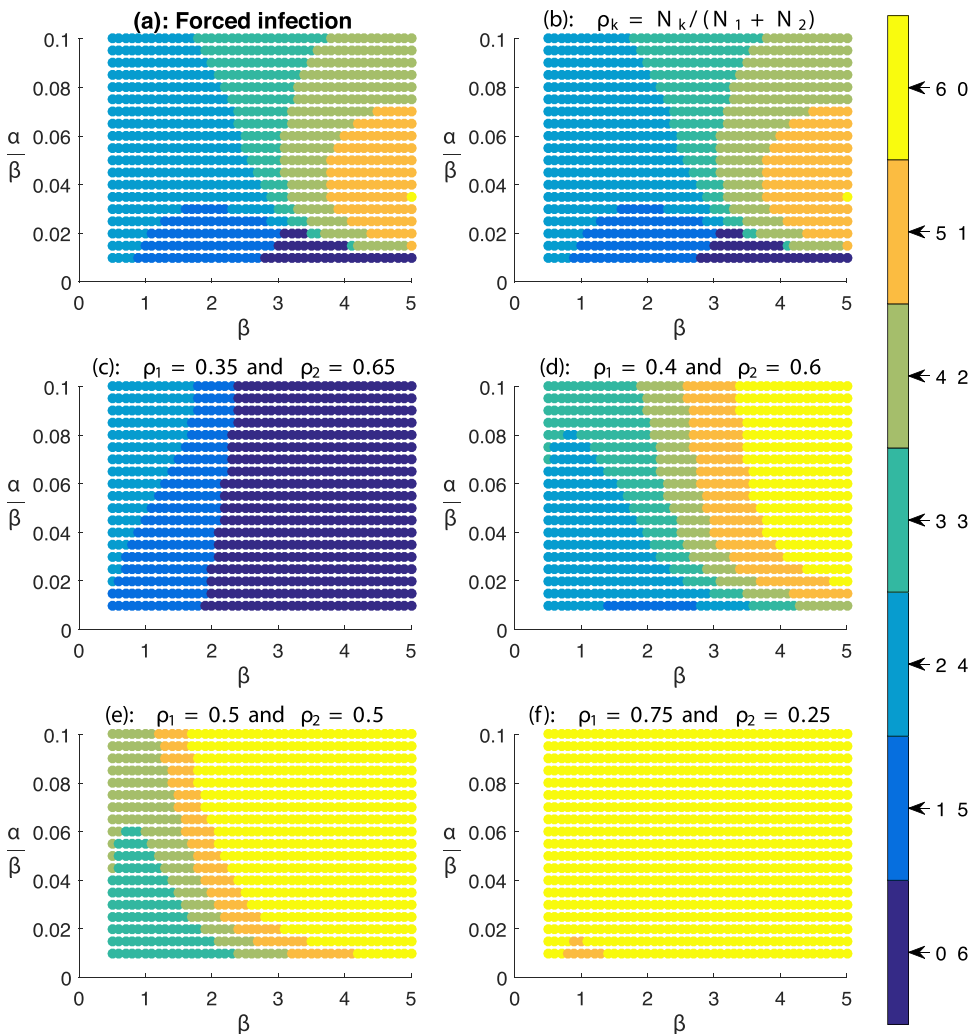


Fig. 1. Optimal strategies for a two-patch example with  $N_1 = 9$ ,  $N_2 = 15$ ,  $V = 6$  vaccines,  $\gamma = \frac{1}{2}$ , within-patch infection rate  $\beta \in [0.5, 5]$  and cross-patch infection rate  $\alpha \in [0.01, 0.1] \times \beta$  for different probabilities of attempted import. (a): Forced infection, (b): Import proportional to population size  $\rho_k \propto N_k$  for  $k = 1, 2$ , (c):  $\rho_1 = 0.35$  and  $\rho_2 = 0.65$ , (d):  $\rho_1 = 0.4$  and  $\rho_2 = 0.6$ , (e):  $\rho_1 = 0.5$  and  $\rho_2 = 0.5$ , (f):  $\rho_1 = 0.75$  and  $\rho_2 = 0.25$ . The colour labels represent the number of vaccines allocated to each patch,  $(V_1, V_2)$ . This figure shows that the relative rates of import affect the optimal strategies.

attempted import of infection. Further, we also incorporate a *delay* between the first onset of infection in the population and vaccination. In a real-world scenario, it is unlikely for public health officials to know instantaneously about the first onset of infection in the population and then be able to apply intervention measures immediately. Hence, there is a delay between these events. To account for this in the modified BDP algorithm, we let  $t_1$  be the time until vaccination can occur after infection is present in the population, measured in days. Then, the transition probabilities of moving from one state to another in time  $t_1$  can be calculated by  $P_1 := P(t_1) = e^{Qt_1}$ . In the example considered in Section 3.1, we vary the delay between first infection and vaccination,  $t_1$ , to explore how prophylactic vaccination compares with vaccinating after infection is present in the population. For a more detailed explanation of the standard backward dynamic programming algorithm and how it relates to our problem, we refer the reader to Appendix A.1.

Another important aspect of the modified BDP algorithm is how vaccines are allocated during an epidemic. As backward dynamic programming determines the best action to take at each time step,  $t'$ , we simplify the modelling, without loss of generality, and limit the number of possible actions by only allowing a single vaccine to be allocated at each time step  $t'$ . Then, given the current state of the process,  $\mathbf{S}_b$ , and the chosen action,  $a_b$ , the process moves to the next state  $\mathbf{s}'$  according to the probability  $P(\mathbf{s}'|\mathbf{S}_b, a_b)$ , which can be obtained from the transition probability matrix,  $P(t') = e^{Qt'}$ . Note that this transition probability matrix is different from the matrix  $P_1 = e^{Qt_1}$  and is used after the delay between first infection and vaccination,  $t_1$ . Hence, the process first evolves according to the transition probabilities in  $P_1$  for time  $t_1$ . After which, the process evolves according to  $P(t') = e^{Qt'}$ .

Next, we consider a simple example and apply the modified BDP algorithm to compare between prophylactic vaccination and reactive vaccination.

### 3.1. Results

Consider an example with 40 individuals split between two patches, and there are  $V = 10$  vaccines available to be allocated to the population over a fixed time-horizon of  $T = 12$  days. For this problem, we have chosen to fix  $\gamma = \frac{1}{2}$  and  $\beta = 1$ . Further, we also assume that the cross-patch infection rates,  $\alpha_{kj}$  for  $j \neq k$  and  $k, j = 1, 2$ , are identical and set to  $\alpha_{12} = \alpha_{21} = \alpha = 0.01$ . Recall that backward dynamic programming determines the best action to take at each time step,  $t'$ , corresponding to at most one vaccine allocated at each time step. To allow for a fairer comparison between prophylactic vaccination and reactive vaccination, we have chosen to let  $t' = \frac{1}{100}$ , as this essentially means that all vaccines are able to be allocated almost simultaneously after the delay for the reactive vaccination. For larger values of  $t'$ , reactive vaccinations are allocated at a slower rate and the disease could spread to more people.

Under this set up, we consider four ways to split the population of 40 individuals into two patches of size  $N_1$  and  $N_2$ :

1.  $[N_1, N_2] = [2, 38]$ ,
2.  $[N_1, N_2] = [10, 30]$ ,
3.  $[N_1, N_2] = [15, 25]$ , and
4.  $[N_1, N_2] = [20, 20]$ .

Then, for each  $[N_1, N_2]$ , we vary the time until vaccination can occur after infection is present in the population,  $t_1 \in \left\{ \frac{1}{3}, \frac{1}{2}, 1, 2 \right\}$ , which corresponds to a delay of eight hours, twelve hours, one and two days, respectively.

We apply the modified BDP algorithm (Appendix A.1 Algorithm 1) to each problem to obtain the minimum mean final epidemic size for the vaccination schemes of interest: strict-prophylactic vaccination; predominantly-prophylactic vaccination; and, reactive vaccination. Fig. 2 displays the results obtained from the modified BDP algorithm for each problem and each delay  $t_1$ .

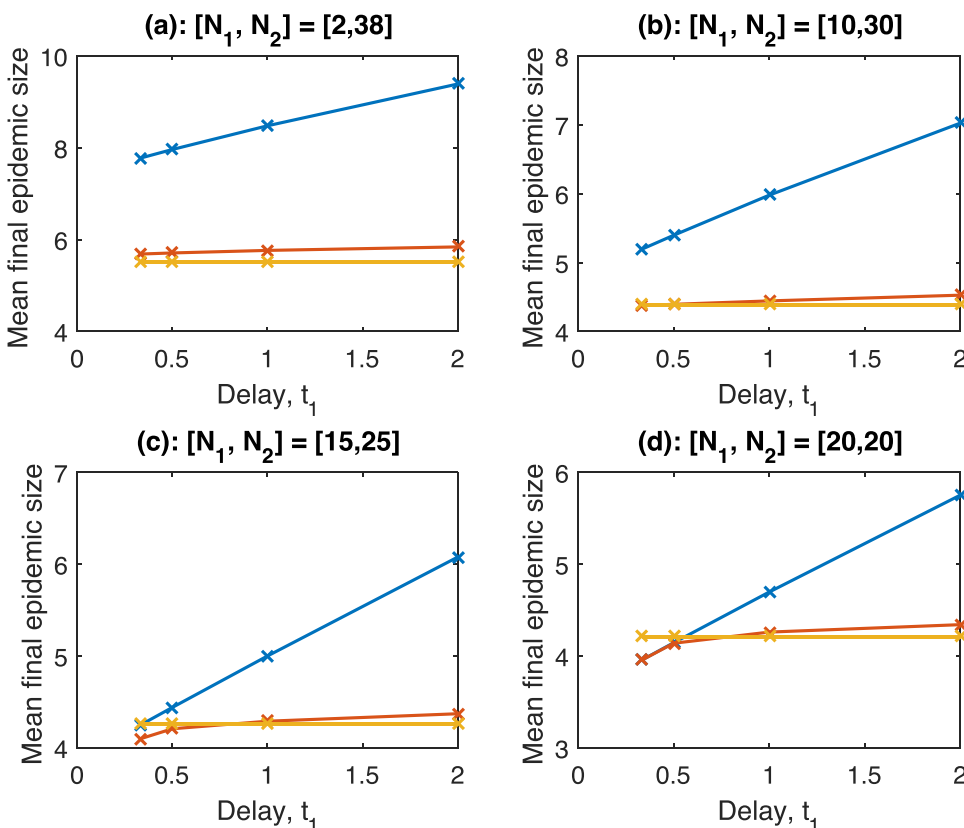


Fig. 2. Mean final epidemic size comparison of the three types of vaccination schemes, reactive vaccination (blue), predominantly-prophylactic vaccination (red) and strict-prophylactic vaccination (yellow), for a total population size of 40 individuals and  $V = 10$  vaccines by varying  $t_1$ , the delay between the first onset of infection in the population and vaccination. The following parameters are assumed fixed: the recovery rate  $\gamma = \frac{1}{2}$ , the within-patch infection rate  $\beta = 1$ , the cross-patch infection rate  $\alpha = 0.01$  and the time-horizon  $T = 12$  days. The subplots correspond to different patch sizes, (a):  $[N_1, N_2] = [2, 38]$ , (b):  $[N_1, N_2] = [10, 30]$ , (c):  $[N_1, N_2] = [15, 25]$ , (d):  $[N_1, N_2] = [20, 20]$ . Note that the y-axis does not start at zero, but instead has been chosen to make the signal more visible. This shows that, in general, strict-prophylactic vaccination is preferred as it results in smaller mean final epidemic sizes than withholding vaccines until after infection is present in the population, especially with realistic delays.

