



## Markov Chain Monte Carlo Methods for Switching Diffusion Models

John C. Liechty; Gareth O. Roberts

*Biometrika*, Vol. 88, No. 2. (Jun., 2001), pp. 299-315.

Stable URL:

<http://links.jstor.org/sici?sici=0006-3444%28200106%2988%3A2%3C299%3AMCMCMF%3E2.0.CO%3B2-0>

*Biometrika* is currently published by Biometrika Trust.

---

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/bio.html>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

---

The JSTOR Archive is a trusted digital repository providing for long-term preservation and access to leading academic journals and scholarly literature from around the world. The Archive is supported by libraries, scholarly societies, publishers, and foundations. It is an initiative of JSTOR, a not-for-profit organization with a mission to help the scholarly community take advantage of advances in technology. For more information regarding JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

# Markov chain Monte Carlo methods for switching diffusion models

BY JOHN C. LIECHTY

*Department of Marketing, Smeal College of Business Administration,  
Pennsylvania State University, University Park, Pennsylvania 16802-3007, U.S.A.*

jcl12@psu.edu

AND GARETH O. ROBERTS

*Department of Mathematics and Statistics, Lancaster University, Lancaster LA1 4YF, U.K.*

g.o.roberts@lancaster.ac.uk

## SUMMARY

Reversible jump Metropolis–Hastings updating schemes can be used to analyse continuous-time latent models, sometimes known as state space models or hidden Markov models. We consider models where the observed process  $X$  can be represented as a stochastic differential equation and where the latent process  $D$  is a continuous-time Markov chain. We develop Markov chain Monte Carlo methods for analysing both Markov and non-Markov versions of these models. As an illustration of how these methods can be used in practice we analyse data from the New York Mercantile Exchange oil market. In addition, we analyse data generated by a process that has linear and mean reverting states.

*Some key words:* Change-point model; Reversible jump Markov chain Monte Carlo; Variable dimension time-series model.

## 1. INTRODUCTION

Recently, reversible jump Markov chain Monte Carlo methods (Green, 1995) have been presented for time series models where the observed process  $X$ , conditional on the latent or hidden process  $D$ , is Markov; see a University of Nottingham technical report by F. Ball, Y. Cai and A. O’Hagan for a discussion where  $D$  is a continuous-time Markov chain, and see Hodgson (1999) and Hodgson & Green (1999) for a discussion where  $D$  represents a more general Markov process. In this paper, we will extend this methodology to a rather general class of switching diffusion models, and apply our results to data from the New York oil markets.

Both Hodgson (1999) and the report by Ball et al. present reversible jump methods that are similar to ours. Hodgson (1999) discusses ‘shift’, ‘birth’ and ‘death’ reversible jump moves for a two-state latent process, and the Ball et al. report presents similar moves for multiple-state latent processes. The main distinction between this paper and other related work is the class of models that we consider. We show how our reversible jump methods can be used to analyse both Markov and non-Markov versions of these models, and we discuss convergence diagnostics and model choice tools for these models.

In the Markov case, we restrict our attention to the class of models where the observed

process  $X = \{X_t; 0 \leq t \leq T\}$  can be represented as a stochastic differential equation with a drift term that is determined by a hidden Markov chain  $D$ :

$$dX_t = \mu(D_t, X_t, t) dt + \sigma(X_t, t) dW_t, \quad (1.1)$$

where  $W$  is a Wiener process and  $D$  is a continuous-time Markov chain. In the non-Markov case, we consider models where the drift of  $X$  is a complex function of the past sample path and  $D$ . In the case where, conditional on the  $D$ ,  $X$  is Markov, as in (1.1), filtering methods can also be used to sample from the hidden process  $D$ . However, our methods extends beyond the Markov case to provide a flexible general purpose technique that can be used to sample from  $D$  for both the Markov and non-Markov cases.

For all our models, the hidden Markov chain  $D$  can be characterised in terms of a jump chain  $(i_0, i_1, \dots)$  and associated waiting times  $(t_0, t_1, \dots)$ . Conditional on the fact that a jump has occurred, the probability that  $D$  changes from state  $i_j$  to  $i_{j+1}$  is given by the jump chain transition matrix  $P_j$  which, together with the inverses of the average waiting times  $\lambda_i$ , is summarised in an infinitesimal generator or a 'Q' matrix. The Q matrix for  $D$  is given by

$$(Q)_{i,j} = \lambda_{ij} = \begin{cases} p_{ij} \times \lambda_i, & \text{if } i \neq j, \\ -\lambda_i, & \text{if } i = j, \end{cases}$$

where  $p_{ij}$  is the  $(i, j)$  element of  $P_j$ .

Since  $X$  is usually observed over a fixed time interval  $[0, T]$ , the number of parameters needed to describe a particular realisation of  $D$  varies with the number of times that this realisation changes state, so that we require reversible jump Markov chain Monte Carlo methods.

The outline for this paper is as follows. In § 2 we give an overview of appropriate Markov chain Monte Carlo methods and convergence diagnostics, and we discuss model choice tools. The acceptance probabilities for the reversible jump proposals are given in Appendix 2. In § 3 we use these methods to analyse data from the New York Mercantile Exchange oil market, and in § 4 we show how these methods can be extended for non-Markov state space models, giving the methods an advantage over filtering-based Markov chain Monte Carlo methods.

## 2. MODELS, METHODS, DIAGNOSTICS AND MODEL CHOICE

### 2.1. Linear hidden Markov chain models

We will begin by specifying a linear hidden Markov chain model. Although these linear models are relatively simple, they present a clear framework for discussing the challenging task of updating  $D$  within the Markov chain Monte Carlo approach.

A linear hidden Markov chain model has a different drift term for each state of the hidden Markov chain:

$$dX_t = \mu(D_t) dt + \sigma dW_t \equiv \left( \sum_{i=1}^m u_i I_{\{D_t=i\}} \right) dt + \sigma dW_t,$$

where  $I_{\{\cdot\}}$  represents the indicator function and where each  $\mu_i$  is a constant. To avoid problems with identification, we require that

$$\mu_1 < \dots < \mu_m.$$

As an example, we will specify a two-state model. The prior density for the drift term is

a constrained bivariate normal density:

$$\mu_1, \mu_2 \sim N \left\{ \begin{pmatrix} \tilde{\mu}_1 \\ \tilde{\mu}_2 \end{pmatrix}, \begin{pmatrix} \tau_1 & 0 \\ 0 & \tau_2 \end{pmatrix} \middle|_{\mu_1 < \mu_2} \right\}.$$

The prior densities for the off-diagonal elements of the  $Q$  matrix of  $D$  are exponential densities with mean  $1/\beta$ . For all the models we considered, we used hyperparameters that corresponded to relatively vague prior specifications. For example, a priori, the drift terms were centred at zero,  $\tilde{\mu}_1 = \tilde{\mu}_2 = 0$ , and had variances three times the size of the variance of the observed process, using the assumption that the process followed a Brownian motion with a drift equal to zero, that is  $\tau_1 = \tau_2 = 3\hat{\sigma}^2$ , where

$$\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \frac{(X_{t_i} - X_{t_{i-1}})^2}{t_i - t_{i-1}} \tag{2.1}$$

and  $n + 1$  is the number of observations of  $X$ . For  $D$  the a priori average waiting time in each state was one-third of the duration of the observed process,  $1/\beta = T/3$ , and a priori  $D$  was given equal probability of starting in either state 1 or state 2.

