Epidemics 11 (2015) 7–13

Contents lists available at ScienceDirect

Epidemics

journal homepage: www.elsevier.com/locate/epidemics

Optimal prophylactic vaccination in segregated populations: When can we improve on the equalising strategy?

Matt J. Keeling[a,b,c,*], J.V. Ross[d]

[a] WIDER Centre, University of Warwick, UK
[b] Mathematics Institute, University of Warwick, UK
[c] School of Life Sciences, University of Warwick, UK
[d] School of Mathematical Sciences, University of Adelaide, Australia

ARTICLE INFO

Article history:
Received 30 September 2014
Received in revised form 13 January 2015
Accepted 14 January 2015
Available online 24 January 2015

Keywords:
Vaccination
Optimal control
Households
Metapopulation

ABSTRACT

One of the fundamental problems in public health is how to allocate a limited set of resources to have the greatest benefit on the health of the population. This often leads to difficult value judgements about budget allocations. However, one scenario that is directly amenable to mathematical analysis is the optimal allocation of a finite stockpile of vaccine when the population is partitioned into many relatively small cliques, often conceptualised as households. For the case of SIR (*susceptible–infectious–recovered*) dynamics, analysis and numerics have supported the conjecture that an equalising strategy (which leaves equal numbers of susceptible individuals in each household) is optimal under certain conditions. However, there exists evidence that some of these conditions may be invalid or unsuitable in many situations. Here we consider how well the equalising strategy performs in a range of other scenarios that deviate from the idealised household model. We find that in general the equalising strategy often performs optimally, even far from the idealised case. However, when considering large subpopulation sizes, frequency-dependent transmission and intermediate levels of vaccination, optimality is often achieved through more heterogeneous vaccination strategies.

Introduction

Mathematical modelling has had a profound influence on public health associated with infectious diseases; most public-health decisions are now supported by detailed mathematical predictions that quantify the incremental costs and benefits of any new policy. This is particularly true for changes to vaccination programs (including the introduction of new vaccines) where there are potentially many subtle non-linearities between the distribution of vaccine and the public-health benefits (Anderson and May, 1983; Bansal et al., 2006; van Hoek et al., 2011). In principle the aim of this modelling for vaccination is relatively simple: to find a strategy that produces the maximum reduction in cases (and in particular severe health outcomes) for a given cost (Woodhall et al., 2009; Baguelin et al., 2010; Klepac et al., 2011; Brown and Jane White, 2011). Yet despite this apparent simplicity, determining the optimal policy is highly computationally intensive due to the vast range of strategies that can be investigated (Dushoff et al., 2007; Hall et al., 2007; Keeling and White, 2011). In addition when there are multiple desirable outcomes, it is generally impossible to optimise all of them simultaneously and a careful definition of the objective is required (Hollingsworth et al., 2011).

The ground-breaking work of Ball et al. (1997), Ball and Lyne (2002) is seen as offering one of the few explicit and rigorous results in this complex field. In this work it was demonstrated that an equalising strategy was optimal for control of an SIR-type infection in a population segregated into households (or component subpopulations) of just 2, 3 and 4 individuals. Further, it was conjectured, supported by extensive numerics, that this result holds for all subpopulation sizes. Here an equalising strategy is one which leaves an equal number of individuals susceptible in each household (or subpopulation) irrespective of the size of the household (e.g. all households of size 3 or more are left with just 3 susceptible individuals). However, the results are more precise and constrained than usually appreciated (Ball et al., 1997; Ball and Lyne, 2002). Firstly, they only strictly apply to density-dependent transmission, which in the household context means that the risk of transmission between any two members of the household is independent of household size. However, data from detailed household

* Corresponding author at: Mathematics Institute, University of Warwick, UK. Tel.: +44 02476 574832; fax: +44 02476 524182.
E-mail address: M.J.Keeling@warwick.ac.uk (M.J. Keeling).

http://dx.doi.org/10.1016/j.epidem.2015.01.002

studies of influenza suggest that the rate of transmission between any two household members decreases monotonically with household size (Cauchemez et al., 2004, 2009; House et al., 2012). Secondly, optimality refers to maximising the reduction in early household-to-household transmission (the between-household reproductive number or $R_\star$) for a given supply of vaccine; equivalently, this allows the calculation of the minimal amount of prophylactic vaccine required to reduce the reproductive number to one, and hence prevent a large-scale outbreak of infection. Therefore the equalising strategy does not necessarily hold for on-going vaccination during an outbreak, nor does it necessarily protect the most people from infection when there is insufficient resources to reach the elimination threshold (Hollingsworth et al., 2011).

Although the equalising strategy and results in this paper are generally expressed in terms of household-based transmission, the implications are more wide ranging. Our findings apply to any population that can be modelled in a metapopulation format: multiple distinct relatively small subpopulations with strong transmission within the subpopulations but weaker transmission between them. The only other condition is that there must be a large number of these subpopulations such that we can take expectations of their behaviour. As such the models formulated here equally apply to human populations aggregated into schools, hospitals and local communities, livestock aggregated into farms, or wildlife that can be spatially aggregated into regions of suitable habitat. Therefore, although for brevity and historical consistency, we refer to households throughout this paper the findings hold for any appropriate grouping or subpopulation.