From Fig. 2, we see that, in general, withholding vaccines until after infection is present in the population results in a larger mean final epidemic size than strict-prophylactic vaccination. Hence, in general, strict-prophylactic vaccination is optimal. The only instances where there might be some benefit in withholding vaccines (some or all) until after infection is present in the population occurs when the delay between infection and vaccination is less than 1 day and when the patch sizes are relatively even. The same analysis was performed for different values of within-patch,  $\beta$ , and cross-patch,  $\alpha$ , infection rates. For larger  $\beta$  values or larger  $\alpha$  values, we found that prophylactic vaccination resulted in a smaller mean final epidemic size than when vaccines were withheld. Note that the inherent time scale of the model considered is the mean infectious period,  $\frac{1}{\gamma} = 2$  days in our examples. Therefore, the above results can be interpreted more generally by thinking of  $t_1 = \left\{ \frac{1}{6}, \frac{1}{4}, \frac{1}{2}, 1 \right\}$  mean infectious periods as opposed to days. In this language, we have that strict-prophylactic vaccination is optimal unless both the delay between infection and vaccination is less than half a mean infectious period and the patch sizes are relatively even.

Hence, given the typical mean infectious period of a respiratory virus and anticipated delays in responding, there appears to be no benefit to withholding vaccines until after the infection is present for the basic metapopulation model considered. If vaccination could be applied very rapidly after the first onset of infection, then some minor benefit exists. In general for this metapopulation with minimal structure, we find that prophylactic vaccination results in a smaller mean final epidemic size and hence is the preferred scheme. Next, we focus on the problem of prophylactic vaccination and determining the best way of allocating a fixed number of vaccines to such a metapopulation.

#### 4. Optimal prophylactic vaccination

Consider the metapopulation model described in Section 2. We know from the previous section that the preferred vaccination scheme is to allocate vaccines before the first onset of infection in the population. Hence, given a fixed number of vaccines, we consider all possible ways of allocating the available vaccines to the population and the optimal strategy is the one that results in the smallest mean final epidemic size. We first discuss how to solve this problem exactly before explaining the need for an approximation to the exact solution and considering other strategies found in the literature.

##### 4.1. Optimal strategy

As the optimal strategy corresponds to the allocation with the smallest mean final epidemic size, we need to calculate the mean final epidemic size for all possible allocations of available vaccines to the population. Let  $V_k$  be the number of vaccines allocated to patch  $k$  with the total number of vaccines,  $V = \sum_{k=1}^m V_k$ , and let  $u_k = N_k - V_k$  denote the number of unvaccinated susceptible individuals in patch  $k$ . Recall from Section 2 that the number of recovered (including vaccinated) individuals at the end of the epidemic for every possible initial state  $\mathbf{s}$ ,  $\zeta(\mathbf{s})$ , can be obtained simultaneously by solving,

$$Q \cdot \zeta = -\mathbf{f},$$

and extracting the appropriate element from  $\zeta$ . Then, for the initial state of interest  $\mathbf{u} = (u_1, 0, \dots, u_k, 0, \dots, u_m, 0)$ , the mean final epidemic size given the infection started in patch  $k$  is  $[\zeta(\mathbf{u} + \text{inf}_k) - V]$ . By multiplying this with the probability of an individual in patch  $k$  becoming infected,  $\rho_k \times \left(\frac{u_k}{N_k}\right)$ , and summing over all patches, we obtain the mean final epidemic size for the initial state of interest  $\mathbf{u}$ ,

$$\zeta'(\mathbf{u}) = \sum_{k=1}^m \rho_k \times \left(\frac{u_k}{N_k}\right) [\zeta(\mathbf{u} + \text{inf}_k) - V].$$

The problem of determining the optimal strategy corresponding to the minimum mean final epidemic size can then be written as the following minimisation problem,

minimise  $\zeta'(\mathbf{u})$ ,

subject to

$$0 \leq V_k \leq N_k, \quad \text{where } u_k = N_k - V_k, \quad \text{for } k = 1, \dots, m,$$

$$V = \sum_{k=1}^m V_k.$$

We note that this optimisation problem is relatively easy to solve once the vector  $\zeta$  is calculated as it simply involves selecting the strategy corresponding to the smallest mean final epidemic size  $\zeta'(\mathbf{u})$ , which is itself a fairly trivial calculation given  $\zeta$ .

The difficulty with the optimisation problem lies with the calculation of the vector  $\zeta$  as it is exact and requires solving a system of linear equations involving the  $Q$  matrix which is of order  $|\mathbb{S}| \times |\mathbb{S}|$ . Further, the size of the state space grows with the population size as

$$|\mathbb{S}| = \prod_{k=1}^m \frac{(N_k + 1)(N_k + 2)}{2}.$$

Hence, as the population size increases, obtaining the mean final epidemic size using this method becomes more computationally intensive and for large population sizes, this method is computationally intractable. This indicates a need to develop a more computationally-efficient method of estimating the mean final epidemic size to sufficient accuracy in order to develop an approximate strategy.

##### 4.2. Approximate strategy

Here, we detail our approximation method which consists of two approximations. The first we call the *average initial infection rate approximation* and the second the *weakly-coupled final sizes approximation*. Also, involved in our method is a rule for determining which approximation to use for a given parameter set.

First, consider the average initial infection rate approximation. To best control the epidemic, we are interested in determining the vaccine allocation that minimises the average initial infection rate. For an initial state  $\mathbf{u} = (u_1, 0, \dots, u_k, 0, \dots, u_m, 0)$ , where  $u_\ell = N_\ell - V_\ell$  for  $\ell=1, \dots, m$ , the average initial infection rate,  $r$ , can be calculated by

$$r = \sum_{k=1}^m \rho_k \times \left(\frac{u_k}{N_k}\right) \times \left[ \frac{\beta(u_k - 1)}{N_k - 1} + \sum_{j \neq k}^m \left( \frac{\alpha_{kj} u_j}{N_j} + \frac{\alpha_{jk} u_j}{N_k} \right) \right].$$

For details of its derivation, we refer the reader to Appendix A.2. Then, we evaluate  $r$  for all feasible combinations of  $u_k$ , for  $k = 1, \dots, m$ , and the combination corresponding to the smallest value of  $r$  is chosen as the optimal initial state.

Next, we consider the weakly-coupled final sizes approximation which is suitable when the cross-patch infection rate is sufficiently small relative to the within-patch infection rate. Recall from Section 4.1 that the problem with obtaining the optimal strategy exactly is that the calculation of the mean final epidemic size  $\zeta^*(\mathbf{u} + \text{inf}_k)$ , where  $\zeta^*(\mathbf{u} + \text{inf}_k) = \zeta(\mathbf{u} + \text{inf}_k) - V$ , can be computationally intensive, or even intractable, for large population sizes. Hence, for this approximation, we develop a less computationally-intensive weakly-coupled method to approximate  $\zeta^*(\mathbf{u} + \text{inf}_k)$ . We refer the reader to Appendix A.3 for more details on this approximation and provide the idea behind this approximation below.

As we assume that the cross-patch infection rate is sufficiently small relative to the within-patch infection rate, the probability of infection spreading between patches is small. Therefore, we approximate the mean final epidemic size for each patch using single population models. The basis of this approximation is to calculate the mean final epidemic size of the initially infected patch and then add to it the mean final epidemic size of each patch which has not yet experienced infection, weighted by the probability of each of those patches becoming infected. Note that we consider all possible pathways that the infection can spread from the initially infected patch to all fully susceptible patches. From this, we can obtain the approximate mean final epidemic size for each strategy and the optimal strategy is the one corresponding to the smallest approximated mean final epidemic size.

As the average initial infection rate approximation and the weakly-coupled final sizes approximation each return an optimal strategy for a given parameter set, we need a method to determine which approximation, and hence corresponding strategy, is most appropriate for that parameter set. Recall that the weakly-coupled final sizes approximation is restricted to the region where the cross-patch infection rate is sufficiently small relative to the within-patch infection rate. On the other hand, there is no restriction on the average initial infection rate approximation. Hence, our method needs to determine, for a given parameter set, if the weakly-coupled final sizes approximation can be applied or the average initial infection rate approximation should be used instead.

Let us consider the transition rates of the model (described in Section 2) and focus specifically on the infection events. The rate of infection in patch  $k$ , for  $k = 1, \dots, m$ , is,

$$\frac{\beta s_k i_k}{N_k - 1} + \sum_{j \neq k}^m \left( \frac{\alpha_{kj}}{N_j} + \frac{\alpha_{jk}}{N_k} \right) s_k i_j = s_k \left[ \frac{\beta i_k}{N_k - 1} + \sum_{j \neq k}^m \left( \frac{\alpha_{kj}}{N_j} + \frac{\alpha_{jk}}{N_k} \right) i_j \right]$$

As the weakly-coupled final sizes approximation assumes that the cross-patch infection rate is sufficiently small relative to the within-patch infection rate, we want, for each  $k = 1, \dots, m$ ,

$$\sum_{j \neq k}^m \left( \frac{\alpha_{kj}}{N_j} + \frac{\alpha_{jk}}{N_k} \right) i_j \ll \frac{\beta i_k}{N_k - 1}.$$

Note that the above expression is state-dependent as it depends on the values of  $i_j$  and  $i_k$ . Instead, we want our expression to be state-independent and so we bound these terms by the largest possible value they can take,  $N_j$  and  $N_k$  (Teo, 2017). Thus, the expression of interest to us can be written, for each  $k = 1, \dots, m$ , as

$$\sum_{j \neq k}^m \left( \frac{\alpha_{kj}}{N_j} + \frac{\alpha_{jk}}{N_k} \right) N_j \ll \frac{\beta N_k}{N_k - 1}.$$

For large  $N$ , this expression can be approximated by,

$$\sum_{j \neq k}^m \left( \alpha_{kj} + \frac{N_j}{N_k} \alpha_{jk} \right) \ll \beta.$$

We need this condition to hold for all patches  $k = 1, \dots, m$ , for the weakly coupled final sizes approximation to be a suitable approximation. Hence, for some chosen cut-off value  $c \ll 1$ , if

$$\frac{1}{\beta} \times \max_k \left\{ \sum_{j \neq k}^m \left( \alpha_{kj} + \frac{N_j}{N_k} \alpha_{jk} \right) \right\} < c,$$

then the weakly-coupled final sizes approximation is applied. After investigation of potential choices for the cut-off value  $c$  (see (Teo, 2017)), a cut-off value of  $c = 0.175$  was chosen as a robust choice (for a range of population sizes, number of patches and parameter values).

### 4.3. Other strategies in the literature

We compare our approximate strategy to other strategies found in the literature, in particular the *equalising* strategy (Ball et al., 1997; Ball and Lyne, 2002), the *deterministic* strategy, obtained from a modified model of Keeling and Shattock (Keeling and Shattock, 2012), and a *fair (pro-rata)* strategy (Azman and Lessler, 2014), used by many as a basis for comparison.

The *equalising* strategy involves distributing vaccines in such a way that an equal, or close to equal, number of susceptible individuals remain in each household after vaccination. We note that the main difference between the household model used for the equalising strategy and the model we have used is in their focus. Our metapopulation model considers a few large groups whilst the household model considers many small groups. Although there is a difference in their focus, we believe that the equalising strategy is a good comparison to make. This strategy was shown to be *optimal* for a stochastic SIR model with an infinitely large population of households of size two, three and four and was conjectured to hold for larger household sizes, with numerical investigations supporting this conjecture. Their definition of optimality was to determine the minimal number of vaccines and an optimal allocation scheme that brings the basic household reproduction number,  $R_*$ , below 1. It was also assumed that vaccination rendered an individual immune and that the transmission of infection is density dependent, that is, the chance of transmission between any two individuals in the population is independent of the household size. Fig. 3 illustrates this strategy for a simple example consisting of three households of size  $N_1 = 2, N_2 = 4, N_3 = 6$  with  $V = 6$  vaccines.