The likelihood for  $X$ ,

$$L(X|D, \mu_1, \mu_2) = \exp \left\{ \frac{1}{\sigma^2} \int_0^T (I_{\{D_s=1\}}\mu_1 + I_{\{D_s=2\}}\mu_2) dX_s - \frac{1}{2\sigma^2} \int_0^T (I_{\{D_s=1\}}\mu_1 + I_{\{D_s=2\}}\mu_2)^2 ds \right\},$$

is given by the Cameron–Martin–Girsanov formula; see for example Øksendal (1980, p. 123).

The full conditional density for  $\mu_1$  and  $\mu_2$  is given by

$$\mu_1, \mu_2 | D, \lambda_1, \lambda_2, X \sim N \left\{ \begin{pmatrix} a_1/b_1 \\ a_2/b_2 \end{pmatrix}, \begin{pmatrix} 1/b_1 & 0 \\ 0 & 1/b_2 \end{pmatrix} \middle|_{\mu_1 < \mu_2} \right\},$$

where

$$a_i = \frac{1}{\sigma^2} \int_0^T I_{\{D_s=i\}} dX_s + \frac{\tilde{\mu}_i}{\tau_i^2}, \quad b_i = \frac{1}{\sigma^2} \int_0^T I_{\{D_s=i\}} ds + \frac{1}{\tau_i^2}.$$

The full conditional density for  $\lambda_{ij}$ , with  $i, j = 1, 2$  and  $i \neq j$ , is given by

$$\lambda_{ij} | D, \mu_1, \mu_2, \lambda_j, X \sim \text{Ga} \left( \vartheta_{ij} + 1, \int_0^T I_{\{D_s=i\}} ds + \beta_i \right), \tag{2.2}$$

where  $\vartheta_{ij}$  is the number of times that the Markov chain  $D$  jumps from state  $i$  to state  $j$  on  $[0, T]$ . The full conditional density for  $D$  is nonstandard and is given by

$$\begin{aligned} \pi(D|X, \mu_1, \mu_2, \lambda_{12}, \lambda_{21}) &\propto \lambda_{12}^{\vartheta_{12}} \lambda_{21}^{\vartheta_{21}} \exp \left\{ \int_0^T \left( \frac{I_{\{D_s=1\}}}{\sigma^2} \mu_1 + \frac{I_{\{D_s=2\}}}{\sigma^2} \mu_2 \right) dX_s \right\} \\ &\times \exp \left[ - \int_0^T \left\{ I_{\{D_s=1\}} \left( \lambda_{12} + \frac{1}{2\sigma^2} \mu_1^2 \right) \right. \right. \\ &\quad \left. \left. + I_{\{D_s=2\}} \left( \lambda_{21} + \frac{1}{2\sigma^2} \mu_2^2 \right) \right\} ds \right]. \end{aligned}$$

2.2. *Markov chain Monte Carlo methods*

We shall use Metropolis–Hastings algorithms to simulate asymptotically from the posterior distribution; see Tierney (1998) for an introduction to the algorithm at the level of generality that we shall need to adopt. We shall require Metropolis–Hastings rules which update the unknown parameters, and the hidden Markov sample path.

The basic Markov chain Monte Carlo algorithm for the scalar parameters of a linear model is straightforward. The drift terms can be updated using rejection sampling and the jump rates can be updated by sampling directly from (2.2). Since the full conditional density of  $D$  is nonstandard, non-Gibbs algorithms need to be used to update  $D$ .

We use three different methods for updating  $D$ . The first method is an independence sampler which ignores the current realisation of  $D$  and proposes realisations of  $D$  that are considerably different, in terms of the posterior density. This sampler tends to result in large but infrequent moves. The other two methods make small modifications to the current realisation of  $D$ . The second method is a refinement sampler in which one of the jump times of  $D$  is changed, and the third method is a birth-death sampler in which an interval of  $D$  is either created or removed. These two samplers tend to result in small but frequent moves. The independence sampler has obvious advantages when the posterior distribution is multimodal or when a poor initial value of  $D$  has been chosen, whereas the refinement sampler and birth-death sampler have the advantage of more efficiently exploring the modes of the posterior distribution.

To illustrate how we use all three samplers, we give the basic Markov chain Monte Carlo algorithm for the two-state model and then give an overview of the three samplers. A description of the birth-death sampler algorithm is given in Appendix 1 and the acceptance probabilities for all three samplers are given in Appendix 2. It should be noted that in practice we chose our initial parameter values, for the basic algorithm, by drawing a random sample from their posterior distributions.

**ALGORITHM 1** (*The basic Markov chain Monte Carlo algorithm*).

*Step 1. Select a random starting point,  $\theta^0 = \{\mu_1^0, \mu_2^0, \lambda_{12}^0, \lambda_{21}^0, D^0\}$ .*

*Step 2. Generate a sample for  $(\mu_1, \mu_2)$ , with  $\mu_1 < \mu_2$ , using rejection sampling.*

*Step 3. Generate samples for  $\lambda_{12}^1$  and  $\lambda_{21}^1$  from (2.2).*

*Step 4. Generate a proposed realisation  $Y$  using one of the following sub-algorithms:*

*use the independence sampler with probability  $p_{IS}$ ,*

*use the refinement sampler with probability  $p_{RS}$ ,*

*use the birth-death sampler with probability  $p_{BDS}$ ,*

*where  $p_{IS} + p_{RS} + p_{BDS} = 1$  and  $p_{IS}, p_{RS}, p_{BDS} > 0$ .*

*Step 5. Calculate the acceptance probability  $\alpha$ ; see Appendix 2.*

*Step 6. With probability  $\alpha$  let  $D^1 = Y$ ; otherwise let  $D^1 = D^0$ .*

*Step 7. Repeat Steps 2–6.*

The independence sampler is straightforward and is based on the prior density of  $D$ .

**SUB-ALGORITHM 1** (*The independence sampler*). *Generate a proposed realisation at random from the prior density of  $D$  conditioned on  $Q$ :*

$$Y \sim f_{MC}(Q),$$

*where  $f_{MC}$  is the standard density for a continuous-time Markov chain.*

Before detailing the refinement sampler, we need to define the jump times and label the intervals of  $D$ . The  $i$ th jump time for  $D$  is the  $i$ th time that  $D$  changes state and is given

by  $T_i = \sum_{j=1}^i t_j$ . If  $D$  has at least one jump time, then  $[0, T_1]$  is the first interval and  $[T_n, T]$  is the last interval. The first and last intervals are called external intervals; the remaining intervals are called internal intervals.

The refinement sampler generates a proposed realisation  $Y$  by changing the location of one of the jump times of the current  $D$ .

**SUB-ALGORITHM 2** (*The basic refinement sampler*).

*Step 1. Randomly select a point uniformly from the interval  $[0, T]$ :  $T' \sim \text{Un}(0, T)$ .*

*Step 2. Change one of the jump times closest to  $T'$  as follows:*

*if  $T'$  falls on the first interval, let  $T_1 = T'$ ;*

*if  $T'$  falls on the last interval, let  $T_n = T'$ ;*

*if  $T'$  falls on an interval,  $[T_i, T_{i+1}]$ , let  $T_i = T'$  with probability  $\frac{1}{2}$ ; otherwise let  $T_{i+1} = T'$ .*

The birth-death sampler sub-algorithm generates a proposed realisation  $Y$  by creating a new interval, a birth, in the current realisation of  $D$ , or by removing an interval, a death, from the current realisation of  $D$ .

See Appendix 1 for a detailed description of the birth-death sampler algorithm and for a description of the types of birth and death moves that are allowed.

### 2.3. Convergence diagnostics

To help assess when our Markov chains have converged in distribution, we implemented the commonly adopted Gelman & Rubin statistic (Gelman & Rubin, 1992) and the  $L^1$  statistic, which is a diagnostic that we introduce specifically for this context to assess the mixing of the hidden Markov chain. For a general review of convergence diagnostics see Cowles & Carlin (1996) or Brooks & Roberts (1999).