It is worth stressing that the equalising strategy and any improvements outlined in this paper make the simplifying assumption that all individuals in the population are equal apart from their household composition. In practice, for human populations, both age and underlying health status dominate the consequences of infection and hence the need to protect by vaccination. Therefore it should be stressed that strategies purely based on household size are idealisations and alternative targeting should often take priority. However, there are at least two scenarios with this understanding could be practically useful. Firstly, once the most vulnerable or high risk members of the population are protected, the equalising strategy may provide a means of slowing or containing epidemic spread when the number of vaccine doses are limited. Secondly, for livestock infections individual-level heterogeneity is generally less of a consideration, so targeting based purely on animal numbers may be effective especially when dealing with a costly vaccine.

Here we examine both analytically and numerically the generality of the equalising strategy. We first review the previous work and methodology (Ball et al., 1997; Ball and Lyne, 2002) before considering the relevance of the equalising strategy for populations that do not obey density-dependent transmission, while maintaining the same condition for optimality. We then numerically explore the use of the equalising strategy under alternative optimisation criteria.

The traditional equalising strategy

We first define the stochastic *SIR* model in an infinitely large population of households to set the nomenclature and parameters of the system. Throughout this paper we formulate and simulate models that are Markovian in nature (i.e. transitions occur as stochastic rates that only depend on the current state of the system, such that there is no historical knowledge), whereas the original work on the equalising strategy held for any form of transmission dynamics (Ball et al., 1997; Ball and Lyne, 2002). While this explicit decision about the nature of the system is necessary to produce our quantitative comparison of vaccination priorities, we believe that the qualitative findings will hold more generally. We define the model in terms of the transitions between states and the rates at which these transitions occur. There are three possible transitions within a household of size n:

External Infection

$$(S,I,R) \rightarrow (S-1,I+1,R) \quad \text{Rate} = \alpha_n \bar{I} S$$

Internal Infection

$$(S,I,R) \rightarrow (S-1,I+1,R) \quad \text{Rate} = \beta_n I S$$

Recovery

$$(S,I,R) \rightarrow (S,I-1,R+1) \quad \text{Rate} = \gamma_n I$$

where S, I and R ($S+I+R=n$) refer to the number of susceptible, infectious and recovered/resistant individuals in a household, while α, β and γ capture the rates of external transmission, internal transmission and recovery; $\bar{I}$ is the proportion of the population that is infected, which is calculated as the weighted average over all households. We further define h_n to be the proportion of households containing n individuals. Throughout we make the natural assumptions that α and γ are not dependent on household size, although we retain the dependence in the equations as much as possible. We note here that the action of vaccination is to successfully immunise susceptible individuals, effectively turning them into recovered individuals; hence as we are considering prophylactic vaccination (before an outbreak) we assume the action of vaccination is to begin an epidemic with a mixture of susceptible and recovered individuals in each household.

Here we have allowed the fundamental rates to be functions of the household size, n; this is in contrast to the earlier modelling studies where these rates where assumed independent of the household size (Ball et al., 1997; Becker and Starczak, 1997; Ball and Lyne, 2002). We now need to introduce some epidemiological notation; when $\beta_n=\beta$ then transmission increases with the number of individuals in the household and such transmission is known as density dependent (even though it arises when the parameter is independent of population size), in contrast when β_n is a function of the household size n the transmission is referred to as frequency dependent. The independence of parameters from household size (and hence the assumption of density-dependent transmission) made in previous work has an important epidemiological consequence: the rate of transmission between susceptible and infected individuals does not depend on the number of recovered individuals in the household. Therefore rather than considering the behaviour of a household of type (S, I with $R=n-S-I$), in the limited case where independence is assumed, it is sufficient to consider a smaller household of size $S+I$ without any recovered individuals.

Following the work of Ball et al. (1997), Ball and Lyne (2002), the following statements must hold for the equalising strategy to be optimal:

(i) given two households of size n and two doses of vaccine, it is better to vaccinate one individual in each household rather than two individuals in a single household, for all n; and,
(ii) given a household of size n, a household of size $n+1$, and single dose of vaccine, it is better to vaccinate an individual in the household of size $n+1$, for all n.

Here, 'better' and 'optimal' are defined with respect to minimising the expected number of secondary households infected. To generate the mathematics equivalent to these two verbal conditions, requires us to consider the household reproductive number,

$R_\star$ – defined as the expected number of secondary households infected by an average infected household early in the outbreak (Ball et al., 1997; Ross et al., 2010). The fact that we are considering early outbreak dynamics means that it is reasonable to assume that each household has experienced at most one external transmission event as $\bar{I}$ is small. The household reproductive number ($R_\star$) is a sum over all household sizes:

$$R_\star = \sum_n \left(\frac{\alpha_n n h_n}{\sum_k \alpha_k k h_k} \right) \mu_n. \quad (1)$$

The first term in this summation is the probability that an externally infected individual resides in a household of size n (where h_k is the distribution of household sizes). The second term (μ_n) is the average number of external secondary cases generated by a newly infected household of size n, and is proportional to the expected number of cases within the household Z_n (House et al., 2012):

$$\mu_n = \frac{\sum_k \alpha_k k h_k}{\sum_k k h_k} Z_n. \quad (2)$$

For notational convenience, we also define $R_\star = \sum_n R^n_\star$, where:

$$R^n_\star = \left(\frac{\alpha_n n h_n}{\sum_k \alpha_k k h_k} \right) \mu_n$$

is the size-biased contribution to $R_\star$ of a household of size n.

