Constructing the effect of alternative intervention strategies on historic epidemics

A. R. Cook^{1,2,3,*}, G. J. Gibson^{2,3}, T. R. Gottwald⁴ and C. A. Gilligan¹

¹*Department of Plant Sciences, University of Cambridge, Cambridge CB2 3EA, UK*

²*Department of Actuarial Mathematics and Statistics, and ³Maxwell Institute, Heriot-Watt University, Edinburgh EH14 4AS, UK*

⁴*US Department of Agriculture, Agricultural Research Service, US Horticultural Research Laboratory, Fort Pierce, FL 34945, USA*

Data from historical epidemics provide a vital and sometimes under-used resource from which to devise strategies for future control of disease. Previous methods for retrospective analysis of epidemics, in which alternative interventions are compared, do not make full use of the information; by using only partial information on the historical trajectory, augmentation of control may lead to predictions of a paradoxical increase in disease. Here we introduce a novel statistical approach that takes full account of the available information in constructing the effect of alternative intervention strategies in historic epidemics. The key to the method lies in identifying a suitable mapping between the historic and notional outbreaks, under alternative control strategies. We do this by using the Sellke construction as a latent process linking epidemics. We illustrate the application of the method with two examples. First, using temporal data for the common human cold, we show the improvement under the new method in the precision of predictions for different control strategies. Second, we show the generality of the method for retrospective analysis of epidemics by applying it to a spatially extended arboreal epidemic in which we demonstrate the relative effectiveness of host culling strategies that differ in frequency and spatial extent. Some of the inferential and philosophical issues that arise are discussed along with the scope of potential application of the new method.

Keywords: Bayesian inference; citrus canker; common cold; epidemic control; intervention strategies; stochastic epidemics

1. INTRODUCTION

During the outbreak of an epidemic, decisions are taken on how to intervene in order to mitigate its impact (Anderson *et al.* 2004). The time scale in which a decision must be taken and the paucity of information on key epidemiological parameters early in the epidemic make the choice of intervention or control difficult (Ferguson *et al.* 2001a,b). The decision is often controversial (Cunningham *et al.* 2002; Kitching *et al.* 2006; Wingfield *et al.* 2006), particularly so if the control strategy effected involves pro-active culling of non-symptomatic animals or crops or, for human diseases, travel restrictions. Following the cessation of the outbreak, one question that naturally arises is: *was the right choice of control made?* To answer this, we need to determine what would

have occurred had an alternative, mooted choice of control been implemented. Moreover, we need to do this while taking account of uncertainty in both the parameters governing the process and the underlying, partially observed trajectory of the outbreak (definitions of terminology used in the paper are provided in table 1).

With a little thought, it becomes apparent that the question is ill-posed, since there are at least two interpretations, as follows. Was the best choice made given what was known at the *time of the decision*? Was the best choice made in the light of what is known *now*? We call these *prospective* and *retrospective* questions of the appropriateness of the choice of control, respectively. The prospective situation is the common one that confronts epidemiologists and decision makers in an emerging epidemic. Since there is little information to inform choices made at an early stage, though, the prospective question amounts merely to querying the professional competence of the *person or persons* who made the decision.

It is the retrospective appropriateness that is of most general interest when looking back at an historic epidemic. The retrospective question addresses the

*Author and address for correspondence: Department of Plant Sciences, University of Cambridge, Cambridge CB2 3EA, UK (a.r.cook@ma.hw.ac.uk).

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rsif.2008.0030> or via <http://journals.royalsociety.org>.

Table 1. Terminology used in the paper.

an <i>intervention strategy</i> is any action that may change the way an epidemic invades a population. Non-intervention is also considered to be an intervention strategy
the <i>actual</i> or <i>historic epidemic</i> is the one that did occur in a given place and time under the <i>actual</i> or <i>historic intervention strategy</i>
a <i>mooted</i> or <i>alternative intervention strategy</i> is one that may differ from the actual intervention strategy (it may also be the same)
a <i>notional epidemic</i> is one that did not occur but might have occurred had a mooted intervention strategy been effected rather than the actual one
<i>parametric information</i> encapsulates the uncertainty in the parameters of a model. A single vector of parameter values is capable of generating multiple <i>epidemic trajectories</i> or <i>realizations</i> (though not all are equally likely), of which only one occurs in one temporal and spatial locality
<i>trajectory information</i> encapsulates the uncertainty in a single epidemic trajectory. A single trajectory could have been generated by multiple parameter values, though again not all are equally likely. Note that this information does not have to provide a complete representation of the outbreak
a <i>prospective</i> analysis of the effect of a mooted intervention strategy uses current parametric and trajectory-based information to determine what the possible future effects would be, and often is undertaken while the outbreak is at an early stage
a <i>retrospective</i> analysis of the effect of a mooted intervention strategy uses current parametric and trajectory-based information to determine what the possible past effects would have been, and might be carried out once the outbreak has ceased
a <i>semi-retrospective</i> analysis of the effect of a mooted intervention strategy uses up-to-date parametric information and partial trajectory information to determine what the possible past effects would have been in an ensemble of realizations of alternative realities, which have trajectories that, once the first change in the intervention occurs, are independent of the historic trajectory. Such an analysis also might be carried out once the outbreak has ceased

ultimate effectiveness of the actual and mooted *actions* in the light of what has passed (cf. table 1). It is the retrospective question, too, that interests those who have been personally affected by an intervention. A farmer whose entire herd or crop is culled will want to know whether swifter implementation of control measures might have saved the enterprise, not in a long-run population of outbreaks, but in the specific, actual outbreak in which the losses occurred.

