Analysis of the sensitivity properties of a model of vector-borne bubonic plague

Megan Buzby¹, David Neckels², Michael F. Antolin³ and Donald Estep^{1,4,*}

¹*Department of Mathematics, ³Department of Biology, and ⁴Department of Statistics, Colorado State University, Fort Collins, CO 80523, USA*

²*Computational and Information Systems Laboratory, National Center for Atmospheric Research, PO Box 3000, Boulder, CO 80307-3000, USA*

Model sensitivity is a key to evaluation of mathematical models in ecology and evolution, especially in complex models with numerous parameters. In this paper, we use some recently developed methods for sensitivity analysis to study the parameter sensitivity of a model of vector-borne bubonic plague in a rodent population proposed by Keeling & Gilligan. The new sensitivity tools are based on a variational analysis involving the adjoint equation. The new approach provides a relatively inexpensive way to obtain derivative information about model output with respect to parameters. We use this approach to determine the sensitivity of a quantity of interest (the force of infection from rats and their fleas to humans) to various model parameters, determine a region over which linearization at a specific parameter reference point is valid, develop a global picture of the output surface, and search for maxima and minima in a given region in the parameter space.

Keywords: vector-borne disease; model sensitivity analysis; variational analysis

1. INTRODUCTION

Keeling & Gilligan (2000) analysed a deterministic model of vector-borne bubonic plague en route to developing a stochastic metapopulation model. The goal is to use the disease and population dynamics of rats and fleas to explain how sporadic epidemics can occur in human populations. The behaviour and sensitivity of the deterministic model with respect to its parameters constitute an important part of Keeling & Gilligan's considerations. Their sensitivity analysis is essentially based on the properties of the linearization of the model at a single reference value for the parameters, which is claimed to represent the behaviour of the model over a large region in parameter space. We explore the underlying basis for the Keeling & Gilligan analysis and the consequences of their conclusions.

Adjoint-based analysis (Lanczos 1996; Marchuk *et al.* 1996) is a classic way to obtain derivatives of a quantity of interest computed from the solution of a differential equation with respect to parameters in the model. In recent work (Estep & Neckels 2006, 2007), we used the derivative information obtained from the adjoint problem to develop computationally efficient sensitivity analysis tools. In this paper, we apply adjoint-based analysis and the sensitivity analysis tools from Estep & Neckels (2006, 2007) to Keeling & Gilligan's model. We determine the sensitivity of a specific quantity of interest to model parameters, determine a region over

which linearization at a given parameter value is valid, develop a global picture of the output surface and search for maxima and minima in the quantity of interest in a given region in the parameter space.

1.1. A model of plague

We consider the Keeling & Gilligan (2000) SIRNF model that describes the dynamics of outbreaks of plague caused by the bacterium *Yersinia pestis* in populations of rats and humans. We follow their description rather closely. They begin with an SIR model for the rat population that gives the changes in the number of susceptible (S_R), infectious (I_R) and resistant (R_R) rats

$$\begin{cases} \dot{S}_R = r_R S_R \left(1 - \frac{T_R}{K_R}\right) + r_R R_R (1 - p) - d_R S_R \\ \quad - \beta_R \frac{S_R}{T_R} F(1 - \exp(-a T_R)), \\ \dot{I}_R = \beta_R \frac{S_R}{T_R} F(1 - \exp(-a T_R)) - (d_R + m_R) I_R, \\ \dot{R}_R = r_R R_R \left(p - \frac{T_R}{K_R}\right) + m_R g_R I_R - d_R R_R, \end{cases} \quad (1.1)$$

where $\dot{\cdot}$ denotes the time derivative and $T_R = S_R + I_R + R_R$ is the total size of the rat population. The net reproduction rate of susceptible and infected rats is

*Author and address for correspondence: Department of Mathematics, Department of Statistics, Colorado State University, Fort Collins, CO 80523, USA (estep@math.colostate.edu).

Table 1. Reference values for the parameters, derived from experiments or field observations or otherwise set within biologically realistic bounds.

parameter	significance	value
$\mu_1 = r_R^*$	net rat reproductive rate	5 yr^{-1}
$\mu_2 = p^*$	proportion of offspring that inherit the disease	0.975
$\mu_3 = K_R$	rat carrying capacity	2500 km^{-2}
$\mu_4 = d_R^*$	natural rat death rate	0.2 yr^{-1}
$\mu_5 = \beta_R^*$	transmission rate to rats from fleas	4.7 yr^{-1}
$\mu_6 = m_R$	rate that rats leave the infected class	20 yr^{-1}
$\mu_7 = g_R^*$	fraction of rats that become resistant	0.02
$\mu_8 = a$	searching efficiency of fleas	4×10^{-3}
$\mu_9 = r_F$	net flea reproductive rate	20 yr^{-1}
$\mu_{10} = d_F^*$	natural flea death rates	10 yr^{-1}
$\mu_{11} = K_F^*$	flea carrying capacity	6.57

given by r_R with carrying capacity K_R . The proportion of offspring that inherit resistance to the disease from their parents is given by p . The natural death rate of all rats is d_R . The number of free infected fleas looking for a host is F . The infection term, $(1 - \exp(-aT_R))$, corresponds to infected fleas randomly searching for a new rat host for some given time period (Nicholson & Bailey 1935). If they find a host and it is susceptible, then the rat becomes infected with a given probability. Thus, β_R is the transmission rate from fleas to rats and a is a measure of the searching efficiency of the fleas. Rats leave the infected class at a rate of m_R and a fraction g_R of these survive to become resistant; the remainder die and release their infected fleas back into the environment.