Using the implications that parameters are independent of household size, and the definition of $R_\star$ above, it can be shown that both of the verbal conditions (i) and (ii) are satisfied if:

$$D_n \equiv n\mu_n - 2(n-1)\mu_{n-1} + (n-2)\mu_{n-2} \geq 0 \qquad \forall n.$$

That is, if $n\mu_n$ is an increasing convex function of n. When β is independent of n, this has been proven for small subpopulations and conjectured to hold for all subpopulation sizes (Ball et al., 1997; Ball and Lyne, 2002). Therefore, in the case where all epidemiological parameters (α, β and γ) are independent of n, the equalising strategy is conjectured to hold; in which case, equalising susceptibility is optimal for reducing $R_\star$ and therefore also optimal for determining the lowest number of vaccine doses required to achieve elimination or prevent invasion, which both occur at $R_\star = 1$.

The equalising strategy and alternative forms of within-household transmission

We now allow β_n to depend on the household size, and in particular consider transmission of the form $\beta_n \propto (n-1)^{-q}$, so that pure frequency-dependent transmission $q = 1$ as well as alternative scalings with household size can be included. (Note that $q = 0$ refers to the case of density-dependent transmission discussed in Section "The traditional equalising strategy"). We retain the assumption that the vaccine confers perfect immunity; the impact of imperfect vaccines have already been shown to break the equalising strategy (Ball et al., 2004). In this case the two verbal conditions necessary for the equalising strategy to hold require slight modification leading to a change in the associated mathematical conditions. In contrast to the model with density-dependent transmission outlined above, when frequency-dependent transmission (or any transmission with $q > 0$) is assumed then recovered (or vaccinated) individuals play a key role in blocking the within-household transmission; as β_n reduces with household size, therefore recovered individuals increase household size without playing an active role in transmission. It is therefore no longer true that a household with recovered individuals can be treated as a smaller household without recovered individuals.

We now define two new terms $\mu_{n,s}$ and $R^{n,s}_\star$ which correspond to the definitions above, but assume that only s out of n individuals in the household are susceptible. This is to account for vaccination shifting susceptibles to the immune class. Therefore, $\mu_{n,s}$ is the expected number of external secondary cases (outside the household) generated by a newly infected household of size n when s individuals are initially susceptible, assuming all individuals outside the household are susceptible. $R^{n,s}_\star$ is the size-biased contribution to $R_\star$ from a household of size n (with s susceptibles); and is comparable to our earlier definition but accounting for both household size and the number of susceptibles. Again, we make the simplifying assumption throughout our numerical investigations that transmission outside the household is independent of household size (that is $\alpha_n = \alpha$); however we retain the notation in the initial formulae to indicate how this heterogeneity would enter the conditions. The two necessary conditions for the equalising strategy to be optimal for all levels of vaccination now become:

(i) given two households of sizes n and m with S susceptibles in each and two doses of vaccine, it is better to vaccinate one individual in each household rather than two individuals in a single household, $\forall(S, n, m)$. Mathematically this can be expressed as:

$$R^{n,S-1}_\star + R^{m,S-1}_\star = \left(\frac{\alpha_n n h_n}{A}\right)\left(\frac{S-1}{n}\right)\mu_{n,S-1} + \left(\frac{\alpha_m m h_m}{A}\right)\left(\frac{S-1}{m}\right)\mu_{m,S-1}$$
$$\leq \left(\frac{\alpha_n n h_{n-1}}{A}\right)\left(\frac{S}{n}\right)\mu_{n,S} + \left(\frac{\alpha_m m h_m}{A}\right)\left(\frac{S-2}{m}\right)\mu_{m,S-2} = R^{n,S}_\star + R^{m,S-2}_\star$$

where $A = \sum_k \alpha_k k h_k$. Hence we require $E_{S,n,m} \geq 0$, where

$$E_{S,n,m} = (S)\mu_{n,S} + (S-2)\mu_{m,S-2} - (S-1)\mu_{n,S-1} - (S-1)\mu_{m,S-1}$$

(ii) given two households of sizes n and m with S susceptibles in one and $S+1$ susceptibles in the other, and a single dose of vaccine, it is better to vaccinate an individual in the household with $S+1$ susceptibles $\forall(S, n, m)$. Writing this mathematically:

$$R^{n,S}_\star + R^{m,S}_\star = \left(\frac{\alpha_n n h_n}{A}\right)\left(\frac{S}{n}\right)\mu_{n,S} + \left(\frac{\alpha_m m h_m}{A}\right)\left(\frac{S}{m}\right)\mu_{m,S}$$
$$\leq \left(\frac{\alpha_n n h_n}{A}\right)\left(\frac{S-1}{n}\right)\mu_{n,S-1} + \left(\frac{\alpha_m m h_m}{A}\right)\left(\frac{S+1}{m}\right)\mu_{m,S+1}$$
$$= R^{n,S-1}_\star + R^{m,S+1}_\star$$

Hence our second condition becomes $F_{S,n,m} \geq 0$, where

$$F_{S,n,m} = (S-1)\mu_{n,S-1} + (S+1)\mu_{m,S+1} - (S)\mu_{n,S} - (S)\mu_{m,S}$$

We now consider a simple limit (one household becoming large) to demonstrate a case where the equalising strategy must fail. It is intuitive that whenever we move away from the linear density-dependent transmission ($q = 0$) and set $q > 0$, then for finite S but household size m tending to infinity the expected number of cases in the household tends to one (that is no within-household transmission occurs), and hence $\mu_{m,S}$ tends to a constant $\mu^* = \mu_{1,1}$. Considering the case where $n = S > 2$ and $m \to \infty$ and examining conditions (i) and (ii) it is clear that $E_{n,n,m} = -F_{n,n,m}$ and hence in this limit it is impossible to satisfy both criteria and the equalising strategy fails.

