

Neighbourhood control policies and the spread of infectious diseases

L. Matthews^{1*}, D. T. Haydon², D. J. Shaw¹, M. E. Chase-Topping¹,
M. J. Keeling^{3,4} and M. E. J. Woolhouse¹

¹Centre for Tropical Veterinary Medicine, University of Edinburgh, Easter Bush, Roslin EH25 9RG, UK

²Department of Zoology, University of Guelph, Guelph, Ontario N1G 2W1, Canada

³Mathematics Institute, and ⁴Department of Biological Sciences, University of Warwick, Gibbet Hill, Coventry CV4 7AL, UK

We present a model of a control programme for a disease outbreak in a population of livestock holdings. Control is achieved by culling infectious holdings when they are discovered and by the pre-emptive culling of livestock on holdings deemed to be at enhanced risk of infection. Because the pre-emptive control programme cannot directly identify exposed holdings, its implementation will result in the removal of both infected and uninfected holdings. This leads to a fundamental trade-off: increased levels of control produce a greater reduction in transmission by removing more exposed holdings, but increase the number of uninfected holdings culled. We derive an expression for the total number of holdings culled during the course of an outbreak and demonstrate that there is an optimal control policy, which minimizes this loss. Using a metapopulation model to incorporate local clustering of infection, we examine a neighbourhood control programme in a locally spreading outbreak. We find that there is an optimal level of control, which increases with increasing basic reproduction ratio, R_0 ; moreover, implementation of control may be optimal even when $R_0 < 1$. The total loss to the population is relatively insensitive to the level of control as it increases beyond the optimal level, suggesting that over-control is a safer policy than under-control.

Keywords: epidemiology; basic reproduction number; foot-and-mouth disease virus; control strategies; culling theory; optimal control

1. INTRODUCTION

Recent years have seen major epidemics of locally spreading infectious diseases—for example classical swine fever in The Netherlands (Dijkhuizen 1999; Stegeman *et al.* 1999) and foot-and-mouth disease in the UK (Gibbens *et al.* 2001). The need for a quantitative framework, which allows the expected impacts of different control options to be explored in a rigorous manner, has recently been highlighted (The Royal Society 2002).

The control of any infectious disease can, in principle, be achieved in two ways: first, by reducing transmission from an infected to each susceptible individual (isolation), and, second, by limiting the number of susceptible individuals (vaccination). The literature contains numerous comparisons of alternative control strategies for a given disease outbreak (Garner & Lack 1995; Nielen *et al.* 1999; Leslie 2000), but less attention has been focused on the optimal implementation of a given control policy. Though there exists some work on how best to isolate infected individuals (Wickwire 1975; Greenhalgh 1988), much of the literature (which is primarily concerned with human health issues) focuses on the second option and considers the optimal design of vaccination strategies (May & Anderson 1984; Agur *et al.* 1993; Muller 1998).

For livestock diseases, ‘stamping-out’ disease by culling livestock on infected holdings is an option that has been

used extensively in the control of a number of economically important diseases such as foot-and-mouth disease, classical swine fever and Newcastle disease (Elbers *et al.* 1999; Gibbens *et al.* 2001; Hutchinson 1975). A further measure (Gibbens *et al.* 2001) is the pre-emptive culling of livestock on holdings that have not yet been identified as infected but are deemed to be at enhanced risk of infection. This strategy prevents further transmission from these holdings, but may also result in the removal of some uninfected holdings. This leads to a trade-off: increased levels of culling produce a greater reduction in transmission, but increase the number of uninfected holdings culled. We focus on the optimal level of pre-emptive control.

The efficacy of a control policy is usually measured in terms of its ability to reduce the basic reproduction ratio, R_0 , below a value of 1. R_0 is the average number of new infections produced by one infected individual introduced into a fully susceptible population. If $R_0 > 1$ then, on average, the number of new infections will grow, whereas if $R_0 < 1$, new infections will, on average, decline and a major outbreak cannot occur (Anderson & May 1991).

There are, however, further potential requirements of a control policy—for example, spatial containment of the outbreak, a reduction in the outbreak duration, minimization of overall losses to the population or a combination of these factors. Here, we shall consider control strategies that minimize total losses to the population.

We consider a control policy that consists of two parts: first, the culling of infected holdings when they are

* Author for correspondence (louise.matthews@ed.ac.uk).

discovered, and, second, the pre-emptive culling of holdings whose history (for example recent contact with or proximity to an infected holding) suggests an enhanced risk of exposure. Since the pre-emptive policy cannot directly identify infected holdings, its implementation will involve the removal of some uninfected holdings. We aim to optimize the pre-emptive control policy in order to minimize the overall losses to the population. In the rest of this paper, we will refer to this pre-emptive culling as 'control' or 'pre-emptive control' and assume that the culling of holdings identified as having been infected proceeds at a fixed rate. The unit of investigation here is always the livestock holding, rather than the individual animal.

We develop simple mathematical models to capture the effects of the pre-emptive removal of at-risk holdings, such as those in the immediate neighbourhood of an infected holding, taking parameter values from the UK 2001 foot-and-mouth disease epidemic to illustrate model behaviour.