The  $L^1$  statistic monitors the behaviour of  $D$  for  $m \geq 2$  independent realisations of the Markov chain Monte Carlo Markov chain. When the Markov chain components from all independent Markov chain Monte Carlo Markov chains begin to exhibit similar behaviour, there is evidence that the Markov chain component,  $D$ , may have converged in distribution. Our diagnostic is based on

$$\delta(D_1, D_2) = \int_0^T I_{\{D_{1,s} \neq D_{2,s}\}} ds,$$

defined as the time that two Markov chains  $D_1$  and  $D_2$  are in different states. Given  $\delta$ , the  $L^1$  statistic,

$$L^1 = \frac{2 \sum_{j < i} \delta(D_i, D_j)}{m(m-1)},$$

is the average distance between the Markov chain elements. When the  $L^1$  statistic begins to exhibit stationary behaviour, we conclude that the marginal densities of  $D$  have converged in distribution.

### 2.4. Model choice

For model choice we calculate Bayes factors and other tools that are based on an in-sample estimation of  $X$ .

Given a set of data  $X = \{X_{t_i} : T = t_n > \dots > t_0 = 0\}$  and results from a Markov chain

Monte Carlo analysis,  $\hat{X}$  is calculated as follows:

$$\hat{X}_{t_j} = \sum_{i=1}^j \delta \hat{X}_{t_i} + X_{t_0},$$

where  $\delta \hat{X}_{t_i}$  is the expected change in the process  $X$  over the interval  $[t_{i-1}, t_i]$ , taking the expectation with respect to the posterior distribution of the parameters together with the incremental distribution of  $X$ . For a two-state model the expected change is given by

$$\delta \hat{X}_{t_i} = \{\text{pr}(D_t = 1 | \theta_{-D}, X) \hat{\mu}_1 + \text{pr}(D_t = 2 | \theta_{-D}, X) \hat{\mu}_2\} (t_i - t_{i-1}),$$

where  $\theta_{-D}$  represents all of the model parameters except for  $D$ . A plot of  $X$  and  $\hat{X}$ , which we call the visual goodness of fit, offers a simple way of assessing a model.

We also calculate a traditional  $R^2$  statistic as a crude measure of the amount of variation in the data explained by the model and we calculate the average  $L^1$ -distance,  $\bar{L}^1$ , between  $\hat{X}$  and  $X$ , where

$$R^2 = 1 - \frac{\sum_{i=1}^n (X_{t_i} - \hat{X}_{t_i})^2}{\sum_{i=1}^n (X_{t_i} - \bar{X}_{t_i})^2}, \quad \bar{L}^1 = \frac{1}{n} \sum_{i=1}^n |X_{t_i} - \hat{X}_{t_i}|.$$

### 3. DYNAMICS OF THE NEW YORK MERCANTILE EXCHANGE OIL MARKET

#### 3.1. *The data and model assumptions*

We now use our methods to investigate the dynamics of the price of a New York Mercantile Exchange future contract for crude oil to be delivered to New York harbour in June 1995; see Fig. 1(a). The price was recorded at ten-minute intervals during trading hours from 4th–23rd May 1995, and was reported in U.S. cents per barrel. The data reveal price changes over ten-minute intervals, intra-day price movements, and price changes overnight or over a weekend.

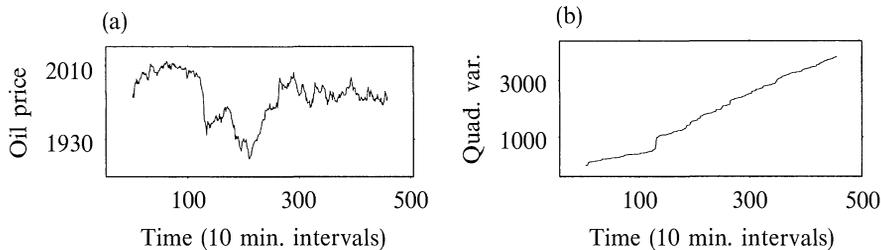


Fig. 1. U.S. oil market example. (a) shows oil price 4–23 May 1995, observed every 10 minutes and reported in U.S. cents per barrel. (b) shows quadratic variation process for the data in (a).

We analyse the intra-day data using several linear hidden Markov chain models. For each analysis we assume that the price at the beginning of a trading session is the same as the price at the end of the last session or we ignore the overnight price movements. This assumption is reasonable for this dataset because the overnight price movements are of the same order of magnitude as the intra-day price movements. When the overnight price movements are significant, several approaches can be used to model these price movements. For example, overnight movements could be modelled as normal increments with a mean and variance that depends on the hidden Markov chain  $D$ .

Our other major assumptions are that the instantaneous volatility is constant and that

it is reasonable to estimate  $\sigma^2$  using (2.1). This estimate of the volatility is calculated a priori and then used as a covariate. The assumption of constant volatility is reasonable over relatively short time periods; see Polson & Roberts (1994) for further discussion of issues related to heteroscedasticity and estimation of volatility from discretely observed pairs. In particular, they introduced the following informal diagnostic of heteroscedasticity. Consider  $Q_t = \int_0^t \sigma_t^2 dt$ , which is approximated by

$$\hat{Q}_{t_j} = \sum_{i=1}^j \frac{(X_{t_i} - X_{t_{i-1}})^2}{t_i - t_{i-1}}.$$

This estimate of the quadratic variation process can be used as a visual diagnostic for the assumption of constant variance, under which we expect  $Q_t$  to be approximately linear. The quadratic variation for the oil data is given in Fig. 1(b); except for a brief period where the oil price drops sharply, the volatility appears relatively constant. During the period when the slope of  $\hat{Q}_{t_j}$  increases significantly, between observations 124 and 136, all three of our models identify a negative drift that is strong enough to inflate  $\hat{Q}_t$  artificially. If  $\hat{Q}_t$  is adjusted for this drift term over the period in question then the slope of  $\hat{Q}_t$  becomes essentially constant over the entire dataset.

The assumption that the observations of  $X$  are dense enough to justify using (2.1) can be investigated by calculating the ratio of the estimate of the volatility over the estimate of the drift squared. If this ratio is large, when compared to 1, we assume that (2.1) gives a reasonable estimate of the volatility. For this dataset (2.1) results in  $\hat{\sigma}^2 = 8.35797$ . The ratio of volatility over each drift is given in Table 1. This ratio is less than one for the strongest decreasing drift for each model. Since the price tends to stay in this decreasing state for a short period of time, between 80 and 160 minutes depending on the model, we assume that the impact of the first drift is unimportant. The ratio for the remaining drifts is greater than one, which suggests that these drifts have little or no impact on the volatility estimate. These ratios, together with the estimates of the quadratic variation process in Fig. 1, suggest that it is reasonable to estimate the volatility using (2.1).

Table 1. *U.S. oil market example. Ratio of estimates of volatility over (drift)<sup>2</sup>,  $\hat{\sigma}^2/\hat{\mu}_i^2$ . As this ratio increases, the influence of the drift on the estimate of the volatility decreases*

Ratio	Two-state model	Three-state model	Four-state model
$\hat{\sigma}^2/\hat{\mu}_1^2$	0.13	0.14	0.12
$\hat{\sigma}^2/\hat{\mu}_2^2$	7256.94	109.11	8.60
$\hat{\sigma}^2/\hat{\mu}_3^2$	—	8.26	457.72
$\hat{\sigma}^2/\hat{\mu}_4^2$	—	—	4.06

### 3.2. Model choice

The Bayes factors comparing the two-state to the five-state linear hidden Markov chain model with the random walk model, Brownian motion without a drift, were calculated using the fourth and final sampling based estimator proposed by Newton & Raftery (1994). The logarithms of these Bayes factors are reported in Table 2.