In assessing the retrospective question, it is imperative that full use is made of the information provided by the historic epidemic, of which there are two types: parametric, i.e. information about the process, and trajectory based, i.e. information encapsulated in the epidemic's trajectory (figure 1 for schematic). Until now, however, it has seemed to many researchers that studying populations of outcomes independent from reality is the only way to address notional controls. This *semi-retrospective* approach was taken by Riley *et al.* (2003) for the SARS epidemic of 2003 and by Keeling *et al.* (2001) for the

Figure 1. Diagram representing the different approaches described in the paper. In this simple scenario, the actual intervention (control) ρ is implemented at time t_{actual} , lasts until the end of the outbreak and results in epidemic trajectory X . The choice of ρ makes use of parametric (θ) and trajectory-based (X) information until time t_{actual} . After the outbreak is over, an alternative mooted intervention ρ' is considered which would have been implemented at time t_{mooted} . This would have resulted in the notional trajectory X' , which has a distribution reflecting our uncertainty in the parameters and trajectory. In the semi-retrospective approach, all available parametric information is used, but the only trajectory information used is that occurring before the intervention. In the (fully) retrospective approach described in this paper, all available information (on X and θ) is used.

foot-and-mouth disease (FMD) outbreak in the UK in 2001. Riley *et al.* (2003) noted the difficulty in interpreting the results, describing the notional epidemics of SARS generated with parameters estimated from Hong Kong cases as representing cities with 'Hong Kong-like characteristics' rather than Hong Kong *per se*. One of the authors of the foot-and-mouth paper later pointed out a paradox of the approach, namely that implementing control measures more swiftly than happened in reality leaves 'a significant probability of a worse outcome than was actually observed' (Woolhouse 2003). The cause of this paradox is illustrated by analogy to the following example. Suppose that one mooted intervention during the 2001 FMD epidemic in Great Britain (Ferguson *et al.* 2001a,b), in addition to the controls implemented in reality, was to vaccinate all livestock on the Isle of Man (itself unaffected by the epidemic despite its proximity to the heavily affected areas of Dumfries, Galloway and Cumbria). It is reasonable to believe that there would be no effect of this alternative intervention strategy on the actual outbreak. This should be reflected in predictions that the notional epidemic be the same as the actual one. Ignoring the information content of the observed trajectory of the epidemic, however, leads to a distribution of possible outcomes following the spurious vaccination, some of which have greater and some less disease than the true outbreak, and which is not necessarily centred on the actual outcome. This paradox and the awkwardness of interpretation can be rectified by incorporating trajectory-based information in analyses to compare alternative control strategies.

A heuristic, retrospective solution has been proposed by Haydon *et al.* (2003) and applied to FMD. This involves constructing an *epidemic tree* to summarize the course of infection by linking each infected farm to the donor that infected it. Contact tracing was used to identify some links, while other links were unknown and were constructed using ad hoc data-driven rules, such as the nearest potential donor, or a randomly selected farm from the set of potential donors weighted by a function of distance to the recipient. The notional effect of mooted interventions, such as swifter implementation of the national movement ban, could then be considered by removing branches of the tree. The approach enabled Haydon *et al.* (2003) to obtain estimates of disease levels under mooted interventions. Haydon *et al.* (2003) recognized these to be underestimates, reflecting an inherent bias in their methods, which allows only a single incoming branch to any recipient, so that when a mooted intervention results in the removal of a branch, the notionally non-infected host unit is considered safe from further infection, regardless of the infective pressure exerted upon it by other farms. Note that although the approach of Haydon *et al.* (2003) does eliminate the paradox of potential increase in disease levels under greater control, this is partly due to the inherent restriction that branches are capable of being removed but not of being inserted.

The current paper formalizes the heuristic approach of Haydon *et al.* (2003) by introducing and testing a novel framework for evaluating mooted interventions on historic outbreaks. The new method makes full use of all information available in a statistically coherent fashion that overcomes the bias of previous work and is applicable to a very broad class of models. The framework treats mooted interventions consistently regardless of whether the difference between them and the actual intervention be slight or major. Section 2 introduces the background for the approach and the way it differs from the standard method for prospective analysis. We follow this by introducing the methodology that allows historic and notional epidemics to be coupled. This is done by matching the latent processes generating epidemics that differ only through the effect of the intervention strategies used. The new methodology is illustrated with two examples—one simple, the other more complex—based on historic data: the general SIR model applied to common cold data and a spatio-temporal SI model applied to data on a spatially extended arboreal disease. In the latter, interventions take the form of physical removal of symptomatic and, possibly, also asymptomatic hosts. In the concluding discussion, we compare the approach with previous attempts and discuss some of the inferential issues that arise.

2. PRELIMINARIES: DISTINGUISHING RETROSPECTIVE FROM PROSPECTIVE APPROACH

The question of the prospective appropriateness of differing control strategies is, in principle at least, easy to answer using statistical decision theory (Berger 1993). We denote the unknown parameters by θ , the actual epidemic process by X , the original choice of

intervention by ρ and observed data by $D(X)$. The utility function, characterizing the costs and benefits of the intervention ρ , is denoted $U(X, \rho)$. We also consider a mooted alternative intervention ρ' and the resulting notional epidemic process X' ; there are natural generalizations to more than two alternatives. Throughout we work within the Bayesian paradigm (Lee 2004) and use p to denote both probability mass and density.

The best prospective choice of intervention (ρ or ρ') given what is known at a time T_1 during the outbreak is the one maximizing the expected utility conditional on this knowledge

$$\begin{aligned} & \mathbf{E}\{U(X_{t>T_1}, \rho) | D(X_{t<T_1})\} \\ &= \int U(x_{t>T_1}, \rho) p(x_{t>T_1} | D(X_{t<T_1})) dx_{t>T_1}, \end{aligned} \quad (2.1)$$

$$\begin{aligned} & \mathbf{E}\{U(X'_{t>T_1}, \rho') | D(X_{t<T_1})\} \\ &= \int U(x'_{t>T_1}, \rho') p(x'_{t>T_1} | D(X_{t<T_1})) dx'_{t>T_1}. \end{aligned} \quad (2.2)$$

Although the problem is easily posed, carrying out this integration may be computationally challenging. One approach utilizes Monte Carlo simulation—draw values of θ from $p(\theta | D(X_{t<T_1}))$, use these to generate samples from $p(X_{t>T_1} | D(X_{t<T_1}))$ and $p(X'_{t>T_1} | D(X_{t<T_1}))$, evaluate utilities and take averages. The expected effect of different strategies can then be compared.