We explore the robustness of the optimal control policy to variation in control effort and examine its sensitivities to the basic reproduction ratio, R_0 , and the ratio of short- (i.e. local) to long-range transmission. Finally, we examine the role of targeted control programmes that additionally take into account differing susceptibilities and transmissibilities of holdings within the at-risk population, which might occur, for example, as a result of different species compositions or differing biosecurity levels.

2. THE MODEL

We initially consider a susceptible-exposed-infectious-removed (SEIR) model of infection dynamics within the population of livestock holdings. Holdings are assumed to fall into one of the following categories: susceptible (S); exposed to infection but not yet infectious (E); infectious and to be removed on discovery of clinical signs (I); or removed by the control programme (R). The transmission rate is given by β , the rate of progression from the exposed to the infectious class is given by γ and the rate of removal by the control programme of holdings identified as infected is given by σ . Thus, exposed holdings arise at a rate $\beta SI/N$ (where $N = S + E + I + R$) and progress to the fully infectious class at a rate γE ; they are removed from the infectious class at a rate σI . This leads to a basic reproduction ratio, R_0 , for an infectious holding of β/σ (Anderson & May 1991).

In addition to the removal of identified infections, we also include a pre-emptive control that removes, at a rate c , holdings that have not been identified as infected, but whose histories suggest that they have been at enhanced risk of exposure (thus c is a measure of pre-emptive control effort). Implementation of the pre-emptive control programme is assumed to remove a proportion, f , of the exposed holdings resulting from a given infectious holding, thus preventing those holdings from progressing to the infectious stage. It is natural to assume that increased levels of control will lead to a greater proportion of exposed holdings being removed prior to reaching the infectious stage; we therefore allow the proportion, f , to depend on the removal rate, c .

Since a proportion, f , of exposed holdings never progress to the infectious stage, a given level of control, c , can




Figure 1. The proportion of infectious cases prevented, $f(c)$, by the pre-emptive removal of holdings at enhanced risk of infection, for a given pre-emptive culling rate or control effort, c (the unit of time is the infectious lifetime of one holding). Data (squares) are derived from the distribution of distances of infectious cases from their sources for the UK 2001 foot-and-mouth disease epidemic. If control is assumed to prevent exposed holdings within a certain distance of their source from becoming infectious, then the proportion of exposed holdings resulting from and lying within a given distance of their source provides an estimate for f . The number of holdings removed from that region (determined by the size of the region and the average density of holdings) corresponds to the parameter κ , from which the culling rate is derived using $c = \kappa\sigma$. The solid line is a curve fitted to the data.

be thought of as producing an effective reduction in the transmission rate from β to $\beta(1 - f(c))$. We therefore define the effective reproduction ratio, R_e , for an infectious holding as $R_e = \beta(1 - f)/\sigma = (1 - f)R_0$. Note that standard approaches to the implementation of control would aim to achieve a value of R_e below 1. We shall see that alternative strategies exist when the aim is to minimize losses to the population.

The number of holdings removed by the pre-emptive control policy per identified infection, during its infectious lifetime ($1/\sigma$), is defined as κ , which we refer to as the culling ratio, and is given by $\kappa = c/\sigma$. Because exposed holdings cannot be directly identified, the pre-emptive culling policy results in the removal of some uninfected holdings. For every κ holdings removed, $R_0 f$ infections are prevented from progressing to the infectious stage; thus, the proportion of the removed holdings that are expected not to have been exposed is given by $1 - (R_0 f/\kappa)$.

Figure 1 illustrates a typical shape for $f(c)$, derived using data on the distribution of distances of infectious cases from their sources during the UK 2001 foot-and-mouth disease epidemic. Since a rational policy will prioritize the removal of those holdings at the greatest risk of exposure, we expect the curve $f(c)$ to be steep for low values of c , where small increase in control can lead to a large change in the effective transmission rate, and to flatten off at large values of c , where a small increase in the level of control has a relatively small effect on the effective transmission rate.

We now argue that a fraction $c + \beta(1 - f)$ of susceptible and exposed holdings change status per unit time per infectious holding, of which c go straight to the removed category R and $\beta(1 - f)$ will go to R via E and I . We also assume that exposed holdings progress rapidly to the

infectious stage (i.e. γ is large). In this case, the exposed population remains small, and is almost in equilibrium with the infectious population, and so its dynamics may be neglected. This is a standard approach, which is frequently used to simplify the analysis of transmission dynamics (Anderson & May 1991). Making this approximation, we obtain the following set of equations:

$$\begin{aligned}\frac{dS}{dt} &= -\beta(1 - f(c))SI/N - cSI/N, \\ \frac{dI}{dt} &= \beta(1 - f(c))SI/N - \sigma I, \\ \frac{dR}{dt} &= \sigma I + cSI/N.\end{aligned}\quad (2.1)$$

(a) Total loss to the population during the course of the outbreak

The set of equations (2.1) is readily solved (by dividing the first equation by the second and integrating with respect to I) to provide an expression for $T(c)$, the total fraction removed from the population for a given control effort, c , over the course of the outbreak. For a population of N holdings with I_0 index infections, we obtain:

$$T(c) = 1 - \left(1 - \frac{I_0}{N}\right)e^{-R_c T(c) - \frac{\kappa}{N} I_0}. \quad (2.2)$$

