



A coupled hidden Markov model for disease interactions

Chris Sherlock and Tatiana Xifara,
Lancaster University, UK

Sandra Telfer
University of Aberdeen, UK

and Mike Begon
University of Liverpool, UK

[Received March 2012. Final revision January 2013]

Summary. To investigate interactions between parasite species in a host, a population of field voles was studied longitudinally, with presence or absence of six different parasites measured repeatedly. Although trapping sessions were regular, a different set of voles was caught at each session, leading to incomplete profiles for all subjects. We use a discrete time hidden Markov model for each disease with transition probabilities dependent on covariates via a set of logistic regressions. For each disease the hidden states for each of the other diseases at a given time point form part of the covariate set for the Markov transition probabilities from that time point. This allows us to gauge the influence of each parasite species on the transition probabilities for each of the other parasite species. Inference is performed via a Gibbs sampler, which cycles through each of the diseases, first using an adaptive Metropolis–Hastings step to sample from the conditional posterior of the covariate parameters for that particular disease given the hidden states for all other diseases and then sampling from the hidden states for that disease given the parameters. We find evidence for interactions between several pairs of parasites and of an acquired immune response for two of the parasites.

Keywords: Adaptive Markov chain Monte Carlo sampling; Forward–backward algorithm; Gibbs sampler; Hidden Markov models; Zoonosis

1. Introduction

1.1. *Motivating problem*

In natural populations, animals are likely to be infected by a variety of pathogens, either simultaneously or successively. Interactions between these pathogens, which can be synergistic or antagonistic, can affect infection biology (e.g. the intensity of one or both infections), or host susceptibility to infection, or may impact on the host's morbidity or/and mortality. However, the biological processes that are involved are often too complex to allow clear-cut predictions regarding the outcome of such interactions. To explore potential interactions, a longitudinal study was undertaken by recording the sequences of infection events for different parasites in

Address for correspondence: Tatiana Xifara, Department of Mathematics and Statistics, Fylde College, Lancaster University, Lancaster, LA1 4YF, UK.
E-mail: t.xifara@lancaster.ac.uk

Reuse of this article is permitted in accordance with the terms and conditions set out at <http://wileyonlinelibrary.com/onlineopen#OnlineOpenTerms>.

four spatially distinct populations of field voles (*Microtus agrestis*). The data are records of six pathogens: three species of *Bartonella* bacteria (*B. taylorii*, *B. grahamii* and *B. doshiae*), cowpox virus, the bacterium *Anaplasma phagocytophilum* and the protozoan *Babesia microti*. Aside from their intrinsic interest as a community of pathogens, *Bartonella*, *Anaplasma*, *Babesia* and cowpox virus infections may also be zoonotic: capable of being transmitted from animals to humans and causing disease.

As in most capture–mark recapture studies, a different set of voles was caught at each session leading to incomplete profiles for all subjects. The data set therefore contains many missing observations; for example a profile for a given vole and a given disease from the first to last observation times for that vole might be NPxxPNxP, where x, N and P respectively indicate a missing observation, a negative response and a positive response. Inference on incomplete data in longitudinal and capture–recapture studies is a major problem; for examples see Daniels and Hogan (2008) and Pradel (2005). Previous analyses of our and related data sets (see Telfer *et al.* (2010) and Begon *et al.* (2009)) have examined all pairs of observations for a given vole that occurred exactly 1 lunar month apart and for which the first of the two observations was an N. The influence of each covariate on the probability of contracting a disease is then ascertained through logistic regression. In this paper we offer a more realistic model and a more powerful analysis methodology for investigating the effects of previous infections for each disease on the other diseases. We use a hidden Markov model (HMM) for each disease (Section 2.1) and perform inference via a Gibbs sampler; this allows us to use all of the data set and to infer covariate effects on a given disease, even when these covariates are the (potentially missing or hidden) states of the other five diseases.

1.2. Data

We analyse data collected between March 2005 and March 2007 from field voles in Kielder Forest, which is a man-made forest on the English–Scottish border. The voles were trapped at four grassy clear-cut sites within the forest, with each site at least 3.5 km from the nearest neighbouring site. Individuals were trapped within a 0.3 ha live trapping grid comprising 100 traps set at 5 m intervals, with trapping taking place every 28 days from March to November, and every 56 days from November to March. Begon *et al.* (2009) have provided further details of the study area and the trapping design.

Captured voles were marked with a unique identifying passive transponder tag to be recognized in later captures. At each capture, a 20–30- μ l blood sample was taken for pathogen diagnostic tests. Polymerase chain reaction assays were used to test directly for evidence of infection with *Anaplasma phagocytophilum*, *Babesia microti* and the three *Bartonella* species (see Courtney *et al.* (2004), Bown *et al.* (2008) and Telfer *et al.* (2008)). Antibodies to cowpox virus were detected by immunofluorescence assay (see Chantrey *et al.* (1999)). A brief description of the observed and derived variables is given in Table 1.

After some processing (which is described in detail in Xifara (2012)) our data set contains 4344 captures of 1841 voles. Only voles that have been caught at least twice are directly informative about transition probabilities (see Section 2.1), although voles that have been captured only once still contribute to inference for the initial distribution of each hidden Markov chain (see Section 3.1.1).

The data set contains a substantial fraction of missing data: almost half of the voles are not captured at every lunar month between the first and last times that they were observed. Thus, even for many of the voles that were observed at least twice, not all of the covariates are available, either because the vole was not caught in a given lunar month, or sometimes because the vole

Table 1. Description of variables in the data set and their possible outcomes

<i>Variable</i>	<i>Description</i>
Tag	Unique number that identifies each vole
Site	Identifier for the capture site (four-level factor)
Sex	Male or female
Lm	Capture time point in whole lunar months (1–27; integer)
Weight	Weight in grams rounded to the nearest 0.5 g
Sin	$\sin(2\pi \text{ Lm}/13)$
Cos	$\cos(2\pi \text{ Lm}/13)$
Tay	<i>B. taylorii</i> , N (negative) or P (positive)
Grah	<i>B. grahamii</i> , N (negative) or P (positive)
Dosh	<i>B. doshiae</i> , N (negative) or P (positive)
Cow	Cowpox, N (negative) or P (positive)
Ana	<i>Anaplasma</i> , N (negative) or P (positive)
Bab	<i>Babesia</i> , N (negative) or P (positive)

Table 2. Frequency of missing values per vole as a function of the number of lunar months from the first capture time to the last capture time

<i>Lunar months from first to last capture</i>	<i>Frequencies for the following numbers of values missing:</i>				
	<i>0</i>	<i>1</i>	<i>2</i>	<i>3</i>	<i>≥ 4</i>
0	832	—	—	—	—
1	275	—	—	—	—
2	132	74	—	—	—
3	75	55	15	—	—
4	30	49	33	9	—
5	21	24	34	5	3
6	7	7	33	15	2
7	1	4	25	16	9
>7	0	2	27	12	15

was caught but a given variable was not ascertained. Table 2 shows the frequency of missing values derived from the first cause. The number of additional missing values, where it was not possible to ascertain the status of a particular disease, despite the vole being captured, is given in Table 3, which also shows the frequency of positive (P) and negative (N) records for each disease.

1.3. Statistical challenges

We aim to investigate potential interactions between the six pathogens of the study. In particular, for each disease *d*, we wish to evaluate the way in which the presence or absence of each of the other diseases (and perhaps further information such as whether or not any infection is in its first month) affects the probability that vole contracts *d*. Additionally where applicable we are interested in how other diseases affect the probability of recovery from *d*.

We could model each disease as a two-state discrete time Markov chain, where state 1 corresponds to no disease and state 2 to presence of disease; however, this two-state model imposes

Table 3. Summary information for the six diseases; number of missing values despite the vole being captured and numbers of negative N and positive P responses

Disease	Number of additional missing values	Number of N	Number of P
<i>B. doshiae</i>	46	3583	715
<i>B. grahamii</i>	44	3468	832
<i>B. taylorii</i>	32	3139	1173
<i>Babesia</i>	0	2354	1990
Cowpox	85	1408	2851
<i>Anaplasma</i>	6	4107	231

a very specific structure. For example the length of any infection is geometrically distributed; however, it might be that the probability of remaining infected when a disease is in its first month (an acute phase) is different from that in subsequent months (chronic phase). It has also been found (e.g. Telfer *et al.* (2010)) that acute and chronic phases of a disease d_1 can have different effects on the probability that a vole contracts disease d_2 . A two-state semi-Markov model (see, for example, Guédon (2003)) could account for the first effect, at the expense of extra complexity, but not the second. To represent both the dynamics and the influence of each disease with minimal extra complexity adequately, therefore, in this analysis the dynamics of all except one of the diseases is modelled as a Markov chain with more than two states. Section 2.1 details the model for each disease.