It may initially appear that a similar approach can be used retrospectively to find the best decision at time T_1 based on what we know at the present (time T_2 , say), replacing $D(X_{t<T_1})$ by $D(X_{t<T_2})$ in the algorithm above. This is inappropriate, though, if the distribution of $X'_{t>T_1}$ is in any way dependent on $X_{t>T_1}$, which will be the case if both parametric and trajectory-based information are to be fully utilized. We term this the *semi-retrospective* approach, since it conditions on some information but disregards other information.

An alternative approach is needed to make full use of all information at our disposal. This requires that pairs of epidemics be *coupled* so that the distribution of the effect of one intervention conditional on that of another can be determined.

3. COUPLING EPIDEMICS BY MATCHING LATENT PROCESSES

Imagine for a moment that the ‘interventions’ are just alternative ways of observing the system for the purposes of collecting data for inference. We assume that collecting the data has no bearing on the epidemic outcome (no pathogens are inadvertently spread by the collectors, for example) so that the actual epidemic process X is maintained regardless of whether ρ or ρ' is carried out, although $D(X)$ and $D'(X)$ differ. This scenario holds when we attempt to devise retrospectively optimal designs of observation schemes (Cook *et al.* in press). Then the distribution of $D'(X)$ conditioned on $D(X)$ is

$$p(D'(X) | D(X)) = \int p(D'(X) | x) p(x | D(X)) dx, \quad (3.1)$$

which may be sampled using Markov chain Monte Carlo (MCMC) and data augmentation (Gibson & Renshaw

Figure 2. The difference between the Poisson and Sellke constructions in terms of the conditional distribution of t_B and t'_B . In the observed epidemic, background sources infect host A at time t_A ; B is then infected at time t_B either by background sources or by A. In the notional epidemic, A is removed immediately after infection and so it is unable to infect B. Under the Sellke construction, B is infected later (green dot). Under the Poisson construction, B may be infected at the same time as before (if it was infected by background sources in the observed epidemic, orange dot); alternatively its infection time follows a shifted exponential distribution (orange line) if it was infected by A in the observed epidemic.

1998; O'Neill & Roberts 1999) or may be deterministic if $D'(X)$ is fully specified by $D(X)$.

By analogy, one way to evaluate the effect of a notional intervention strategy on an historic epidemic is to seek the joint distribution of two epidemics X and X' that differ in some sense only through the effect of their intervention strategies ρ and ρ' . The distribution of X' conditioned on our knowledge of X then gives our best prediction of what would have happened had the notional intervention been chosen rather than the actual one.

The joint distribution $p(X, X')$ cannot, however, ever be validated empirically, since it is impossible to observe both X and X' , although both marginal distributions can be found by repeated sampling. The impossibility of observing the effects of two treatments on the same sampling unit has long been noted (Rubin 1974) and has been called the *fundamental problem of causal inference* (Holland 1986). Some assumption is therefore necessary to circumvent the problem. Causal inference underlies the approach of Haydon *et al.* (2003), who assumed that inferred branches of their epidemic tree that were not removed by a mooted intervention would be invariant to that intervention, and therefore would be maintained in the notional outbreak. This is also the approach we take in this paper, although our approach allows potential contacts that did not cause infection in reality also to be invariant to changes in the intervention strategy.

Suppose, therefore, that we can identify a latent or underlying stochastic process Z whose sample path is unaffected by the choice of intervention and which, together with the intervention, determines the outcome of the epidemic

$$X = g(Z, \rho), \quad (3.2)$$

$$X' = g(Z, \rho'). \quad (3.3)$$

It then follows that the marginal for the notional outcome conditioned on the actual epidemic is

$$p(X'|X) = \int_{\mathcal{Z}} p(X'|z)p(z|X) dz, \quad (3.4)$$

where $\mathcal{Z}$ is the space of Z s consistent with X and X' . The problem is that there are many different choices of Z , and in general they give different results. In selecting a Z process to match different interventions, we propose the following desiderata:

(D1) Z should represent something we might reasonably expect to be invariant to changes in ρ .

(D2) If $\rho' = \rho$ then Z should give $X' = X$.

(D3) If ρ' is 'close' to ρ , then so too should be X' to X .

(D4) Since the value of Z is unknown in practice, we should be able to evaluate in a straightforward way its distribution conditional on the observed part of X , perhaps numerically.

The most important of these is the first one, i.e. the validity of the reasoning that the physical nature of Z should be maintained for differing interventions.

3.1. Poisson construction

One common way of modelling an epidemic is as a modified Poisson process. If hosts mix homogeneously and contact sufficient for disease to spread occurs at a constant rate β , say, between each pair of hosts, then the occurrence of contacts in the population is a collection of Poisson processes of rate β . Non-homogeneous mixing of hosts may also be accounted for by allowing β to vary with the distance (in space or social space) between hosts (as in the example in §5), for example. Infection and hence disease is assumed to spread across a contact if at that time one host is infectious and the other susceptible. Interventions may take the form of actively removing hosts before they spread infection, or reducing the number of contacts. Under the Poisson construction, Z is the infinite set of contact times and the hosts involved.

3.2. Sellke's construction

The infinite nature of the Poisson Z process is computationally undesirable, so we seek an alternative that is more manageable and yet functionally similar. We therefore consider a construction due to Sellke (1983), which is an equivalent way of formulating standard stochastic epidemic processes. (The approach is connected to the idea of non-centred parameterizations (Papaspiliopoulos *et al.* 2003), since there is a one-to-one relationship between an individual host's Sellke threshold and the cumulative distribution function of the infection time of that host.)