The proportion of all removed holdings (other than the index cases) that pass through the infectious stage is given by $R_c/(\kappa + R_c)$. Therefore, taking into account the index cases, the fraction of the total population that comprises infectious cases, $T_I(c)$, is given by

$$T_I(c) = (T(c) - I_0/N) \frac{R_c}{\kappa + R_c} + \frac{I_0}{N}. \quad (2.3)$$

(b) Optimizing the pre-emptive control policy

The optimal level of control will minimize the total fraction of the population removed, $T(c)$. To determine the location of this optimum (which might of course be zero control), we take a given level of control, c , and ask whether it could be improved upon. By differentiating equation (2.2) (see Appendix A) we can show that the total fraction removed, $T(c)$, decreases with increasing control effort, c , whenever

$$\beta \frac{df}{dc} > \frac{I_0}{NT(c)}. \quad (2.4)$$

This condition has a straightforward biological interpretation. The left-hand side, $\beta df/dc$, represents, at a given level of control, the ratio of the number of additional exposed holdings that could be prevented from becoming infectious to the number of additional holdings that would have to be culled to achieve that reduction. The right-hand side of the expression gives the ratio of initial infections to total removals at a given level of control. Obviously, a higher level of control is a sensible strategy only if, in order to prevent the occurrence of one additional infectious holding, we pre-emptively cull fewer holdings than would anyway have been infected and culled during the course of the outbreak as a result of that infectious holding.

The gradient, df/dc , is expected to decline with increased levels of control (see, for example, figure 1). Consequently, if $\beta df/dc$ does not exceed $I_0/NT(c)$ when $c = 0$, then in fact the inequality can never be satisfied and $T(c)$ will never fall below $T(0)$. In this case, zero pre-emptive control represents the optimal strategy.

If, however, $\beta df/dc$ does exceed $I_0/NT(c)$ when $c = 0$, then $T(c)$ will decline as levels of control increase until c reaches the optimal level of control at which $\beta df/dc = I_0/NT(c)$. Clearly, increasing c beyond the optimal level will result in an increase in $T(c)$.

In summary, the value of $f(c)$ at a given level of control, c , determines the total loss to the population, but, importantly, the gradient, df/dc , determines whether increased levels of control can further reduce losses.

(c) Targeted control programmes

The control programme described above assumes that the policy is applied equally to all holdings thought to be at an enhanced risk of infection. Here, we describe modifications to the model to take into account differences in susceptibility or transmissibility between holdings and adjust the policy accordingly.

Consider first the possibility that holdings are known to have different susceptibilities to infection. Here, we assume that a proportion, p_h , of holdings that would have been subject to control under the non-targeted scheme are at low risk of infection, and will therefore be omitted from the targeted programme. The remaining proportion, $p_h = 1 - p_h$, is at high risk of infection and subject to control. For risks of infection of s_h and s_l for high- and low-susceptibility holdings, respectively, a removal rate, c , results in an effective transmission rate of $\beta(1 - p_h s_h f(c/p_h))$, which can be substituted into equation (2.1).

A second possibility is that holdings are known to have different transmissibilities. We assume that a proportion, q_h , of holdings have high transmissibilities and a proportion, $q_l = 1 - q_h$, have low transmissibilities. A greater level of control (corresponding to a control effort c_h) is applied in response to identified infection in a high transmitter and a lesser level of control (c_l) is applied in response to identified infection in a low transmitter, giving an overall removal rate of $c = q_h c_h + q_l c_l$. For transmissibilities of t_h and t_l for high and low transmitters, respectively, we obtain an effective transmission rate of $\beta(1 - q_h t_h f(c_h) - q_l t_l f(c_l))$. In this case, seeking the optimal control policy involves identifying two removal rates c_h and c_l that jointly minimize losses to the population.

In both cases, for ease of analysis, we have assumed that the relative proportions of holdings (p_h , p_l , q_h and q_l) remain constant during the course of the outbreak. Such an approximation best reflects reality for small outbreaks. For a full analysis when this assumption is relaxed see Keeling *et al.* (2001).

(d) Metapopulation model

The simplicity of the model presented above permits insight into the effect of the control programme. However, the assumption of a homogeneously mixing population will over-estimate the 'contact' rate between susceptible and infected holdings in a locally spreading outbreak. In reality, we expect to find clustering of infection and for

local availability of susceptibles to play a key role in the dynamics of the outbreak.

Since ‘contact’ can occur in a variety of ways, for example, through nearness to an infectious holding, movement of livestock or mutual use of a market, the concept of ‘locality’ need not necessarily imply spatial proximity. With this in mind, we propose a simple extension of the homogeneous-mixing model, which, though not an explicitly spatial model, begins to encompass the effects of clustering of infection.

We break down the susceptible population into clusters of holdings within which short-range transmission occurs, and between which long-range transmission occurs. Correspondingly, the transmission rate, β , is broken down into a short-range transmission rate, β_s , corresponding to infections generated within the cluster, and a long-range transmission rate, β_l , corresponding to infections generated outside the cluster in question.