The Bayes factors give very strong support for all the hidden Markov chain models when compared with the random walk model. Given the Bayes factors for the competing

Table 2. U.S. oil market example. Bayes factors, BF, on a log scale with respect to the random walk model and  $R^2$  and  $\bar{L}^1$  statistics for the two-state–five-state models

Model	log(BF)	$R^2$	$\bar{L}^1$
Two-state	18.89	0.65	15.37
Three-state	21.21	0.82	10.56
Four-state	24.01	0.89	8.49
Five-state	24.41	0.91	7.67

linear models, we conclude that the four-state model is the one to choose. The four-state model is chosen over the five-state model because, although the Bayes factor of 1.5 between the two models supports the five-state model, in our view this is not strong enough to justify the increased complexity of the five-state model. This conclusion is supported by the visual goodness of fit,  $R^2$  and the  $\bar{L}^1$  diagnostics; see Fig. 2 and Table 2.

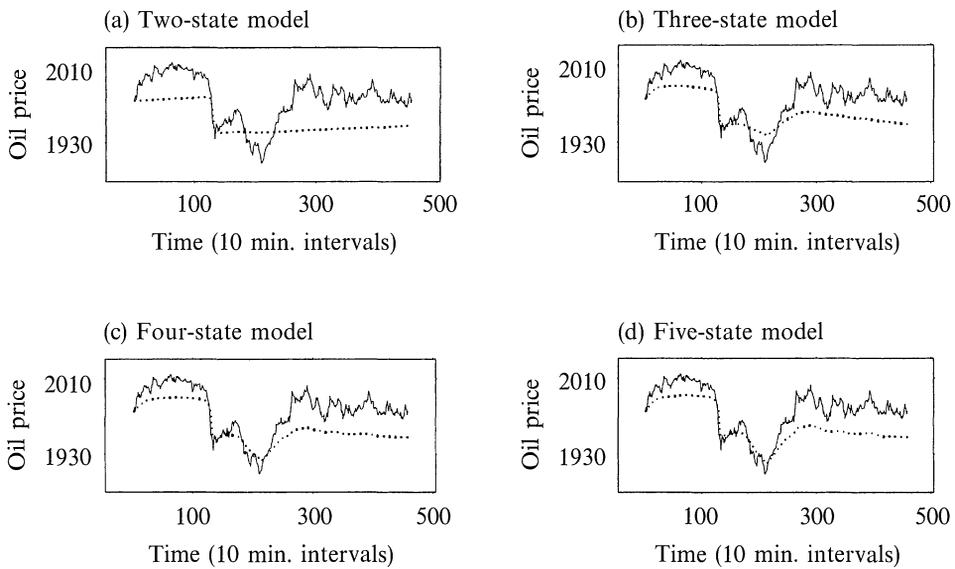


Fig. 2. U.S. oil market example. Visual goodness of fit for the two-state–five-state models. Solid lines, oil price in U.S. cents per barrel; dotted lines, visual goodness of fit.

### 3.3. Parameter estimates for the four-state model

The posterior density for the hidden four-state Markov chain  $D$  is summarised by plotting, for each point in time, the probability of being in each state; see Fig. 3.

As with all the linear hidden Markov chain models, the four-state model identifies a strong decreasing state that lasts for a short period of time; see Table 3 for estimates of the drifts and waiting times. These data also exhibit a weak decreasing state, a flat state and an increasing state. Given that a jump occurred, the probabilities of jumping between each of these states were estimated and are reported in Table 4.

This study offers interesting insights into the dynamics of the crude oil market over the period that the data were observed. The four-state model indicates that, after the market experiences a negative correction, a crash, the market has roughly an equal chance of

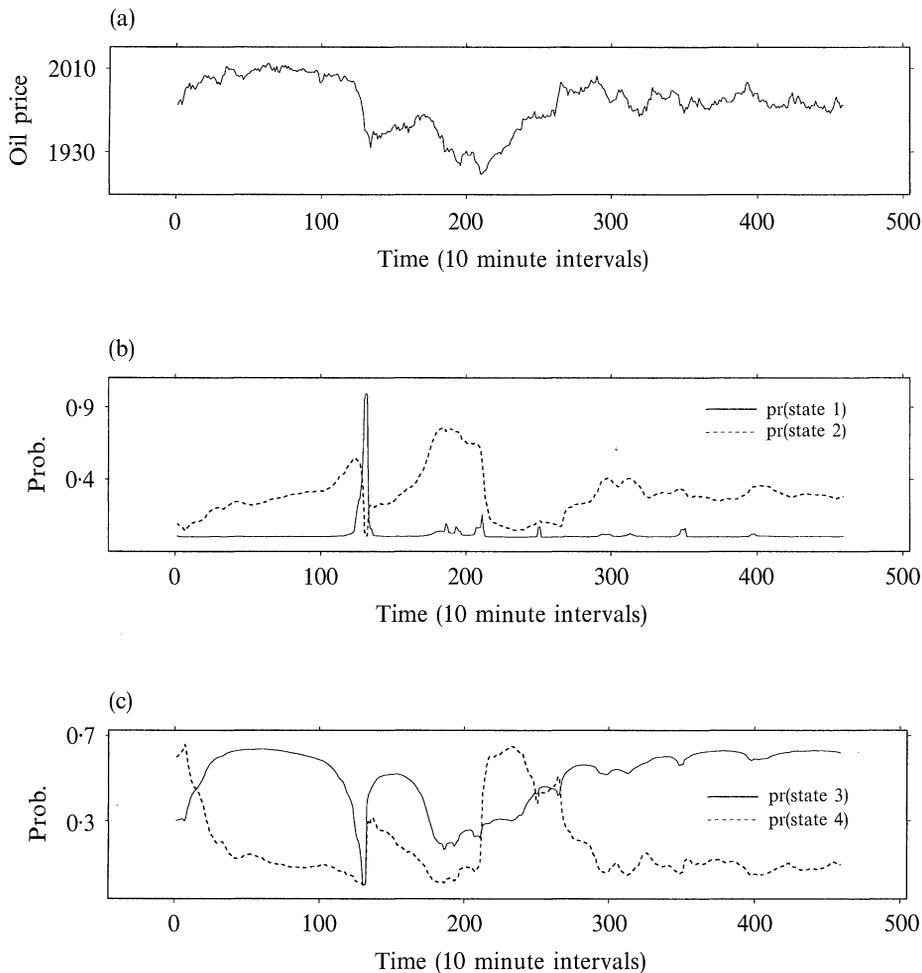


Fig. 3. U.S. oil market example. (a) displays oil prices, observed every 10 minutes and reported in U.S. cents per barrel. For the four-state model, (b) shows the probabilities of being in the decreasing states, states 1 and 2, and (c) shows the probabilities of being in the increasing states, states 3 and 4.

Table 3. U.S. oil market example. Summary statistics for the four-state model slopes or drifts,  $\mu_i$ , given in U.S. cents per 10 minutes, and inverses of the expected waiting times,  $\lambda_{ii}$ , given in (10 minutes)<sup>-1</sup>

State	Point estimate of $\mu_i$	Standard deviation	Point estimate of $\lambda_{ii}$	Standard deviation
1	-8.221	2.719	0.123	0.060
2	-0.986	0.973	0.033	0.019
3	0.135	0.404	0.022	0.014
4	1.435	1.441	0.074	0.047

jumping to any of other states. Alternatively, when an increasing state, a rally, finishes, the market tends to 'land softly' or tends to not change to the strong decreasing state. This soft landing behaviour holds for both the flat and increasing states.