Sellke's construction assigns to each individual j in the population a threshold or resistance to infection $Z_j \sim \text{Exp}(1)$ that must be overcome before j becomes infected. The threshold is overcome by the accumulation of infective pressure—if the rate of infection of j from all sources at time u is $\phi_j(u)$ (that may vary according to host heterogeneity and the evolving contact structure, cf. Cook *et al.* 2007), then the time t_j at which j is infected is the solution of $\int_0^{t_j} \phi_j(u) du = Z_j$. (When no solutions exist because $Z_j > \int_0^\infty \phi_j(u) du$, then there is insufficient infective pressure to infect j and the host escapes infection.) In the general stochastic epidemic model with homogeneous mixing of hosts, the rate of

Figure 3. (a–d) The difference between the Poisson and Sellke constructions in terms of the retrospective joint distribution of the number of infective and removed hosts at time 7 under the two interventions. (a) Retrospective: Poisson and (b) retrospective: Sellke. Also shown is the joint distribution taking realizations under the two interventions to be independent after times $t=0$ and 3, corresponding to what we termed the semi-retrospective approach. (c) Semi-retrospective: fully independent and (d) semi-retrospective: independent post intervention. Marginal distributions are shown in the margins of the plots. These have the same distribution regardless of the method used to generate them.

infection of j at time t is

$$\phi_j(t) = \sum_i \beta \mathbf{1}\{i \in \mathcal{I}(t)\} \mathbf{1}\{j \in \mathcal{S}(t)\}, \quad (3.5)$$

where β is the rate of infection from one infectious host to one susceptible host, $\mathbf{1}\{A\} = 1$ if A is true and 0 otherwise, and $\mathcal{S}(t)$ and $\mathcal{I}(t)$ are the sets of susceptible and infective hosts at time t , respectively. Sellke (1983) shows that this is equivalent to the standard formulation of the infection process of the general stochastic epidemic model, such as the Poisson process approach described previously or Gillespie's (1977) algorithm. Under the Sellke construction, Z is the set of thresholds, with one threshold per host.

3.3. Comparison of Poisson and Sellke constructions

Denoting distributions under the Poisson construction as p_P and under the Sellke as p_S , then although $p_P(X) = p_S(X)$ and $p_P(X') = p_S(X')$ (i.e. the distribution of an epidemic X is identical under the two approaches), it is not generally true that $p_P(X, X') = p_S(X, X')$. This is

illustrated with a simple example (figure 2). Two hosts (A and B) are infected by background sources at rate β . Once one host is infected, it infects the other also at rate β . Suppose that A is infected first. The two interventions considered are ρ : do nothing, and ρ' : remove the first host to become infected immediately upon its infection. The distribution of the time A is infected is $t_A \sim \text{Exp}(\beta)$. Under ρ , $(t_B - t_A | t_A) \sim \text{Exp}(2\beta)$ and under ρ' , $(t_B - t_A | t_A) \sim \text{Exp}(\beta)$.

Now suppose that we have observed X completely and thus know t_A and t_B . Under the Sellke construction, there is a one-to-one mapping between the latent process Z and infection times given the parameters, and so the notional infection time for B takes a point mass at $t'_B = t_A + 2(t_B - t_A)$. Under the Poisson construction, however, B was infected by the background source with probability 1/2, in which case its infection is unaffected by the removal strategy and $t'_B = t_B$; otherwise it was infected by A in reality and so its notional distribution is thus $(t'_B - t_B | t_B) \sim \text{Exp}(\beta)$, since there is no information on the next infectious contact from the background source to B. The difference is illustrated in figure 2.

In practice, it is justifiable to use the Sellke as an approximation to the Poisson construction. Infection times are not observed precisely in reality; instead they are typically censored, recorded to the nearest day, for example. In the presence of censoring, event times assume a joint distribution conditional on what is observed ($D(X)$ rather than X), and the two approaches give very similar conditional distributions $p(X'|D(X))$ (not shown).

This similarity of behaviour is illustrated through the following simulation example based upon the general SIR epidemic (Bailey 1975). Motivated by the case study in §4, consider a homogeneously mixing population of size $N=262$ with rate of infection per S–I pair (as in equation (3.5)) $\beta=0.003$ and rate of recovery per infected individual $\gamma=0.66$. We start with 10 individuals infected at time 0. (The high initial number of infectives is for convenience of representation and does not affect the generality of the comparison.)

A total of 100 000 realizations of Z is generated for both Sellke and Poisson constructions (see appendix A in the electronic supplementary material). These are used to obtain the distribution of the number of infective and removed hosts at time $t=7$ under the interventions ρ : do nothing, and ρ' : halve β from time $t=3$ onwards (which we implement in the Poisson construction using Rényi's splitting theorem (Rényi 1964; Srivastava 1971) to discard each contact with probability $1/2$). These are shown in figure 3, both jointly and marginally. Also shown is the relationship taking the notional disease trajectory to be independent of the true one after times $t=0$ and 3. For this case, the Sellke construction provides an excellent approximation to the joint distribution of the number of infectives $I(7)$ and $I'(7)$ generated under the Poisson construction, with the two quantities strongly correlated. This provides *prima facie* justification for using the more tractable Sellke construction in practice. Additional simulations (not shown) indicate that conditional on a complete realization of the epidemic, the resulting notional trajectory under Sellke matches the mean of the distribution of Poisson trajectories closely. In contrast, the mean using the semi-retrospective approach matches the Poisson trajectories well only when a major intervention occurs so that little information from the actual outbreak is relevant in constructing the notional trajectory.

4. EXAMPLE: THE COMMON COLD ON TRISTAN DA CUNHA

We compare the semi-retrospective approach with the (fully) retrospective approach by applying both to historic data on the common cold on the remote island of Tristan da Cunha during an outbreak starting January 1965 (Hammond & Tyrrell 1971; Shibli *et al.* 1971) and considering the effectiveness of two alternative interventions.