Only knowledge of the presence or absence of the disease is available to us. In general, this equates to knowledge of a subset of the state space in which the true state must lie, but not to the exact state of the chain. For example, for all except one disease, states 2 and 3 both correspond to the presence of the disease. In disease modelling, HMMs arise when the Markov model for disease progression has several stages, or states, but these are not directly observed (e.g. Guihenneuc-Jouyaux *et al.* (2000) and Chadeau-Hyam *et al.* (2010)). Often the relationship between the state of the Markov chain and the observation is stochastic, although in our case no stochasticity is involved, but the state of the Markov chain is nonetheless hidden. Furthermore, observations are only available to us when the vole has been captured. The forward-backward algorithm (see Section 2.3) can be applied to any discrete time HMM with a finite state space and addresses both of these issues.

We consider $D = 6$ diseases, and hence six interacting (or coupled) HMMs. It is possible to consider the coupled Markov chains for each disease together as a single Markov chain on an extended state space. In this case the likelihood function is straightforward to evaluate by using the forward-backward algorithm (see for example Zucchini and MacDonald (2009)) and a Bayesian analysis can then be performed by using Markov chain Monte Carlo (MCMC) methods. In our particular scenario the state spaces have size 4, 4, 4, 3, 3 and 2, which would lead to an extended state space of size $4^3 \times 3^2 \times 2 = 1152$. Since the forward-backward algorithm applied to an HMM with n states takes $O(n^2)$ operations, a naive implementation of the algorithm applied to the extended state space would have a complexity of $O(1152^2)$ compared with a complexity of $O(3 \times 4^2 + 2 \times 3^2 + 2^2)$ for six coupled chains; equivalently, 100000 iterations of an algorithm which deals with each chain separately would take approximately the same central processor unit time as five or six iterations of the single-chain algorithm. In our specific scenario, but certainly not in generality, some of the transition probabilities in each individual chain are 0, and (in our scenario) only 32768 elements of the extended transition matrix would

be non-zero. The use of sparse matrix routines could therefore reduce the efficiency ratio to approximately 468. Such a reduction in computational efficiency would only be justified if the fraction of missing data were very close to 1 so that the mixing of our Gibbs sampler would be extremely slow.

Pradel (2005) analysed capture–recapture data by using an HMM, and incorporation of covariate information within this framework via an appropriate link function is straightforward (see Lachish *et al.* (2011) and Zucchini and MacDonald (2009) (section 8.5.2)). However, the methodology does not allow the use of multiple HMMs nor, therefore, can it use the state of each HMM as a covariate for the other HMMs. We require six HMMs (one for each disease) and we wish to use covariate information such as the time of year and weight of the vole. Furthermore we wish the covariate set for each disease to include the states of the HMMs for the other diseases. For each disease d , we shall represent the probability of each possible state change through a logistic regression. However, some of the covariates, the states of the other HMMs, are unknown. Our solution is a Gibbs sampler which employs the forward–backward algorithm and adaptive random-walk Metropolis steps to sample from the true posterior distribution of all of the HMMs and the covariate parameters jointly.

1.4. Outline

The remainder of this paper is organized as follows. Section 2 describes the model which was used for each disease, gives its likelihood function and outlines the imputation of missing weight values and the other fixed covariate values. The MCMC algorithm is described in Section 3 and we present our results, including the sensitivity study, in Section 4. The paper concludes with a discussion.

The data that are analysed in the paper can be obtained from

<http://www.blackwellpublishing.com/rss>

2. Modelling the hidden and missing data

2.1. Hidden Markov models and notation

HMMs are used when observations are influenced by a Markov process but the state of the Markov process itself cannot be determined exactly from the observations. Usually the relationship between the Markov process and the observation process is stochastic, but (as in our application) this need not be so. For various examples and applications of HMMs see, for example, Zucchini and MacDonald (2009). Purely to simplify our subscript notation we consider each vole to have been first observed at a local (to the vole) time of 1 and last observed at (local) time T . For disease d ($d = 1, \dots, D$), let the state space for the Markov chain be $\mathbb{S}^{[d]}$ and the state space for the observation process be $\mathbb{Y} := \{\mathbb{N}, \mathbb{P}\}$. For a given vole and for disease d , the unobserved Markov chain and the observations are respectively

$$\mathbf{X}^{[d]} := (X_1^{[d]}, X_2^{[d]}, \dots, X_T^{[d]}) \in (\mathbb{S}^{[d]})^T$$

and

$$\mathbf{Y}^{[d]} := (Y_1^{[d]}, Y_2^{[d]}, \dots, Y_T^{[d]}) \in (\mathbb{Y} \cup \{\text{missing}\})^T.$$

Note that $Y_t^{[d]}$ is conditionally independent of $X_1^{[d]}, \dots, X_{t-1}^{[d]}, X_{t+1}^{[d]}, \dots, X_T^{[d]}$ given $X_t^{[d]}$. The observed process $\{Y^{[d]}\} \in \mathbb{Y}$ is related to the state of the hidden process $\{X^{[d]}\}$ by a likelihood vector $\mathbf{I}_t^{[d]}$, which has elements

$$l_{i,i}^{[d]} = \begin{cases} \mathbb{P}(Y^{[d]} = y_i^{[d]} | x_i^{[d]} = i) & \text{if } y_i^{[d]} \in \{\mathbf{N}, \mathbf{P}\}, \\ 1 & \text{if } y_i^{[d]} = \text{missing}. \end{cases}$$

This vector is defined for each disease in Section 2.5.

We take the discrete time interval of each Markov chain to be 1 lunar month. Since trapping sessions in winter took place every 2 lunar months (see Section 1.2) this inevitably leads to missing observations for any vole caught several times over the winter, even if it is caught at every trapping session. For each unknown transition probability (see Section 2.5) we have a logistic regression model; for example the probability $p_{11,t}^{[d]}$ that a given vole will be in state 1 (disease d absent) at time $t + 1$ given that it is in state 1 at time t is given by

$$\text{logit}(p_{11,t}^{[d]}) = (\mathbf{z}_t^{[d]})^T \boldsymbol{\beta}_{11}^{[d]}.$$

Here $\mathbf{z}_t^{[d]}$ is the vector of covariates at time t , which for all models includes the states of the other diseases at time t , $\mathbf{x}_t^{[-d]}$, as well as a deterministic covariate vector $\mathbf{z}_t^{*[d]}$. This deterministic vector was chosen via forward fitting of logistic regression models that were very similar to those of Telfer *et al.* (2010) and Begon *et al.* (2009) (see Section 1.1). However, whereas Telfer *et al.* (2010) and Begon *et al.* (2009) allowed both the current covariates and covariates 1 lunar month into the future to influence the response 1 lunar month into the future, we allow only the current covariates to influence the future response; further details are available in Xifara (2012). For all diseases the deterministic vector consists of a time trend (lunar month Lm as a continuous covariate), a seasonal cycle in the form of \sin and \cos , and sex, weight and site. The covariate vector for cowpox also includes a different trend with lunar month for each site, and for all other diseases it allows for a different seasonal cycle for each sex (see Table 1 for detailed covariate descriptions).

We denote the transition probability matrix from time t to $t + 1$ for disease d by $P^{[d]}(\boldsymbol{\beta}^{[d]}, \mathbf{x}_t^{[-d]}, \mathbf{z}_t^{[d]})$, i.e. $P_{i,j}^{[d]}(\boldsymbol{\beta}^{[d]}, \mathbf{x}_t^{[-d]}, \mathbf{z}_t^{[d]}) = \Pr(X_{t+1}^{[d]} = j | X_t^{[d]} = i)$, and let the initial distribution for the hidden chain be $\boldsymbol{\pi}^{[d]}$. Fig. 1 depicts a simplification of our scenario, where there are just two diseases. Note that the states of all chains at time $t + 1$, $X_{t+1}^{[1]}, \dots, X_{t+1}^{[D]}$, are independent conditional on the states of all chains at time t .