The flea population dynamics are modelled by the flea index N , which is the average number of fleas living on a rat (each rat is assumed to have the same number of fleas) and the number of free infectious fleas F that are searching for a new host

$$\begin{cases} \dot{N} = r_F N \left(1 - \frac{N}{K_F}\right) + \frac{d_F}{T_R} F (1 - \exp(-aT_R)), \\ \dot{F} = (d_R + m_R(1 - g_R))I_R N - d_F F. \end{cases} \quad (1.2)$$

The net reproduction rate of the fleas is given by r_F with a carrying capacity of K_F , while the death rate is d_F .

The human population dynamics do not affect the number of new human cases since a subsequent transmission between humans is rare (Gage & Kosoy 2005). The potential for human cases is directly related to the number of infected fleas that fail to find a rat host and instead choose a human host

$$\mathcal{F}_I = F \exp(-aT_R). \quad (1.3)$$

The first *quantity of interest*, $\mathcal{F}_I$ (labelled λ_H in Keeling & Gilligan 2000), represents the upper bound or the *potential* force of infection to humans.

Keeling & Gilligan used a stochastic model derived from (1.1) and (1.2) to explain how plague can lay dormant but then break out into an epidemic in the human population. Their explanation depends on consideration of plots of $\mathcal{F}_I$ versus F and T_R as well as plots of $\mathcal{F}_I$ and I_R/m_R versus time. These show that their model predicts that the highest potential for a human epidemic occurs after a major epizootic in the rat population, which leaves a large number of infected fleas looking for a host. This is discussed in terms of the reproductive ratio of the disease, $R_0 = (\beta_R K_F / d_F) (1 - \exp(-aK_R))$.

1.2. Four assertions

Keeling & Gilligan made assertions about the behaviour of the model (1.1) and (1.2) over a large region in the parameter space centred at specific reference values (table 1).

Our analysis addresses four assertions that depend on claims about the sensitivity of the model with respect to the parameters. As a measure of sensitivity of an output variable V with respect to a parameter λ , Keeling & Gilligan used $S = \mu / V(\mu) \times \partial V / \partial \lambda$, where μ is the reference value. The four assertions are as follows.

- (i) From multiple simulations, we note that, even when parameters are changed by a factor of 2, the essential pattern of sensitivity ... remains, showing that $S = (\mu / V(\mu))(\partial V / \partial \lambda)$ is a robust measure of the effects of parameter change.
- (ii) Only K_F , the carrying capacity of fleas per rat, has any effect (on $\mathcal{F}_I$ and I_R/m_R) that is much stronger than linear, although r_R , $1 - p$, K_R , and a all have effects on the number of rat or human cases that are close to linear.
- (iii) ... it is sufficient to consider whether results are robust to changes in the parameter a .
- (iv) The basic reproductive ratio of the disease, R_0 , is related to the parameters via aK_R in the following way: $R_0 = (\beta_R K_F / d_F)(1 - \exp(-aK_R))$. From standard theory, rat epizootics can occur when $R_0 > 1$, which corresponds to $aK_R > 0.39$. It can also be seen that a large human outbreak is possible when $0.5 < aK_R < 20$.

These claims depend on the underlying fact that the linearization of the model with respect to the parameters at the reference values describes the behaviour of the model over a large region in the parameter space. We point out that there is no *a priori* reason to believe that Taylor's theorem can be used in this way.

1.3. A quick overview of sensitivity analysis

Describing response to variations in model input, e.g. initial and boundary data and parameters, is key to understanding a mathematical model of a complex system. Inputs into a model are subject to many sources of uncertainty, variability and measurement error. The model itself may be subject to uncertainty arising from incomplete information or poor understanding of the physical processes and driving forces. Sensitivity

analysis is the organized study of the way in which the output of a model responds to different sources and kinds of variations in the input into the model and in the model itself.

We consider the problem of computing a quantity of interest

$$q(\mathbf{x}(\boldsymbol{\lambda})) = q(\boldsymbol{\lambda}) = \int_0^T \langle \mathbf{x}(s; \boldsymbol{\lambda}), \boldsymbol{\psi}(s) \rangle ds, \quad (1.4)$$

where $\boldsymbol{\psi}$ is a function of time corresponding to the desired information and $\mathbf{x} \in \mathbb{R}^n$ solves the initial value problem

$$\begin{cases} \dot{\mathbf{x}}(t; \boldsymbol{\lambda}) = \mathbf{f}(\mathbf{x}(t; \boldsymbol{\lambda}); \boldsymbol{\lambda}_1), & t > 0, \\ \mathbf{x}(0; \boldsymbol{\lambda}) = \boldsymbol{\lambda}_0. \end{cases} \quad (1.5)$$

Here, $\mathbf{f} : \mathbb{R}^{n+p} \rightarrow \mathbb{R}^n$ is the model and $\boldsymbol{\lambda} = (\boldsymbol{\lambda}_1, \boldsymbol{\lambda}_0)^\top \in \mathbb{R}^d$ ($d = p + n$) are the parameters, with $\boldsymbol{\lambda}_1 \in \mathbb{R}^p$ representing model parameters and $\boldsymbol{\lambda}_0 \in \mathbb{R}^n$ the initial conditions. The quantity of interest could be a statistical quantity, values of the solution at a given time or a more complicated function determined from the solution values.