The control policy is assumed to act locally—preventing transmission only to other holdings within the cluster. We use the homogeneous-mixing model to describe transmission and control within a cluster.

We assume that the outbreak is initiated by a single infected holding, which generates a cluster of infected holdings. The cluster of infected holdings has the potential to create further clusters via long-range transmission events. The basic reproduction ratio for a cluster, R_C , is the product of the number of long-range transmission events per holding and the total number of infectious holdings in a cluster, thus $R_C = (\beta_l/\sigma)T_1N_H$, where N_H is the number of holdings in a cluster and T_1 is the fraction of the holdings that has been infectious. To determine the total number of affected clusters, we approximate the infection process by applying susceptible–infectious–removed (SIR) type dynamics to the population of clusters. The fraction of infected clusters is therefore given by equation (2.2) with control set to zero.

Under this formulation, reinfection of clusters cannot occur, but simulation shows that the effect of re-infection is expected to make only a small contribution to model output. A similar approach, that of considering within-cluster dynamics in order to calculate a basic reproduction number for a cluster and thus approximate the global infection dynamics, was taken by Muller *et al.* (2000).

We do not attempt an analytic investigation of the existence of an optimal control policy for this model, but exploratory simulation analysis leads us to expect broad qualitative agreement with the homogeneous-mixing model.

In § 3, we compare the behaviour of the homogeneous-mixing model with that of the metapopulation model, for a specific example of a control problem, and examine in detail the optimal control policy for the metapopulation model.

(e) *An example of a locally spreading outbreak*

Data from the UK 2001 foot-and-mouth disease epidemic allow us to construct an example of a pre-emptive control policy based on the removal of holdings within the neighbourhood of an infectious holding.

Data on the distribution of distances of infected cases from their sources are used to derive $f(c)$ as shown in figure 1. In this paper, we do not remove holdings explicitly

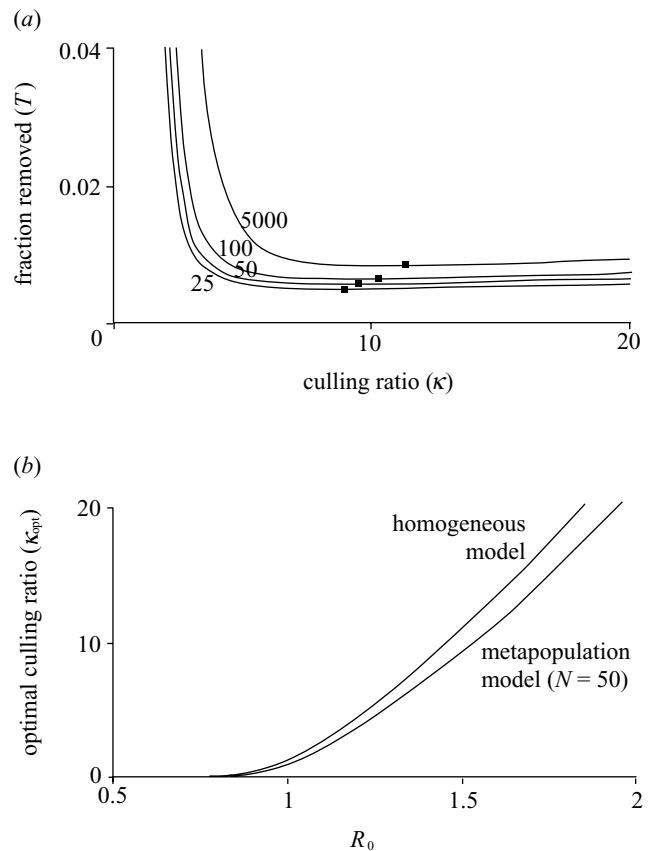


Figure 2. (a) Total fraction removed from the population as a function of culling ratio, κ , for cluster sizes of 25, 50, 100 and 5000, and $R_0 = 1.5$. Squares mark the optimal control effort for each cluster size. (b) Dependence of the optimal value of the culling ratio, κ_{opt} , on the basic reproduction ratio, R_0 , for cluster sizes of 50 and 5000 (the homogeneous mixing model).

on the basis of distance from an infectious holding, but remove a selected number of holdings (κ) at random from a cluster containing an infectious holding, and use $f(c)$ to determine the proportion of exposed holdings that has been removed.

The effect of varying degrees of clustering on model output is investigated. We set the total population size, N , to be 5000, which we divide (in four alternative models) into clusters of size 25, 50, 100 and 5000 (which corresponds to the homogeneous-mixing model). The breakdown of the transmission rate into short- and long-range components, β_s and β_l , is guided by the data on the distribution of distances of infected cases from their sources. Accordingly, we choose the proportions of total transmission occurring within clusters of 25, 50, 100 and 5000 to be 0.65, 0.75, 0.85 and 1.0, respectively. Values for the default transmission rate, β , and the rate of removal of identified infections, σ , are chosen to give a default value for R_0 of 1.5.