Next, we consider a *fair (pro-rata)* strategy which allocates vaccines in proportion to the number of susceptible individuals in each patch. Due to vaccines being discrete, it is possible for there to be more than one fair strategy; herein we compare all such strategies. Fig. 3 illustrates this strategy for a simple example consisting of three households of size  $N_1 = 2, N_2 = 4, N_3 = 6$  with  $V = 6$  vaccines.


The final strategy we consider is the *deterministic* strategy obtained using a variation of the deterministic model presented by Keeling and Shattock (Keeling and Shattock, 2012). We made slight changes to the notation used by Keeling and Shattock in (Keeling and Shattock, 2012), incorporated a single attempted import of infection into their model and derived analogous formulae to their equations for the final epidemic size and reproduction number. The details of our modifications to their formulae are presented in Appendix A.4. Here, we present our method for obtaining the deterministic strategy. The final size for patch  $\ell$ , for  $\ell = 1, \dots, m$ , where infection has been seeded in patch  $k$ , can be calculated using the final size equation

$$Z_\ell^\infty(k) = (N_\ell - V_\ell) - S_\ell^0(k) \exp \left( - \sum_{j=1}^m \frac{R_0^{j\ell}}{N_\ell} Z_j^\infty(k) \right)$$


This can be solved using simple recursion as follows,

$$Z_\ell^{t+1}(k) = (N_\ell - V_\ell) - S_\ell^0(k) \exp \left( - \sum_{j=1}^m \frac{R_0^{j\ell}}{N_\ell} Z_j^t(k) \right), \quad \text{for } \ell = 1, \dots, m,$$


Equalising strategy:



Household 1




Household 2




Household 3


Fair strategy:



Household 1



Household 2



Household 3

**Fig. 3.** Example illustrating the allocation of vaccines using the equalising strategy and the fair strategy for a simple example with three households of size  $N_1 = 2$ ,  $N_2 = 4$ ,  $N_3 = 6$  with  $V = 6$  vaccines. Here, blue represents a susceptible individual and orange represents a vaccinated individual.

with

$$S_\ell^0(k) = \begin{cases} N_k - V_k - 1, & \text{if } \ell = k, \\ N_\ell - V_\ell, & \text{otherwise,} \end{cases}$$

and initial condition,

$$Z_\ell^1(k) = \begin{cases} 1, & \text{if } \ell = k, \\ 0, & \text{otherwise,} \end{cases}$$

where  $R_0^{j\ell}$  is the expected number of secondary cases produced in patch  $\ell$  by a single infected individual in patch  $j$ , assuming all individuals in patch  $\ell$  are initially susceptible, and

$$\frac{R_0^{j\ell}}{N_\ell} = \begin{cases} \frac{\beta}{\gamma(N_\ell - 1)}, & \ell = j, \\ \frac{\alpha_j}{\gamma N_\ell} + \frac{\alpha_\ell}{\gamma N_j}, & \ell \neq j. \end{cases}$$

Then, the final epidemic size for a given initial state  $\mathbf{u}(k) = (u_1, 0, \dots, u_k - 1, 1, \dots, u_m, 0)$ , where infection has been seeded in patch  $k$ , is

$$\xi^{(D)}(\mathbf{u}(k)) = \sum_{\ell=1}^m Z_\ell^\infty(k),$$

where  $Z_\ell^\infty(k)$  is the final size for patch  $\ell$  obtained through recursion.

This gives us the final epidemic size when infection has been seeded in patch  $k$ . Hence, for a given parameter set, the mean final epidemic size for a given strategy  $\mathbf{u} = (u_1, 0, \dots, u_k, 0, \dots, u_m, 0)$  is,

$$\zeta^{(d)}(\mathbf{u}) = \sum_{k=1}^m \rho_k \times \left( \frac{u_k}{N_k} \right) \times \xi^{(D)}(\mathbf{u} + \text{inf}_k)$$

The optimal deterministic strategy is the one corresponding to the smallest  $\zeta^{(d)}(\mathbf{u})$ , calculated out of all possible strategies  $\mathbf{u}$ .

#### 4.4. Results

##### 4.4.1. Small populations

We first consider a small example with three patches of sizes  $N_1 = 6$ ,

$N_2 = 12$  and  $N_3 = 18$ , with  $V = 9$  vaccines available to the population. For this example, we make the assumption that the cross-patch infection rate is the same for all patches, i.e.,  $\alpha_{kj} = \alpha$  for  $k, j = 1, 2, 3$  and  $j \neq k$ . To allow for easy presentation of the results, we fix the recovery rate  $\gamma = \frac{1}{2}$  and vary the within-patch infection rate  $\beta \in [0.5, 5]$  and the cross-patch infection rate  $\alpha \in [0.01, 0.1] \times \beta$ . Fig. 4a consists of the optimal strategies and Fig. 4b consists of the approximate strategies for this example. Note that as there are a large number of possible vaccine allocation schemes, we have chosen a colour map that only considers the required vaccine allocations.

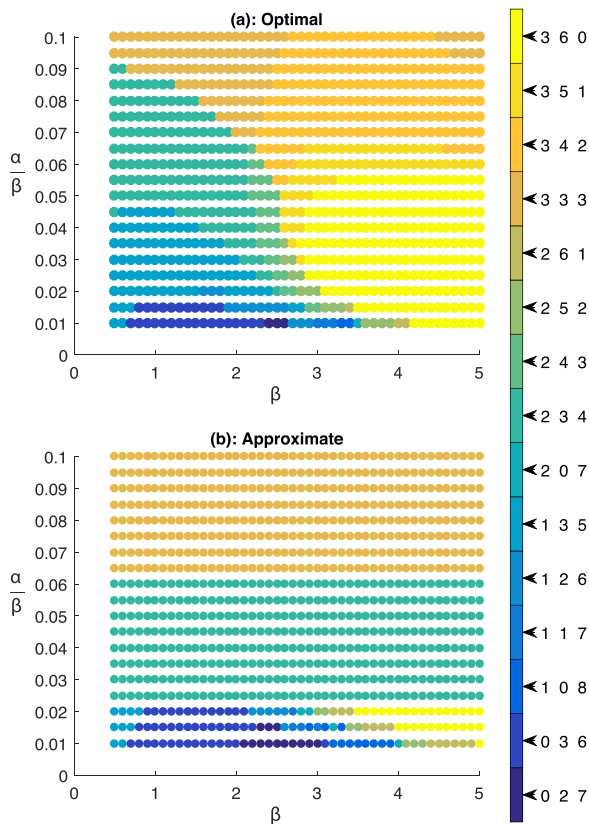
We can see from this figure that there are some differences between the approximate and optimal strategies. However, we note that comparing the strategies may not be the best approach as strategies which differ by a single vaccinated individual may not result in a large difference in the mean final epidemic size. Instead, we use the relative difference in mean final epidemic size as a measure to determine how well a proposed strategy performs in comparison with the optimal strategy.

Let  $\zeta_{opt}$  denote the mean final epidemic size for the optimal strategy and  $\zeta_{prop}$  denote the mean final epidemic size for the proposed strategy. Note that the mean final epidemic sizes,  $\zeta_{opt}$  and  $\zeta_{prop}$ , are calculated exactly using the method in Section 4.1. Then, the relative difference is given by,

$$\frac{|\zeta_{prop} - \zeta_{opt}|}{\zeta_{opt}}$$

Note that this is the relative difference for a single parameter set. If we have a range of parameters and are interested in comparing between various proposed strategies, then a single number which describes how well a proposed strategy performs over the range of parameters is preferred. Hence, we have chosen to use the average and maximum relative differences as our measures for comparison between proposed strategies.

Fig. 5 consists of the relative difference between each proposed strategy (approximate, deterministic, equalising and fair) and the optimal strategy and Table 1 consists of the average and maximum relative differences. Note that in Table 1, we consider two parameter regions. The first region consists of the entire parameter range and the



**Fig. 4.** Three-patch example with  $N_1 = 6$ ,  $N_2 = 12$ ,  $N_3 = 18$ ,  $V = 9$  vaccines,  $\gamma = \frac{1}{2}$ , within-patch infection rate  $\beta \in [0.5, 5]$  and between-patch infection rate  $\alpha \in [0.01, 0.1] \times \beta$ . (a): Optimal strategies, (b): Approximate strategy. The colour labels represent the number of vaccines allocated to each patch,  $(V_1, V_2, V_3)$ . Only the optimal allocations are shown in this figure. This figure shows that the optimal and approximate strategies match reasonably well especially for small  $\frac{\alpha}{\beta}$  values.

second region has the additional condition of  $\beta \in [0.5, 0.8]$  which corresponds to influenza-like parameters, where the basic reproduction number,  $R_0 = \beta/\gamma \in [1, 1.6]$ .

We observe from both Fig. 5 and Table 1 that the approximate strategy performs considerably better than the deterministic and equalising strategies and is also slightly better than the fair strategy for both parameter regions considered. Further, we also note that the average and maximum relative differences for the approximate strategy are very small in magnitude. Hence, this gives us confidence that the approximate strategy performs very well when compared to other strategies.

Another important point to note is the amount of time required to obtain the optimal strategies. Both the exact method and the approximate method were run on a Lenovo NeXtScale system consisting of 120 nodes, where one core of an Intel® 2.3 GHz Xeon E5-2698v3 node with 12 GB of RAM was used. The exact method required 3313 s to obtain the optimal strategies whilst the approximate method took only 1.74 s to obtain the approximate strategies. This considerable speed up in computational time combined with the small relative differences enhances the attractiveness of the approximate method. Further, as the size of the patches increases, obtaining the optimal strategies exactly becomes computationally infeasible, making an approximation method necessary to determine the best strategies.

#### 4.4.2. Larger populations

Next, we explore an example of a larger problem where it is computationally infeasible to obtain the optimal strategies exactly. Consider an example with 1800 individuals split into three patches of

sizes  $N_1 = 300$ ,  $N_2 = 600$  and  $N_3 = 900$ , and  $V = 450$  vaccines. Note that this is exactly the example used in the previous section, but scaled up by a factor of 50. As with the previous example, we fix the recovery rate for all patches ( $\gamma = \frac{1}{2}$ ) and assume that the within-patch infection rate is the same for all patches and the cross-patch infection rate is also the same for all patches. We also choose to vary the within-patch infection rate  $\beta \in [0.5, 5]$  and the cross-patch infection rate,  $\alpha \in [0.01, 0.1] \times \beta$ . Figs. 6a and 6b consist of the approximate and deterministic strategies for each  $(\alpha, \beta)$  pair respectively. From these figures, we observe that the approximate strategy differs from the deterministic strategy for  $\beta > 1.5$ .

Unlike the previous example, it is computationally infeasible to calculate the mean final epidemic size exactly for this problem. Hence, we use Sellke’s simulation method (Sellke, 1983) to obtain the final epidemic size for a given  $(\alpha, \beta)$  pair and a given strategy (initial state). As this only provides us with a single realisation of the final epidemic size, we chose to use  $10^6$  simulations to obtain a reasonable estimate for the mean final epidemic size.