Table 4. *U.S. oil market example. Estimate of the jump chain transition matrix for the four-state model*

State	State			
	1	2	3	4
1	0.0000	0.3004	0.3553	0.3443
2	0.3053	0.0000	0.3354	0.3593
3	0.2282	0.4442	0.0000	0.3276
4	0.1962	0.3418	0.4620	0.0000

We realise that this study is based on a small sample and that these observations about the dynamics of the oil market may or may not apply to larger sets of data.

3.4. *Convergence diagnostics*

We calculated the convergence diagnostics mentioned in § 2.3, based on five replications of the Markov chain Monte Carlo Markov chain for each analysis. The convergence diagnostics were similar for each analysis, and we only report the convergence diagnostics for the four-state model.

Plots of the Gelman & Rubin (1992) statistics and the  $L^1$  statistic, given in Fig. 4, suggest that the four-state model appears to have converged by around 15 000 iterations.

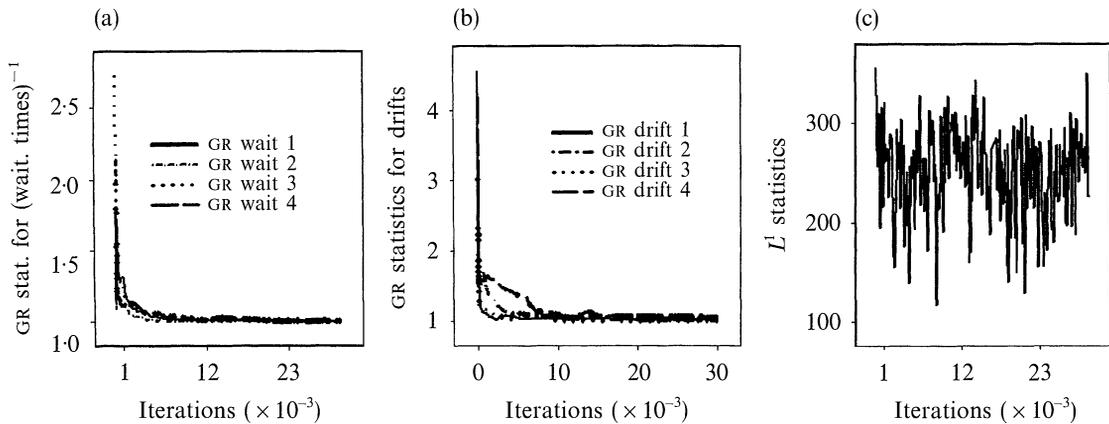


Fig. 4. *U.S. oil market example. For the four-state model: (a) shows Gelman & Rubin (1992), GR, statistics for the inverse of the waiting times, GR wait 1–GR wait 4; (b) shows Gelman & Rubin statistics for the drift terms, GR drift 1–GR drift 4; (c) shows the  $L^1$  statistic.*

In addition, the acceptance rates for the two-, three- and four-state models are given in Table 5. As expected, the acceptance rates for the independence sampler are very low as

Table 5. *U.S. oil market example. Acceptance rates for the two-state, three-state and four-state models*

	Acceptance rates		
	Two-state model	Three-state model	Four-state model
Independence sampler	0.000211	0.000026	0.000189
Refinement sampler	0.006346	0.062231	0.220441
Birth-death sampler	0.003830	0.026512	0.098163

a consequence of the algorithm's ambitious global proposals. However, the refinement sampler and birth-death sampler exhibit higher and reasonable acceptance rates, even though the proposed moves can still substantially alter the hidden Markov sample path in one iteration.

#### 4. NON-MARKOV MODELS

##### 4.1. Model description

Our methods can be extended and used to analyse a wide range of hidden Markov models and hidden non-Markov models. For example, these methods can be extended for a hidden non-Markov model, where  $D$  is a continuous-time Markov chain and  $X$  is either a linear diffusion or an Ornstein–Uhlenbeck process with a level of attraction that depends on a past level of  $X$ ; see Øksendal (1980, p. 100) for a discussion of the Ornstein–Uhlenbeck process.

The Itô diffusion associated with this linear Ornstein–Uhlenbeck model is given by

$$dX_t = \left\{ \sum_{i=1}^{m-1} \mu_i I_{\{D_t=i\}} + \beta I_{\{D_t=m\}} (X_{t_m(t)} - X_t) \right\} dt + \sigma dW_t,$$

where  $D$  is an  $m$ -state Markov chain and where

$$t_m(t) = \max_{s \leq t} \{s : dD_s \neq 0, D_r = m (s < r \leq t)\}$$

is the most recent time that  $D$  jumped to the Ornstein–Uhlenbeck state. The linear Ornstein–Uhlenbeck model switches between a collection of linear diffusions and an Ornstein–Uhlenbeck process where the level of attraction for the Ornstein–Uhlenbeck state is the value of  $X$  at the time that  $D$  enters the Ornstein–Uhlenbeck state.

The linear Ornstein–Uhlenbeck model is non-Markov because  $X_{t_m(t)}$  is a function of  $X$  at the last time that  $D$  enters the Ornstein–Uhlenbeck state, which is a random time in the past.

##### 4.2. Markov chain Monte Carlo methods

The linear Ornstein–Uhlenbeck model can be analysed using a modification of the basic algorithms presented in § 2. The modifications to the algorithm and each of the sub-algorithms are straightforward. The only substantial modification required concerns the formula for calculating the acceptance probability,  $\alpha$ , associated with modified versions of the refinement sampler and birth-death sampler sub-algorithms; see Appendix 2. The acceptance probabilities need to be modified to accommodate proposed changes to  $D$  where the beginning of an Ornstein–Uhlenbeck interval is changed. As an example we give the acceptance probability for a proposed refinement move where the beginning of an Ornstein–Uhlenbeck interval is modified:

$$\alpha(D, Y) = \min \left\{ \frac{\pi(Y|X, \theta_{-Y})}{2\pi(D|X, \theta_{-D})}, 1 \right\},$$

where  $\theta_{-D}$  stands for all parameters except for the hidden Markov chain  $D$  and

$$\begin{aligned} \frac{\pi(Y|X, \theta_{-Y})}{\pi(D|X, \theta_{-D})} = \exp \left\{ \frac{1}{\sigma^2} \int_{T'}^{T_2} \beta(X_{T'} - X_s) dX_s - \frac{1}{2\sigma^2} \int_{T'}^{T_2} \beta^2(X_{T'} - X_s)^2 ds \right. \\ \left. + \frac{1}{\sigma^2} \int_{T_1}^{T_2} \beta(X_{T_1} - X_s) dX_s - \frac{1}{2\sigma^2} - \frac{1}{\sigma^2} \int_{T_1}^{T_2} \mu_1 dX_s \right. \\ \left. + \frac{1}{2\sigma^2} \int_{T_1}^{T_2} \mu_1^2 ds + (\lambda_1 - \lambda_m)(T_1 - T') \right\}. \end{aligned}$$

Appropriate modifications to the acceptance probability for a modified birth-death sampler sub-algorithm can be arrived at using a similar approach.

### 4.3. Analysis of simulated data

We generated three-state linear Ornstein–Uhlenbeck data where  $D$  enters an Ornstein–Uhlenbeck state at three different times. We then analysed these data using a three-state Ornstein–Uhlenbeck model and a three-state linear model, showing that the linear model has difficulty in identifying the mean reverting state.