Given that we know the equalising strategy fails when households are allowed to become infinitely big (for $q > 0$), it is natural to assess the parameter range for which the strategy holds when there is an upper bound on the household size. Given that solution of the two conditions is not analytically tractable, we rely on fast numerics to evaluate E and F across a range of parameters (Ross et al., 2010; House et al., 2012). Numerical investigation shows that it is

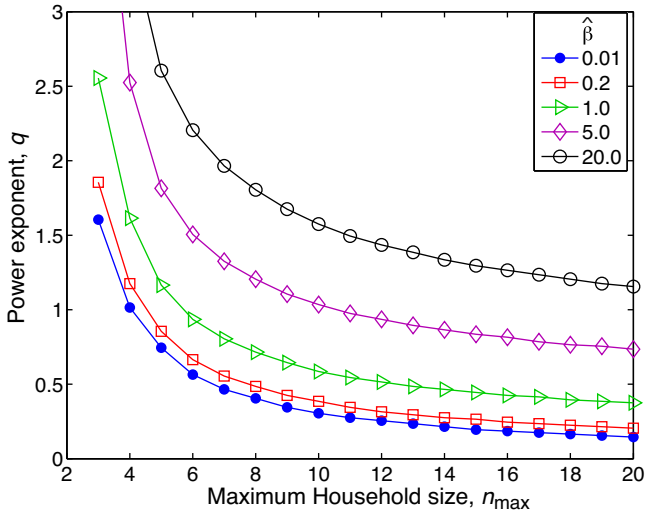

Fig. 1. Contours separating regions of parameter space in which the equalising strategy is always optimal (below the curve) from where it is not (above the curve). We assume throughout that the within-household transmission rate is given by $\beta_n = \widehat{\beta}(n-1)^{-q}$ and vary the maximal household size $n_{\max}$, the value of the exponent q and the transmission scalar $\widehat{\beta}$. Curves correspond to particular values of $\widehat{\beta}$. Note that under the assumption that all individuals experience equal force of infection from outside the household ($\alpha_n = \alpha$) these results are independent of the value of α; although α should be sufficiently large that an outbreak occurs in an unvaccinated population (Throughout we set $\gamma_n = 1$, and assume Markovian dynamics for computational simplicity).

sufficient to consider $E_{S,S,n_{\max}}$ and $F_{S,n_{\max},S}$ for all $S \le n_{\max}$, where $n_{\max}$ is the largest household size.

In Fig. 1, we assume that $\beta_n = \widehat{\beta}(n-1)^{-q}$ (in addition to $\alpha_n = \alpha$ and $\gamma_n = 1$), and investigate when both E and F are positive (and hence the equalising strategy always holds) for a range of $\widehat{\beta}$ and q values and a range of maximal household sizes, $n_{\max}$. We note that our findings are independent of the value of α and γ, although we require α sufficiently large and γ sufficiently small to allow transmission to be sustained before vaccination. We consistently find that the equalising strategy holds for smaller household sizes ($n_{\max}$), larger within-household transmission rates ($\widehat{\beta}$) and smaller exponents (q). As an example, if we assume pure frequency-dependent transmission ($q = 1$), and state that the largest household size is $n_{\max} = 10$, then the equalising strategy holds and is optimal only if $\widehat{\beta} > 4.51\gamma$; however when the maximum household size is assumed to be $n_{\max} = 6$, then the strategy is necessarily optimal for all $\widehat{\beta} > 1.29\gamma$. (Note, $\widehat{\beta} = 1.29\gamma$ corresponds to the situation where the probability of transmission between an infected-susceptible pair of individuals within a household are approximately 56%, 39%, 30%, 24% and 21% in households of size 2–6 respectively, and is comparable with parameters estimated for influenza House et al., 2012)

From the general behaviour of Fig. 1, we conclude that for any two of the three parameters ($n_{\max} > 2$, $\widehat{\beta} > 0$, $q > 0$) a value of the remaining parameter can be found for which the equalising strategy does not necessarily hold. However, for relatively small households, there are substantive regions of epidemiologically realistic parameters for which the equalising strategy is optimal in reducing between household transmission.

The optimal strategy for minimising $R_\star$

The equalising strategy seeks to minimise the household reproductive number $R_\star$ for a given level of immunisation. We note that $R_\star$ depends linearly on the number of households in a given state and hence the question of optimal vaccination (or minimising $R_\star$ for a fixed level of vaccination) becomes a linear programming problem, which can be solved with considerable efficiency (Becker and Starczak, 1997; Ball and Lyne, 2002). Assuming an infinite number of households, this problem is expressed as:

$$\begin{aligned} \underset{h_{n,s}}{\text{minimise}} \quad & R_\star = \sum_n \sum_s \left(\frac{\alpha_n n h_n}{\sum_k \alpha_k k h_k} \right) \left(\frac{s}{n} \right) \mu_{n,s} \quad \left(\frac{h_{n,s}}{h_n} \right) \\ \text{such that} \quad & \sum_s h_{n,s} = h_n, \forall\, n \\ & \sum_n \sum_s \left(\frac{n-s}{n} \right) h_{n,s} \le \nu \\ & h_{n,s} \ge 0, \ \forall (n,s) \end{aligned} \tag{3}$$

where the decision variables are $h_{n,s}$ (the number of households of size n vaccinated such that s individuals remain susceptible); and ν is the proportion of the population for which vaccine is available, that is the *vaccination coverage*. We solve this linear program using `linprog` in MATLAB®, and assess the performance relative to the equalising strategy.