2.2. Likelihood function

We now provide full detail of the likelihood for a given vole. The likelihood for the data is simply the product of these likelihoods over all 1841 voles. Let $\boldsymbol{\beta}^{[1:D]} := (\boldsymbol{\beta}^{[1]}, \boldsymbol{\beta}^{[2]}, \dots, \boldsymbol{\beta}^{[D]})$, $\mathbf{y}^{[1:D]} = (\mathbf{y}^{[1]}, \mathbf{y}^{[2]}, \dots, \mathbf{y}^{[D]})$ and $\mathbf{x}^{[1:D]} = (\mathbf{x}^{[1]}, \mathbf{x}^{[2]}, \dots, \mathbf{x}^{[D]})$. The conditional independence structure leads to a complete-data likelihood of

$$L(\mathbf{y}^{[1:D]}, \mathbf{x}^{[1:D]}; \boldsymbol{\beta}^{[1:D]}, \boldsymbol{\pi}^{[1:D]}) = \prod_{d=1}^D \pi_{x_1^{[d]}}^{[d]} l_{1,x_1^{[d]}}^{[d]} \prod_{t=1}^{T-1} P_{x_t^{[d]}, x_{t+1}^{[d]}}^{[d]}(\boldsymbol{\beta}^{[d]}, \mathbf{x}_t^{[-d]}, \mathbf{z}_t^{[d]}) l_{t+1,x_{t+1}^{[d]}}^{[d]}.$$

The observed data likelihood for the vole is then

$$L(\mathbf{y}^{[1:D]}; \boldsymbol{\beta}^{[1:D]}, \boldsymbol{\pi}^{[1:D]}) = \sum_{\mathbf{x}^{[1:D]} \in \mathbb{S}_*^T} L(\mathbf{y}^{[1:D]}, \mathbf{x}^{[1:D]}; \boldsymbol{\beta}^{[1:D]}, \boldsymbol{\pi}^{[1:D]}), \tag{1}$$

where $\mathbb{S}_* := \mathbb{S}^{[1]} \times \dots \times \mathbb{S}^{[D]}$. For a single chain the summation over the hidden states can be written as a matrix product; this simplification is not possible for coupled chains as the transition matrix for each disease depends on the state of each of the other diseases. Our Bayesian analysis requires prior distributions for $\boldsymbol{\beta}^{[1:D]}$ and $\boldsymbol{\pi}^{[1:D]}$, which are detailed in Section 3.1. The product of the observed data likelihoods over all voles multiplied by the prior distribution for $\boldsymbol{\beta}^{[1:D]}$ and $\boldsymbol{\pi}^{[1:D]}$ gives, up to a constant of proportionality, the joint posterior for $\boldsymbol{\beta}^{[1:D]}$, $\boldsymbol{\pi}^{[1:D]}$ and $\mathbf{x}^{[1:D]}$.




Fig. 1. Directed graph of a realization of two parallel HMMs where $y_t^{[d]}$, $t = 1, \dots, 5$, where present, are the observed values of disease d ($d = 1, 2$) and $X_t^{[d]}$, $t = 1, \dots, 5$, are the states of the hidden Markov chain for d , arising from (unknown) initial distribution $\pi^{[d]}$; the nodes from $\beta^{[d]}$ reflect dependence of the transition matrix on the (unknown) covariate parameters; to simplify the presentation and to focus on the coupled HMM we omit the deterministic covariate vectors $\mathbf{z}_t^{[d]}$.

2.3. The forward-backward algorithm

The forward-backward algorithm that was developed by Baum *et al.* (1970) (see also Zucchini and MacDonald (2009), Scott (2002), Rabiner (1989) and Chib (1996), for example) may be applied to any discrete time HMM with a finite state space and provides us with two useful tools. The *forward recursion* is a computationally efficient algorithm for calculating the likelihood of the observed data, whereas the *backward recursion* provides us with the distribution of each hidden state X_t , given the state at the next time point, x_{t+1} , and all of the observations. Both will form part of the Gibbs sampling scheme that will be described in detail in Section 3.2.

2.4. Other missing covariates

As mentioned in Section 1.3, for many voles not all of the covariates are available. For a given vole, the covariates sex, site and Lm clearly carry over to the missing records. The unobserved disease states will be treated dynamically and will be sampled from the conditional distribution as part of the Gibbs sampling scheme (see Section 3.2). Such sampling could perhaps also be performed for weight. However, here we adopt a simpler approach whereby each missing weight value is imputed once via linear interpolation between the two nearest observed values for that vole. The robustness of inference to other sensible imputed weight values obtained by using a growth model is investigated in Section 4.3.

2.5. Details of the Markov models for individual diseases

The remainder of this section gives a brief description of each disease in the study and describes the HMM that is used to model it. All transition probabilities are time dependent since some of the covariates are time dependent; however, for ease of notation we drop any explicit reference to this time dependence. A more detailed description of the host resources that are required by these parasites and a discussion about host immune responses can be found in Telfer *et al.* (2008).

2.5.1. *Bartonella* species

Bartonella is a genus of bacteria that infects mammals (including humans), which is usually transmitted by arthropods. The species that are investigated here are transmitted by fleas (Bown *et al.*, 2004). We assume that the effect of other diseases and covariates on the probability that a vole will recover from a particular *Bartonella* species after the second (third, fourth etc.) lunar month is the same as for the effect on the probability of recovery after the first month; there are no grounds for assuming otherwise. However, since the majority of *Bartonella* infections last for 1 month and only a few last more than 2 (Birtles *et al.*, 2001; Telfer *et al.*, 2008) the overall probabilities of recovering after the first and second month are likely to be different. Additionally a vole’s chance of contracting a particular *Bartonella* species for the first time might be different from the chance of contracting it again after recovery from it in the past, although, again, there is no reason to assume that the effects of other diseases and covariates on this are likely to be different. This suggests that each *Bartonella* species could be sensibly modelled by using a Markov chain with four states: 1, no infection; 2, new infection; 3, old infection; 4, uninfected but has had a past infection. However, the observed sequence indicates either negative N or positive P status. In particular, an observation of $y_i^{[d]} \equiv N$ corresponds to a hidden process of $X_i^{[d]} = 1$ or $X_i^{[d]} = 4$ with likelihood vector $\mathbf{l}_i^{[d]} = (1, 0, 0, 1)$, and an observation of $y_i^{[d]} \equiv P$ corresponds to $X_i^{[d]} = 2$ or $X_i^{[d]} = 3$ with $\mathbf{l}_i^{[d]} = (0, 1, 1, 0)$. The time inhomogeneous transition probability matrix from time t to time $t + 1$ for this Markov chain is

$$P = \begin{pmatrix} 1 - p_{12} & p_{12} & 0 & 0 \\ 0 & 0 & 1 - p_{24} & p_{24} \\ 0 & 0 & 1 - p_{34} & p_{34} \\ 0 & p_{42} & 0 & 1 - p_{42} \end{pmatrix}. \tag{2}$$

Each transition probability is governed by a logistic regression as follows:

$$\text{logit}(p_{12}) = \beta_{0,12} + \mathbf{z}^T \boldsymbol{\beta}_{\text{contract}}, \tag{3}$$

$$\text{logit}(p_{24}) = \beta_{0,24} + \mathbf{z}^T \boldsymbol{\beta}_{\text{recover}}, \tag{4}$$

$$\text{logit}(p_{34}) = \beta_{0,34} + \mathbf{z}^T \boldsymbol{\beta}_{\text{recover}}, \tag{5}$$

$$\text{logit}(p_{42}) = \beta_{0,42} + \mathbf{z}^T \boldsymbol{\beta}_{\text{contract}}. \tag{6}$$

As justified above, we use the same vector of covariate effects $\boldsymbol{\beta}_{\text{contract}}$ for the two probabilities related to contracting the particular *Bartonella* species. Similarly we use the same covariate effects for the two probabilities relating to recovery from the disease, $\boldsymbol{\beta}_{\text{recover}}$; we allow only the intercepts to differ. This assumption prevents a further increase in the, already large, number of

parameters to be estimated. For example, the logistic model for the probability of contracting *B. taylorii* for the first time at lunar month $Lm + 1$ will be

$$\begin{aligned} \text{logit}(p_{12}^{\text{tay}}) = & \beta_{0,12} + \beta_{\text{wt}} \text{weight} + \beta_{Lm} Lm + \beta_{\text{sex}} I(\text{male}) + \beta_{s2} I(\text{site2}) + \beta_{s3} I(\text{site3}) + \beta_{s4} I(\text{site4}) \\ & + \beta_{\text{sin}} \sin + \beta_{\text{cos}} \cos + \beta_{\text{sex:sin}} I(\text{male}) \sin + \beta_{\text{sex:cos}} I(\text{male}) \cos + \beta_{\text{grah2}} I(\text{grah2}) \\ & + \beta_{\text{grah3}} I(\text{grah3}) + \beta_{\text{grah4}} I(\text{grah4}) + \beta_{\text{dosh2}} I(\text{dosh2}) + \beta_{\text{dosh3}} I(\text{dosh3}) \\ & + \beta_{\text{dosh4}} I(\text{dosh4}) + \beta_{\text{bab2}} I(\text{bab2}) + \beta_{\text{bab3}} I(\text{bab3}) + \beta_{\text{cow2}} I(\text{cow2}) + \beta_{\text{cow3}} I(\text{cow3}) \\ & + \beta_{\text{ana2}} I(\text{ana2}). \end{aligned}$$

Here and elsewhere $I(\cdot)$ denotes the indicator function, and $[\text{disease}]x$ is a statement that the hidden chain for $[\text{disease}]$ is in state x .