For a range of parameter sets, we numerically solve expression (2.2) (and the related expressions for the metapopulation model) for the total fraction removed, T , as control effort, c , increases. In § 3, we use this method to determine the optimal level of control and to examine its sensitivities to the basic reproduction ratio, R_0 , and the level of local clustering. We also compare the optimal targeted-control policies with the optimal non-targeted policy.

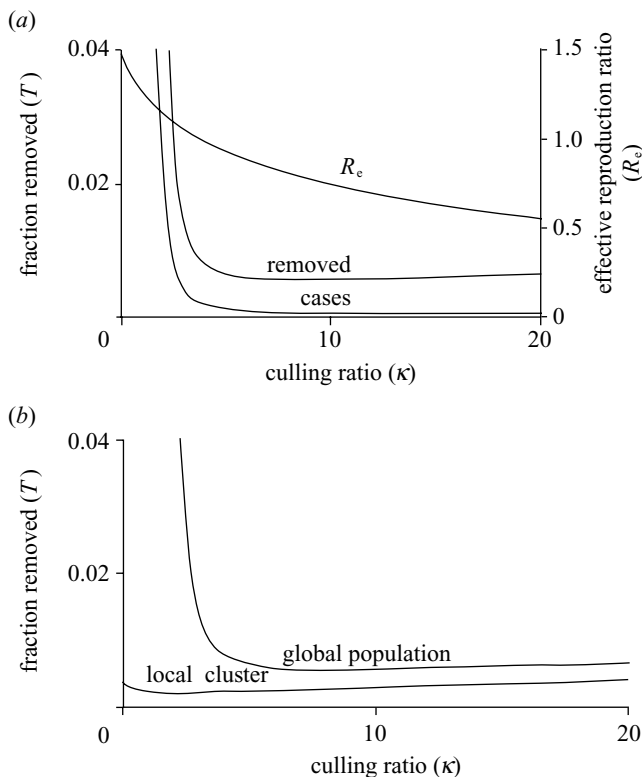


Figure 3. (a) Impact of control policy on the total fraction removed from the population, total fraction of the population infected (cases) and the effective reproduction ratio, R_e , for a cluster size of 50, $R_0 = 1.5$. (b) Total fraction of the global population removed and fraction of the global population removed within the single (local) cluster as a function of the culling ratio, κ , for a cluster size of 50, $R_0 = 1.5$.

3. RESULTS

We first examine the sensitivity of the total fraction removed over the course of the outbreak to the control effort applied. Figure 2a shows the fraction removed as a function of the culling ratio for both the homogeneous-mixing and the metapopulation models (for the three different sizes of local cluster). There are three key features of each curve: (i) the implementation of control can produce a substantial reduction in the total fraction removed from the population; (ii) there exists an optimal level of control that minimizes the total loss to the population (shown as squares); and (iii) the fraction removed is relatively insensitive to the level of control as the optimal level is exceeded.

In figure 2b, we examine the dependence of the optimal culling ratio on the basic reproduction number, R_0 . The level of control required to minimize the fraction removed rises with increasing R_0 . An important observation is that a non-zero level of control may optimize the outcome of the epidemic even when $R_0 < 1$. A greater level of control is required under the homogeneous description of the outbreak dynamics than under the metapopulation description. However, the differences in the predicted optimal culling ratio between the models are small compared with the range of control efforts that would produce a near-optimal result for either of the models.

For the remainder of this section we restrict our atten-

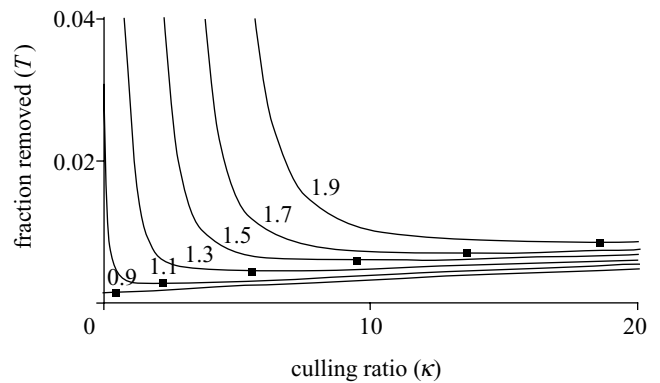


Figure 4. Total fraction removed as a function of culling ratio, κ , for $R_0 = 0.9, 1.1, 1.3, 1.5, 1.7$ and 1.9 and a cluster size of 50. Squares mark the optimal control effort.

tion to the dynamics of the metapopulation model for the intermediate cluster size of 50. In figure 3a, we compare the total fraction removed as a function of the culling ratio with the corresponding fall in the effective reproduction ratio, R_e , and the fraction of the population with identified infections (cases). R_e falls gradually with increasing control effort while the effect on the total size of the outbreak is much more dramatic. The figure clearly demonstrates the key point that, if the criterion for the implementation of a control policy is to minimize losses to the population, the aim should not be the greatest reduction in R_e , but to seek an optimal value to which R_e should be reduced (in this example to *ca.* 0.75).