Two fundamental tools for sensitivity analysis are *density estimation* and computing partial derivatives with respect to the parameters. Density estimation is concerned with computing the probability distribution associated with the quantity of interest $q(\mathbf{x}(\boldsymbol{\lambda}))$ computed from a solution of a model, where the model inputs, e.g. parameters $\boldsymbol{\lambda}$, are assumed to be random variables with associated distributions. Density estimation describes how the model output varies as the parameters vary over the *entire* parameter space. Parametric density estimation is concerned with the case in which the output distribution is known, in which case the goal is to determine best values for the parameters defining the distribution. Non-parametric density estimation is used when the output distribution is unknown or not recognizable, as often happens, e.g. for nonlinear models. In this case, random samples are drawn from the parameter space, the model is solved and the quantity of interest is computed for each sample; finally, a histogram is computed from the values to form an approximate distribution. Various smoothing techniques are often used to reduce the effects of the discrete binning associated with a histogram.

There are two classic approaches to computing partial derivatives of model output with respect to parameters in the model. In forward linearization analysis (Caswell 2007), equations for the partial derivatives of the solution of a model with respect to parameters are determined by differentiating the model and using the Chain Rule

$$\frac{\partial}{\partial t} \left(\frac{\partial \mathbf{x}}{\partial \boldsymbol{\lambda}_1} \right) = D_{\mathbf{x}} \mathbf{f}(\mathbf{x}(t); \boldsymbol{\lambda}_1) \frac{\partial \mathbf{x}}{\partial \boldsymbol{\lambda}_1} + D_{\boldsymbol{\lambda}_1} \mathbf{f}(\mathbf{x}(t); \boldsymbol{\lambda}_1),$$

where $\partial \mathbf{x} / \partial \boldsymbol{\lambda}_1$ is the gradient of $\mathbf{x}$ with respect to parameter $\boldsymbol{\lambda}_1$, and $D_{\mathbf{x}} \mathbf{f}$ and $D_{\boldsymbol{\lambda}_1} \mathbf{f}$ are the Jacobians of $\mathbf{f}$ with respect to $\mathbf{x}$ and $\boldsymbol{\lambda}_1$, respectively. The resulting (large) system of equations for the solution and its partial derivatives is integrated.

The second classic approach uses the mathematical notions of duality and adjoint operators (Lanczos 1996; Marchuk et al. 1996; Estep & Newton in preparation).

This approach is not as easy to describe as forward linearization analysis. The most familiar example is provided by the technique of Green's functions in differential equations (Lanczos 1996). In the adjoint-based approach, a specific adjoint problem

$$\begin{cases} -\dot{\boldsymbol{\phi}}(t) - D_{\mathbf{x}} \mathbf{f}(\mathbf{x}(t); \boldsymbol{\lambda}_1)^\top \boldsymbol{\phi}(t) = \boldsymbol{\psi}(t), & T \geq t \geq 0, \\ \boldsymbol{\phi}(T) = \mathbf{0} \end{cases} \quad (1.6)$$

is defined for each quantity of interest $q(\boldsymbol{\lambda})$ to be computed from the model solution. This adjoint problem yields partial derivatives of the quantity of interest with respect to model parameters relatively cheaply (see appendix A).

Asymptotically, values computed from the two approaches are equal. There are circumstances under which each approach is preferred. The forward linearization technique yields values of partial derivatives of the solution at *all* points in time. From this, one can compute partial derivatives of particular quantities of interest using the Chain Rule. On the other hand, the augmented system is very large and expensive to solve. If the values of the partial derivatives at all points in time are not required, then much of the computational work is wasted. In case that only partial derivatives of particular quantities of interest are desired, the adjoint-based technique is computationally more efficient.

In a recent work, Estep and co-workers have developed new computational tools for sensitivity analysis using the adjoint equation. This work uses the fact that partial derivatives of a quantity of interest with respect to parameters are computed relatively cheaply by solving the adjoint equation. In Estep & Neckels (2006, 2007), efficient numerical methods for density estimation have been developed. The derivative information provided by the adjoint equation is used to either create a higher order approximation or derive adaptive sampling. This adaptive sampling is not probability based. In appendix A, we describe this work briefly.

The computations in this paper were performed with an earlier version of globally accurate adaptive sampling package (GAASP; Estep et al. 2006). GAASP provides computational tools for solving ordinary differential equations with quantitative estimates of numerical accuracy and tools for analysing the parameter and data sensitivity of differential equation models.

2. ANALYSIS OF THE MODEL

We use adjoint-based sensitivity analysis and the new computational tools described in appendix A to analyse Keeling & Gilligan's model in several ways.

- We find a (relatively small) region in the parameter space over which the linearization at the reference value provides a reasonably accurate description of the behaviour of the model.
- We construct a piecewise constant approximation of the quantity of interest over a large region in the parameter space that is close in size to the region

Figure 1. Plots $\mathcal{F}_I$ (solid line) and I_R/m_R (dashed line) versus time, where $\mathcal{F}_I$ is the force of infection to humans and I_R/m_R is the ratio of the number of infected rats relative to the rate of recovery. (a) Reference parameter values with $a=0.004$ (a is the search area of fleas). (b) Values on the ‘upper’ boundary of $\mathcal{R}_L$ with $a=0.006$ and $p=1.0$ (p is the probability that rats inherit resistance). (c) Values in one corner of $\mathcal{R}_S$ with $a=0.004125$. Note that the behaviour at this point matches that at the reference value shown in the plot on the left very closely.