The general SIR stochastic epidemic model (Bailey 1975) is fitted to the daily numbers of removed individuals (i.e. those whose symptoms have ceased and have moved from the I to the R class) using standard data augmentation and MCMC integration techniques to obtain the joint posterior distribution of parameters and unobserved event times (e.g. Gibson & Renshaw 1998; O'Neill & Roberts

1999, and appendix B in the electronic supplementary material). The model has infection as in equation (3.5), i.e. infection occurs at rate β per S–I pair in the absence of intervention. The control strategies whose effectiveness we wish to investigate take the form of reductions to β 3 days after the first removal, perhaps due to warnings issued by the island's physician to reduce contact with other islanders. Letting $\beta\xi$ be the rate after 3 days, the two strategies we consider are ρ_1 : $\xi=0.5$ and ρ_2 : $\xi=0.9$. Infected individuals recover, leaving $\mathcal{I}(t)$ and entering the set of removed individuals, $\mathcal{R}(t)$, at rate γ . The data are $\{R(t)=|\mathcal{R}(t)|: t=0, \dots, 19\}$, the population size of $N=262$ and the fact that no subsequent infections occurred in the outbreak. The initial infection is assumed to arise by some other, unmodelled, process. Flat priors are taken for β and γ on the region $[0, 100]^2$.

As part of the parameter estimation routine, the set of Sellke thresholds is calculated and used to estimate the distribution of the number of removals with time, under the two control strategies of interest. Non-infected hosts are given randomly generated thresholds, conditional on being non-infected, making use of the memoryless property of the exponential distribution. Details of the implementation may be found in the electronic supplementary material, appendix B.

Posterior medians and credible regions (Lee 2004) are shown in figure 4*b,d*. This figure also shows predictions using the semi-retrospective approach (figure 4*a,c*), with infection trajectories diverging from the actual one on day 3, when the intervention strategy begins. The problems with the latter approach are clear when one considers the predicted effect of a small reduction to the rate of infection—paradoxically, the expected amount of disease increases (figure 4*a*). Under our fully retrospective approach, however, we would instead predict a small decrease in the numbers of infections (figure 4*b*). When the change is more marked ($\xi=0.5$, figure 4*c,d*), the expected behaviour under the two approaches is similar, but taking the semi-retrospective approach inflates the variance considerably relative to the fully retrospective approach, yielding considerably poorer predictions.

5. EXAMPLE: CITRUS CANKER IN FLORIDA

Citrus canker is an economically important disease of citrus trees caused by the bacterium *Xanthomonas axonopodis* pv. *citri* (Graham *et al.* 2004). It was found to have been reintroduced to Florida, USA, in 1995, following a successful eradication programme from 1986 to 1992. From October 1997 to July 1999, Gottwald *et al.* (2002) collected spatio-temporal data on the locations and disease status of citrus trees in five residential districts of Florida. A detailed analysis will appear elsewhere; in this paper, we use the data to evaluate how effective a strategy of regular monitoring and removal of symptomatic trees, and potentially non-symptomatic trees in their vicinity, would have been at reducing the number of trees infected.

Let the rate of infection of j , a susceptible tree, be

$$\phi_j(t) = \epsilon + \sum_i \beta f(d_{ij}, \alpha) \mathbf{1}\{i \in \mathcal{I}(t)\} \mathbf{1}\{j \in \mathcal{S}(t)\}, \quad (5.1)$$

where $d_{i,j}$ is the distance in kilometres separating trees i and j and $f(d, \alpha)$ is a dispersal kernel. We use the exponential kernel $f(d, \alpha) = \exp(-d/\alpha)$; results are robust to sensible choices of kernel (not shown). Rates of host-to-host infection are also governed by β , while ϵ is the *per capita* rate of infection from external sources outwith the study area and potentially from anywhere within the infected area of the state.

We analyse disease progress maps made at 30 day intervals over 360 days from 26 October 1997 to 20 October 1998 in a 10.3 km² area of Miami, Dade county (site D1 in Gottwald *et al.* 2002). The trees in the study area were sampled extensively, yielding a very rich dataset, with 1124 of 6056 trees becoming infected. See Gottwald *et al.* (2002) for a fuller description of the data collection and some caveats regarding their interpretation. Note that the epidemic was ultimately interrupted by a prolonged period of dry conditions; prior to that, infected trees were effectively continuously infectious. During this year, removal efforts focused on clearing a backlog from elsewhere in the state and so no trees were removed from the epidemic from the population in question. We investigate the effect of notional removal strategies in which, at intervals of Δ days, all trees are assessed for infection. Infection is detected with probability q (including $q=1$): trees found to be infected are removed immediately along with all other trees within a circle of radius r metres centred on the infected tree.

In reality, tree disease status was assessed by a team of phytopathologists (Gottwald *et al.* 2002). In the mooted interventions, we allow for the possibility of non-detection accounting for less formal disease assessment. For the purposes of illustrating the methodology, here we focus exclusively on a simple criterion for the effectiveness of intervention: the total number of trees removed relative to the numbers infected in the historic epidemic. This does not take account of the force of infection generated by the population, nor of the costs of surveying and removing trees.

We used an MCMC routine to sample the posterior distribution of $(\alpha, \beta, \epsilon)$ (taking uniform priors over the region of interest) and the infection times $\{t_i : i = 1, \dots, 1124\}$. In figure 5, we present mean and 95% credible regions for $\sum_i \mathbf{1}\{i \in \mathcal{I}(t)\} + \mathbf{1}\{i \in \mathcal{R}(t)\}$ (i.e. the number of trees lost to disease and/or removal) against time, under three interventions: removals every 60 days ($\Delta=60$, $q=1$, $r=0$), removals every 120 days ($\Delta=120$, $q=1$, $r=0$) and 40% removals every 120 days ($\Delta=120$, $\pi=0.4$, $r=0$). In table 2 we present the ratio of trees lost for these and several other removal strategies during a 360 day period relative to the actual number infected during this interval. This ratio is greater than 1 if pre-emptive removal of susceptible trees is carried out ($r>0$); otherwise the maximum the ratio may take is unity, as all other intervention strategies lead to a strict non-increase in the infective pressure on all hosts.