Our inability to calculate the mean final epidemic size exactly for this large population means that we are unable to use the previous method of comparison where we calculated the relative difference in the exact mean final epidemic size between the optimal strategy and the proposed strategies (approximate, deterministic, equalising and fair). Instead, as we have used Sellke’s simulation method to estimate the mean final epidemic size, we perform statistical tests to determine if the estimated mean final epidemic size between the strategies are significantly different. This can then tell us how well the approximate strategy performs in comparison to the deterministic, equalising and fair strategies. We highlight some key points of our method of comparison and refer the reader to Appendix A.5 for more details.

The first statistical test we perform is the one-way ANOVA test, which tests if the means for all strategies are the same. For our problem, we have chosen to perform the one-way ANOVA test at a 5% significance level using the aov function in R (R Core Team, 2015). If we fail to reject the null hypothesis for the one-way ANOVA test, then we conclude that there is insufficient evidence that the estimated mean final epidemic sizes of the strategies are significantly different. Otherwise, the null hypothesis is rejected and we perform pairwise comparisons using Dunnett’s test (Dunnett, 1955, 1964), at the 5% significance level with the glht function in R (Hothorn et al., 2008), to determine which strategy or strategies differ from the approximate strategy. If the null hypothesis for Dunnett’s test is rejected, then we conclude that there is a difference in the estimated mean final epidemic size between the approximate strategy and the other strategy of interest. Further, if the difference (other – approximate) in the estimated mean final epidemic sizes is positive, then the approximate strategy is better than the other strategy. Otherwise, the other strategy is better.

We perform this testing procedure for the  $(\alpha, \beta)$  pairs considered, each with  $10^6$  simulations. Fig. 7 displays the results of the one-way ANOVA test. We observe from this figure that for all  $(\alpha, \beta)$  pairs, at least one of the mean final epidemic sizes is significantly different from the other means at the 5% significance level. Hence, we perform pairwise comparisons, using Dunnett’s test, to determine whether the approximate strategy performs better than the other strategies. Fig. 8 consists of the difference in sample means between the deterministic, equalising and fair strategies with the approximate strategy, as well as the  $p$ -values from Dunnett’s test for each pairwise comparison. By showing both the difference in sample means as well as the  $p$ -values, we can use these figures to determine regions where the means are significantly different (for any choice of significance level) and also determine if the other strategies outperform the approximate strategy. Fig. 9 combines the results of Dunnett’s test at a 5% significance level with the difference in sample means for each pairwise comparison. It indicates where the other strategies are significantly different from the approximate strategy as well as how different the sample means of the strategies are through the intensity of the colour. Fig. 10 combines the results of



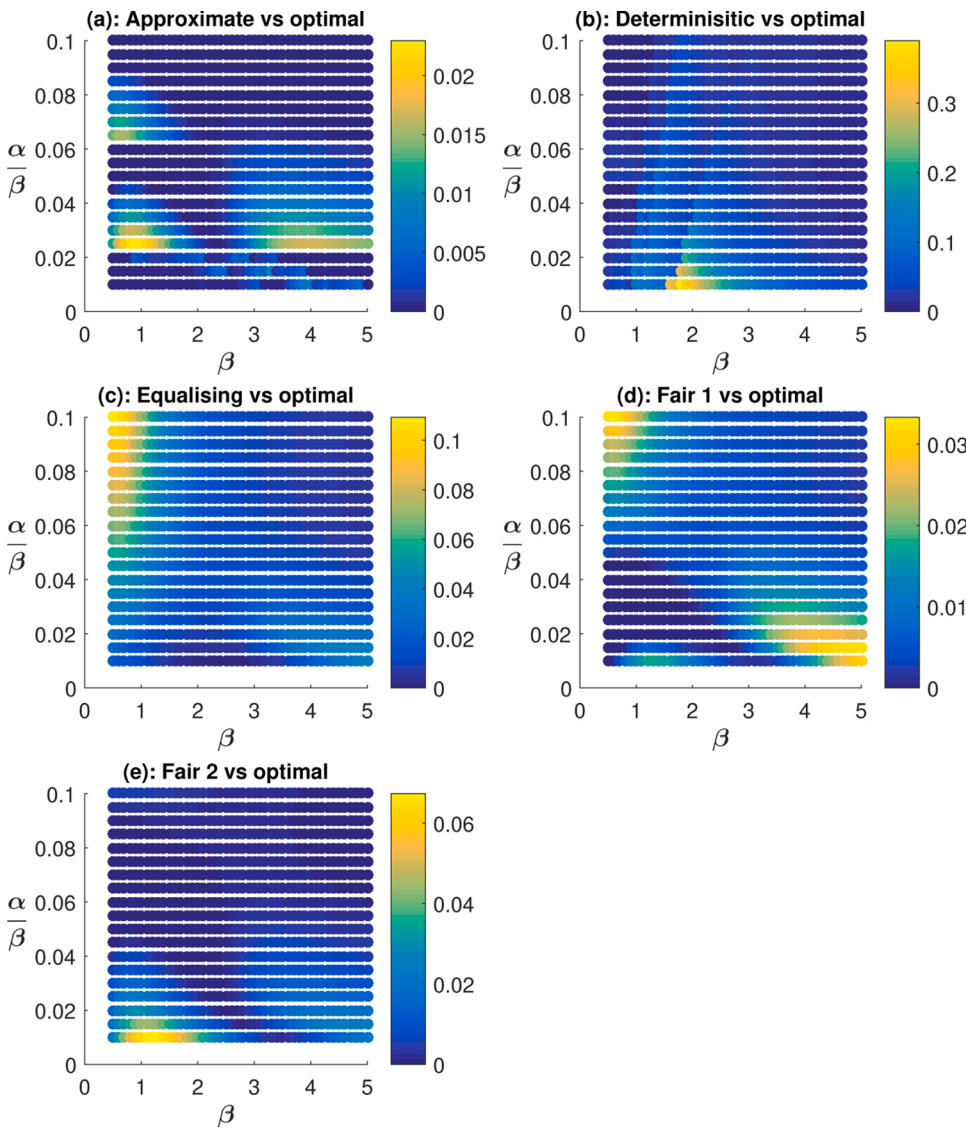


Fig. 5. Relative difference in mean final epidemic size between the different strategies and the optimal for the three-patch example with  $N_1 = 6, N_2 = 12, N_3 = 18, V = 9$  vaccines,  $\gamma = \frac{1}{2}$ , within-patch infection rate  $\beta \in [0.5, 5]$  and between-patch infection rate  $\alpha \in [0.01, 0.1] \times \beta$ . (a): Approximate strategy versus optimal, (b): Deterministic strategy versus optimal, (c): Equalising strategy  $(V_1, V_2, V_3) = (0, 2, 7)$  versus optimal, (d): Fair strategy  $(V_1, V_2, V_3) = (1, 3, 5)$  versus optimal, (e): Fair strategy  $(2, 3, 4)$  versus optimal. Note each subfigure has its own colour scale. Further, note that two fair strategies are possible due to the rounding of proportions/individuals. This shows that the approximate strategy compares well with the optimal strategy and performs better than the other strategies considered.

**Table 1**  
Average and maximum relative difference for each proposed strategy for the three-patch example with  $N_1 = 6, N_2 = 12, N_3 = 18, V = 9$  vaccines and  $\gamma = \frac{1}{2}$ .

Proposed strategy	Region	Relative Difference	
		Average	Maximum
Approximate	Full	0.0027	0.0229
Deterministic	Full	0.0363	0.3895
Equalising	Full	0.0212	0.1093
Fair (1,3,5)	Full	0.0074	0.0333
Fair (2,3,4)	Full	0.0058	0.0673
Approximate	Influenza-like	0.0047	0.0222
Deterministic	Influenza-like	0.0127	0.0509
Equalising	Influenza-like	0.0610	0.1093
Fair (1,3,5)	Influenza-like	0.0098	0.0333
Fair (2,3,4)	Influenza-like	0.0074	0.0509

the individual Dunnett’s test at a 5% significance level into a single figure. From these figures, we observe that the approximate strategy performs well when compared to the equalising and fair strategies.


When comparing with the deterministic strategy, we observe from the figures that the deterministic strategy appears to perform better than

the approximate strategy in the region where  $\frac{\alpha}{\beta} = 0.01$ . However, when we consider Fig. 9a, which combines the results of Dunnett’s test at a 5% significance level with the difference in sample means, we see that when the approximate strategy performs better than the deterministic strategy, the difference in sample means is larger than the difference when the deterministic strategy performs better. The approximate strategy is better, or almost as good, over the parameter space considered; hence, the approximate strategy appears robust to the precise values of  $(\alpha, \beta)$ .


### 5. Discussion

The question of how best to allocate a limited supply of vaccines to a population to minimise the expected number of people that become infected over the course of an epidemic is of great importance to public health officials. We addressed this question in two parts, when to allocate and where to allocate. For both parts, we utilised a stochastic SIR metapopulation model and assumed that the entire population is initially susceptible with infection occurring through a single attempted import into the population.

We first explored the question of when best to allocate vaccines, before infection or after infection is present in the population, and considered three different allocation strategies: strict-prophylactic vaccination, predominantly-prophylactic vaccination and reactive



**Fig. 6.** Large three-patch example with  $N_1 = 300$ ,  $N_2 = 600$ ,  $N_3 = 900$ ,  $V = 450$  vaccines,  $\gamma = \frac{1}{2}$ , within-patch infection rate  $\beta \in [0.5, 5]$  and between-patch infection rate  $\alpha \in [0.01, 0.1] \times \beta$ . (a): Approximate strategy, (b): Deterministic strategy. The colour labels represent the number of vaccines allocated to each patch, ( $V_1$ ,  $V_2$ ,  $V_3$ ). This figure shows that the approximate and deterministic strategies differ for  $\beta > 1.5$ .



**Fig. 7.** Comparison to determine if the means of all strategies are the same using a one-way ANOVA test for the large three-patch example with  $N_1 = 300$ ,  $N_2 = 600$ ,  $N_3 = 900$ ,  $V = 450$  vaccines and  $\gamma = \frac{1}{2}$ . (a): Indicates whether all means are the same or at least one mean is different at a 5% significance level, (b):  $p$ -values for the one-way ANOVA test.

vaccination. Through the use of a modified backward dynamic programming algorithm which included a delay between the first onset of infection and vaccination, we found that in general, strict-prophylactic vaccination resulted in a smaller mean final epidemic size. The strength of this result from our investigation using a simple SIR

metapopulation model suggests that this might hold for any metapopulation which lacks significant structure which might be exploited. Part of the reason for the success of the prophylactic scheme here is its accounting for import of infection in the objective for determining the allocation. Hence, for the next question of where best to allocate a