Fig. 2 shows a detailed comparison between the true optimal vaccination strategy, the equalising strategy and random vaccination, in terms of their effect on the household reproductive number, $R_\star$. Two particular illustrative examples are chosen: a relatively rapidly transmitted infection in human households (using household size distributions from 2001 U.K. Census Office for National Statistics, 2001) and a livestock-based infection transmitted within and between cattle farms (using cattle numbers from the 2007 U.S.D.A. agricultural census United States Department of Agriculture, 2007). In both cases pure frequency-dependent transmission ($q = 1$) is assumed to hold, and parameters correspond to a region where the equalising strategy does not hold for all levels of vaccination. It is clear that in both cases the equalising strategy produces a substantial drop in $R_\star$ in comparison to random vaccination. (We note that if q is substantially greater than one, then scenarios exist where random vaccination produces a lower $R_\star$ than the equalising strategy; however such values of q are unlikely to be epidemiologically realistic.) For the small household sizes associated with human infections, the true optimal distribution of vaccine is only a marginal improvement over the equalising strategy, reducing $R_\star$ by less than one percent. However, for the larger aggregated subpopulations associated with livestock farming, the percentage improvement can be substantive suggesting that it would be advantageous to consider optimising prophylactic vaccination strategies for livestock outbreaks.

The equalising strategy and other measures of optimality

The equalising strategy is defined as optimal in terms of minimising the household reproductive number, $R_\star$, and hence is the strategy that prevents a major epidemic ($R_\star < 1$) for the least amount of vaccine. Therefore in terms of purchasing sufficient vaccine to prevent a major future outbreak, the equalising strategy (when it holds) provides an optimal distribution of vaccine.

However, in many scenarios the converse problem holds, there is an existing stockpiled supply of vaccine, which must be used optimally to protect the population. Here the challenge is to determine a strategy (distribution of vaccine between households) that minimises the total number of cases (known as the final epidemic size, R_∞) for a fixed number of vaccine doses. Unfortunately, no simple set of strategic rules appears to be optimal in this context, and we need to exhaustively consider the outcomes of the many potential distributions of vaccine numerically. Therefore, for ease of numerical tractability, in this study we restrict our focus to homogeneous populations where all individuals exist in households of

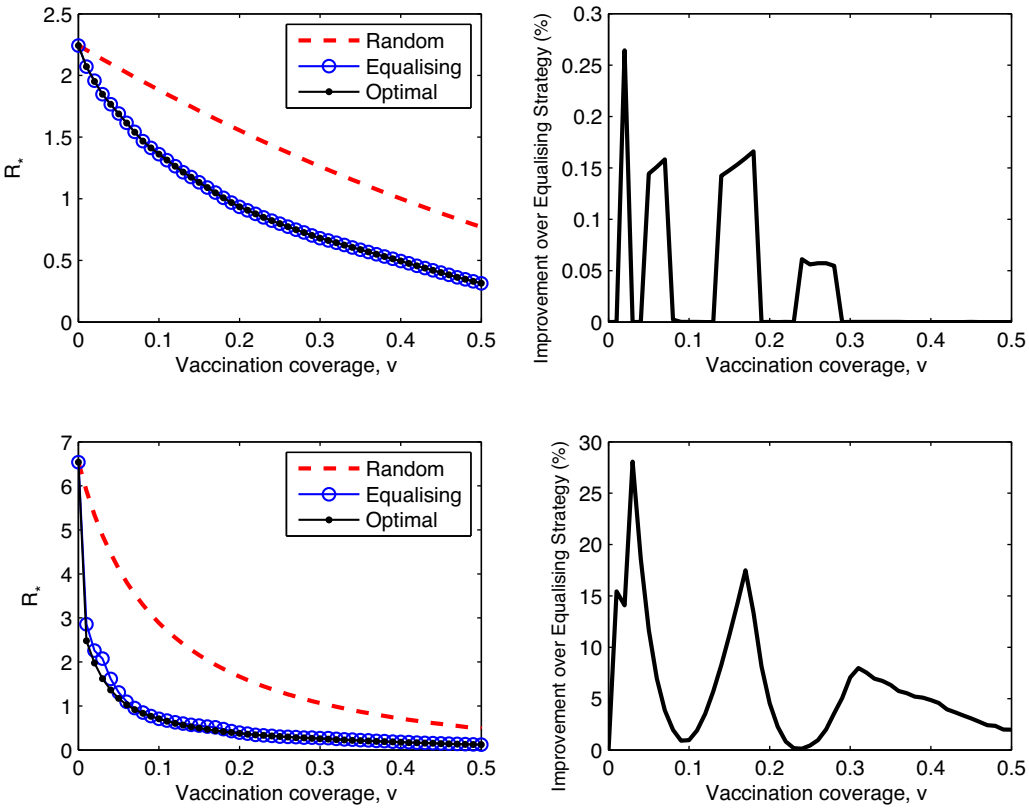

Fig. 2. Comparison of random vaccination (red dashed), equalising strategy (blue circles) and optimal vaccination (black dots) for minimising $R_\star$, as determined by linear programming (3); the dynamics are assumed Markovian for computational simplicity. The left-hand graphs give the value of the household reproductive number, $R_\star$, for the three strategies, while the right-hand graphs show the relative improvement of the true optimal distribution of vaccination (*opt*) compared to the equalising strategy (*ES*): $(R_\star^{ES} - R_\star^{opt})/R_\star^{ES}$. The top row of graphs correspond to influenza-like parameters $\beta = 2(n-1)^{-1}$, $\alpha = 1$, $\gamma = 1$ and household size distributions for the U.K. in 2001 (Office for National Statistics, 2001). The lower row of graphs use lower transmission parameters $\beta = (n-1)^{-1}$, $\alpha = 0.5$, $\gamma = 1$ and distributions for the number of cattle on livestock farms in the USA (United States Department of Agriculture, 2007). In both cases pure frequency-dependent transmission ($q = 1$) is assumed within the household or farm. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

a fixed size, and vary four key quantities: the within-household transmission parameter, $\widehat{\beta}$ (the type of transmission, frequency or density dependent is irrelevant here as we are considering a single household size); the household reproductive number ($R_\star$); the supply of vaccine; and the number of individuals in a household, n. Fig. 3 shows the maximal reduction in the final epidemic size (R_∞) for any distribution of vaccine compared to the equalising strategy; the final epidemic size is calculated using machine-precision Markov Chain methods (House et al., 2012).