For both analyses we used appropriate versions of the independence sampler, refinement sampler and birth-death sampler sub-algorithms to update  $D$ . The visual goodness of fit and the probabilities of being in the decreasing state and the Ornstein–Uhlenbeck state for the three-state linear Ornstein–Uhlenbeck model, and the visual goodness of fit and probabilities of being in the increasing and decreasing states for the three-state linear model are given in Fig. 5. In addition the model choice tools for both models are given in Table 6.

It is clear from the visual goodness of fit plot, in Fig. 5, that the three-state linear model has difficulty distinguishing between the Ornstein–Uhlenbeck state, which ideally should be treated as a state with drift equal to zero, and the decreasing state. This is not surprising, given that the process that generated the data exhibits both linear trending behaviour and mean reverting behaviour. While the generating process is in the mean reverting state the process occasionally moves away from and then back towards a fixed level of attraction, creating behaviour which could look like a collection of short-term linear trends. Since the linear model does not have a mean reverting component, the short term trends from the Ornstein–Uhlenbeck state compete with the longer trends of the decreasing and increasing states. This ‘competition’ can make it difficult for the linear model to identify the Ornstein–Uhlenbeck state properly.

In contrast, the linear Ornstein–Uhlenbeck model clearly identifies all three of the states, see Fig. 5, and, as measured by our model choice tools in Table 6, there is very strong evidence, with a Bayes factor of  $\exp(32.35)$ , for choosing the linear Ornstein–Uhlenbeck model over the linear model.

In conclusion, if data are generated by a process that has linear and mean reverting states, a linear hidden Markov chain model will fit the data better than a model with a constant drift, but a linear Ornstein–Uhlenbeck model will be better at identifying the mean reverting state and will describe the data better than either a model with a constant drift or a linear hidden Markov chain model.

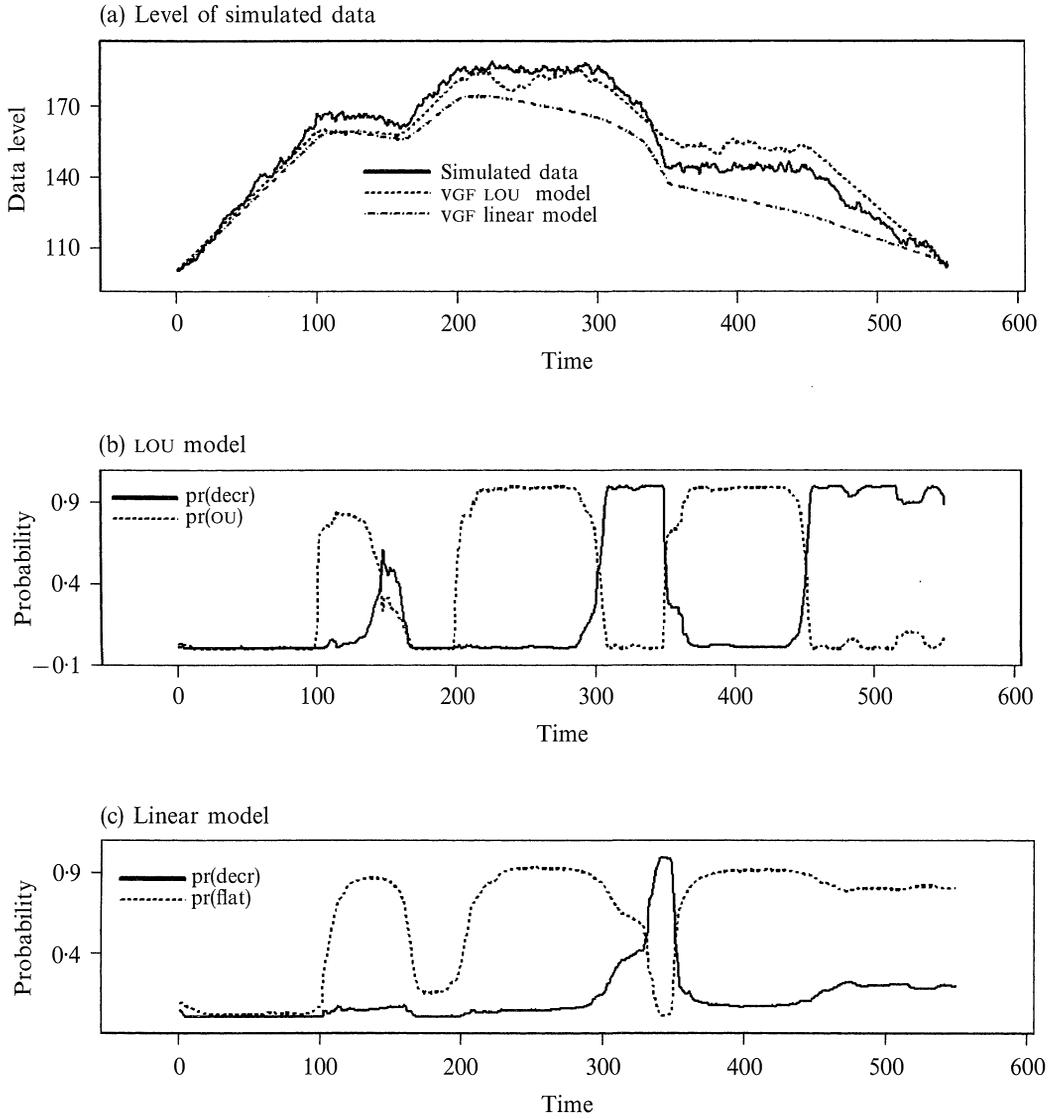


Fig. 5. Simulated data example using three-state linear Ornstein-Uhlenbeck, LOU, model and linear model. (a) shows levels of simulated data and visual goodness of fit, vGF. (b) and (c) show probabilities of being in a decreasing state,  $pr(decr)$ , flat state,  $pr(flat)$  or Ornstein-Uhlenbeck state,  $pr(ou)$ , (b) for three-state linear Ornstein-Uhlenbeck model, and (c) for three-state linear model.

Table 6. Simulated data example. Bayes factor, BF, on a log scale with respect to the random walk model, and  $R^2$  and  $\bar{L}^1$  statistics for the three-state linear Ornstein-Uhlenbeck, LOU, model and for the three-state linear model

Model	$\log(BF)$	$R^2$	$\bar{L}^1$
Three-state LOU	62.62	0.945	5.28
Three-state linear	30.37	0.805	9.92

## APPENDIX 1

*The birth-death sampler algorithm*

First let  $\mathfrak{g}_D = \sum_{i \neq j} \mathfrak{g}_{ij}$  be the number of times a Markov chain  $D$  changes state over an interval  $[0, T]$ . As mentioned in § 2, the birth-death sampler sub-algorithm generates a proposed realisation  $Y$  by creating a new interval in the current realisation of  $D$ , or by removing an interval from the current realisation of  $D$ . Three different types of birth are allowed, a left birth, a right birth and a middle birth, and three different types of death, a left death, a right death and a middle death; see Fig. 6.

For a two-state model, most of the birth modifications are the same as a refinement modification. For example, a left birth on any interval except for the first interval and a right birth on any interval except for the last interval reduce to a refinement modification. For models with more than two states, some of the birth modifications are the same as a refinement modification. For example, a left birth reduces to a refinement modification when the state of the new interval is the same as the state of the interval immediately to its left. A similar statement can be made for right births.

In order to distinguish clearly between the refinement sampler and the birth-death sampler sub-algorithms, we restrict the birth-death sampler so that it cannot propose births that would result in a refinement modification. This restriction will in turn result in two different versions of the birth-death sampler sub-algorithm, one version for models with two states and another version for models with more than two states. Since these sub-algorithms are similar we will only give the version for models with more than two states.

SUB-ALGORITHM 3 (*The birth-death sampler for three or more states*).