Figure 2. Output distributions of $\partial \mathcal{F}_I / \partial \lambda_i$, $i=4, 5, 6, 7$, for a uniform input distribution on $\mathcal{R}_L$. Variance: (a) 209, (b) 3285, (c) 0.19268 and (d) 48 908.

over which Keeling & Gilligan wanted to draw conclusions about the model behaviour.

- We implement a method of steepest ascent and descent to find the local maxima and minima in a large region in the parameter space.

We then draw several conclusions about the model sensitivity and examine the validity of assertions (i)–(iv).

We make some initial observations. First, Keeling & Gilligan stated that their conclusions hold for parameter values that vary by as much as a factor of 2. This is imprecise since several parameters have physical limitations. Instead, we consider a large rectangle $\mathcal{R}_L$ with sides $[\mu_i/2, 3\mu_i/2]$ for each i (if either endpoint goes past a physical bound, we use the physical bound instead). Even a region of that size allows a wide range of behaviours (figure 1). Second, $\mathcal{F}_I$ depends on the final time T . Keeling & Gilligan used $T=100$. However, (1.1) and (1.2) are used to construct a stochastic model that, Keeling & Gilligan noted, has a different long-time behaviour. Indeed, it is the transient behaviour of (1.1) and (1.2) which is relevant to the stochastic model. Therefore, we use a smaller time, $T=10$. This affects the conclusions quantitatively but not qualitatively.

2.1. Linearization at the reference value

Sampling $\mathcal{F}_I$ over the large rectangle $\mathcal{R}_L$ shows that the linearization at the reference point μ does not describe the behaviour of the model over $\mathcal{R}_L$. For

example, if we use a uniform distribution on $\mathcal{R}_L$, the partial derivatives $\partial \mathcal{F}_I / \partial \lambda_i$ have approximately normal distributions centred at 0 for $i=4, 5, 6, 7$ but have variances 210, 3300, 0.19, 49 000, respectively (figure 2). This reflects the fact that, when we sample over the entire region $\mathcal{R}_L$, varying more than one parameter at a time results in a more complicated behaviour than that observed when varying each parameter individually and that $\mathcal{F}_I$ depends non-linearly on some parameters. We provide evidence in figure 3. In addition, it demonstrates that there are enormous scale differences in the responses to the various parameters.

Since the quantity of interest depends smoothly on the parameters, Taylor’s theorem shows that linearization at the reference value provides a good approximation on all *sufficiently small* rectangles centred at the reference value. The goal is to try to determine the largest possible rectangle on which linearization produces a reasonable approximation. To find such a region, we successively refine $\mathcal{R}_L$ one-parameter dimension at a time until we find a smaller rectangle, $\mathcal{R}_S$, over which linearization appears to be valid.

One way to do this would be to assume a uniform distribution on the parameters, then, holding all parameters fixed but one, compute a least-squares linear approximation to the observed values of the output as that one parameter varies and then checking goodness of fit using the R^2 -statistic. We can decrease the size of the rectangle in that parameter dimension until we achieve a reasonable fit.

Figure 3. (a) Plot of $\mathcal{F}_I$ versus λ_4 when λ_4 alone is varied and when both λ_4 and λ_7 are varied simultaneously, demonstrating that varying more than one parameter complicates the response (dots, λ_4 and λ_7 vary; pluses, λ_4 alone varies). (b) Plot of $\mathcal{F}_I$ versus λ_8 holding the other parameters fixed, which demonstrates a nonlinear dependence. Parameter values are taken from $\mathcal{R}_L$.

However, this approach does not provide a way to deal with scale differences in the parameters. We wish not only to determine a relatively large region on which linearization is valid but also to determine a region over which the *relative* variation in parameters is roughly the same. Our criteria for deciding whether linearization is valid are based on the observation that a one-dimensional linear transformation maps a normal input distribution to another normal distribution. We fix all the parameters at the reference values except one and place a normal distribution on the one parameter that is allowed to vary. The normal distribution is centred at the reference value and the variance is computed so that the range for the parameter in the current rectangle represents a 95% CI. The normal is then truncated to the current rectangle. We then compute the corresponding output distribution of $\mathcal{F}_I$ using a kernel density estimator based on (A 5). If the output distribution is not approximately normal (which can be checked crudely using a least-squares line fit and R^2 statistic after changing variables), we refine the rectangle by halving the size in the dimension of the parameter in question, and try again. Once we determine a range of valid linearization for a parameter, we proceed to the next parameter. After working through all the parameters, we obtain a smaller rectangle, $\mathcal{R}_S$.

While $\mathcal{R}_S$ is a region on which the linearization of $\mathcal{F}_I$ provides an accurate approximation, the variances of the distributions of $\mathcal{F}_I$ corresponding to the distributions on the different parameters on $\mathcal{R}_S$ still vary over a very large range. We refine $\mathcal{R}_S$ further in order to reduce the variances corresponding to the different parameters to all having relative size 1, where we scale each output distribution corresponding to each parameter by the size of the reference value for that parameter. We stop after 16 iterations to obtain $\mathcal{R}_S$ with side lengths $[31\mu_i/32, 33\mu_i/32]$. Further refinement did not lower the variances for λ_2 and λ_8 . However, sampling on $\mathcal{R}_S$ suggests that the time behaviour of $\mathcal{F}_I$ and I_R/m_R is relatively constant over $\mathcal{R}_S$ (figure 1). On the other hand, if we increase the dimensions of $\mathcal{R}_S$ by 50% in each dimension, this is no longer true.