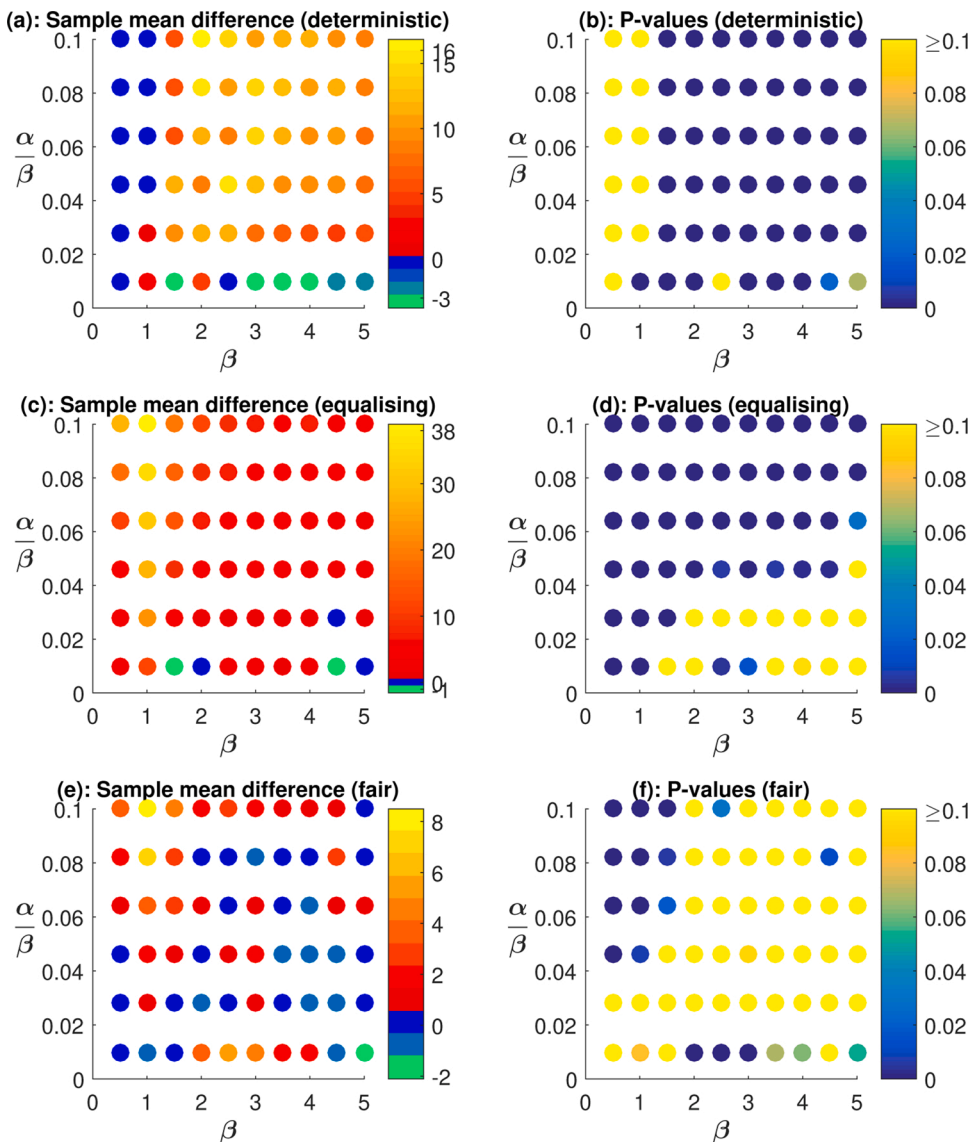


Fig. 8. Comparison to determine whether the other strategies are better than the approximate strategy using Dunnett's test for the large three-patch example with  $N_1 = 300$ ,  $N_2 = 600$ ,  $N_3 = 900$ ,  $V = 450$  vaccines and  $\gamma = \frac{1}{2}$ . (a): Difference in sample means (deterministic – approximate), (b):  $p$ -values for Dunnett's test, (c): Difference in sample means (equalising – approximate), (d):  $p$ -values for Dunnett's test, (e): Difference in sample means (fair – approximate), (f):  $p$ -values for Dunnett's test.

limited supply of vaccines to a metapopulation, we focused on the case of strict-prophylactic vaccination.

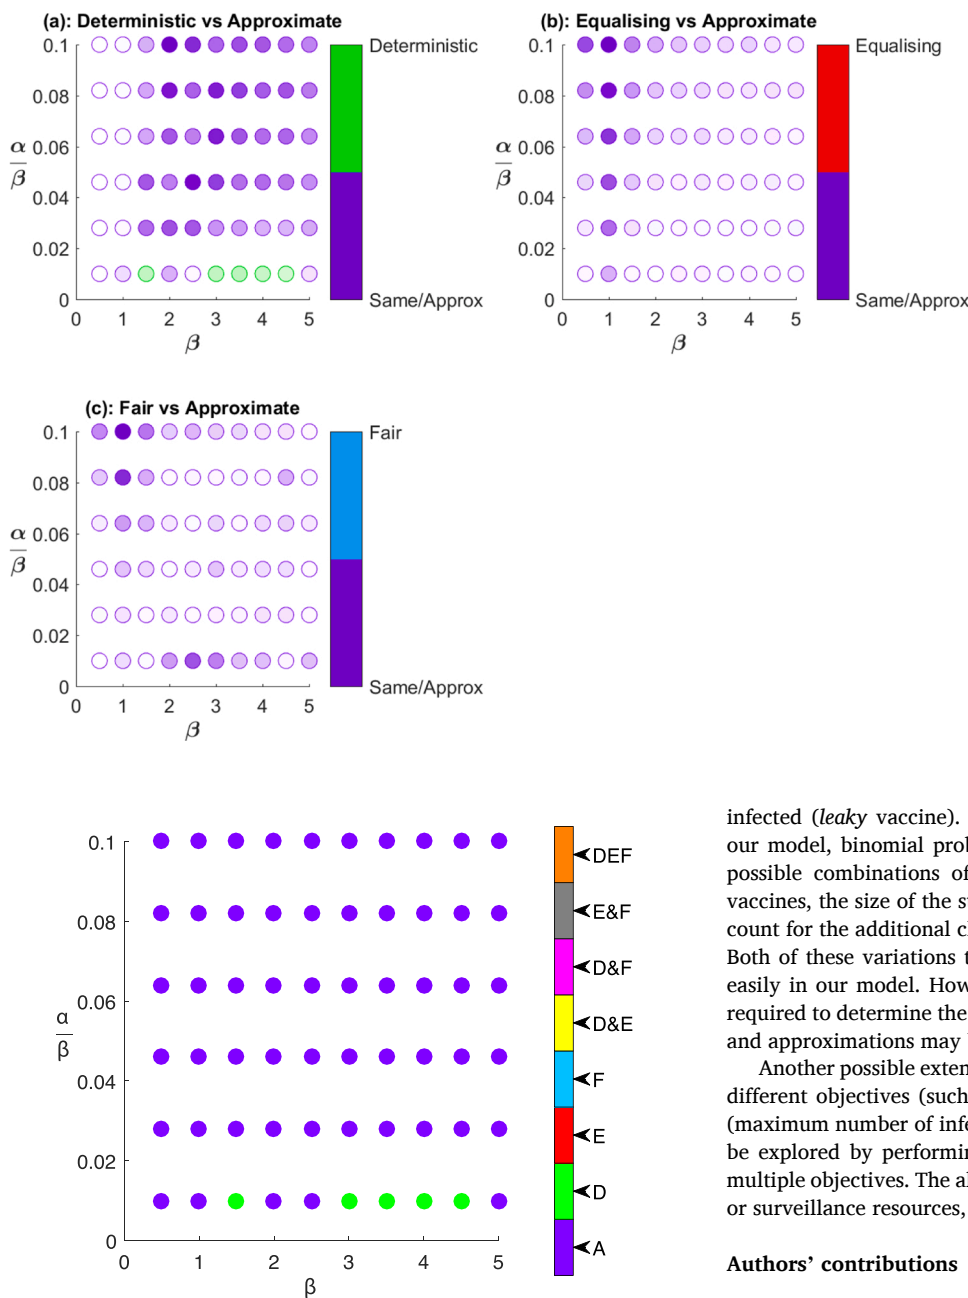
Determining the optimal strategy for a given set of parameters can be a computationally intensive task especially as the population size increases; for large population sizes, this becomes computationally intractable. Hence, we developed a more computationally-efficient method to obtain an approximately optimal strategy. Our approximation method consists of two approximations, the average initial infection rate approximation and the weakly-coupled final sizes approximation, as well as a rule for determining which approximation to use for a given parameter set. We considered two examples, a small population size and a larger population size, and compared the performance of our approximate strategy with other strategies in the literature (deterministic, equalising and fair). In general, we observed that our approximate strategy performed well when compared to the other strategies and also appeared to be more robust than the other strategies for a range of parameters.

In reality, during an outbreak of an infectious disease, it is difficult to know the values of disease parameters such as transmission ( $\beta$ ,  $\alpha$ ) or recovery ( $\gamma$ ) rates. Similarly, the probability of a single attempted import ( $\rho$ ) or the rate of import ( $\epsilon$ ) may also be unknown. However, these parameters might be able to be approximated and/or estimated. For example, the probability of a single attempted import or the rate of

import can be informed by the international arrivals into cities, and the disease parameters can be informed through historical values and updated via first few hundred studies (Black et al., 2017). The uncertainty in the parameters emphasises the importance of robustness of the strategies to a range of parameter values. We observed, in the examples considered, that our approximate strategy returns optimal allocations which performed better than the optimal allocations obtained from other strategies for a range of ( $\alpha$ ,  $\beta$ ) pairs.

In addition to the robustness to parameter values, these examples also provide us with confidence that the approximate strategy performs well across examples of different sizes. Overall, we note that whenever the optimal strategy can be obtained, it should, of course, be used. For patch sizes over a thousand, we are unable to obtain the optimal strategy and so the approximate strategy we developed should be used instead. Then, if we consider even larger population sizes (patch sizes in the millions), there appears to be some benefit in using the deterministic strategy.

This work can be extended in a number of ways. One possible extension is to consider other types of vaccination. Here, we have assumed that vaccination renders an individual immune to the disease. However, this may not be entirely realistic. Some vaccines may not be completely effective and so only a proportion of vaccinated individuals are immune to the disease (*all-or-nothing* vaccination), or perhaps,



**Fig. 9.** Comparison of the approximate strategy with other strategies using Dunnett's test at a 5% significance level for the large three-patch example with  $N_1 = 300$ ,  $N_2 = 600$ ,  $N_3 = 900$ ,  $V = 450$  vaccines and  $\gamma = \frac{1}{2}$ . (a): Indicates the deterministic strategy is better (green) or the approximate strategy is better or there is no difference between strategies (purple), (b): Indicates the equalising strategy is better (red) or the approximate strategy is better or there is no difference between strategies (purple), (c): Indicates the fair strategy is better (blue) or the approximate strategy is better or there is no difference between strategies (purple). The intensity of the colour increases with the difference in sample means.

**Fig. 10.** Combines individual tests between the approximate strategy and other strategies using Dunnett's test at a 5% significance level for the large three-patch example with  $N_1 = 300$ ,  $N_2 = 600$ ,  $N_3 = 900$ ,  $V = 450$  vaccines and  $\gamma = \frac{1}{2}$ . The legend indicates which strategies are better than the approximate strategy, where A, D, E and F represent the approximate strategy, the deterministic strategy, the equalising strategy and the fair strategy, respectively. Note that purple indicates either the approximate strategy performed better or that of the other strategies.

vaccination might only reduce an individual's chance of becoming

### Appendix A

#### A.1 Modified BDP algorithm

The modified BDP algorithm utilises backward dynamic programming (Powell, 2011; Puterman, 1994) to determine an optimal vaccination strategy. Here, we discuss the terms in the backward dynamic programming algorithm and how they relate to our problem. Note that we have chosen

infected (*leaky* vaccine). To incorporate all-or-nothing vaccination in our model, binomial probabilities would need to be included and all possible combinations of each strategy considered. While for leaky vaccines, the size of the state space of our model would increase to account for the additional class of susceptible but vaccinated individuals. Both of these variations to vaccination can be incorporated relatively easily in our model. However, they increase the computational effort required to determine the optimal solution and so, other simplifications and approximations may be needed for useful results to be obtained.