From this analysis we find that parameter space naturally splits into three regimes dependent upon the amount of available vaccine, v. When there are sufficient doses of prophylactic vaccine available, the infection cannot generate a large-scale epidemic and multiple distributions of vaccine will all lead to negligible final epidemic sizes (white regions in Fig. 3). For levels of vaccine that are just below the critical amount needed to prevent an epidemic (coloured grey) the equalising strategy is optimal, for the simple reason that it minimises the household reproductive number to a value just above one and hence generates a relatively small epidemic size. These regions of parameter space in which the equalising strategy is optimal are greatest: for lower values of the household reproductive number ($R_\star$), for lower within-household transmission rates ($\widehat{\beta}$) and for smaller household sizes (n). For other regions (coloured blue to red in Fig. 3) the optimal prophylactic vaccine distribution that minimises the final epidemic size is heterogeneous – protecting some households while leaving others unvaccinated or poorly protected, even though all households are identical. More precisely, in this case of a single fixed household size, in extensive numerical searches all optimal strategies have been found to take a particular form: a proportion p of households have ν_1 individuals vaccinated while the remaining $1-p$ of households have ν_2 individuals vaccinated, where p, ν_1 and ν_2 are parameters that need to be determined, although as expected for optimality all the vaccine doses will be used generating a relationship between the three parameters ($v = p\nu_1 + (1-p)\nu_2$). These findings compare well with deterministic two-population models (Keeling and Shattock, 2012) and again suggest that when the aggregated subpopulations (households) are sufficiently large there is good reason to consider optimal deployment of vaccination beyond the equalising strategy.

Using the UK household sizes as given in the 2011 national census when 80% of households are occupied by three people or less, limits the strategy space and allows a numerical investigation. For plausible epidemiological rates ($\beta_n = \widehat{\beta}(n-1)^{-q}$, $\widehat{\beta} \le 5$, $q \le 1, R_\star \le 5$), our numerical investigations suggest that many future human pandemics are likely to lie within the regions of parameter space where the equalising strategy is optimal. That is, for any available level of vaccine coverage, the equalising strategy will both minimise the early between-household transmission and minimise the total number of future cases. However, for larger aggregations of susceptible hosts, such as farms, schools, or large social cliques we expect there to be substantial savings in optimally targeting vaccination.