We now explore the link between the behaviour of the outbreak at a global level and the dynamics within the local clusters. In figure 3b, we compare the total fraction of the population removed across all clusters with the fraction (of the global population) removed within a single (local) cluster. Though the implementation of control measures has relatively little effect on the total losses within a single infected cluster, the consequence of increasing control is a reduction in the *infected* fraction (not shown) within the cluster. This substantially reduces the number of infected clusters that may be produced by long-range transmission events, and hence greatly reduces the total losses to the global population.

In figure 4, the fraction removed versus the culling ratio is shown for a range of values of R_0 . Both the level of control required to minimize the total loss to the population and the minimum loss attained increase as R_0 increases. The curves diverge substantially at low levels of control effort. In this region of parameter space small increases in R_0 can lead to substantial increases in the fraction removed. At higher levels of control, the curves are much more closely spaced for different values of R_0 . In this region of parameter space, variation in R_0 translates into relatively small changes in the fraction removed. Furthermore, if there are perturbations to R_0 , and the level of control remains constant, the fraction removed remains close to its optimal value.

The results shown so far correspond to a non-targeted culling policy, which treats the herds as homogeneous. We now consider two targeted control programmes (again for a cluster size of 50): the first omits holdings of low suscep-

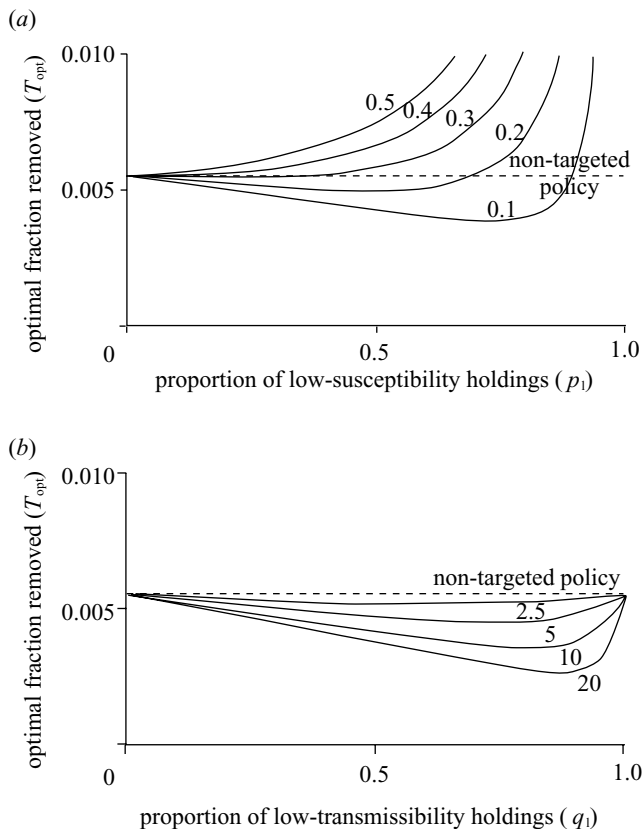


Figure 5. (a) Optimal total fraction, T_{opt} , removed for the non-targeted culling policy (dashed line) and a targeted culling policy directed at high-susceptibility holdings (solid lines), as a function of the proportion, p_l , of low-susceptibility holdings for relative susceptibilities of the low-risk holdings (s_l/s_h) of 0.1, 0.2, 0.3, 0.4 and 0.5. Cluster size of 50; $R_0 = 1.5$. (b) Optimal total fraction, T_{opt} , removed for the non-targeted culling policy (dashed line) and a targeted control policy that directs differing levels of control at holdings with high and low transmissibilities (solid lines), as a function of the proportion, q_l , of low-transmissibility holdings for relative transmissibilities of the high-transmissibility holdings (t_h/t_l) of 2.5, 5, 10 and 20. Cluster size of 50; $R_0 = 1.5$.

tibility from the culling programme (figure 5a), and the second applies a greater level of control in response to the detection of holdings of high transmissibility than for holdings of low transmissibility (figure 5b).

In figure 5a, for a range of relative susceptibilities of low-risk holdings, s_l/s_h , we show the variation in the optimal fraction removed with the proportion of low-risk holdings, p_l . For each set of parameter values the overall reproduction ratio, R_0 , is maintained at a value of 1.5. For a given proportion of low-risk holdings, the benefits of the targeted policy are greatest when the relative susceptibility of the low-risk holdings is lowest. For low proportions of low-risk holdings, the selective policy differs little from the non-targeted policy and so confers little benefit. At high proportions of low-risk holdings, the contribution of these holdings to the overall dynamics of the outbreak is sufficiently great that selectively omitting them from the cull reduces the overall efficacy of the programme and increases the fraction removed. Importantly, if the difference in susceptibilities between the groups is small, the

optimal fraction removed for the targeted policy will always exceed that of the non-targeted policy.

In figure 5b, for a range of relative transmissibilities of high-risk holdings, t_h/t_l , we show the variation in the optimal fraction removed with the proportion of low-transmissibility holdings, q_l . As above, for each set of parameters, the overall reproduction ratio, R_0 , is maintained at a value of 1.5. Targeting the control effort can reduce the optimal fraction removed. The reductions are greatest at high relative transmissibilities and when the proportion of low-transmissibility holdings, q_l , is high (though, as q_l approaches 1, the low-transmissibility holdings comprise such a substantial proportion of the epidemic that the control policy again approaches the non-targeted policy).