2.2. Adaptive sampling and dimension reduction

Recall that, in the adaptive sampling algorithm, we add points in the parameter direction in which the partial derivative of $\mathcal{F}_I$ is the largest. As the sampling proceeds, the number of sample points added in each parameter direction indicates the sensitivity of $\mathcal{F}_I$ with respect to the parameters. In figure 4, we plot the value of the parameters at the centre of each division in fast adaptive parameter sampling (FAPS) against the number of samples used to create the piecewise constant approximation at each iteration. This shows that $\mathcal{F}_I$ is highly sensitive to λ_1 , λ_3 , λ_5 , λ_{10} and λ_{11} and not very sensitive to λ_2 , λ_4 , λ_6 , λ_7 , λ_8 and λ_9 on $\mathcal{R}_L$. We emphasize that the same conclusion can be made from a FAPS computation with just 50 sample points.

If $F(x, \omega)$ is a target distribution and $\hat{F}_n(x)$ is an approximation to $F(x)$ computed using n samples, the standard Kolmogorov–Smirnov (K–S) statistic

$$D_n(\omega) = \sup_x |\hat{F}_n(x, \omega) - F(x, \omega)|$$

is used to check the accuracy. In figure 5, we plot the K–S statistic for FAPS versus the number of sample points. By way of comparison, we plot the K–S statistic for the error of the mean of many densities determined by computing histograms of the Monte Carlo samples drawn from uniform random sampling of the parameter space along with the asymptotic upper bound given by the law of the iterated logarithm (Estep & Neckels 2006). Note that the error of a single Monte Carlo density has a much larger variance than that of the average of many Monte Carlo computations, as can be seen in the plot.

For a small number of sample points, FAPS converges at the same rate as the *average* of a number of Monte Carlo computations as well as the theoretical bound. Restricting the addition of sample points to the directions with the largest derivatives leads to the staircase appearance of the plots and the levelling out of the statistic for moderate sample sizes.

In figure 5, we also plot the K–S statistic for FAPS when λ_2 , λ_4 , λ_6 , λ_7 , λ_8 and λ_9 are held fixed at the reference value. This confirms the observation from

Figure 4. Plots of parameter values versus number of points used in FAPS. This indicates that $\mathcal{F}_I$ is sensitive to $\lambda_1 = r_R$, $\lambda_3 = K_R$, $\lambda_5 = \beta_R$, $\lambda_{10} = d_F$ and $\lambda_{11} = K_F$ over $\mathcal{R}_L$. Keeling & Gilligan predicted the most sensitive dependence on $\lambda_1 = r_R$, $\lambda_2 = p$, $\lambda_3 = K_R$, $\lambda_8 = a$ and $\lambda_{11} = K_F$. (a) λ_1 , (b) λ_2 , (c) λ_3 , (d) λ_4 , (e) λ_5 , (f) λ_6 , (g) λ_7 , (h) λ_8 , (i) λ_9 , (j) λ_{10} and (k) λ_{11} .

Figure 5. Plots of the K-S statistic. (a) FAPS results, the upper bound given by the law of the iterated logarithm, and the mean of several Monte Carlo samples. Solid line, FAPS; dashed line, mean of Monte Carlo results; dotted line, theoretical bound. (b) The individual K-S statistic for the Monte Carlo computations. Dotted line, Monte Carlo results; solid line, mean of Monte Carlo results. (c) Results for FAPS when λ_2 , λ_4 , λ_6 , λ_7 , λ_8 and λ_9 are held fixed at the reference value. Curves described are the same as in (a).

figure 4 that $\mathcal{F}_I$ is sensitive to the parameters λ_1 , λ_3 , λ_5 , λ_{10} and λ_{11} , and we can safely reduce the dimension of the model by considering variations in these parameters alone.

2.3. Searching for extreme values in parameter space

As part of predicting human epidemics, Keeling & Gilligan examined R_0 at extremal values of a and K_R while fixing the other parameters at the reference values. There is a very narrow window, $0.39 < aK_R < 0.5$, for which plague can affect the rat population significantly without causing a human epidemic. Their approach is reasonable if the model is insensitive to parameter changes as they claim, given that $\partial R_0 / \partial a = \partial R_0 / \partial \lambda_8$ is very large at the reference value. However, we have

seen that the model is sensitive to five parameters and also to simultaneous parameter changes.

Alternatively, we search for local maxima and minima in the parameter domain using a gradient search algorithm. Starting at a point, we can take a small step in the direction of the gradient (or minus the gradient) and recompute new values before stepping again. In this way, we create a path of linked sample points along which we eventually find either an extremal value where the gradient is zero or a point where the path leaves the parameter domain. The final stopping point depends on the initial point as, in general, we expect to find many local extrema.

We use this in two ways. First, we start at the reference value and conduct searches on both $\mathcal{R}_S$ and on a larger rectangle $\mathcal{R}_L$ with side lengths $[\mu_i/2, 2\mu_i]$. Keeling & Gilligan's claimed amounts to stating that any search starting from the reference value for the

parameters should leave $\tilde{\mathcal{R}}_L$ along the $\lambda_8 = a$ axis. Leaving the reference value, the first few steps of the gradient search are roughly parallel to the $\lambda_8 = a$ axis, as $\partial R_0 / \partial a$ is the dominant component in the gradient of R_0 near the reference value. If $\partial R_0 / \partial a$ remains the dominant component in the gradient of R_0 at all points, then the gradient search will always move roughly parallel to the $\lambda_8 = a$ axis. If, however, R_0 becomes more sensitive to changes in other parameters away from the reference value, the gradient search will change direction.