From table 2 it is clear that removal of symptomatic trees would have substantially reduced the amount of trees being infected in the study area. The more frequent the surveying and removal, the less disease would have resulted, although losses depend nonlinearly upon

surveying frequency. These results seem to suggest that pro-active removal of asymptomatic trees in the vicinity of a known infective is not an effective strategy, since it is predicted to increase total losses; indeed, the broadest removal radius considered by us was predicted to result in almost all trees being removed. We consider this ostensible inefficacy to be an artefact of analysing a non-isolated population. The study area considered here forms part of a greater population of citrus statewide (Gottwald *et al.* 2002) and pro-actively culling trees only in some areas allows disease to be reintroduced from elsewhere. However, this and the other study areas were unusual in that disease was allowed to increase without intervention by the regulatory agencies specifically so that the epidemic could be studied to aid the development of intervention strategies. Within all other locations throughout the state intense culling of symptomatic and surrounding trees was practised. From a statewide perspective, culling is desirable, for although culling results in heavy to complete losses of tree populations locally, the practice effectively lowers ϵ at the geographic scale.

6. DISCUSSION

This paper introduced a novel approach to analysing retrospectively the effects of mooted interventions on historic epidemics. Coupling epidemics by matching their latent processes allows us to make full use of both parametric and trajectory-based information. By so doing, we avoid paradoxical results that may often occur when information about the epidemic trajectory is only partially used; these include predictions of a high probability of a more severe epidemic occurring with more effective controls than were used in reality (figure 4). Our method is easy to implement as part of a parameter estimation routine using standard MCMC techniques. The choice of latent process to match is subjective, but may be guided by desiderata, which the Sellke construction satisfies.

There is a logical inconsistency in using information from the whole of an outbreak to parametrize the model but only part of the information on the trajectory of the outbreak: that coming before a mooted intervention. This inconsistency is most evident when we wish to assess the effect of an incremental change to an actual strategy. Our fully retrospective approach is a natural alternative that preserves patterns and pathologies in the data. It preserves patterns wholly when the mooted intervention has effects identical to the actual one, and partly and decreasingly so when the mooted intervention changes more and more from the actual one. When the mooted intervention is drastically different from the actual one, all trajectory information is lost, and our approach gives predictions that match those using the semi-retrospective approach.

The mechanism used to couple epidemics was to match latent processes that we *assume* to be unaffected by control. We did this using the Sellke construction to transform the effect of one intervention strategy (including non-intervention) to another. The choice of matching-process is subjective. Even the semi-

Figure 4. Effect of two notional intervention strategies on an historic epidemic of the common cold on the island of Tristan da Cunha. These take the form of scaling the infection rate by ξ after $t=3$; time has been translated so that the first removal occurs on day 1. (a,c) The semi-retrospective distributions are what we would obtain using the posterior for the model parameters based on the whole of the epidemic but assuming that the trajectory of the epidemic after $t=3$ is independent of the actual epidemic. (b,d) The fully retrospective distributions show the distribution of the effect of the interventions on the epidemic conditional on the actual epidemic. Predictions take the form of credible regions.

retrospective approach rests on an assumption: that the notional invasion trajectory is independent of the actual one, i.e. that the occurrence of events in the observed epidemic has no correlation with their occurrence in the notional outbreak. This is one extremum of the set of possible assumptions regarding the relationship between two epidemics. Indeed, *any* division of unobservables into ‘parameters’ and ‘outcomes’ is inherently arbitrary.

The main philosophical issue with any approach to evaluating the benefits provided by the actual intervention compared with the outcome of an alternative strategy is that the relationship between a notional and the actual epidemic can never be verified. This is not a new issue in modelling: in making any predictions based on a model, we implicitly trust that the model provides a reasonable description of reality and may be extrapolated to future or alternative conditions. Indeed, non-verifiability of cause and effect is an old philosophical issue dating back to Hume. The theory of causal inference (e.g. Rubin 1974, 2007; Holland 1986; Cox 1992; Greenland & Brumback 2002) has been developed to overcome (at least partially) this obstacle and is frequently used in medicine, the social sciences and econometrics when randomization of sampling units is not possible. Most of the causal inference literature focuses on scenarios in which treatments are applied to a sample of units that respond independently. Clearly this is inappropriate

for contagious infections since the infectious status of individuals in the population are not independent (Halloran & Struchiner 1995). Our approach differs in that we are interested in the effect of ‘treatment’ on a population rather than an individual, and only one ‘treated’ population is observed. We therefore tackled the problem by using the Sellke construction to decompose the epidemic into simpler components and then again to reconstruct the notional outbreak under a mooted intervention. As in causal inference (see Cox 1992) the assumptions we make can be justified from first principles but cannot be independently tested.

We have applied the method in two ways. First, we used temporal data and a simple and accessible model. Here we made simplifying assumptions, assuming no latent period (cf. Arruda *et al.* 1997; Heikkinen & Järvinen 2003) and homogeneous mixing (cf. Becker & Hopper 1983) and susceptibility (cf. Heikkinen & Järvinen 2003), for example. The validity and tractability of the method are not, however, dependent on these simplifying assumptions. To illustrate the generality of the approach, we also used a spatio-temporal model to analyse the effect of removal of trees on an economically important disease of citrus in Florida.

The most pressing extension of the work is to incorporate economic factors such as treatment costs, in order to identify economically optimal strategies

Figure 5. Retrospective effect of three removal strategies on lost citrus trees in an area of Florida. In (a), solid lines are posterior means and dashed lines 95% credible intervals. Points mark the observed number of infections (there were no removals in reality). Colours distinguish three notional control strategies: ρ'_1 : sample the population every 60 days and remove all infected trees (corresponding to (b)); ρ'_2 : sample every 120 days and remove all infected trees (c); and ρ'_3 : sample every 120 days but only detect and remove infection with probability 40% (d). The three lower panels show maps of the areas, with coloured symbols indicating the posterior mean probability that the corresponding tree would have been infected by time 360 had the intervention taken place (legend in top panel).

(Forster & Gilligan 2007). This may be effected within the current framework by suitably adapting the utility function. This remains the subject of future work.