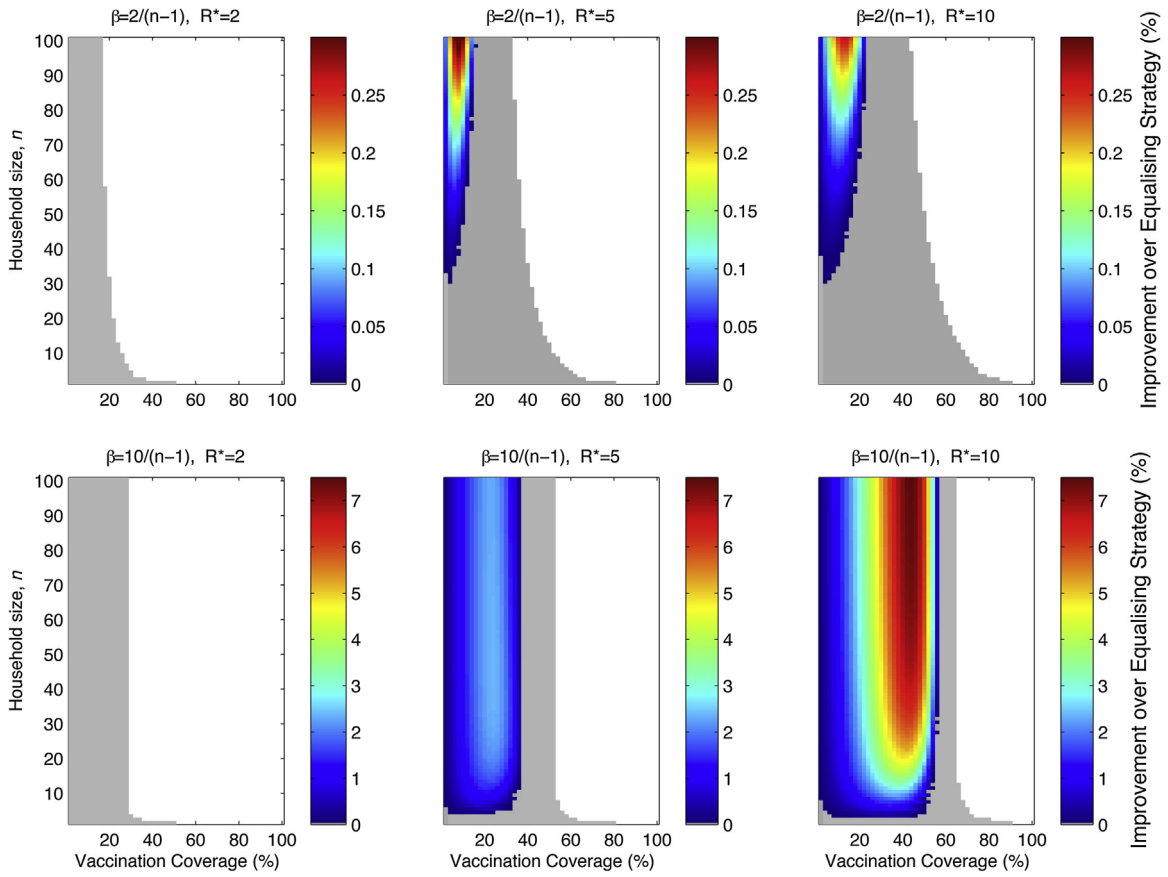

Fig. 3. A large sweep over parameter space showing the maximal improvement in the final epidemic size (R_∞) that can be gained over the equalising strategy by choosing an optimal prophylactic distribution of vaccine. The x-axis gives the available stockpile of vaccine in terms of the proportion of the population that can be vaccinated at the start of the outbreak, the y-axis is the size of household, while the different panels refer to different within and between household transmission. Regions coloured white correspond to sufficient vaccine to eliminate the infection. Grey regions are where the equalising strategy is optimal in minimising the final epidemic size. For regions coloured blue to red we show the reduction in final epidemic size relative to the final epidemic size for the equalising strategy: $(R_\infty^{ES} - R_\infty^{opt})/R_\infty^{ES}$. Panels from left to right have increasing household reproduction number ($R_\star = 2, 5, 10$) in the absence of vaccination; panels from top to bottom show increasing within-household transmission (Throughout we set $\gamma = 1$, and assume Markovian dynamics for computational simplicity; note that the colour-scale changes between top and bottom sets of panels). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Conclusion

The equalising strategy was proposed by Ball et al. (1997) and was demonstrated to minimise the between-household reproductive number when the within-household transmission is density dependent (i.e. transmission scales linearly with the number of susceptible and infectious individual in the household, $q = 0$). Here we have extended this concept and shown that while the equalising strategy holds in many epidemiologically important scenarios, it does not necessarily hold true in all cases of interest. For populations with a heterogeneous mix of household sizes and for a range of frequency-dependent transmission relations ($q > 0$), we were able to examine whether the equalising strategy was optimal in terms of reducing the between household reproductive number (as in Ball et al., 1997). Large maximal household sizes or small within-household transmission rates combined with frequency-dependent transmission break the optimality of the equalising strategy (Fig. 1). When our optimisation criterion is to minimise the final outbreak size following vaccination we were only able to consider populations with a single fixed household size; here the equalising strategy proves to be sub-optimal for large household sizes and high within and between household transmission. In all cases where the equalising strategy fails, the optimal strategy is necessarily more heterogeneous with a proportion of households well protected while the remainder receiving less vaccine, despite the fact that all households are identical. Intuitively this arises from the non-linear impact of vaccination, such that the gains from additional protection in some households is greater than the losses for leaving others unprotected. What these results show is that in many situations careful targeting of vaccination could result in substantial increases in efficacy of control. However, both theoretical and practical questions remain.

When the equalising strategy fails there is no simple rule for determining the optimal distribution of vaccine. If we are interested in minimising the between-household reproductive number then highly efficient linear programming can be utilised to find the true optimum (Fig. 2), and in certain circumstances the form of the optimal policy may be determined explicitly (Ball and Lyne, 2002). However, if we wish to minimise a more complex quantity, such as the final epidemic size or early growth rate of infection, then the numerics become far more computationally demanding. When minimising final epidemic size for households of a fixed size (Fig. 3), determining the optimal strategy is matter of scanning over a two-dimensional parameter space, which is numerically intensive but remains computationally feasible. However, in populations of mixed household sizes the task of finding the true optimal distribution becomes far more challenging due to the increased number of plausible strategies that need to be explored. Using the U.K. household-size distribution, from the 2011 census, we found that for any available level of vaccine coverage, the equalising strategy will both minimize the early growth rate and minimize the final epidemic size. However, for populations with larger

aggregations of susceptible hosts, for example when considering farms and schools, considerable benefit may come from determining the optimal allocation strategy.

Finally, and of more applied importance, is the key question of how such results could be used in practice. Obtaining the precise degree of targeted vaccination, within and between households, as required to reach optimality is clearly unworkable; although proxy measures, such as vaccinating children, may provide close approximations (House and Keeling, 2009). In addition, for human infections considerations such as age-structured mixing and risk factors for severe illness are also major contributing factors (Keeling and White, 2011), such that the aim is often to maximise vaccination of such high-risk groups. However, the findings of this paper may be of considerable benefit once these high-risk individuals are protected, suggesting optimal patterns of vaccination to slow the spread of infection until sufficient vaccination is produced to fully control an outbreak. Alternatively, these optimality results may have greater applicability when attempting to control outbreaks of livestock infections. With a limited stock-pile of vaccine – such as those held by the National Veterinary Stockpile in the USA or by International Governmental Vaccine Banks – the targeting of vaccine during the early stages of an outbreak may be highly advantageous in terms of maintaining control. In particular this work suggests that careful allocation of vaccine stocks for prophylaxis based on animal numbers within a region could provide significant benefits to the livestock industry.