Second, we choose five points near the reference value and follow those paths to see if they stay close to the reference trajectory.

On $\mathcal{R}_S$, the searches starting from all six initial points exit the rectangle on the face perpendicular to the λ_8 axis. However, the exit points vary greatly over this face because variations in the other parameters have an important effect. This affects the gradient searches as they move in the directions that are not roughly parallel to the λ_8 axis. Therefore, Keeling & Gilligan's claim is partially correct on $\mathcal{R}_S$. On the larger rectangle $\tilde{\mathcal{R}}_L$, the searches starting at the reference value leave the larger rectangle $\tilde{\mathcal{R}}_L$ on the lower face perpendicular to λ_7 and the upper face perpendicular to λ_2 . In other words, the gradient search does not follow the λ_8 axis as derivatives with respect to other parameters come to dominate the gradient. The searches that begin from the five other starting points all reach different local minima inside $\tilde{\mathcal{R}}_L$. On the other hand, all of the searches corresponding to following the negative gradient leave on the face perpendicular to λ_2 , though at six different points.

3. CONCLUSIONS

We begin by addressing Keeling & Gilligan's assertions (i)–(iv).

- (i) Regarding assertion (i), the linearization at the reference values simply does not represent the behaviour of the nonlinear model over the large rectangle $\mathcal{R}_L$. In particular, there is an enormous variance in the sensitivity on $\mathcal{R}_L$; simultaneously varying multiple parameters on $\mathcal{R}_L$ leads to unpredictable results and there are some definite nonlinear dependencies. The linearization at the reference values does accurately represent the behaviour of the model over a much smaller rectangle $\mathcal{R}_S$.
- (ii) Regarding assertion (ii), we find that the quantity of interest $\mathcal{F}_I$ is most sensitive to variations in the parameters $\lambda_1 = r_R$, $\lambda_3 = K_R$, $\lambda_5 = \beta_R$, $\lambda_{10} = d_F$ and $\lambda_{11} = K_F$ in contrast to the findings of Keeling & Gilligan, which predict the most sensitive dependence on $\lambda_1 = r_R$, $\lambda_2 = p$, $\lambda_3 = K_R$, $\lambda_8 = a$ and $\lambda_{11} = K_F$. That the model is sensitive to three parameters ($\lambda_1 = r_R$, $\lambda_3 = K_R$ and $\lambda_{11} = K_F$) under both analyses is no surprise, as these directly control growth rates and sizes of the rat and flea populations, and thus have the greatest potential for altering disease dynamics. It is also not surprising that FAPS found

sensitivity to $\lambda_5 = \beta_R$, the transmission rate from fleas to rats, and $\lambda_{10} = d_F$, the death rate of fleas. Both of these parameters influence the number of infected fleas and thus $\mathcal{F}_I$. A recent model of bubonic plague dynamics by Webb *et al.* (2006) is also found to be sensitive to transmission rates and infectious periods.

- (iii) Regarding assertion (iii), we cannot predict the behaviour of the model, even on the smallest rectangle $\mathcal{R}_S$, by fixing all the parameters except for a at the reference values and considering the effect of varying a . Moreover, there is no apparent correlation between the direction of the gradient at the reference point and nearby points and the location of extremal values along paths given by the method of steepest ascent and descent. We conclude that the model is sensitive to a number of parameters in addition to a , and that in some regions of the parameter space it is more sensitive to parameters other than a .
- (iv) Regarding assertion (iv), when predicting the behaviour of the model for the values of aK_R , we have to consider the range for the other parameter values. Thus, predicting plague outbreaks in rats that could spill over into human populations depends upon more parameters than just a , the search area by fleas and K_R , the carrying capacity for the rat population (Gage & Kosoy 2005).

As a practical consequence of our analysis of the Keeling & Gilligan model, we conclude that the general danger of large-scale outbreaks of plague in humans predicted by the model was overstated. The reason is that the force of infection from rats and their fleas to humans, $\mathcal{F}_I$, depends strongly on more parameters than just the rat carrying capacity, K_R , and searching efficiency of fleas, a . In particular, Keeling & Gilligan concluded that the reproductive ratio of disease, R_0 , is predicted to be greater than 1 for even fairly low carrying capacities of rat populations provided the searching efficiency of fleas a is high. On the other hand, our analysis shows that, in large parts of the parameter space, the model is relatively insensitive to a and that $R_0 > 1$ and human risk depends more on large rat populations than on flea behaviour. Large populations of rodents still provide the risk of human outbreaks, but details of epidemiology within the rat and flea plague reservoir, especially flea survival and transmission rates, must also be specified (see Gage & Kosoy 2005).

We emphasize that our observations do not necessarily contradict the scientific conclusions of Keeling & Gilligan (2000) because their main results are based on a stochastic model, not (1.1) and (1.2). However, our results do raise questions about the connection between the stochastic model analysed by Keeling & Gilligan and the underlying deterministic model that was used to create the stochastic model. The properties of the stochastic and deterministic models are clearly different.