2.5.2. *Babesia*

Babesia microti can cause haemolytic anaemia in infected hosts. It is a chronic infection, which is to say that once a host has been infected it is never again free of the disease. The effect of a *Babesia* infection on the probabilities of contracting or recovering from one of the other diseases may depend on whether the *Babesia* infection is acute (in its first month) or chronic.

We therefore model *Babesia* by using a Markov chain with the following three states: 1, no infection; 2, acute infection; 3, chronic infection. Here the likelihood vector that connects the states with the observations is analogous to that for *Bartonella* species but ignoring state 4. The transition matrix is

$$P = \begin{pmatrix} 1 - p_{12} & p_{12} & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{pmatrix}.$$

As in the previous section, a logistic regression relates p_{12} to the covariates, including the states of the other diseases.

2.5.3. *Anaplasma*

Anaplasma phagocytophilum is a tick-borne bacterium that causes the disease granulocytic ehrlichiosis in humans. In the data set there are relatively few positive records for *Anaplasma* and thus little power to ascertain transition probabilities and covariate effects from a third state of, for example, ‘currently uninfected but was previously infected’. Therefore, we use a two-state Markov chain with transition probability matrix

$$P = \begin{pmatrix} 1 - p_{12} & p_{12} \\ p_{21} & 1 - p_{21} \end{pmatrix},$$

and separate logistic regressions relating p_{12} and p_{21} to the covariates, including the states of the other diseases. This therefore is the only disease for which the underlying Markov model is not hidden.

2.5.4. *Cowpox*

In voles and other wild rodents, infection with cowpox virus is known to last for approximately 4 weeks (Bennett *et al.*, 1997). The diagnostic test, however, detects antibodies to the virus, not the virus itself. Antibodies appear approximately 2 weeks after contracting the infection but then remain present in the bloodstream of a vole for the rest of its life (Bennett *et al.*, 1997). Since the disease lasts for approximately 1 month we model the progression as a Markov chain with

three states: 1, antibodies absent and disease absent; 2, antibodies present and disease present; 3, antibodies present and disease absent. Therefore, the form of the transition matrix and the relationship between the states and the response is identical to that for *Babesia*. The difference is in the interpretation: here state 3 corresponds to a positive response but absence of the disease, whereas for *Babesia* state 3 corresponds to a positive response which means that the disease is present.

3. Bayesian approach

3.1. Choice of prior

3.1.1. Initial probability distribution

The likelihood (Section 2.2) and the forward–backward algorithm (Section 2.3) require, for each disease, the initial distribution $\pi^{[d]}$ of the Markov chain on the set of states for that disease at the first observation time for each vole. Our time inhomogeneous Markov chains admit no limiting distribution and so the popular choice of setting the initial distribution to the limiting distribution of the chain is not available to us.

We choose then to estimate this distribution for each disease through our Gibbs sampler in Section 3.2. We choose independent and relatively vague Dirichlet priors $\pi^{[d]} \sim \text{Dir}(\alpha^{[d]})$, with $\alpha^{[d]}$ set to a vector of 1s with length equal to the cardinality of state space $\mathbb{S}^{[d]}$. This is equivalent to a uniform prior on each $\pi^{[d]}$.

3.1.2. Prior distributions for the regression parameters

A similar longitudinal data set to the one that we analyse was also available to us. This additional data set arises from an earlier, 3-year study which was conducted using the same sampling design, but where the response for *Bartonella* was a single indicator for presence and absence, rather than an indicator for each species. For each of *Babesia*, *Anaplasma* and cowpox we could therefore fit a logistic regression to a subset of the additional data set as briefly described in Sections 1.1 and 2.1, except that the three indicator covariates for presence or absence of each *Bartonella* species were replaced with a single indicator covariate for presence or absence of at least one *Bartonella* species. Parameter estimates from these analyses were used to inform our choice of prior for similar parameters in our main analysis.

Since the additional data set does not distinguish the *Bartonella* species, there is not an exact correspondence between parameters from the simple analyses and the parameters in our main model, and some of the parameters in our main model have no counterpart in the simple analyses. The priors for each $\beta^{[d]}$ ($d = 1, \dots, D$) in our main analysis are independent and Gaussian with covariance matrices that we denote by $V^{[d]}$. For *Babesia*, *Anaplasma* and cowpox, where parameters do approximately correspond, we set the prior mean for the parameter in our main analysis to the maximum likelihood estimate for the corresponding parameter in the simple analysis of the additional data set and the block of $V^{[d]}$ that is associated with these parameters to nine times the analogous block from the simple analysis. Where no corresponding maximum likelihood estimate exists, the prior mean was set equal to 0 and the block of $V^{[d]}$ was a diagonal matrix where the diagonal elements were set to 9.

3.2. Adaptive Metropolis-within-Gibbs algorithm

In our data set, the target parameter β can be naturally partitioned into six subblocks: one for each disease. In the Gibbs scheme we wish to update the covariate parameters for each given disease d , $\beta^{[d]}$, by using a *random-walk Metropolis* (RWM) step (see for example Gilks *et al.*

(1996)); however, the efficiency of a given RWM algorithm depends heavily on the choice of the variance of the proposal jump, $\Sigma^{[d]}$. Both Roberts and Rosenthal (2001) and Sherlock and Roberts (2009) suggested that an RWM algorithm might achieve near optimal efficiency when $\Sigma^{[d]}$ correctly represents the general shape of the target distribution, e.g. if it is proportional to the variance of $\beta^{[d]}$ or the inverse curvature at the mode. We therefore generalize the *adaptive RWM* algorithm that was described in Sherlock *et al.* (2010) to an *adaptive Metropolis-within-Gibbs* algorithm on D subblocks.

Let $\beta^{[1:D]} = (\beta^{[1]}, \beta^{[2]}, \dots, \beta^{[D]})$, where $\beta^{[d]}$ is the parameter set for the d th subblock. A single iteration starts from an initial value $\beta^{[1:D]}$, cycles through all the subblocks updating each in turn and finishes with $\beta^{[1:D]} = (\beta^{[1]}, \beta^{[2]}, \dots, \beta^{[D]})$. In the update for the d th block the current and proposed values are respectively

$$\begin{aligned} \beta^{[1:D]} &:= \beta^{[1]}, \dots, \beta^{[d-1]}, \beta^{[d]}, \beta^{[d+1]}, \dots, \beta^{[D]}, \\ \beta^{*[1:D]} &:= \beta^{[1]}, \dots, \beta^{[d-1]}, \beta^{[d]} + \varepsilon, \beta^{[d+1]}, \dots, \beta^{[D]}. \end{aligned} \tag{7}$$

Here the proposal jump ε in expression (7) for the d th subblock at the n th iteration is sampled from a mixture distribution:

$$\varepsilon \sim \begin{cases} N\{\mathbf{0}, (m_n^{[d]})^2 \Sigma_n^{[d]}\} & \text{with probability } 1 - \delta, \\ N\{\mathbf{0}, (m_0^{[d]})^2 \Sigma_0^{[d]}\} & \text{with probability } \delta. \end{cases} \tag{8}$$

Here δ is a small positive constant, which we set to 0.1, $\Sigma_0^{[d]}$ is a fixed covariance matrix and $m_0^{[d]}$ is set to the theoretically derived value $2.38/k^{1/2}$ (Roberts and Rosenthal, 2001), where k is the dimension of $\beta^{[d]}$. The matrix $\Sigma_n^{[d]}$ is the estimated variance of $\beta^{[d]}$ by using the sample from the Markov chain to date. The scaling factor $m_n^{[d]}$ for the adaptive part is initialized to $m_0^{[d]}$. Sherlock *et al.* (2010) detailed the adaptation of the scaling factor at each iteration, which leads to an equilibrium acceptance rate of 25%, which is close to the optimal acceptance rate (approximately 23%) that was derived by Roberts and Rosenthal (2001).

The conditional likelihood for $\beta^{[d]} | \mathbf{y}^{[d]}, \mathbf{x}^{[-d]}, \boldsymbol{\pi}^{[d]}$ is calculated by using the forward part of the forward–backward algorithm for disease d . This provides the conditional posterior for $\beta^{[d]}$ and hence the acceptance probability for the proposed value of $\beta^{[d]}$. The states of the hidden Markov chain for disease d can then be sampled from their conditional distribution given the observed data for that disease, $\boldsymbol{\pi}^{[d]}, \beta^{[d]}$, and the states of the other diseases using the backward part of the forward–backward algorithm.

Given the states of the hidden chain for a particular disease d and, in particular, the initial state for each vole, the conditional conjugacy of the Dirichlet distribution allows straightforward sampling from the conditional posterior for $\boldsymbol{\pi}^{[d]}$.