Step 1. Select a birth with probability  $\rho = (\frac{1}{2})I_{\{\mathfrak{g}_D > 0\}} + (1)I_{\{\mathfrak{g}_D = 0\}}$ ; otherwise select a death.

Step 2. Select the interval to be modified by generating a random time,  $T' \sim \text{Un}(0, T)$ . The interval that contains  $T'$  is the interval that will be modified.

Step 3. If a birth is selected in Step 1 use (a); otherwise use (b)–(d).

(a) If a birth is selected, with equal probabilities create the birth interval using a middle birth, a left birth or a right birth. For a middle birth generate  $T'' \sim \text{Un}(T_i, T_{i+1})$ , and let  $[\min\{T', T''\}, \max\{T', T''\}]$  become the birth interval. For a left birth let  $[T_i, T']$  become the birth interval. For a right birth let  $[T', T_{i+1}]$  become the birth interval.

(b) If the first interval is to be removed, absorb it into the second interval.

(c) If the last interval is to be removed, absorb it into the next to last interval.

(d) If an internal interval is to be removed, with equal probabilities absorb it into the interval immediately preceding or immediately following.

Step 4. If a birth is selected in Step 1, choose a new state by excluding the state of the interval selected in Step 2 and by excluding any states which would result in a refinement type move, and then select one of the remaining states with equal probabilities.

## APPENDIX 2

*Acceptance probabilities*

This appendix contains the acceptance probabilities for Sub-algorithms 1 and 2 presented in § 2 and for Sub-algorithm 3 presented in Appendix 1. In each case the acceptance probability  $\alpha$  is a function of the current values of the scalar parameters, the data and the competing realisations of the hidden Markov chain.

*The independence sampler.* The acceptance probability for Sub-algorithm 1 is given by

$$\alpha(D, Y) = \min \left\{ \frac{L(X|Y, \mu_1, \mu_2)}{L(X|D, \mu_1, \mu_2)}, 1 \right\}.$$

*The refinement sampler.* Recall that we call the first and last intervals of a Markov chain  $D$

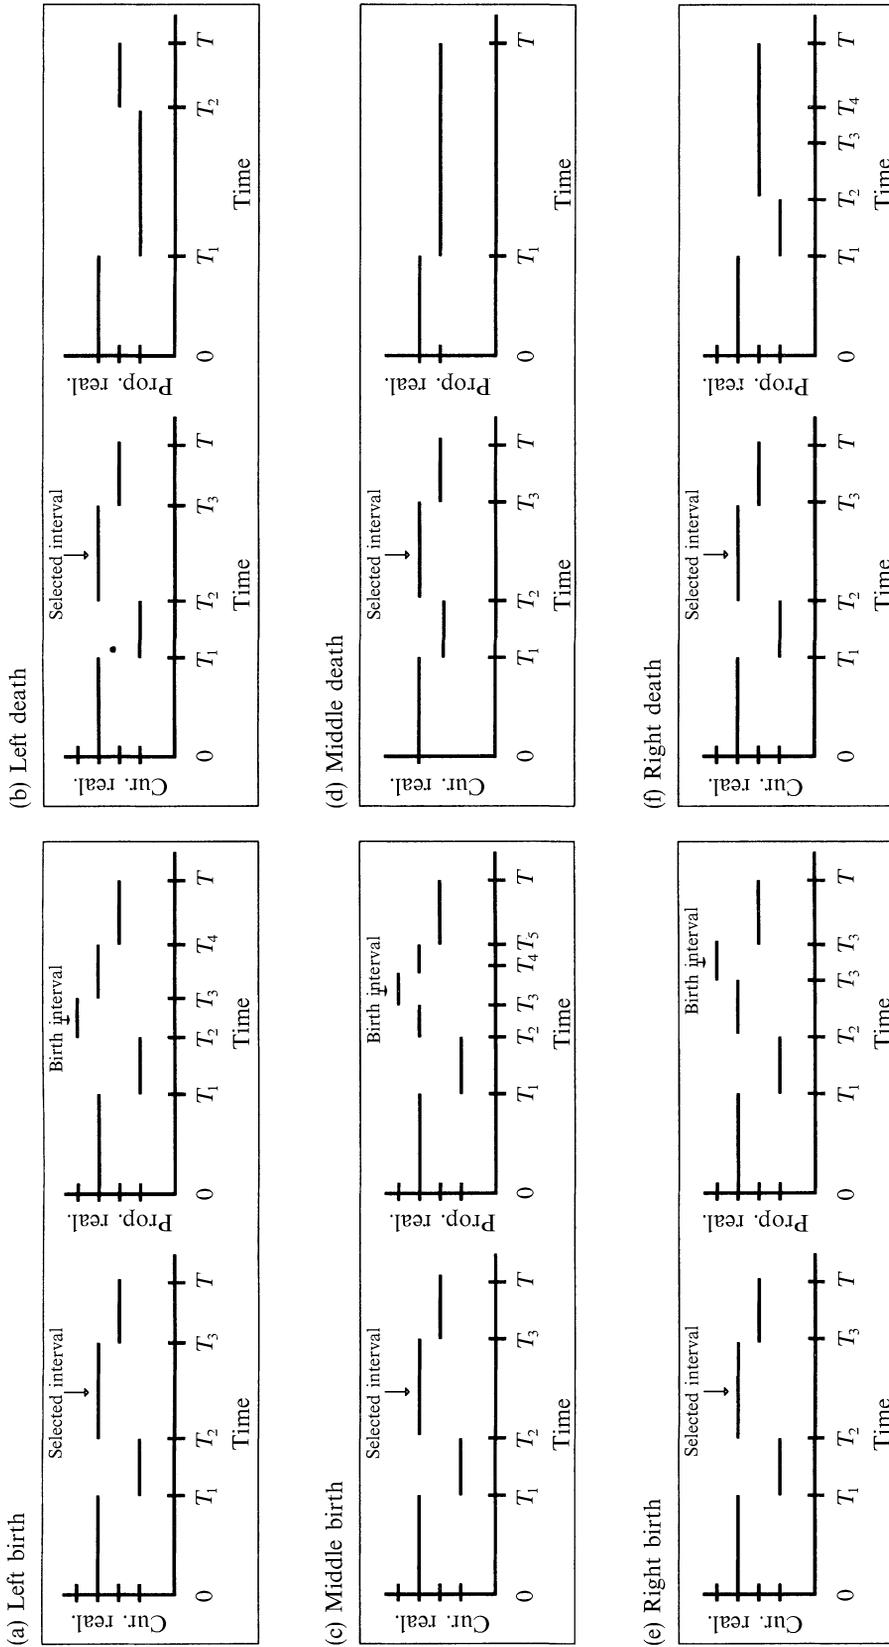


Fig. 6. Possible births and deaths. Current realisations and proposed realisations plotted against time.

external intervals and we call any remaining intervals internal intervals. In addition, recall that  $\vartheta_D$  represents the number of times  $D$  changes state on the interval  $[0, T]$  and that  $T_i$  represents the time when the  $i$ th jump occurs.

The acceptance probability for Sub-algorithm 2 is given by

$$\alpha(D, Y) = \min \left\{ a \frac{\pi(Y|X, \theta_{-Y})}{\pi(D|X, \theta_{-D})}, 1 \right\},$$

where  $a = \frac{1}{2}$  if an external interval is selected and  $\vartheta_D > 1$ ,  $a = 2$  if an internal interval is selected and  $T_1$  or  $T_n$  is modified,  $a = 1$  for all other cases, and  $T_1$  is the first jump time of  $D$  and  $T_n$  is the last jump time of  $D$ .