Acknowledgments

This research was funded by The Royal Society, through the International Exchanges Scheme, and supported under the Australian Research Council's Discovery Project and Future Fellowship funding schemes (DP110102893 and FT130100254; JVR) and by the Department of Health (MJK). We thank Thomas House, Lorenzo Pellis and Andrew Black for their thoughts and comments on versions of this manuscript.

References

Anderson, R.M., May, R.M., 1983 May. Vaccination against rubella and measles: quantitative investigations of different policies. J. Hyg. 90 (02), 259–325.

Baguelin, Marc, Van Hoek, Albert Jan, Jit, Mark, Flasche, Stefan, Peter, J., White, W., Edmunds, John, 2010 March. Vaccination against pandemic influenza A/H1N1v in England: a real-time economic evaluation. Vaccine 28 (12), 2370–2384.

Ball, F., Lyne, O., 2002. Optimal vaccination policies for stochastic epidemics among a population of households. Math. Biosci. 177, 333–354.

Ball, F., Mollison, D., Scalia-Tomba, G., 1997. Epidemics with two levels of mixing. Ann. Appl. Probab. 7, 46–89.

Ball, Frank, Britton, Tom, Lyne, Owen, 2004. Stochastic multitype epidemics in a community of households: estimation and form of optimal vaccination schemes. Math. Biosci. 191 (1), 19–40.

Bansal, Shweta, Pourbohloul, Babak, Meyers, Lauren Ancel, 2006. A comparative analysis of influenza vaccination programs. PLoS Med. 3, 1816–1825.

Becker, N.G., Starczak, D.N., 1997. Optimal vaccination strategies for a community of households. Math. Biosci. 139, 117–132.

Brown, V.L., Jane White, K.A., 2011. The role of optimal control in assessing the most cost-effective implementation of a vaccination programme: HPV as a case study. Math. Biosci. 231, 126–134.

Cauchemez, S., Carrat, F., Viboud, C., Valleron, A.J., Boelle, P.Y., 2004. A Bayesian MCMC approach to study transmission of influenza: application to household longitudinal data. Stat. Med. 23, 3469–3487.

Cauchemez, Simon, Donnelly, Christl A., Reed, Carrie, Ghani, Azra C., Fraser, Christophe, Kent, Charlotte K., Finelli, Lyn, Ferguson, Neil M., 2009. Household transmission of 2009 pandemic influenza A (H1N1) virus in the United States. N. Engl. J. Med. 361, 2619–2627.

Dushoff, Jonathan, Plotkin, Joshua B., Viboud, Cecile, Simonsen, Lone, Miller, Mark, Loeb, Mark, Earn, David J.D., 2007. Vaccinating to protect a vulnerable subpopulation. PLoS Med. 4, 921–927.

Hall, I.M., Egan, J.R., Barrass, I., Gani, R., Leach, S., 2007. Comparison of smallpox outbreak control strategies using a spatial metapopulation model. Epidemiol. Infect. 135, 1133–1144.

Hollingsworth, T.D., Klinkenberg, D., Heesterbeek, H., Anderson, R.M., 2011. Mitigation strategies for pandemic influenza A: balancing conflicting policy objectives. PLoS Comput. Biol. 7, e1001076.

House, T., Keeling, M.J., 2009. Household structure and infectious disease transmission. Epidemiol. Infect. 137 (5), 654–661.

House, T., Ross, J.V., Sirl, D., 2012. How big is an outbreak likely to be? Methods for epidemic final-size calculation. Proc. R. Soc. A: Math. Phys. Eng. Sci. 469 (2150), 20120436.

House, Thomas, Inglis, Nadia, Ross, Joshua V., Wilson, Fay, Suleman, Shakeel, Edeghere, Obaghe, Smith, Gillian, Olowokure, Babatunde, Keeling, Matt J., 2012. Estimation of outbreak severity and transmissibility: influenza A(H1N1)pdm09 in households. BMC Med. 10 (1), 117.

Keeling, Matt J., Shattock, Andrew, 2012 June. Optimal but unequitable prophylactic distribution of vaccine. Epidemics 4 (2), 78–85.

Keeling, Matt J., White, Peter J., 2011 May. Targeting vaccination against novel infections: risk, age and spatial structure for pandemic influenza in Great Britain. J. R. Soc. Interface 8 (58), 661–670.

Klepac, Petra, Laxminarayan, Ramanan, Grenfell, Bryan T., 2011. Synthesizing epidemiological and economic optima for control of immunizing infections. Proc. Natl. Acad. Sci. U. S. A. 108, 14366–14370.

Office for National Statistics, 2001. 2001 Census: Aggregate Data (England and Wales). UK Data Service Census Support.

Ross, J.V., House, T., Keeling, M.J., 2010. Calculation of disease dynamics in a population of households. PLoS ONE 5 (3), e9666.

United States Department of Agriculture, 2007. 2007 Census of Agriculture. United States Summary and State Data, AC-07-A0-51.

van Hoek, Albert Jan, Melegaro, Alessia, Zagheni, Emelio, Edmunds, W. John, Gay, Nigel, 2011. Modelling the impact of a combined varicella and zoster vaccination programme on the epidemiology of varicella zoster virus in England. Vaccine 29 (13), 2411–2420.

Woodhall, Sarah C., Jit, Mark, Cai, Chun, Ramsey, Tina, Zia, Sadique, Crouch, Simon, Birks, Yvonne, Newton, Robert, Edmunds, W. John, Lacey, Charles J.N., 2009 August. Cost of treatment and QALYs lost due to genital warts: data for the economic evaluation of HPV vaccines in the United Kingdom. Sex Transm. Dis. 36 (8), 515–521.