We therefore simulate from the joint posterior distribution of the coefficients of the logistic regressions for the transition probabilities, the hidden disease states and the initial probability distribution of the hidden states with the following MCMC algorithm.

At the start of the current iteration of the chain let the covariate parameters be $\beta^{[1:D]}$, let the hidden states be $\mathbf{X}^{[1:D]}$ and the initial distribution of the hidden states be $\boldsymbol{\pi}^{[1:D]}$; denote their values at the start of the next iteration as $\beta^{[1:D]}$, $\mathbf{X}^{[1:D]}$ and $\boldsymbol{\pi}^{[1:D]}$ respectively.

Each iteration of the Gibbs sampler is as follows.

- Step 1: perform an adaptive RWM update according to $\beta^{[1]} | \mathbf{y}^{[1]}, \beta^{[1]}, \mathbf{x}^{[2:D]}, \boldsymbol{\pi}^{[1]}$.
- Step 2: simulate the hidden states for the first disease from $\mathbf{X}^{[1]} | \mathbf{y}^{[1]}, \beta^{[1]}, \mathbf{x}^{[2:D]}, \boldsymbol{\pi}^{[1]}$.
- Step 3: simulate the initial probability distribution of the chain for the first disease $\boldsymbol{\pi}^{[1]} | \mathbf{x}^{[1]}$.
- Step 4: perform an adaptive RWM update according to $\beta^{[2]} | \mathbf{y}^{[2]}, \beta^{[2]}, \mathbf{x}^{[1]}, \mathbf{x}^{[3:D]}, \boldsymbol{\pi}^{[2]}$.
- Step 5: simulate the hidden states for the second disease from $\mathbf{X}^{[2]} | \mathbf{y}^{[2]}, \beta^{[2]}, \mathbf{x}^{[1]}, \mathbf{x}^{[3:D]}, \boldsymbol{\pi}^{[2]}$.

Step 6: simulate $\pi^{[2]}|\mathbf{x}^{[2]}$.

⋮

Step 3D – 2: perform an adaptive RWM update according to $\beta^{[D]}|\mathbf{y}^{[D]}, \beta^{[D]}, \mathbf{x}^{[1:D-1]}, \pi^{[D]}$.

Step 3D – 1: simulate $\mathbf{X}^{[D]}|\mathbf{y}^{[D]}, \beta^{[D]}, \mathbf{x}^{[1:D-1]}, \pi^{[D]}$.

Step 3D: simulate $\pi^{[D]}|\mathbf{x}^{[D]}$.

The adaptive RWM step requires the fixed covariance matrix $\Sigma_0^{[d]}$. For each disease, a separate non-adaptive RWM step was performed for the logistic regressions coefficients that are associated with the HMM for this disease that are not associated with the other diseases, e.g. weight and sex. The block of $\Sigma_0^{[d]}$ that is associated with these covariates was estimated directly from this run. Each of the remaining β -coefficients was given a small proposal variance and was assumed to be uncorrelated with any of the other coefficients. Also to ensure a sensible non-singular $\hat{\Sigma}_n^{[d]}$, for each disease, proposals from the adaptive part were only allowed once at least 1000 proposed jumps had been accepted.

4. Analysis and results

4.1. Convergence of the algorithm and model diagnostics

All the computationally intensive parts of the algorithm were coded in C within an R (R Development Core Team, 2012) wrapper. On a computer with an Intel Nehalem 2.26 GHz central processor unit, 100000 iterations of the algorithm took approximately 3 h.

Three independent Markov chains of length 350000 were generated from the algorithm in Section 3.2; each chain was started from a different position. Six of the 233 trace plots from one of the chains are reproduced in Fig. 2. For most of these, over the first few tens of thousands of the iterations the variance of the proposal increases as the adaptive algorithm learns the shape of the posterior; this was so in many of the 233 trace plots.

The Gelman–Rubin statistic (Gelman and Rubin, 1992) R was calculated from the three chains for each of the 233 components of β . Fig. 3 shows the mean of the estimated R -statistics, the maximum of the estimated R -statistics along with the maximum of the 97.5% quantiles of R -statistics, plotted against iteration number. The plot suggests that a burn-in period of 150000 iterations should be more than sufficient. Inference is therefore performed using the final 200000 iterations from each of the three runs combined.

To assess model fit we examine the posterior predictive distribution of the data (see Robert *et al.* (1999)). We chose, at random, 100 captures where all six diseases were observed and created an alternative data set where all diseases for these captures were marked as missing. We refitted the model and estimated the posterior probabilities \hat{p} that these artificially removed observations are positive. A Hosmer–Lemeshow test (e.g. Collett (2003)) for each disease provides a p -value for the null hypothesis that each of the true (binary) observations arises from a Bernoulli trial with the given posterior probability. For the six diseases we obtained p -values of 0.443 (*B. doshiae*), 0.188 (*B. grahamii*), 0.061 (*B. taylorii*), 0.141 (*Babesia*), 0.56 (*Anaplasma*) and 0.322 (cowpox).

4.2. Posterior inference

We are interested in interactions between diseases, e.g. in whether or not the presence or absence of disease d_1 affects the probability of a change of state for disease d_2 . In each logistic regression for each transition matrix, we therefore examine the coefficients that correspond to the states of the other diseases. We are also interested (for *Bartonella*) in whether or not the status of an

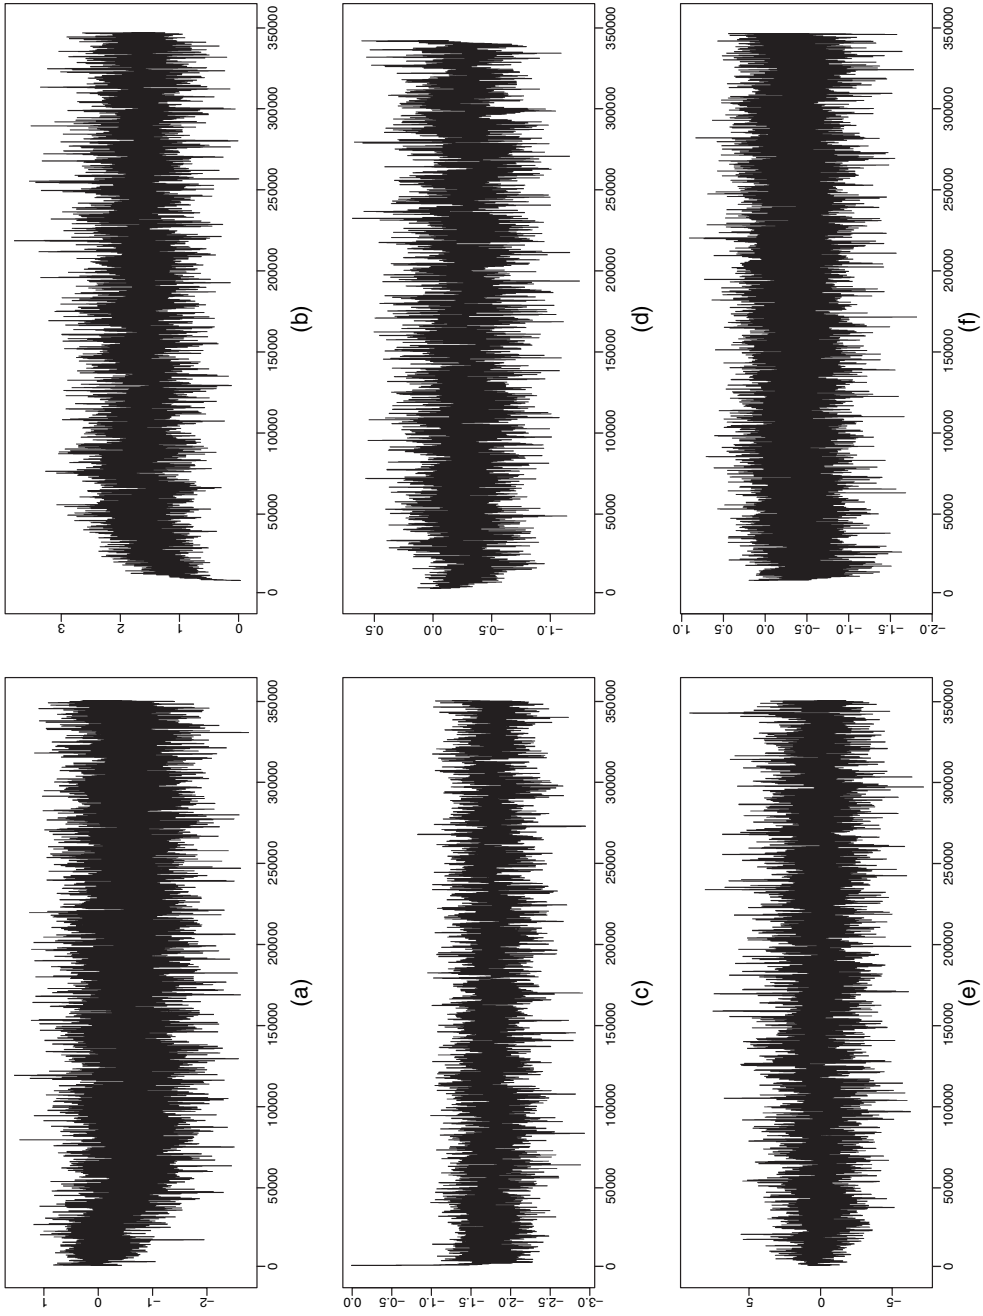


Fig. 2. Trace plots of the β -coefficients for (a) *ana21.grah3*, (b) *ana21.grah3*, (c) *ana21.grah3*, (d) *ana21.grah3*, (e) *ana21.grah3* and (f) *ana21.grah3*