*The birth-death sampler.* The following notation helps simplify the presentation:  $\vartheta_D$  is the number of times  $D$  changes state;  $t_s$  is the length of the selected interval;  $t_b$  is the length of the new birth interval;  $t_d$  is the length of the interval that results from a death;  $T_{s-}$  and  $T_{s+}$  are the times at which the selected interval stops and starts;  $T_{b-}$  and  $T_{b+}$  are the times at which the birth interval stops and starts;  $i$  is the state of the selected interval;  $j$  is the state of the interval preceding the selected interval;  $k$  is the state of the interval following the selected interval;  $l$  is the state of the birth interval; and  $m$  is the number of elements in the state space of  $D$ .

The acceptance probability  $\alpha$  for the birth-death sampler algorithm can be written in terms of the product of ratios of likelihoods, prior densities and proposal measures. This acceptance probability is given by

$$\alpha(D, Y) = \min \left\{ \frac{q_p(Y, D)\pi_Y(Y|X, \theta_{-Y})}{q_p(D, Y)\pi_D(D|X, \theta_{-D})}, 1 \right\},$$

where

$$\frac{q_p(Y, D)\pi_Y(Y|X, \theta_{-Y})}{q_p(D, Y)\pi_D(D|X, \theta_{-D})} = \frac{q_p(Y, D)}{q_p(D, Y)} \times \frac{p(Y|Q)}{p(D|Q)} \times \frac{L(X|Y, \theta_{-Y})}{L(X|D, \theta_{-D})}.$$

We will proceed by defining each of these ratios for models with three or more states. First we will give the ratio of proposal densities. For birth modifications,

$$\frac{q_p(Y, D)}{q_p(D, Y)} = \begin{cases} \frac{3}{2}t_b(m-1), & \text{for } \vartheta_D = 0, \text{ left or right birth,} \\ \frac{3}{2}t_b(m-2), & \text{for } \vartheta_D > 0, \text{ left or right birth of an internal interval,} \\ 3t_b(m-1), & \text{for } \vartheta_D > 0, \text{ birth of an external interval,} \\ \frac{3}{2}t_s t_b(m-1), & \text{for } \vartheta_D = 0, \text{ middle birth,} \\ 3t_s t_b(m-1), & \text{for } \vartheta_D > 0, \text{ middle birth.} \end{cases}$$

For death modifications,

$$\frac{q_p(Y, D)}{q_p(D, Y)} = \begin{cases} \frac{2}{3}\{t_s(m-1)\}^{-1}, & \text{for } \vartheta_D = 1, \text{ left or right birth,} \\ \frac{2}{3}\{t_s(m-2)\}^{-1}, & \text{for } \vartheta_D > 1, \text{ left or right death of an internal interval,} \\ \frac{2}{3}\{t_s(m-1)\}^{-1}, & \text{for } \vartheta_D > 1, \text{ death of an external interval,} \\ \frac{2}{3}\{t_d t_s(m-1)\}^{-1}, & \text{for } \vartheta_D = 2, \text{ middle death,} \\ \frac{1}{3}\{t_d t_s(m-1)\}^{-1}, & \text{for } \vartheta_D > 2, \text{ middle death.} \end{cases}$$

Next we will give the ratio of the prior densities for the two competing Markov chains. For birth modifications,

$$\frac{P(Y|Q)}{P(D|Q)} = \begin{cases} \lambda_{jl}\lambda_{li}\lambda_{ji}^{-1} \exp\{(\lambda_i - \lambda_l)t_b\}, & \text{for left birth not the first interval,} \\ \lambda_{li} \exp\{(\lambda_i - \lambda_l)t_b\}, & \text{for birth of the first interval,} \\ \lambda_{il}\lambda_{li} \exp\{(\lambda_i - \lambda_l)t_b\}, & \text{for middle birth,} \\ \lambda_{il}\lambda_{ik}\lambda_{ik}^{-1} \exp\{(\lambda_i - \lambda_l)t_b\}, & \text{for right birth not the last interval,} \\ \lambda_{il} \exp\{(\lambda_i - \lambda_l)t_b\}, & \text{for birth of the last interval.} \end{cases}$$

For death modifications,

$$\frac{P(Y|Q)}{P(D|Q)} = \begin{cases} \lambda_{jk} \lambda_{ji}^{-1} \lambda_{ik}^{-1} \exp\{(\lambda_i - \lambda_j)t_s\}, & \text{for left death not the last interval,} \\ \lambda_{ji}^{-1} \exp\{(\lambda_i - \lambda_j)t_s\}, & \text{for death of the last interval,} \\ \lambda_{ji}^{-1} \lambda_{ij}^{-1} \exp\{(\lambda_i - \lambda_j)t_s\}, & \text{for middle death,} \\ \lambda_{jk} \lambda_{ji}^{-1} \lambda_{ik}^{-1} \exp\{(\lambda_i - \lambda_k)t_s\}, & \text{for right death not the first interval,} \\ \lambda_{ik}^{-1} \exp\{(\lambda_i - \lambda_k)t_s\}, & \text{for death of the first interval.} \end{cases}$$

Finally we will give the ratio of the likelihoods of  $X$  given the two competing Markov chains. For birth modifications,

$$\frac{L(X|Y, \mu_1, \dots, \mu_m)}{L(X|D, \mu_1, \dots, \mu_m)} = \exp \left\{ \int_{T_{b-}}^{T_{b+}} \left( \frac{\mu_l - \mu_i}{\sigma^2} \right) dX_s - \frac{(\mu_l^2 - \mu_i^2)t_b}{2\sigma^2} \right\}.$$

For death modifications,

$$\frac{L(X|Y, \mu_1, \dots, \mu_m)}{L(X|D, \mu_1, \dots, \mu_m)} = \exp \left\{ \int_{T_{s-}}^{T_{s+}} \left( \frac{\mu_\zeta - \mu_i}{\sigma^2} \right) dX_s - \frac{(\mu_\zeta^2 - \mu_i^2)t_s}{2\sigma^2} \right\},$$

where  $\zeta = j$  for a left or middle death and  $\zeta = k$  for a right death.

#### REFERENCES

- BROOKS, S. P. & ROBERTS, G. O. (1999). Assessing convergence of Markov chain Monte Carlo algorithms. *Statist. Comp.* **8**, 319–35.
- COWLES, M. K. & CARLIN, B. P. (1996). Markov chain Monte Carlo convergence diagnostics: a comparative review. *J. Am. Statist. Assoc.* **91**, 883–904.
- GELMAN, A. & RUBIN, D. (1992). Inference from iterative simulation using multiple sequences (with Discussion). *Statist. Sci.* **7**, 457–511.
- GREEN, P. J. (1995). Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* **82**, 711–32.
- HODGSON, M. E. A. (1999). A Bayesian restoration of an ion channel signal. *J. R. Statist. Soc. B* **61**, 95–114.
- HODGSON, M. E. A. & GREEN, P. J. (1999). Investigating Markov model discrimination for ion channels. *Proc. R. Soc. Lond. A* **455**, 3425–48.
- NEWTON, M. A. & RAFTERY, A. E. (1994). Approximate Bayesian inference with the weighted likelihood bootstrap (with Discussion). *J. R. Statist. Soc. B* **56**, 3–48.
- ØKSENDAL, B. (1980). *Stochastic Differential Equations*. New York: Springer Verlag.
- POLSON, N. G. & ROBERTS, G. O. (1994). Bayes factors for discrete observations from diffusion processes. *Biometrika* **81**, 11–26.
- TIERNEY, L. (1998). A note on Metropolis–Hastings kernels for general state spaces. *Ann. Appl. Prob.* **8**, 1–9.

[Received December 1998. Revised June 2000]