Table 2. Posterior effect of various intervention strategies on the number of trees infected or removed during a period of 360 days. (The removal efficacy is labelled q , the removal frequency Δ and the radius of removal r (38.1 m = 125 ft being the original radius around detected infections used to determine asymptomatic trees for removal elsewhere in the state, later extended to 579 m = 1900 ft): if no entry is present in the r column this indicates only the tree with infection detected is removed. The mean ratio of notional to actual losses as well as 95% credible bounds are tabulated.)

q (%)	Δ (days)	r (m)	mean	95% CI
100	60		0.30	(0.28, 0.32)
	120		0.39	(0.35, 0.42)
	240		0.49	(0.48, 0.51)
80	60		0.34	(0.31, 0.37)
	120		0.49	(0.44, 0.53)
	240		0.56	(0.53, 0.59)
	60	579	5.34	(5.30, 5.35)
	60	50	2.41	(2.27, 2.52)
	60	38.1	1.88	(1.77, 1.98)
	60	25	1.32	(1.23, 1.40)
	120	579	5.34	(5.24, 5.37)
	120	50	2.24	(2.08, 2.40)
60	120	38.1	1.78	(1.64, 1.91)
	120	25	1.30	(1.20, 1.39)
	60		0.42	(0.37, 0.47)
	120		0.59	(0.54, 0.64)
	240		0.64	(0.61, 0.68)
	60	579	5.32	(5.23, 5.35)
	60	50	2.36	(2.22, 2.49)
	60	38.1	1.87	(1.74, 1.98)
	60	25	1.34	(1.24, 1.44)
40	120	579	5.30	(5.13, 5.37)
	120	50	2.11	(1.94, 2.27)
	120	38.1	1.69	(1.56, 1.83)
	120	25	1.28	(1.18, 1.37)
	60		0.55	(0.49, 0.60)
	120		0.71	(0.66, 0.76)
	240		0.75	(0.71, 0.79)
	60	579	5.27	(5.11, 5.36)
	60	50	2.24	(2.09, 2.38)
	60	38.1	1.80	(1.67, 1.92)
	60	25	1.34	(1.23, 1.43)
	120	579	5.22	(4.95, 5.37)
	120	50	1.91	(1.74, 2.06)
	120	38.1	1.56	(1.44, 1.68)
	120	25	1.23	(1.14, 1.31)

The scope of potential applications is broad. Vital questions about the effectiveness of varying the timing and intensity of control may be evaluated in diseases of humans (e.g. SARS (Riley *et al.* 2003; Wallinga & Teunis 2004), Spanish influenza (Chowell *et al.* 2006a,b) and Ebola (Lekone & Finkenstädt 2006)) and other animals (e.g. FMD) as well as plants. We note, in particular, that a very nice aspect of the approach of Haydon *et al.* (2003)—their incorporation of known infectious contacts—could easily and consistently be incorporated within our framework also. By using an appropriate latent process such as the Sellke, our method also allows assessments to be made of the risk of inaction following outbreaks in which control measures were actually deployed. The method described in this

paper allows such important issues to be tackled taking full account of all available sources of information.

The data and C++ routines used in the common cold example are also available in the electronic supplementary material.

The authors wish to thank the Biotechnology and Biological Sciences Research Council (BBSRC) for financing the research project (grant no. BB/C007263/1). C.A.G. gratefully acknowledges the support of a BBSRC Professorial Fellowship. Financial support from the United States Department of Agriculture is gratefully acknowledged. A.R.C. carried out some of this research while visiting the University of Tokyo and thanks Prof. Hisashi Inaba for facilitating the visit. We are grateful to two anonymous referees for their helpful comments.