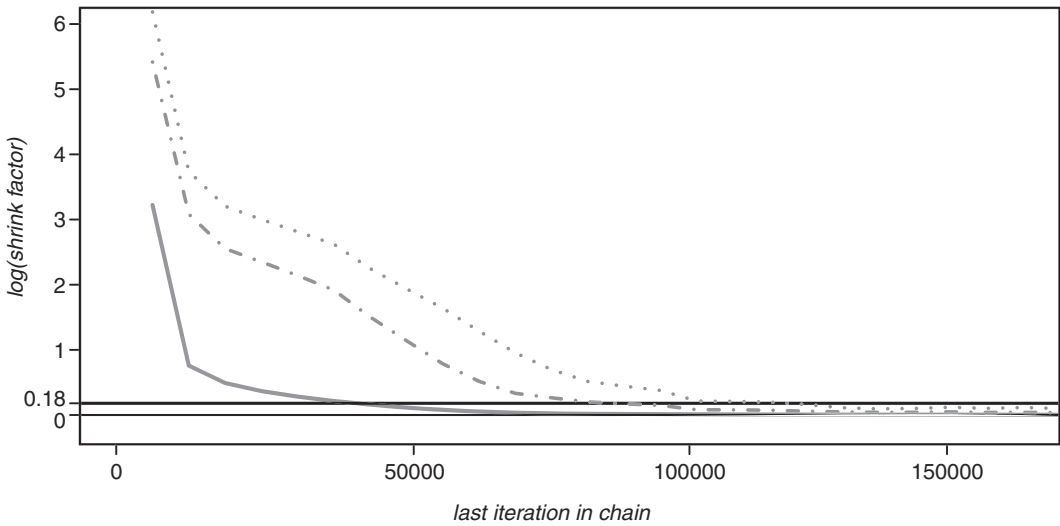


Fig. 3. Combined Gelman–Rubin statistics for all 233 β s, R_1, \dots, R_{233} (the statistics are plotted on a log-scale against iteration number, along with the ideal ratio of $\log(1)$ (0) and the threshold of $\log(1.2)$ (0.18) suggested in Gelman (1996)): —, mean Gelman–Rubin point estimator $(1/233) \sum_i R_i$; - - -, maximum Gelman–Rubin point estimator $\max_i R_i$; ·····, maximum 97.5% quantile estimate $\max_i R_{i,0.975}$

infection (new or old) affects the chance of recovery, and in whether or not a previous infection affects the chance of a new infection with the same species; these correspond respectively to the contrasts $\beta_{0,34} - \beta_{0,24}$ and $\beta_{0,42} - \beta_{0,12}$.

Formal model choice, e.g. via reversible jump MCMC sampling (Green, 1995), is computationally infeasible here. Instead we take a high posterior probability that a given parameter or contrast is positive (or a high probability that it is negative) as indicating a potentially important effect. For an individual parameter we might consider $P(\text{positive}) > 0.975$ or $P(\text{positive}) < 0.025$ as indicating a likely effect. We are, however, interested in a total of 116 parameters and contrasts, which raises a similar problem to that of multiple testing in classical statistics. Although considering probabilities below 0.025 or above 0.975 to indicate a possible interaction, we therefore take probabilities below 0.00025 and above 0.99975 as indicating a very probable interaction.

Table 4 shows those parameters for which the posterior probability of positivity is either above 0.975 or below 0.025. Table 4 shows the posterior median, a 95% credibility interval and the posterior probability that the parameter is positive.

Firstly, and most clearly, the presence of *Babesia* decreases the probability of contracting *Bartonella* and increases the probability of recovery from *Bartonella*. This is true for both chronic and acute *Babesia* infections and for all three species of *Bartonella*. There is no evidence for the reverse interaction, i.e. for the presence of *Bartonella* affecting the chance of contracting *Babesia*.

For two of the three *Bartonella* species (*B. taylorii* and *B. grahamii*) it appears that a vole is less likely to be reinfected following previous exposure whereas it is more likely to recover from an old infection of *B. taylorii* than a new one. Furthermore, there seems to be a decrease in the probability of contracting *B. doshiae* when a vole has been exposed to *B. taylorii* whether or not it is still infected. Finally, infection with *Anaplasma* appears to increase the probability of recovery from *B. grahamii*; there was perhaps some evidence for the same interaction with *B. doshiae* with posterior probability 0.9593. There is no evidence of a change to the probability of recovery from *B. taylorii* (probability 0.416). It is also possible that a current infection with

Table 4. Posterior summaries of model parameters of interest for which the posterior probability of positivity is either above 0.975 or below 0.025†

Transition probability	Covariate	Median	95% credibility interval	Posterior probability
p_{12}^{dosh}	tay2	-1.0124	(-1.8708, -0.1554)	0.0111
	tay3	-1.1077	(-2.0584, -0.1877)	0.0096
	tay4	-1.0304	(-1.8046, -0.2278)	0.0077
	bab2	-1.5636	(-2.4247, -0.8039)	0.0001
p_{24}^{dosh}	bab3	-1.7702	(-2.3095, -1.3235)	0.0000
	bab2	2.9613	(1.6459, 4.5298)	1.0000
	bab3	3.3386	(1.775, 5.5524)	1.0000
	cow2	-1.01	(-2.0937, -0.0691)	0.0171
p_{12}^{grah}	bab2	-1.2579	(-2.0506, -0.5881)	0.0000
	bab3	-2.0294	(-2.6841, -1.5096)	0.0000
p_{24}^{grah}	bab2	2.7818	(1.1981, 5.0649)	0.9999
	bab3	1.7288	(0.8899, 2.754)	1.0000
	ana2	1.0284	(0.0116, 2.1709)	0.9764
p_{12}^{tay}	bab2	-1.6933	(-2.6563, -0.815)	0.0001
	bab3	-1.3346	(-1.8338, -0.8754)	0.0000
p_{24}^{tay}	bab2	1.5824	(0.702, 2.5404)	0.9997
	bab3	1.7513	(1.0015, 2.6279)	1.0000
p_{12}^{bab}	cow3	-0.4939	(-0.9799, -0.0206)	0.0200
<i>Contrasts</i>				
$\beta_{0,42}^{\text{grah}} - \beta_{0,12}^{\text{grah}}$		-1.3248	(-5.6452, -0.2774)	0.0076
$\beta_{0,42}^{\text{tay}} - \beta_{0,12}^{\text{tay}}$		-4.472	(-8.6604, -2.1214)	0.0015
$\beta_{0,34}^{\text{tay}} - \beta_{0,24}^{\text{tay}}$		1.1762	(0.555, 1.8601)	0.9998

†For each parameter the posterior median, a 95% credibility interval and the probability that the parameter is greater than 0 is provided. Each parameter arises from a logistic regression coefficient for a particular transition probability in the HMM for a particular disease. The disease and transition appear in the first column, and the second column indicates the particular disease and state that is influencing the transition probability. State 1 is always taken to be the baseline. The contrasts $\beta_{0,42} - \beta_{0,12}$ and $\beta_{0,34} - \beta_{0,24}$ are defined in Section 4.2.

cowpox virus hinders recovery from *B. doshiae* and previous exposure to cowpox prevents infection with *Babesia*.

4.3. Sensitivity analysis

Three somewhat arbitrary choices were made in the set-up of our model and priors: the interpolation scheme that fills in missing weight values, the prior for the initial distribution for the state of the HMM for each disease and vole, and the exact relationship between parameters that are estimated in the simple analysis of the alternative data set and priors for parameters in the HMMs for the main data set.

An alternative for each of these choices is described below. For each alternative three further chains of length 350000 were created and checked for convergence. Then any sizable changes in the conclusions that would be drawn from the posterior distributions of the parameters were noted.