4. DISCUSSION

This paper presents a model of a control programme for a disease outbreak in a population of livestock holdings. Control of the outbreak is achieved in two ways: by removing infected holdings when they are discovered and by the pre-emptive removal of holdings deemed to be at enhanced risk of infection. Analysis of a homogeneously mixing SIR model of the control programme provides an expression for the total loss (i.e. both infected and uninfected holdings) to the population over the course of the outbreak. Under biologically realistic assumptions, this enables us to: (i) determine a criterion for the implementation of pre-emptive control; and (ii) demonstrate the existence of an optimal level of pre-emptive control that minimizes the total loss. This model underpins a more complex metapopulation model, which is introduced in order to represent the clustering of infection in a locally spreading outbreak. The behaviour of the metapopulation model is relatively robust to the level of clustering of the susceptible population, and broad qualitative agreement with the homogeneous-mixing model is observed.

For a specific example of a locally spreading outbreak, parameterized using data from the UK 2001 foot-and-mouth disease epidemic, we demonstrate the existence of an optimal level of pre-emptive control. This is a key result: total losses to the population are minimized not when transmission is minimized but at an optimal value of the effective reproduction ratio R_e . As would be intuitively expected, an increase in the basic reproduction number R_0 results in an outbreak that is more challenging in two respects: first, a greater control effort is required to minimize total losses to the population, and, second, that minimum loss to the population is greater at higher values of R_0 . This sensitivity highlights the importance of primary management strategies, directed at identified infections, in minimizing transmission and hence reducing R_0 .

In a disease outbreak with complex epidemiology, parameters relevant to the design and implementation of the control programme may be difficult to estimate and may vary both temporally and spatially. Understanding both inter- and intra-herd dynamics could require knowledge of holding size, species composition, source of infection, number of index cases and pathogenesis of the disease. Assessment of the relative proportions of short- and long-range transmission requires accurate assignment of links between sources of infection and their resulting secondary infections. For these reasons, it is essential to establish the

robustness of the control programme to variability or uncertainty in these parameters.

The key points that emerge are that total losses are not highly sensitive to small variations in the control effort around the optimal value, and that losses increase only gradually as control effort increases beyond the optimal value. This suggests that some leeway is acceptable in practice, but that over-control is generally safer than under-control when trying to avoid large losses to the population. Similar arguments apply for variation in R_0 , suggesting that over-control is safer than under-control when there is uncertainty or variability in this parameter.

Long-range transmission events play a key role in determining the overall losses to the population. Though control measures may not substantially affect losses at the local level, they play a crucial role in reducing the overall levels of infection and thus the number of long-range transmission events into new areas of susceptible holdings. This result further highlights the importance of primary control measures such as movement restrictions designed to contain outbreaks locally.

Control policies might take into account differential susceptibilities and transmissibilities of holdings and adjust the management of the outbreak accordingly. Such targeted control programmes can reduce the total losses to the population below the level achieved by the optimal non-targeted policy. However, if holdings regarded as at low risk of infection are to be omitted from the pre-emptive control programme, careful assessment is required to ensure that their omission does not reduce the overall efficacy of the cull programme and result in total losses that exceed those of the non-targeted policy.

Although the model presented here is relatively simple, the robustness of the results to both model structure and variation in parameter values suggests that they will be applicable to a broader class of more complex models. Indeed, good agreement is found between the deterministic results and those from a stochastic version of the homogeneous model (L. Matthews, unpublished observations). Stochastic effects may be especially important where local populations are small. However, we anticipate that a stochastic version of the metapopulation model will support our conclusion that over-control is a safe option, since fade-out of local outbreaks is expected to reduce total losses at high levels of control. Moreover, the same broad patterns are seen when comparing different control policies using a detailed spatially explicit microsimulation model (Keeling *et al.* 2001; M. J. Keeling, unpublished observations).

Finally, though we have explicitly considered 'stamping out' as a control policy, the results presented are more generally applicable to any control policy where the removal of holdings occurs quickly relative to the lifespan of the infection. This would include fast-acting treatments or vaccinations capable of producing rapid immunity to infection, in which case the 'total loss' could be regarded as a measure of the effort required to implement the policy.

Financial support for this work was provided by the Wellcome Trust. The authors are grateful to Defra (formerly MAFF) for making data available from the UK 2001 foot-and-mouth disease epidemic. They thank M. G. Roberts, N. Colegrave, O. P. Judson and two anonymous referees for their especially helpful comments on the manuscript.