Another possible extension is to obtain strategies which are robust to different objectives (such as the duration of an epidemic or the peak (maximum number of infected individuals) of an epidemic). This could be explored by performing multi-criteria optimisation and weighting multiple objectives. The allocation of other resources, such as anti-virals or surveillance resources, could also be investigated in a similar way.

#### Authors' contributions

**Mingmei Teo:** Conceptualization; Formal analysis; Investigation; Methodology; Software; Validation; Visualization; Writing – original draft, review and editing. **Nigel Bean:** Conceptualization; Methodology; Supervision; Writing – review and editing. **Joshua Ross:** Conceptualization; Methodology; Supervision; Writing – review and editing.

#### Conflict of interest

The authors declare no conflict of interest.

to consider the allocation of vaccines as a finite-horizon problem where after a given time  $T$ , no more vaccines are allocated and the process evolves without any intervention. We let  $F_t(S_t, v)$  represent the minimum mean final epidemic size starting in state  $S_t$  at time  $t$  with optimal actions chosen from time  $t$  until the horizon  $T$  and with  $v$  vaccines available. At the time-horizon  $T$ , we know that  $F_T(S_T, v)$  is the mean final epidemic size for state  $S_T$  as no actions are taken from the horizon  $T$ .

As highlighted in Section 3, we have chosen to simplify the modelling and limit the number of possible actions taken at each time step by only allowing at most one vaccine to be allocated at each time step during the infection. That is, the possible actions at time step  $t$  are:

- $a_t = \emptyset =$  no vaccine allocated,
- $a_t = k =$  a single vaccine allocated to patch  $k$ , for  $k = 1, \dots, m$ .

Note that if a vaccine is not allocated, then it is stockpiled for use at a later time step. Hence, for a given state  $S_b$ , the set of possible actions are given by,  $A(S_t) = \{\emptyset\} \cup \bigcup_{k=1}^m \{k | s_k > 0\}$ , as there needs to be at least one susceptible individual in patch  $k$  for vaccination to occur in that patch. After an action is taken, the process evolves to the next time step according to  $P(s' | S_t, a_t)$ , which is the probability of the process moving to state  $s'$ , given the current state  $S_t$  and action  $a_t$  is chosen. This probability can be obtained from the transition probability matrix,  $P(t') = e^{Q t'}$ , which consists of the probabilities of moving from one state to another in time  $t'$ , for all  $t' \geq 0$ .

Recall from Section 3 that our simplification of only allowing at most one vaccine to be allocated at each time step does not restrict the model to only allocating at most one vaccine each day as a time step,  $t'$ , does not necessarily correspond to a day. If we assume that a day corresponds to a time unit of 1, then smaller values of  $t'$  result in the process evolving over smaller time steps (than a day) and so vaccines are allocated at a faster rate. For our example, we let  $t' = \frac{1}{100}$  which essentially means that vaccines are able to be allocated almost simultaneously. In Algorithm 1,  $P_2 = e^{\frac{Q}{100}}$  is the transition probability matrix governing the evolution of the process after the attempted import of infection and when vaccination is allowed to be applied.

The other term we detail is  $C(S_b, a_t)$ , which is the contribution of choosing action  $a_b$  given the process is in state  $S_t$  at time  $t$ . If no vaccine is allocated, that is, action  $a_t = \emptyset$  is chosen, then  $C(S_b, \emptyset) = 0$ , for all  $t$  and  $S_t \in \mathbb{S}$ , as there is no change in the mean final epidemic size. On the other hand, for all other actions  $a_t = k$  for  $k = 1, \dots, m$ , an individual in patch  $k$  is vaccinated and moved to the recovered class. As the method for obtaining the mean final epidemic size in the backward dynamic programming algorithm determines the number of recovered individuals, we need to remove the vaccinated individual from the mean final epidemic size calculation and so  $C(S_b, a_t) = -1$ , for all  $t, S_t \in \mathbb{S}$  and  $a_t \in A(S_t) \setminus \{\emptyset\}$ .

Next, we describe the post-processing step to account for prophylactic vaccination and the attempted import of infection. We note here that unlike vaccination after infection is present in the population, vaccines allocated prophylactically are distributed simultaneously before the attempted import of infection as opposed to a single vaccine each time step. To allow for prophylactic vaccination, we assume that at the initial time  $t = 1$ , the entire population is susceptible. Hence, the initial state  $S_1 = (N_1, 0, \dots, N_k, 0, \dots, N_m, 0)$ . We also assume that the attempted import of infection occurs at time  $t = 1$  and any number of vaccines, up to the maximum available  $V$ , can be allocated to the population before infection. Then, to determine the optimal allocation of these vaccines before infection, we consider all possible ways of allocating the vaccines to the patches as well as the option of withholding the vaccines to be allocated post-infection. We use  $A_{pre}(S_1, V)$  to denote this set of possible actions for state  $S_1$  and vaccines  $V$ . To better understand  $A_{pre}(S_1, V)$ , consider a simple example with two patches, initial state  $S_1$  and  $V = 2$  vaccines. For this example, the set  $A_{pre}(S_1, V)$  consists of the following actions,

- no vaccines pre-allocated and two vaccines available post infection,
- one vaccine pre-allocated to patch 1 or patch 2 and one vaccine available post infection,
- one vaccine pre-allocated to each patch and no vaccines available post infection, and
- two vaccines pre-allocated to patch 1 or patch 2 and no vaccines available post infection.

Note that this set comprises of all three cases, strict-prophylactic vaccination, predominantly-prophylactic vaccination and reactive vaccination, as well as all possible ways of allocating vaccines within each case. Then, to determine the best vaccination scheme, we compare the mean final epidemic sizes of all possible schemes in the set  $A_{pre}(S_1, V)$ .

The general backward dynamic programming algorithm returns the optimal mean final epidemic size for each state in the state space without accounting for the attempted import of infection or prophylactic vaccination. Let  $F_1(S, v)$  be the minimum mean final epidemic size starting in state  $S \in \mathbb{S}$  at time  $t = 1$  with  $v$  vaccines available, for  $v = 0, \dots, V$ , obtained from the general backward dynamic programming algorithm. Then, let  $\hat{F}_1(u, v)$  be the minimum mean final epidemic size of starting in state  $u = (u_1, 0, \dots, u_k, 0, \dots, u_m, 0)$ , at time  $t = 1$  with  $v$  vaccines, for  $v = 0, \dots, V$ , after accounting for the attempted import of infection in the population and prophylactic vaccination. Note that  $u_k = N_k - V_k$  is the number of unvaccinated susceptible individuals in patch  $k$  and  $V_k$  is the number of vaccines allocated to patch  $k$  with  $V - v = \sum_{k=1}^m V_k$ . Thus, we define,

$$\hat{F}_1(u, v) = \sum_{k=1}^m \rho_k \times \left( \frac{u_k}{N_k} \right) \times (F_1(u + \mathbf{inf}_k, v) - (V - v)),$$

for  $v = 0, \dots, V$ , where  $\mathbf{inf}_k = \mathbf{e}_{2k} - \mathbf{e}_{2k-1}$  and  $\mathbf{e}_k$  is a vector of 0s with a 1 in the  $k$ th position. As the method used to determine the mean final epidemic size for the backward dynamic programming algorithm contains vaccinated individuals, we subtract the number of pre-allocated vaccines  $(V - v)$  from the mean final epidemic size.

We can determine the minimum mean final epidemic size under the optimal vaccination scheme, for the initial state  $S_1 = (N_1, 0, \dots, N_k, 0, \dots, N_m, 0)$

and a given number of available vaccines  $V$ , in the following way,

$$\tilde{F}_1(S_1, V) = \min_{a \in A_{\text{pre}}(S_1, V)} \hat{F}_1(S_1', V),$$

where  $S_1'$  is the state of the process and  $V$  is the number of vaccines remaining after allocation  $a \in A_{\text{pre}}(S_1, V)$  has been applied. The optimal vaccination scheme is the allocation  $a$  that solves the above minimisation problem.

Algorithm 1 details the modified BDP algorithm used to determine the optimal vaccination strategy. Note that in Algorithm 1,  $t'$  represents the chosen time step for the problem and  $P_1 = e^{Q_{t_1}}$  is the transition probability matrix containing the probabilities of moving from one state to another in time  $t_1$ , where  $t_1$  is the delay between the first onset of infection in the population and non-prophylactic vaccination.

**Algorithm 1.** Backward dynamic programming algorithm modified to allow for prophylactic vaccination, attempted import of infection, choice in the rate of vaccination post infection and delay between infection and vaccination.

---

1: Initialise the terminal contribution  $F_T(S_T, v)$  for all states  $S_T \in \mathbb{S}$  and all vaccines  $v = 0, \dots, V$ .

2: Set  $t = T - t'$ .

3: **for**  $v = 0, 1, \dots, V$  **do**

4: Calculate:

$$F_t(S_t, v) = \min_{a' \in A(S_t)} \{C(S_t, a') + \sum_{s' \in \mathbb{S}} P_2(s' | S_t, a') F_{t+1}(s', v - \mathbf{1}_{\{a' \neq \emptyset\}})\},$$

for all  $S_t \in \mathbb{S}$ . Set  $a_t(v) =$  value of  $a'$  that solves the minimisation problem. Note that if different actions result in the value of  $C(S_t, a') + \sum_{s' \in \mathbb{S}} P_2(s' | S_t, a') F_{t+1}(s', v - \mathbf{1}_{\{a' \neq \emptyset\}})$  for  $a' \in A(S_t)$  being very close (within  $10^{-10}$ ), then these values are set to be the same. Further, if there are multiple actions that correspond to the minimal value, then the action corresponding to vaccinating the smallest numbered patch is chosen.

5: **end for**

6: If  $t > 1 + t'$ ,  $t = t - t'$  and return to Step 3. Else stop.

7: Set  $t = 1$ . (Post infection, but pre-vaccination period)

8: **for**  $v = 0, 1, \dots, V$  **do**

9: Calculate

$$F_t(S_t, v) = \sum_{s' \in \mathbb{S}} P_1(s' | S_t) F_{t+1}(s', v),$$

$$\hat{F}_t(S_t, v) = \sum_{k=1}^m \rho_k \times \left(\frac{u_k}{N_k}\right) \times (F_t(S_t + \mathbf{inf}_k, v) - (V - v)),$$

align="center"

for all  $S_t \in \mathbb{S}$  and  $\mathbf{inf}_k = \mathbf{e}_{2k} - \mathbf{e}_{2k-1}$ , where  $\mathbf{e}_k$  is a vector of 0s with a 1 in the  $k$ th position.

10: **end for**

11: For initial state  $S_1 = (N_1, 0, \dots, N_j, 0, \dots, N_m, 0)$  and given number of vaccines  $V$ , calculate

$$\tilde{F}_1(S_1, V) = \min_{a \in A_{\text{pre}}(S_1, V)} \hat{F}_1(S_1', V),$$

where  $S_1'$  is the state of the process and  $V$  is the number of vaccines remaining after pre-allocation  $a \in A_{\text{pre}}(S_1, V)$  has been applied. The allocation  $a$  that solves the minimisation problem is the optimal vaccination scheme.