Our analysis emphasizes the fact that nonlinear models like (1.1) and (1.2), even when composed of well-understood mechanisms, can exhibit complicated nonlinear sensitivity to parameters that are not

predictable from linearization at a single point in the parameter space. Keeling & Gilligan's assertions about the model are accurate on a relatively small region in the parameter space surrounding the reference value, but do not hold true in a larger region in the parameter space. Whether this is relevant or not depends on the region in the parameter space on which predictions are needed and how the model results are used. However, in any situation, it is scientifically important to determine the limitations affecting the accuracy of predictions of model behaviour.

When analysing the sensitivity of a model with respect to parameters, the ability to compute gradient information at sample points is invaluable. In this paper, we have used adjoint-based analysis tools to carry out a detailed sensitivity analysis of a nonlinear model of the plague over a large parameter domain. We have also demonstrated that the FAPS and high order parameter sampling numerical methods based on the adjoint problem, variational analysis and the generalized Green's function provide a very efficient way to perform non-parametric density estimation.

The research activities of M.B. were partially supported by the National Science Foundation (DGE-0221595003). The research activities of D.N. were partially supported by the National Science Foundation (DGE-0221595003, MSPA-CSE-0434354). The research activities of M.F.A. were partially supported by the National Science Foundation (EF-0327052 and DEB0217631). D.E.'s work was supported in part by the Department of Energy (DE-FG02-04ER25620, DE-FG02-05ER25699, DE-FC02-07ER54909), the National Aeronautics and Space Administration (NNG04GH63G), the National Science Foundation (DMS-0107832, DMS-0715135, DGE-0221595003, MSPA-CSE-0434354, ECCS-0700559), Idaho National Laboratory (00069249), the Sandia Corporation (PO299784) and the United States Department of Agriculture (58-5402-3-306).

APPENDIX A. THE ANALYTIC TECHNIQUE

To analyse parameter sensitivity, we employ an approach (Eriksson *et al.* 1995, 1996; Estep 1995; Estep *et al.* 2000; Estep & Neckels 2006, 2007) that uses variational analysis involving the generalized Green's function, which solves the (linearized) adjoint problem (1.6) with data specific to the information (1.4) to be computed. Green's function describes how local variation in the model, parameters and data propagate into the computed information. The ability to tailor the analysis to a specific quantity of interest, such as $\mathcal{F}_1$ (1.3), is very important, e.g. the solution of (1.1) and (1.2) is much more sensitive to changes in parameter values than $\mathcal{F}_1$.

We give minimal details about the new approach (see Estep & Neckels (2006, 2007) for information). To model uncertainty and variation in the parameters, we assume that $\lambda = \lambda(\omega)$ is a random vector on a probability space $(\Omega, \mathcal{B}, P)$. This leads naturally to the *density estimation* problem for $q(\omega) = q(\lambda(\omega))$, i.e. approximate the cumulative distribution function $F_q(x)$ of q given an arbitrary input distribution for $\lambda(\omega)$. Using the cumulative distribution function, we can

compute the density as well as any desired statistic, e.g. means and moments.

As a first step, we solve for the deterministic solution $\mathbf{y}(t)$ corresponding to a reference value $\boldsymbol{\mu} = (\boldsymbol{\mu}_1, \boldsymbol{\mu}_0)^\top$ for the parameters and compute the generalized Green's function solving the adjoint problem to the linearized equation

$$\begin{cases} -\dot{\boldsymbol{\phi}}(t) - D_{\mathbf{x}}\mathbf{f}(\mathbf{y}(t); \boldsymbol{\mu}_1)^\top(t)\boldsymbol{\phi}(t) = \boldsymbol{\psi}(t), & T \geq t \geq 0, \\ \boldsymbol{\phi}(T) = \mathbf{0}. \end{cases} \quad (\text{A } 1)$$

A modification (Estep & Neckels 2006, 2007) of the standard variational argument for Green's functions yields the following.

Theorem A.1. *If $\mathbf{f}$ is smooth, then*

$$\begin{aligned} q(\lambda) \approx & \int_0^T \langle \mathbf{y}, \boldsymbol{\psi} \rangle \, ds + \langle \lambda_0 - \boldsymbol{\mu}_0, \boldsymbol{\phi}(0) \rangle \\ & + \int_0^T \langle D_{\lambda_1}\mathbf{f}(\mathbf{y}; \boldsymbol{\mu}_1)(\lambda_1 - \boldsymbol{\mu}_1), \boldsymbol{\phi} \rangle \, ds, \end{aligned} \quad (\text{A } 2)$$

and

$$\left. \begin{aligned} \nabla q(\boldsymbol{\mu})(\mathbf{v}) &= \langle \mathbf{v}_0, \boldsymbol{\phi}(0) \rangle + \int_0^T \langle D_{\lambda_1}\mathbf{f}(\mathbf{y}; \boldsymbol{\mu}_1)\mathbf{v}_1, \boldsymbol{\phi} \rangle \, ds, \\ \mathbf{v} &= (\mathbf{v}_0, \mathbf{v}_1)^\top \in \mathbb{R}^d. \end{aligned} \right\} \quad (\text{A } 3)$$

The theorem implies that derivatives of the quantity of interest can be computed cheaply by solving a *linear* adjoint problem and computing an integral. Note that the first term on the r.h.s. of (A 3) describes the effect of variations in the initial conditions and the second term describes the effect of variations in the model parameters.