In the main analysis, missing weight values were filled in via linear interpolation. As an alternative we considered the logistic growth curve which was proposed in Burthe *et al.* (2009). We assumed Gaussian residuals for the logarithm of weight and allowed the logistic growth

parameters to depend on covariates such as the sex of the vole and the time of year; some of the coefficients were also allowed to include subject-specific random effects. More details are provided in Xifara (2012).

The initial distributions for the states of the HMMs for the diseases are assigned independent Dirichlet priors with the parameter for each disease, $\alpha^{[d]}$, a vector of 1s. As an alternative prior we set each $\alpha^{[d]}$ to be a vector of 2s.

In the main analysis, where there was a rough correspondence between a parameter in the simple analysis of the alternative data set and a parameter in the main analysis, we centred the Gaussian prior in the main analysis on the maximum likelihood estimate from the simple analysis and set the covariance matrix to be nine times the estimated covariance matrix from the simple analysis (Section 3.1.2). As an alternative we use vague but proper Gaussian priors for all parameters.

Parameter estimates with the alternative weight scheme or with the alternative prior distribution of the hidden states were very similar to the estimates from the main set-up. For all the significant covariates none of the posterior probabilities changed by more than 0.005. However, the use of vague priors for the parameters noticeably affected one of the 21 covariates in Table 4. The effect on a vole's probability of contracting *Babesia* when the vole had been exposed to cowpox became apparently unimportant, with posterior probability changing from 0.02 to 0.093. No additional covariates became potentially important (i.e. $p < 0.025$ or $p > 0.975$) in any of the three alternative runs.

5. Discussion

We have described a coupled discrete time HMM for interactions between diseases in a host and used it to analyse data from a longitudinal study of field voles with records of six different pathogens. The Markov model offers a more detailed description than the existing modelling approach that is described in Section 1.1. Furthermore, by explicitly dealing with the missing observations (which comprise approximately 50% of the data set), the inference methodology that we introduce can use more of the data than the existing standard inference methodology.

Inference is performed via a Metropolis-within-Gibbs sampler that cycles through the diseases and, for each disease conditional on the hidden states of all of the other diseases, samples from the parameters of the logistic regressions for the transition probabilities of the HMM using an adaptive RWM step and then from the exact distribution of the hidden states given these parameters. These two steps use respectively the forward and backward parts of the forward-backward algorithm.

The forward-backward Gibbs sampler (e.g. Chib (1996), Scott (2002) and Fearnhead and Sherlock (2006)) also uses the forward-backward algorithm; however, the motivation is different. The forward-backward Gibbs sampler does not use the likelihood from the forward recursion directly, as this would require a Metropolis-Hastings update; instead the backward recursion provides a sample from the posterior distribution of the hidden states given the parameters. Owing to the conditional conjugacy structure of the problems that is targeted by the forward-backward Gibbs sampler it is then possible to sample exactly from the conditional posterior for the parameters given the hidden states and thus to avoid the Metropolis-Hastings step and the associated tuning. Our logistic regression model for the transition probabilities does not allow a simple Gibbs step for updating the parameters conditionally on the hidden states, and so we content ourselves with a Metropolis-Hastings step for the parameters and, for efficiency of mixing, do not condition on the hidden states for the current disease. After the Metropolis-Hastings step we *then* sample from the hidden states for the disease so that these

can be used as covariates for the other diseases; in effect, we therefore sample from the joint conditional distribution of the parameters and the hidden states for the disease. The forward–backward algorithm could be avoided entirely by updating the logistic regression parameters conditionally on the hidden states for *all* diseases, and by sampling from the distribution of each individual hidden state conditionally on all the other hidden states and the transition parameters (e.g. Robert *et al.* (1993) and Robert and Titterton (1998)). However, we believe that the correlation between hidden states and between these states and the parameters would have led to a very poorly mixing MCMC chain; Scott (2002) has discussed the first aspect of this.

Brand (1997), Saul and Jordan (1999), Rezek *et al.* (2000), Zhong and Ghosh (2002) and Natarajan and Nevatia (2007) all examined inference for coupled HMMs. The directed acyclic graph for the HMMs that they considered is the same as in Fig. 1; however, to allow recursions that are similar to those in the forward–backward algorithm all—except Rezek *et al.* (2000)—made the simplifying assumption that the transition probability for a given chain conditional on the others is separable:

$$P(x_{t+1}^d | \mathbf{x}_t^{[1:D]}) \propto \prod_{i=1}^D f_{i,d}(x_{t+1}^d | x_t^{[i]}),$$

for some collection of non-negative functions $f_{i,d}$. Moreover the computational complexity increases with the square of the sum of the number of states in each chain and, in practice, each of Brand (1997), Saul and Jordan (1999), Rezek *et al.* (2000), Zhong and Ghosh (2002) and Natarajan and Nevatia (2007) considered only two chains. We wish to apply logistic regression rather than to assume separability and to consider six chains; furthermore computational complexity of our algorithm increases with the sum of the squares of the number of states in each chain. Rezek *et al.* (2000) performed Bayesian inference via a Gibbs sampling algorithm that is qualitatively similar to our own; the fully conjugate structure that is used does not, however, allow for the effects of any covariates.

We now examine the most major findings of Section 4.2 and briefly discuss the biological insights that they offer. For *B. taylorii*, voles are more likely to recover from an old infection than a new one, which is to be expected given that more complete histories for individuals indicate that most infections last for only 1 month. Previous data also indicate that infections by *B. doshiae* may last longer than infections by *B. taylorii* and *B. grahamii* (Telfer *et al.*, 2007). Also, for *B. grahamii* and *B. taylorii*, previous infection by the species appears to grant some form of immunity to that species, suggesting that hosts can develop an effective acquired immune response. To date, there has been conflicting evidence for such a response in wild populations, suggesting that immune responses may vary between host species and/or *Bartonella* species (Birtles *et al.*, 2001; Kosoy *et al.*, 2004; Bai *et al.*, 2011). Interestingly also infection by *B. taylorii* appears to provide immune cross-protection to infection by *B. doshiae*.

We found, for voles currently infected with *Babesia*, both a reduction in susceptibility to *Bartonella* and an increase in the probability of recovery from *Bartonella* over the next lunar month. We also found no evidence that a current *Bartonella* infection might influence susceptibility to *Babesia* over the next month. Telfer *et al.* (2010) found that the *Babesia* covariates, both at time t_0 and at time t_1 , were significant for predicting the probability of catching *Bartonella* between t_0 and t_1 . The *Bartonella* covariate at time t_1 is also found to be significant in predicting susceptibility to *Babesia*, apparently contradicting our findings. However, as mentioned in Section 2.1, Telfer *et al.* (2010) allowed both the current (t_0) and future (t_1) state of each disease covariate to influence the probability that, for the disease that is being treated as a response, a vole tests positively at time t_1 given that it tests negatively at time t_0 . An assumption of only the negative effects of *Babesia* on *Bartonella* infections that is apparent from inference for our HMM, and no

other dependence between *Bartonella* and *Babesia*, is sufficient to lead to a negative correlation between *Bartonella* and *Babesia* at any given time. Consider two groups of voles: those with *Babesia* at t_0 (group A) and those without *Babesia* at t_0 (group B). Since *Babesia* is a chronic infection, group A voles are more likely to have *Babesia* at t_1 than are voles from group B. However, since *Babesia* impacts negatively on the probability that a vole has *Bartonella*, group A voles are less likely to have *Bartonella* at t_1 than voles from group B. Since the effect of *Babesia* on *Bartonella* is so pronounced (Table 4), it is certainly believable that this negative correlation could be sufficiently strong that each of the two diseases at t_1 appears as an important covariate for the other.

In the application which we have considered, missingness was believed to be independent of disease state; in other scenarios, such as those considered in Pradel (2005), the probability that a given subject will be observed might depend on the states of each of the HMMs. This could be accommodated within our methodology through a further logistic regression for the probability of being observed given the set of hidden states and other covariate information, and several other minor changes as detailed in Pradel (2005).

Acknowledgements

Part of this work was funded through North West Development Agency project N0003235. TX acknowledges also funding from the Engineering and Physical Sciences Research Council and Faculty of Science and Technology at Lancaster University. The fieldwork was supported by funding from the Natural Environment Research Council (grant GR3/13051) and the Wellcome Trust (grants 075202/Z/04/Z and 070675/Z/03/Z). We are grateful to the Joint Editor, the Associate Editor and two referees for their helpful comments.