---

### A.2 Average initial infection rate

Recall, the probability of a randomly chosen individual in patch  $k$  becoming infected is  $\left(\rho_k \times \frac{u_k}{N_k}\right)$ , where  $\rho_k = \frac{e_k}{\sum_{j=1}^m e_j}$ , and if infection is successful, the process moves to the state  $(u_1, 0, \dots, u_k - 1, 1, \dots, u_m, 0)$ . From this state, the process can move to  $m$  possible states of further infection (ignoring the state corresponding to recovery) according to the rates specified in Section 2. Explicitly, the state  $(u_1, 0, \dots, u_k - 1, 1, \dots, u_m, 0)$  can move to the following states with the corresponding rates,

- $(u_1, 0, \dots, u_j, 0, u_k - 2, 2, \dots, u_m, 0)$  at rate

$$\frac{\beta \times (u_k - 1) \times 1}{N_k - 1} + \sum_{\substack{\ell \neq k \\ \ell=1}}^m \left( \frac{\alpha_{\ell k} \times 0 \times (u_k - 1)}{N_k} + \frac{\alpha_{k\ell} \times 0 \times (u_k - 1)}{N_\ell} \right),$$

- $(u_1, 0, \dots, u_j - 1, 1, u_k - 1, 1, \dots, u_m, 0)$  at rate

$$\frac{\beta \times u_j \times 0}{N_j - 1} + \frac{\alpha_{kj} \times 1 \times u_j}{N_j} + \frac{\alpha_{jk} \times 1 \times u_j}{N_k} + \sum_{\substack{\ell \neq j, k \\ \ell=1}}^m \left( \frac{\alpha_{\ell j} \times 0 \times u_j}{N_j} + \frac{\alpha_{j\ell} \times 0 \times u_j}{N_\ell} \right),$$

for  $j = 1, \dots, m$  and  $j \neq k$ .

Hence, the initial infection rate for patch  $k$  is,

$$\frac{\beta \times (u_k - 1) \times 1}{N_k - 1} + \sum_{\substack{j \neq k \\ j=1}}^m \left( \frac{\alpha_{kj} \times 1 \times u_j}{N_j} + \frac{\alpha_{jk} \times 1 \times u_j}{N_k} \right)$$

Then, to obtain the average initial infection rate,  $r$ , for the state of interest  $u = (u_1, 0, \dots, u_k, 0, \dots, u_m, 0)$ , we combine all  $m$  possible cases of initial infection and obtain,

$$r = \sum_{k=1}^m \rho_k \times \left( \frac{u_k}{N_k} \right) \times \left[ \frac{\beta(u_k - 1)}{N_k - 1} + \sum_{\substack{j \neq k \\ j=1}}^m \left( \frac{\alpha_{kj} u_j}{N_j} + \frac{\alpha_{jk} u_j}{N_k} \right) \right]$$

### A.3 Weakly-coupled final sizes approximation

The weakly-coupled final sizes approximation is suitable when the cross-patch infection rate is sufficiently small relative to the within-patch infection rate. Recall from Section 4.1 that the mean final epidemic size for the initial state  $u = (u_1, 0, \dots, u_k, 0, \dots, u_m, 0)$  can be calculated as follows,

$$\zeta^*(u) = \sum_{k=1}^m \rho_k \times \left( \frac{u_k}{N_k} \right) \times \zeta^*(u + \text{inf}_k),$$

where  $\zeta^*(u + \text{inf}_k) = \zeta(u + \text{inf}_k) - V$ . The problem with obtaining the optimal strategy exactly, as identified previously, is that the calculation of the mean final epidemic size  $\zeta^*(u + \text{inf}_k)$  can be computationally intensive, or even intractable, for large population sizes. For this approximation, we develop a less computationally-intensive method to approximate  $\zeta^*(u + \text{inf}_k)$ .

Note that we have mentioned that this approximation is appropriate when the cross-patch infection rate is sufficiently small relative to the within-patch infection rate. Under this condition, the probability of infection spreading between patches is small. Hence, we can approximate the mean final epidemic size for each patch individually, using single population models, and then combine these in an appropriate way to obtain the mean final epidemic size for a given initial state. This reduces the computational complexity of our model, as there is a significant reduction in the state space when using single population models.

Let  $\zeta^*(s_k, i_k)$  denote the mean final epidemic size, excluding vaccinated individuals, for a single population of size  $N_k$  with  $s_k$  susceptible individuals and  $i_k$  infectious individuals. Recall that the entire population is initially susceptible and a randomly chosen individual becomes infected. Assuming that the infected individual belongs to patch  $k$ , the mean final epidemic size for patch  $k$ , obtained from a single population model, is  $\zeta^*(u_k - 1, 1)$ . We then consider each fully susceptible patch  $j$  and add its mean final epidemic size, weighted by the probability of patch  $j$  becoming infected, to the mean final epidemic size for patch  $k$ . Note that we consider all possible ways (that is, different pathways of infection) the infection can spread from the initially infected patch to all fully susceptible patches. As we have assumed that the cross-patch infection rate is sufficiently small relative to the within-patch infection rate, it is unlikely for previously infected patches to be re-infected. Hence, we do not consider re-infection into patches that have already been infected. Further details can be found in (Teo, 2017).

Next, we briefly detail the calculation of the probability of a fully susceptible patch becoming infected. Assume that patch  $k$  is the currently infected patch and patch  $j$  is a fully susceptible patch. Let  $T_I$  denote the set of previously infected patches, excluding patch  $k$ , and let  $T_S$  denote the set of fully susceptible patches. Also, let  $(k \rightarrow j | T_I, T_S)$  denote the infection event from patch  $k$  to patch  $j \in T_S$ , given the previously infected patches in  $T_I$  and the fully susceptible patches in  $T_S$ . Then, we have,

$$P(k \rightarrow j | T_I, T_S) = \left[ 1 - \frac{\gamma}{\gamma + \sum_{\ell \in T_I} \left( \frac{\alpha_{k\ell} u'_\ell}{N_\ell} + \frac{\alpha_{\ell k} u'_\ell}{N_k} \right) + \sum_{j \in T_S} \left( \frac{\alpha_{kj} u_j}{N_j} + \frac{\alpha_{jk} u_j}{N_k} \right)} \right]^{\zeta^*(u_k - 1, 1)} \times \left( \frac{\frac{\alpha_{kj} u_j}{N_j} + \frac{\alpha_{jk} u_j}{N_k}}{\sum_{\ell \in T_I} \left( \frac{\alpha_{k\ell} u'_\ell}{N_\ell} + \frac{\alpha_{\ell k} u'_\ell}{N_k} \right) + \sum_{j \in T_S} \left( \frac{\alpha_{kj} u_j}{N_j} + \frac{\alpha_{jk} u_j}{N_k} \right)} \right),$$

where  $u'_\ell = \lfloor u_\ell - \zeta^*(u_\ell - 1, 1) \rfloor$ , since we require the number of susceptible individuals to be an integer value. This probability is the product of the probability that at least one infectious individual in patch  $k$  infects a susceptible individual in another patch and the conditional probability of one infectious individual in patch  $k$  infecting patch  $j \in T_S$  instead of any other patch, given that cross-infection occurs.

Here, we provide a detailed example of the approximation for  $m=3$  patches. For this case, the quantities of interest to us are  $\zeta^*(u_1 - 1, 1, u_2, 0, u_3, 0)$ ,  $\zeta^*(u_1, 0, u_2 - 1, 1, u_3, 0)$  and  $\zeta^*(u_1, 0, u_2, 0, u_3 - 1, 1)$ . We consider  $\zeta^*(u_1 - 1, 1, u_2, 0, u_3, 0)$  and note that the other two expressions are analogous,

$$\begin{aligned} \zeta^*(u_1 - 1, 1, u_2, 0, u_3, 0) &= \zeta^*(u_1 - 1, 1) \\ &+ P(1 \rightarrow 2 | T_I = \emptyset, T_S = \{2, 3\}) \times [\zeta^*(u_2 - 1, 1) + P(2 \rightarrow 3 | T_I = \{1\}, T_S = \{3\}) \times \zeta^*(u_3 - 1, 1)] \\ &+ P(1 \rightarrow 3 | T_I = \emptyset, T_S = \{2, 3\}) \times [\zeta^*(u_3 - 1, 1) + P(3 \rightarrow 2 | T_I = \{1\}, T_S = \{2\}) \times \zeta^*(u_2 - 1, 1)], \end{aligned}$$

where the probabilities are defined as above for  $P(k \rightarrow j | T_I, T_S)$ .

A.4 Modification to Keeling and Shattock's method

Keeling and Shattock (Keeling and Shattock, 2012) employed a deterministic model with a finite number of patches to determine the optimal allocation of vaccines to a metapopulation with the aim to minimise the final epidemic size. To calculate the final size for patch  $\ell$ , for  $\ell=1, \dots, m$ , they used the following final size equation,

$$Z_\ell^\infty = S_\ell^0 \left[ 1 - \exp \left( - \sum_{j=1}^m \frac{R_0^{j\ell}}{N_\ell} Z_j^\infty \right) \right],$$

where  $Z_\ell^\infty$  is the final size for patch  $\ell$ ,  $S_\ell^0$  is the initial number of susceptible individuals in patch  $\ell$ ,  $N_\ell$  is the size of patch  $\ell$  and  $R_0^{j\ell}$  is the expected number of secondary cases produced in patch  $\ell$  by a single infected individual in patch  $j$ , assuming all individuals in patch  $\ell$  are initially susceptible (the notation is slightly modified from that used in Keeling and Shattock (Keeling and Shattock, 2012)). Then, the authors used simple recursion to attempt to solve the final size equation, given above, in the following way,

$$Z_\ell^{t+1} = S_\ell^0 \left[ 1 - \exp \left( - \sum_{j=1}^m \frac{R_0^{j\ell}}{N_\ell} Z_j^t \right) \right], \text{ for } \ell = 1, \dots, m,$$

with the initial condition,  $Z_\ell^1 = S_\ell^0$ , for  $\ell=1, \dots, m$ .

As we consider the seeding of infection into the population in our model, we adapt their method of calculation of the final epidemic size to suit our problem. To account for the patch in which infection is seeded (patch  $k$ ), we modify the initial condition for the recursive method and also expand the final size equation to replace the first  $S_\ell^0$  term with  $(N_\ell - V_\ell)$ . The initial condition is  $Z_\ell^1(k) = 1$ , if infection is seeded in patch  $k$ , otherwise  $Z_\ell^1(k) = 0$ . In a similar way, we set the second  $S_\ell^0$  term in the expanded final size equation to be  $S_\ell^0(k) = N_k - V_k - 1$ , if infection is seeded in patch  $k$ , otherwise  $S_\ell^0(k) = N_\ell - V_\ell$ . Although there is only a difference of one individual in  $S_\ell^0(k)$  for the infected patch compared with the other patches, for small population sizes, this change can make a substantial difference. Another point of note is that we observed through numerical studies that there was little difference between Keeling and Shattock's method and our revised method for reasonably large values of  $R_0$ ; however, as the value of  $R_0$  decreased, differences arose.