REFERENCES

- Anderson, R. M., Fraser, C., Ghani, A. C., Donnelly, C. A., Riley, S., Ferguson, N. M., Leung, G. M., Lam, T. H. & Hedley, A. J. 2004 Epidemiology, transmission dynamics and control of SARS: the 2002–2003 epidemic. *Phil. Trans. R. Soc. B* **359**, 1091–1105. (doi:10.1098/rstb.2004.1490)
- Arruda, E., Pitkäranta, A., Witek, T. J., Doyle, C. A. & Hayden, F. G. 1997 Frequency and natural history of rhinovirus infections in adults during Autumn. *J. Clin. Microbiol.* **35**, 2864–2868.
- Bailey, N. T. J. 1975 *The mathematical theory of infectious diseases and its applications*, 2nd edn. London, UK: Griffin.
- Becker, N. G. & Hopper, J. L. 1983 Assessing the heterogeneity of disease spread through a community. *Am. J. Epidemiol.* **117**, 362–374.
- Berger, J. O. 1993 *Statistical decision theory and Bayesian analysis*. New York, NY: Springer.
- Chowell, G., Ammon, C. E., Hengartner, N. W. & Hyman, J. M. 2006a Transmission dynamics of the great influenza pandemic of 1918 in Geneva, Switzerland: assessing the effects of hypothetical interventions. *J. Theor. Biol.* **241**, 193–204. (doi:10.1016/j.jtbi.2005.11.026)
- Chowell, G., Nishiura, H. & Bettencourt, L. M. A. 2006b Comparative estimation of the reproductive number for pandemic influenza from daily case notification data. *J. R. Soc. Interface* **4**, 155–166. (doi:10.1098/rsif.2006.0161)
- Cook, A. R., Otten, W., Marion, G., Gibson, G. J. & Gilligan, C. A. 2007 Estimation of multiple transmission rates for epidemics in heterogeneous populations. *Proc. Natl Acad. Sci. USA* **104**, 20 392–20 397. (doi:10.1073/pnas.0706461104)
- Cook, A. R., Gibson, G. J. & Gilligan, C. A. In press. Optimal observation times in experimental epidemic processes. *Biometrics*. (doi:10.1111/j.1541-0420.2007.00931.x)
- Cox, D. R. 1992 Causality: some statistical aspects. *J. R. Stat. Soc. A* **155**, 291–301. (doi:10.2307/2982962)
- Cunningham, I. (chair) *et al.* 2002 *Inquiry into foot and mouth disease in Scotland*. Edinburgh, UK: Royal Society of Edinburgh.
- Ferguson, N. M., Donnelly, C. A. & Anderson, R. M. 2001a Transmission intensity and impact of control policies on the foot and mouth epidemic in Great Britain. *Nature* **413**, 542–548. (doi:10.1038/35097116)
- Ferguson, N. M., Donnelly, C. A. & Anderson, R. M. 2001b The foot-and-mouth epidemic in Great Britain: pattern of spread and impact of interventions. *Science* **292**, 1155–1160. (doi:10.1126/science.1061020)
- Forster, G. A. & Gilligan, C. A. 2007 Optimizing the control of disease infestations at the landscape scale. *Proc. Natl Acad. Sci. USA* **104**, 4984–4989. (doi:10.1073/pnas.0607900104)
- Gibson, G. J. & Renshaw, E. 1998 Estimating parameters in stochastic compartment models using Markov chain methods. *IMA J. Math. Appl. Med. Biol.* **15**, 19–40. (doi:10.1093/imammb/15.1.19)
- Gillespie, D. T. 1977 Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* **81**, 2340–2361. (doi:10.1021/j100540a008)
- Gottwald, T. R., Sun, X., Riley, T., Graham, J. H., Ferrandino, F. & Taylor, E. L. 2002 Geo-referenced spatiotemporal analysis of the urban citrus canker epidemic in Florida. *Phytopathology* **92**, 361–377. (doi:10.1094/PHYTO.2002.92.4.361)
- Graham, J. H., Gottwald, T. R., Cubero, J. & Achor, D. S. 2004 *Xanthomonas axonopodis* pv. *citri*: factors affecting successful eradication of citrus canker. *Mol. Plant Pathol.* **5**, 1–15. (doi:10.1046/j.1364-3703.2004.00197.x)
- Greenland, S. & Brumback, B. 2002 An overview of relations among causal modelling methods. *Int. J. Epidemiol.* **31**, 1030–1037. (doi:10.1093/ije/31.5.1030)
- Halloran, M. E. & Struchiner, C. J. 1995 Causal inference in infectious diseases. *Epidemiology* **6**, 142–151. (doi:10.1097/00001648-199503000-00010)
- Hammond, B. J. & Tyrrell, D. A. J. 1971 A mathematical model of common-cold epidemics on Tristan da Cunha. *J. Hyg. (Camb.)* **69**, 423–433.
- Haydon, D. T., Chase-Topping, M., Shaw, D. J., Matthews, L., Friar, J. K., Wilesmith, J. & Woolhouse, M. E. J. 2003 The construction and analysis of epidemic trees with reference to the 2001 UK foot-and-mouth outbreak. *Proc. R. Soc. B* **270**, 121–127. (doi:10.1098/rspb.2002.2191)
- Heikkinen, T. & Järvinen, A. 2003 The common cold. *Lancet* **361**, 51–59. (doi:10.1016/S0140-6736(03)12162-9)
- Holland, P. W. 1986 Statistics and causal inference. *J. Am. Stat. Assoc.* **81**, 945–960. (doi:10.2307/2289064)
- Keeling, M. J. *et al.* 2001 Dynamics of the 2001 UK Foot and Mouth epidemic: stochastic dispersal in a heterogeneous landscape. *Science* **294**, 813–817. (doi:10.1126/science.1065973)
- Kitching, R. P., Taylor, N. M. & Thrushfield, M. V. 2006 Vaccination strategies for foot-and-mouth disease, and reply by Tildesley, M. *et al. Nature* **445**, E12–E13. (doi:10.1038/nature05604)
- Lee, P. M. 2004 *Bayesian statistics: an introduction*, 3rd edn. London, UK: Arnold.
- Lekone, P. E. & Finkenstädt, B. F. 2006 Statistical inference in a stochastic epidemic SEIR model with control intervention: Ebola as a case study. *Biometrics* **62**, 1170–1177. (doi:10.1111/j.1541-0420.2006.00609.x)
- O'Neill, P. D. & Roberts, G. O. 1999 Bayesian inference for partially observed stochastic epidemics. *J. R. Stat. Soc. A* **162**, 121–129. (doi:10.1111/1467-985X.00125)
- Papaspiliopoulos, O., Roberts, G. O. & Sköld, M. 2003 Non-centred parameterisations for hierarchical models and data augmentation. *Bayesian Stat.* **7**, 307–326.
- Rényi, A. 1964 On two mathematical models of traffic on a divided highway. *J. Appl. Probab.* **1**, 311–320. (doi:10.2307/3211862)
- Riley, S. *et al.* 2003 Transmission dynamics of the etiological agent of SARS in Hong Kong: impact of public health interventions. *Science* **300**, 1961–1966. (doi:10.1126/science.1086478)

- Rubin, D. B. 1974 Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* **66**, 688–701. (doi:10.1037/h0037350)
- Rubin, D. B. 2007 The design versus the analysis of observational studies for causal effects: parallels with the design of randomized trials. *Stat. Med.* **26**, 20–36. (doi:10.1002/sim.2739)
- Sellke, T. 1983 On the asymptotic distribution of the size of a stochastic epidemic. *J. Appl. Probab.* **20**, 390–394. (doi:10.2307/3213811)
- Shibli, M., Gooch, S., Lewis, H. E. & Tyrrell, D. A. J. 1971 Common colds on Tristan da Cunha. *J. Hyg. (Camb.)* **69**, 255–262.
- Srivastava, R. C. 1971 On a characterization of the Poisson process. *J. Appl. Prob.* **8**, 615–616. (doi:10.2307/3212185)
- Wallinga, J. & Teunis, P. 2004 Different epidemic curves for severe acute respiratory syndrome reveal similar impacts of control measures. *Am. J. Epidemiol.* **160**, 509–516. (doi:10.1093/aje/kwh255)
- Wingfield, A., Miller, H. & Honhold, N. 2006 FMD control strategies, and reply by Keeling, M. *et al. Vet. Rec.* **158**, 706–708.
- Woolhouse, M. E. J. 2003 Foot-and-mouth disease in the UK: what should we do next time? *J. Appl. Microbiol.* **94**, 126S–130S. (doi:10.1046/j.1365-2672.94.s1.15.x)