References

- Bai, Y., Calisher, C. H., Kosoy, M. Y., Root, J. J. and Doty, J. B. (2011) Persistent infection or successive reinfection of deer mice with *Bartonella vinsonii* subsp. *arupensis*. *Appl. Environ. Microbiol.*, **77**, 1728–1731.
- Baum, I. E., Petrie, Y., Soules, G. and Weiss, N. (1970) A maximisation technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Statist.*, **41**, 164–171.
- Begon, M., Telfer, S., Burthe, S. J., Lambin, X., Smith, J. M. and Paterson, S. (2009) Effects of abundance on infection in natural populations: field voles and cowpox virus. *Epidemics*, **1**, 35–46.
- Bennett, M., Crouch, A. J., Begon, M., Duffy, B., Feore, S., Gaskell, R. M., Kelly, D. F., McCracken, C. M., Vicary, L. and Baxby, D. (1997) Cowpox in British voles and mice. *J. Compar. Path.*, **116**, 35–44.
- Birtles, R. J., Hazel, S. M., Bennett, M., Bown, K., Raoult, D. and Begon, M. (2001) Longitudinal monitoring of the dynamics of infections due to *Bartonella* species in UK woodland rodents. *Epidem. Infectn*, **126**, 323–329.
- Bown, K. J., Bennett, M. and Begon, M. (2004) Flea-borne *Bartonella grahamii* and *Bartonella taylorii* in Bank Voles. *Emergng Infect. Dis.*, **10**, 684–687.
- Bown, K. J., Lambin, X., Telford, G. R., Ogden, N. H., Telfer, S., Woldehiwet, Z. and Birtles, R. J. (2008) Relative importance of *Ixodes ricinus* and *Ixodes trianguliceps* as vectors for *Anaplasma phagocytophilum* and *Babesia microti* in field vole (*Microtus agrestis*) populations. *Appl. Environ. Microbiol.*, **74**, 7118–7125.
- Brand, M. (1997) Coupled hidden Markov Models for modelling interacting processes. *Technical Report 405*. Massachusetts Institute of Technology Media Laboratory, Cambridge.
- Burthe, S. J., Lambin, X., Telfer, S., Douglas, A., Beldomenico, P., Smith, A. and Begon, M. (2009) Individual growth rates in natural field voles, *Microtus agrestis*, populations exhibiting cyclic population dynamics. *Oecologia*, **162**, 653–661.
- Chadeau-Hyam, M., Clarke, P. S., Guihenneuc-Jouyaux, C., Cousens, S. N., Will, R. G. and Ghani, A. C. (2010) An application of hidden Markov models to the French variant Creutzfeldt–Jakob disease epidemic. *Appl. Statist.*, **59**, 839–853.
- Chantrey, J., Meyer, H., Baxby, D., Begon, M., Bown, K. J., Hazel, S. M., Jones, T., Montgomery, W. I. and Bennett, M. (1999) Cowpox: reservoir hosts and geographic range. *Epidem. Infectn*, **122**, 455–460.
- Chib, S. (1996) Calculating posterior distributions and modal estimates in Markov mixture models. *J. Econometr.*, **75**, 79–97.
- Collett, D. (2003) *Modelling Binary Data*. Boca Raton: Chapman and Hall–CRC.

- Courtney, J. W. L., Kostelnik, M., Zeidner, N. S. and Massung, R. F. (2004) Multiplex real-time PCR for detection of *Anaplasma phagocytophilum* and *Borrelia burgdorferi*. *J. Clin. Microbiol.*, **42**, 3164–3168.
- Daniels, M. J. and Hogan, J. W. (2008) *Missing Data in Longitudinal Data: Strategies for Bayesian Modelling and Sensitivity Analysis*. Boca Raton: Chapman and Hall–CRC.
- Fearnhead, P. and Sherlock, C. (2006) An exact Gibbs sampler for the Markov-modulated Poisson process. *J. R. Statist. Soc. B*, **68**, 767–784.
- Gelman, A. (1996) Inference and monitoring convergence. In *Markov Chain Monte Carlo in Practice* (eds W. R. Gilks, S. Richardson and D. J. Spiegelhalter). London: Chapman and Hall.
- Gelman, A. and Rubin, D. (1992) Inference from iterative simulation using multiple sequences. *Statist. Sci.*, **7**, 457–472.
- Gilks, W. R., Richardson, S. and Spiegelhalter, D. J. (eds) (1996) *Markov Chain Monte Carlo in Practice*. London: Chapman and Hall.
- Green, P. J. (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, **82**, 711–732.
- Guédon, Y. (2003) Estimating hidden semi-Markov chains from discrete sequences. *J. Computnl Graph Statist.*, **12**, 604–639.
- Guihenneuc-Jouyaux, C., Richardson, S. and Longini, I. M. (2000) Modelling markers of disease progression by a hidden Markov process: application to characterising CD4 cell decline. *Biometrics*, **56**, 733–741.
- Kosoy, M., Mandel, E., Green, D., Marston, E. and Childs, J. (2004) Prospective studies of *Bartonella* of rodents: part I, Demographic and temporal patterns in population dynamics. *Vect. Borne Zoonotic Dis.*, **4**, 285–295.
- Lachish, S., Knowles, S. C. L., Alves, R., Wood, M. J. and Sheldon, B. C. (2011) Infection dynamics of endemic malaria in a wild bird population: parasite species-dependent drivers of spatial and temporal variation in transmission rates. *J. Anim. Ecol.*, **80**, 1207–1216.
- Natarajan, P. and Nevatia, R. (2007) Coupled hidden semi Markov models for activity recognition. *Institute of Electrical and Electronics Engineers Wrkshp Motion and Video Computing*.
- Pradel, R. (2005) Multievent: an extension of multistate capture-recapture models to uncertain states. *Biometrics*, **61**, 442–447.
- Rabiner, L. R. (1989) A tutorial on hidden Markov models and selected applications in speech recognition. *Proc. IEEE*, **77**, 257–286.
- R Development Core Team (2012) *R: a Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing.
- Rezek, I., Sykacek, P. and Roberts, S. J. (2000) Learning interaction dynamics with couple hidden Markov models. *IEE Proc. Sci. Measmnt Technol.*, **147**, 345–350.
- Robert, C. P., Celeux, G. and Diebolt, J. (1993) Bayesian estimation of hidden Markov chains: a stochastic implementation. *Statist. Probab. Lett.*, **16**, 77–83.
- Robert, C. P., Rydén, G. and Titterington, D. M. (1999) Convergence controls for MCMC algorithms, with application to hidden Markov chains. *J. Statist. Computn Simuln.*, **64**, 327–355.
- Robert, C. P. and Titterington, D. M. (1998) Reparameterization strategies for hidden Markov models and Bayesian approaches to maximum likelihood estimation. *Statist. Comput.*, **8**, 145–158.
- Roberts, G. O. and Rosenthal, J. (2001) Optimal scaling for various Metropolis-Hastings algorithms. *Statist. Sci.*, **16**, 351–367.
- Saul, K. and Jordan, M. (1999) Mixed memory Markov models: decomposing complex stochastic processes as mixtures of simpler ones. *Mach. Learn.*, **37**, 75–87.
- Scott, S. L. (2002) Bayesian methods for hidden Markov models: recursive computing in the 21th century. *J. Am. Statist. Ass.*, **97**, 337–351.
- Sherlock, C., Fearnhead, P. and Roberts, G. O. (2010) The random walk Metropolis: linking theory and practice through a case study. *Statist. Sci.*, **28**, 172–190.
- Sherlock, C. and Roberts, G. (2009) Optimal scaling of the random walk Metropolis on elliptically symmetric unimodal targets. *Bernoulli*, **15**, 774–798.
- Telfer, S., Begon, M., Bennett, M., Bown, K., Burthe, S., Lambin, X., Telford, G. and Birtles, R. (2007) Contrasting dynamics of *Bartonella* spp. in cyclic field vole populations: the impact of vector and host dynamics. *Parasitology*, **134**, 413–425.
- Telfer, S., Birtles, R., Bennett, M., Lambin, X., Paterson, S. and Begon, M. (2008) Parasite interactions in natural populations: insights from longitudinal data. *Parasitology*, **135**, 767–781.
- Telfer, S., Lambin, X., Birtles, R., Beldomenico, P., Burthe, S. J., Paterson, S. and Begon, M. (2010) Species interactions in a parasite community drive infection risk in a wildlife population. *Science*, **330**, 243–246.
- Xifara, T. D. (2012) A hidden Markov model for disease interactions in field voles. *Technical Report*. Department of Mathematics and Statistics, Lancaster University, Lancaster. (Available from <http://www.math.lancs.ac.uk/~xifara>.)
- Zhong, S. and Ghosh, J. (2002) HMMs and coupled HMMs for multi-channel EEG classification. In *Proc. Int. Jt Conf. Neural Networks*, pp. 1154–1159. New York: Institute of Electrical and Electronics Engineers.
- Zucchini, W. and MacDonald, I. L. (2009) *Hidden Markov Models for Time Series: an Introduction using R*. New York: Chapman and Hall–CRC.