The revised final size equation used to determine the final size for patch  $\ell$ , for  $\ell=1, \dots, m$ , where infection is seeded in patch  $k$ , is,

$$Z_\ell^\infty(k) = (N_\ell - V_\ell) - S_\ell^0(k) \exp \left( - \sum_{j=1}^m \frac{R_0^{j\ell}}{N_\ell} Z_j^\infty(k) \right)$$

Using simple recursion, we can solve the revised final size equation by,

$$Z_\ell^{t+1}(k) = (N_\ell - V_\ell) - S_\ell^0(k) \exp \left( - \sum_{j=1}^m \frac{R_0^{j\ell}}{N_\ell} Z_j^t(k) \right), \text{ for } \ell = 1, \dots, m,$$

with

$$S_\ell^0(k) = \begin{cases} N_k - V_k - 1, & \text{if } \ell = k, \\ N_\ell - V_\ell, & \text{otherwise,} \end{cases}$$

and initial condition,

$$Z_\ell^1(k) = \begin{cases} 1, & \text{if } \ell = k, \\ 0, & \text{otherwise.} \end{cases}$$

Another important point to note is the parameterisation used by Keeling and Shattock in their final size equation. For their model, they chose to use  $R_0^{j\ell}$  whilst our model uses  $\alpha_{j\ell}$ ,  $\beta$  and  $\gamma$ . Hence, we need a mapping between the two parameterisations:

$$\frac{R_0^{j\ell}}{N_\ell} = \begin{cases} \frac{\beta}{\gamma(N_\ell - 1)}, & \ell = j, \\ \frac{\alpha_{j\ell}}{\gamma N_\ell} + \frac{\alpha_{\ell j}}{\gamma N_j}, & \ell \neq j. \end{cases}$$

We have chosen this mapping as it corresponds to the transition rates of our model. Note that  $\ell = j$  represents within-patch infection and we map this case to the within-patch component of the infection rate  $\frac{\beta}{\gamma(N_\ell - 1)}$ . Similarly for  $\ell \neq j$ , we have cross-patch infection and so the cross-patch component of the infection rate is  $\frac{\alpha_{j\ell}}{\gamma N_\ell} + \frac{\alpha_{\ell j}}{\gamma N_j}$ .

As the recursive calculation of the final size for each patch,  $Z_\ell^\infty(k)$ , depends on the knowledge of which patch the infection has been seeded in, this tells us that the final epidemic size is dependent on the initial state. Hence, for a given initial state  $\mathbf{u}(k) = (u_1, 0, \dots, u_k - 1, 1, \dots, u_m, 0)$ , where infection has been seeded in patch  $k$ , the final epidemic size for  $\mathbf{u}(k)$  is,



$$\xi^{(D)}(\mathbf{u}(k)) = \sum_{\ell=1}^m Z_{\ell}^{\infty}(k),$$

where  $Z_{\ell}^{\infty}(k)$  is the final size for patch  $\ell$  obtained through recursion and is dependent on  $\mathbf{u}(k)$  and hence the patch  $k$  where infection is seeded.

#### A.5 Statistical testing procedure

As described in Section 4.4.2, when the optimal strategy and the mean final epidemic size cannot be obtained exactly due to the size of the problem, we are unable to use the relative difference in mean final epidemic size between the optimal strategy and the proposed strategy to determine how well the proposed strategy performs. Instead, to compare between the various proposed strategies (approximate, deterministic, equalising and fair), we simulate the mean final epidemic sizes for each proposed strategy and perform statistical tests to determine if the mean final epidemic size between the strategies are significantly different.

To simulate the mean final epidemic size, we perform many simulations using Sellke's method (Sellke, 1983) to obtain a set of final epidemic sizes for a given  $(\alpha, \beta)$  pair and given strategy. Then, the statistical testing procedure used to compare the performance of our approximate strategy with the other strategies, deterministic, equalising and fair, consists of two statistical tests. The first is the one-way ANOVA test, which tests if the mean final epidemic size for all strategies are the same. The null and alternative hypotheses for this test are,

$$\begin{aligned} H_0 &: \mu_{\text{approx}} = \mu_{\text{deterministic}} = \mu_{\text{equalising}} = \mu_{\text{fair}}, \\ H_a &: \text{not all means are equal,} \end{aligned}$$

respectively, where  $\mu_{\text{strategy}}$  is the mean final epidemic size for a given strategy. Using the `anova` function in R (R Core Team, 2015), we perform the one-way ANOVA test at a 5% significance level. This means that if the  $p$ -value of the test is less than 0.05, then we reject the null hypothesis,  $H_0$ , and conclude that the mean final epidemic size for the strategies are not all equal.

As we are interested in determining if there is a significant difference in the mean final epidemic size between the strategies, if the null hypothesis for the one-way ANOVA test is not rejected, then there is insufficient evidence to conclude that the mean final epidemic sizes between the strategies are significantly different. However, if the null hypothesis is rejected, then not all the mean final epidemic sizes are the same and we are interested in which strategy or strategies are different from the approximate strategy. Hence, we make pairwise comparisons between the approximate strategy and the deterministic, equalising and fair strategies by performing Dunnett's test (Dunnett, 1955, 1964).

Dunnett's test is a multiple comparisons test that tests various treatments with a control. This fits with our problem as the control is the approximate strategy and the other treatments are the deterministic, equalising and fair strategies. Hence, the null and alternative hypotheses are,

$$\begin{aligned} H_0 &: \mu_{\text{approx}} = \mu_{\text{other}}, \\ H_a &: \mu_{\text{approx}} \neq \mu_{\text{other}}, \end{aligned}$$

where  $\mu_{\text{other}}$  is the mean final epidemic size for the other strategy (deterministic, equalising or fair). We use the `glht` function in R (Hothorn et al., 2008) to perform Dunnett's test at a 5% significance level. Hence, if the  $p$ -value returned by the `glht` function is less than 0.05, then we reject the null hypothesis and conclude that there is a difference in the mean final epidemic size of the approximate strategy and the other strategy. Further, we can determine which strategy is better if we consider the difference in the mean final epidemic sizes. If the difference,  $\mu_{\text{other}} - \mu_{\text{approx}}$ , is positive, then the approximate strategy is better than the other strategy. On the other hand, if the difference is negative, then the other strategy is better.

The other important aspect to consider when performing statistical tests is the underlying assumptions made. The assumptions of the one-way ANOVA test are:

- normality of observations in each group,
- equal variance of observations between groups,
- independence of observations within each group, and
- independence of observations between each group.

Dunnett's test has the same assumptions of normality, constant variance and independence.

As each final epidemic size obtained from a single simulation is independent of final epidemic sizes obtained from other simulations, the independence assumption is satisfied. Next, we consider the assumption of constant variance. For this assumption to be reasonable, we require the following condition to be satisfied,

$$\frac{\max SD}{\min SD} < 2,$$

where  $\min SD$  and  $\max SD$  are the minimum and maximum sample standard deviations of the groups, respectively. For all problems considered, this assumption is satisfied except when the within-patch infection rate,  $\beta$ , is 0.5. Further, we also note from (Box, 1954) and (Miller, 1997) that when the size of the groups are equal, moderate differences in the variance of each group does not have a serious impact. Hence, the assumption of constant variance is satisfied for our problem.

Lastly, we consider the normality assumption. For this assumption, we require the set of final epidemic sizes for each strategy of interest to be normally distributed. We know that this is not the case for the final epidemic size as it is either skewed or bi-modal. However, the central limit theorem tells us that for large enough samples, the sample mean is approximately normally distributed. As we perform  $10^6$  simulations to obtain the set of final epidemic sizes, we can apply the central limit theorem to ensure that the sample means are approximately normally distributed (Miller, 1997). Hence, for our problem, all assumptions are satisfied and it is appropriate for us to use the one-way ANOVA and Dunnett's tests.

## References

- Andersson, H., Britton, T., 2000. Stochastic epidemic models and their statistical analysis. Vol. 151 of Lecture Notes in Statistics. Springer New York.
- Azman, A.S., Lessler, J., 2014. Reactive vaccination in the presence of disease hotspots. *Proc. R. Soc. Lond. B: Biol. Sci.* 282 (1798).
- Ball, F.G., Lyne, O.D., 2002. Optimal vaccination policies for stochastic epidemics among a population of households. *Math. Biosci.* 177–178, 333–354.
- Ball, F., Lyne, O., 2006. Optimal vaccination schemes for epidemics among a population of households, with application to variola minor in Brazil. *Stat. Methods Med. Res.* 15 (5), 481–497.
- Ball, F., Sirl, D., 2018. Evaluation of vaccination strategies for SIR epidemics on random networks incorporating household structure. *J. Math. Biol.* 76 (January (1)), 483–530. <https://doi.org/10.1007/s00285-017-1139-0>.
- Ball, F., Mollison, D., Scalia-Tomba, G., 1997. Epidemics with two levels of mixing. *Ann. Appl. Probab.* 7 (1), 46–89.
- Ball, F., Britton, T., Lyne, O., 2004. Stochastic multitype epidemics in a community of households: estimation and form of optimal vaccination schemes. *Math. Biosci.* 191 (1), 19–40.
- Black, A.J., Geard, N., McCaw, J.M., McVernon, J., Ross, J.V., 2017. Characterising pandemic severity and transmissibility from data collected during first few hundred studies. *Epidemics* 19, 61–73.
- Box, G.E.P., 1954. Some theorems on quadratic forms applied in the study of analysis of variance problems, I. Effect of inequality of variance in the one-way classification. *Ann. Math. Stat.* 25 (2), 290–302.
- Duijzer, E., van Jaarsveld, W., Wallinga, J., Dekker, R., 2015. Dose-Optimal Vaccine Allocation Over Multiple Populations. *EI2015-29*. <http://EconPapers.repec.org/RePEc:ems:eureir:79212>.
- Duijzer, E., van Jaarsveld, W., Wallinga, J., Dekker, R., 2016. The most efficient critical vaccination coverage and its equivalence with maximizing the herd effect. *Math. Biosci.* 282, 68–81.
- Dunnnett, C.W., 1955. A multiple comparison procedure for comparing several treatments with a control. *J. Am. Stat. Assoc.* 50 (272), 1096–1121.
- Dunnnett, C.W., 1964. New tables for multiple comparisons with a control. *Biometrics* 20 (3), 482–491.
- Hothorn, T., Bretz, F., Westfall, P., 2008. Simultaneous inference in general parametric models. *Biom. J.* 50 (3), 346–363.
- Keeling, M.J., Ross, J.V., 2015. Optimal prophylactic vaccination in segregated populations: when can we improve on the equalising strategy? *Epidemics* 11, 7–13.
- Keeling, M.J., Shattock, A., 2012. Optimal but unequitable prophylactic distribution of vaccine. *Epidemics* 4 (2), 78–85.
- Medlock, J., Galvani, A.P., 2009. Optimizing influenza vaccine distribution. *Science* 325 (5948), 1705–1708.
- Miller, R.G., 1997. *Beyond ANOVA: Basics of Applied Statistics*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis.
- Pollett, P.K., 2003. Integrals for continuous-time Markov chains. *Math. Biosci.* 182 (2), 213–225. [https://doi.org/10.1016/S0025-5564\(02\)00161-X](https://doi.org/10.1016/S0025-5564(02)00161-X).
- Pollett, P.K., Stefanov, V.T., 2002. Path integrals for continuous-time Markov chains. *J. Appl. Probab.* 39 (4), 901–904. <http://www.jstor.org/stable/3216013>.
- Powell, W.B., 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*. Wiley Series in Probability and Statistics Wiley.
- Puterman, M.L., 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. A Wiley-Interscience publication Wiley.
- R Core Team, 2015. *R: A Language and Environment for Statistical Computing*. Vienna, Austria. <http://www.R-project.org/>.
- Sellke, T., 1983. On the asymptotic distribution of the size of a stochastic epidemic. *J. Appl. Probab.* 20 (2), 390–394.
- Teo, M., 2017. *Optimal Allocation of Vaccines in Metapopulations* [PhD Thesis]. School of Mathematical Sciences, The University of Adelaide. <https://doi.org/10.4225/55/5b2076b298bab>.
- Wallinga, J., van Boven, M., Lipsitch, M., 2010. Optimizing infectious disease interventions during an emerging epidemic. *Proc. Natl. Acad. Sci. USA* 107 (2), 923–928.
- Yuan, E.C., Alderson, D.L., Stromberg, S., Carlson, J.M., 2015. Optimal vaccination in a stochastic epidemic model of two non-interacting populations. *PLOS ONE* 10 (2), 1–25.