APPENDIX A

An optimal level of control, c , minimizes the total fraction of the population removed, $T(c)$. To determine the location of this minimum, we investigate the conditions under which $dT/dc < 0$. Equation (2.2) can be rewritten:

$$1 - T(c) = g(T, c), \quad (\text{A } 1)$$

where

$$g(T, c) = \left(1 - \frac{I_0}{N}\right) e^{-R_c T(c) \frac{K}{N} I_0}.$$

Differentiating equation (A 1) with respect to c , we obtain:

$$-\frac{dT}{dc} = \frac{\partial g/\partial c}{1 + \partial g/\partial T}.$$

We wish to examine the sign of dT/dc . The following analysis shows that the denominator, $1 + \partial g/\partial T$, is always positive. For a fixed value of c , the solution to equation (A 1) lies at the point of intersection of the curves (regarded as functions of T) $1 - T$ and $g(T, c)$. Clearly when $T = 0$, $g(T, c) < 1$ and $1 - T$ is equal to 1. Now, considering the point of intersection of these curves, since the gradient of $1 - T$ is -1 , it is straightforward to see graphically that, at this point, the gradient of $g(T, c)$ must be greater than -1 , i.e. $\partial g/\partial T > -1$ and hence

$$1 + \frac{\partial g}{\partial T} > 0.$$

It follows that the sign of dT/dc is opposite to that of $\partial g/\partial c$. Differentiating $g(T, c)$ with respect to c , we obtain

$$\frac{\partial g}{\partial c} = \frac{Tg(T, c)}{\sigma} \left(\beta \frac{df}{dc} - \frac{I_0}{NT} \right).$$

Therefore, the fraction of the population infected or removed, $T(c)$, declines as control effort increases provided that

$$\beta \frac{df}{dc} > \frac{I_0}{NT(c)}.$$

REFERENCES

- Agur, Z., Danon, Y. L., Anderson, R. M., Cojocaru, L. & May, R. M. 1993 Measles immunization strategies for an epidemiologically heterogeneous population—the Israeli case-study. *Proc. R. Soc. Lond. B* **252**, 81–84.
- Anderson, R. M. & May, R. M. 1991 *Infectious diseases of humans: dynamics and control*. Oxford University Press.
- Dijkhuizen, A. A. 1999 The 1997–1998 outbreak of classical swine fever in The Netherlands: preface. *Prev. Vet. Med.* **42**, 135–137.
- Elbers, A. R. W., Stegeman, A., Moser, H., Ekker, H. M., Smak, J. A. & Pluimers, F. H. 1999 The classical swine fever epidemic 1997–1998 in The Netherlands: descriptive epidemiology. *Prev. Vet. Med.* **42**, 157–184.
- Garner, M. G. & Lack, M. B. 1995 An evaluation of alternate control strategies for foot-and-mouth-disease in Australia—a regional approach. *Prev. Vet. Med.* **23**, 9–32.
- Gibbins, J. C., Sharpe, C. E., Wilesmith, J. W., Mansley, L. M., Michalopoulou, E., Ryan, J. B. M. & Hudson, M. 2001 Descriptive epidemiology of the 2001 foot-and-mouth

- disease epidemic in Great Britain: the first five months. *Vet. Rec.* **149**, 729–743.
- Greenhalgh, D. 1988 Some results on optimal-control applied to epidemics. *Math. Biosci.* **88**, 125–158.
- Hutchinson, H. L. 1975 The control and eradication of Newcastle disease in Northern Ireland. *Vet. Rec.* **96**, 213–217.
- Keeling, M. J., Woolhouse, M. E. J., Shaw, D. J., Matthews, L., Chase-Topping, M., Haydon, D. T., Cornell, S. J., Kapey, J., Wilesmith, J. & Grenfell, B. T. 2001 Dynamics of the 2001 UK foot and mouth epidemic: stochastic dispersal in a heterogeneous landscape. *Science* **294**, 813–817.
- Leslie, J. 2000 Newcastle disease: outbreak losses and control policy costs. *Vet. Rec.* **146**, 603–606.
- May, R. M. & Anderson, R. M. 1984 Spatial heterogeneity and the design of immunization programs. *Math. Biosci.* **72**, 83–111.
- Muller, J. 1998 Optimal vaccination patterns in age-structured populations. *SIAM J. Appl. Math.* **59**, 222–241.
- Muller, J., Schonfisch, B. & Kirkilionis, M. 2000 Ring vaccination. *J. Math. Biol.* **41**, 143–171.
- Nielen, M., Jalvingh, A. W., Meuwissen, M. P. M., Horst, S. H. & Dijkhuizen, A. A. 1999 Spatial and stochastic simulation to evaluate the impact of events and control measures on the 1997–1998 classical swine fever epidemic in The Netherlands. II. Comparison of control strategies. *Prev. Vet. Med.* **42**, 297–317.
- Stegeman, A., Elbers, A. R. W., Smak, J. & de Jong, M. C. M. 1999 Quantification of the transmission of classical swine fever virus between herds during the 1997–1998 epidemic in The Netherlands. *Prev. Vet. Med.* **42**, 219–234.
- The Royal Society 2002 Infectious diseases in livestock: scientific questions relating to the transmission, prevention and control of epidemic outbreaks of infectious disease in livestock in Great Britain. Policy Document 15/02. London: The Royal Society.
- Wickwire, K. H. 1975 Optimal isolation policies for deterministic and stochastic epidemics. *Math. Biosci.* **26**, 325–346.

As this paper exceeds the maximum length normally permitted, the authors have agreed to contribute to production costs.