If we use the reference sample $q(\boldsymbol{\mu})$ as an approximation for the value $q(\lambda)$ at a nearby parameter value λ , theorem A.1 yields a computable error estimate

$$q(\lambda) - q(\boldsymbol{\mu}) \approx \langle \nabla q(\boldsymbol{\mu}), (\lambda - \boldsymbol{\mu}) \rangle. \quad (\text{A } 4)$$

In general, one sample point cannot be used to accurately describe the behaviour of the quantity of interest over a large parameter domain. Therefore, we construct a piecewise local approximation by computing a partition of the parameter space $\mathcal{A}$ into N generalized rectangles $\{R_i\}_{i=1}^N$, choosing a sample point $\boldsymbol{\mu}_i$ inside each R_i , and computing corresponding sample values $q(\boldsymbol{\mu}_i)$ for $i=1, \dots, N$. This yields the piecewise constant approximation

$$\tilde{q}(\lambda) = \sum_{i=1}^N q(\boldsymbol{\mu}_i) \chi_{R_i}(\lambda), \quad (\text{A } 5)$$

where χ_{R_i} is 1 on R_i and 0 elsewhere. The error is estimated by a sum of local contributions

$$q(\omega) - \tilde{q}(\omega) \approx \sum_{i=1}^N \int_{R_i} \langle \nabla q(\boldsymbol{\mu}_i), (\mathbf{z} - \boldsymbol{\mu}_i) \rangle \, d\mu_\lambda(\mathbf{z}). \quad (\text{A } 6)$$

We use (A 6) to construct an adaptive sampling procedure called FAPS (Estep & Neckels 2006). If the estimated error for a given piecewise approximation is too large, we choose additional sample points in the

rectangle with the largest local contribution by adding sample points in the parameter dimension with the largest directional derivative. By adding sampling points in one dimension at a time, we avoid sampling in parameters that are not significantly affecting the quantity of interest.

To apply this approach for the Keeling & Gilligan analysis, we must linearize the nonlinear quantity of interest $\mathcal{F}_I$ before using the representation formula. Since

$$\partial_{\lambda_i} \mathcal{F}_I = \partial_{\lambda_i} F \exp(-aT) - aF \exp(-aT) \partial_{\lambda_i} T,$$

for $i \neq 8$, it suffices to find $\partial_{\lambda_i} F$ and $\partial_{\lambda_i} T$, which are linear functionals of the solution corresponding to data $\psi_1 = \delta(t-T)(0, 0, 0, 0, 1)^\top$ and $\psi_2 = \delta(t-T)(1, 1, 1, 0, 0)^\top$, respectively, where $\delta(s)$ is the delta function at s . For $\lambda_8 = a$, we obtain $\partial_{\lambda_8} \mathcal{F}_I$ in a similar fashion. To calculate these derivatives, we simultaneously integrate two adjoint solutions corresponding to ψ_1 and ψ_2 .

REFERENCES

- Caswell, H. 2007 Sensitivity analysis of transient population dynamics. *Ecol. Lett.* **10**, 1–15. (doi:10.1111/j.1461-0248.2006.01001.x)
- Eriksson, K., Estep, D., Hansbo, P. & Johnson, C. 1995 Introduction to adaptive methods for differential equations. *Acta Numer.* **4**, 105–158.
- Eriksson, K., Estep, D., Hansbo, P. & Johnson, C. 1996 *Computational differential equations*. New York, NY: Cambridge University Press.
- Estep, D. 1995 *A posteriori* error bounds and global error control for approximations of ordinary differential equations. *SIAM J. Numer. Anal.* **32**, 1–48. (doi:10.1137/0732001)
- Estep, D. & Neckels, D. 2006 Fast and reliable methods for determining the evolution of uncertain parameters in differential equations. *J. Comput. Phys.* **213**, 530–556. (doi:10.1016/j.jcp.2005.08.024)
- Estep, D. & Neckels, D. 2007 Fast methods for determining the evolution of uncertain parameters in reaction–diffusion equations. *Comput. Meth. Appl. Mech. Eng.* **196**, 3967–3979. (doi:10.1016/j.cma.2006.10.048)
- Estep, D. & Newton, W. In preparation. Sensitivity analysis and density estimation for population models.
- Estep, D., Larson, M. & Williams, R. 2000 Estimating the error of numerical solutions of systems of reaction–diffusion equations. *Mem. Am. Math. Soc.* **146**, pp. viii + 109.
- Estep, D., Mckeown, B., Neckels, D. & Sandelin, J. 2006 GAASP: globally accurate adaptive sensitivity package. Write to estep@math.colostate.edu for information on obtaining GAASP software.
- Gage, K. L. & Kosoy, M. Y. 2005 Natural history of plague: perspectives from more than a century of research. *Annu. Rev. Entomol.* **50**, 505–528. (doi:10.1146/annurev.ento.50.071803.130337)
- Keeling, M. J. & Gilligan, C. A. 2000 Bubonic plague: a metapopulation model of a zoonosis. *Proc. R. Soc. B* **267**, 2219–2230. (doi:10.1098/rspb.2000.1272)
- Lanczos, C. 1996 *Linear differential operators*. Philadelphia, PA: SIAM.
- Marchuk, G., Agoshkov, V. & Shutyaev, V. 1996 *Adjoint equations and perturbation algorithms in nonlinear problems*. New York, NY: CRC Press.
- Nicholson, A. J. & Bailey, V. A. 1935 The balance of animal populations. *Proc. Zool. Soc. Lond.* **3**, 551–598.
- Webb, C., Brooks, C., Gage, K. & Antolin, M. 2006 Classic flea-borne transmission does not drive plague epizootics in prairie dogs. *Proc. Natl Acad. Sci. USA* **103**, 6236–6241. (doi:10.1073/pnas.0510090103)