Quantifying parameter uncertainty in a coral reef model using Metropolis-Coupled Markov Chain Monte Carlo

Damian Clancy[a], Jason E. Tanner[b], Stephen McWilliam[a], Matthew Spencer*[,c]

[a]*Department of Mathematical Sciences, University of Liverpool, Liverpool, UK*
[b]*SARDI Aquatic Sciences, Henley Beach; School of Earth and Environmental Sciences, University of Adelaide, Australia*
[c]*School of Environmental Sciences, University of Liverpool, Liverpool, UK. Phone +44 (0)151 795 4399, Fax +44 (0)151 795 4404.*

Abstract

Coral reefs are threatened ecosystems, so it is important to have predictive models of their dynamics. Most current models of coral reefs fall into two categories. The first is simple heuristic models which provide an abstract understanding of the possible behaviour of reefs in general, but do not describe real reefs. The second is complex simulations whose parameters are obtained from a range of sources such as literature estimates. We cannot estimate the parameters of these models from a single data set, and we have little idea of the uncertainty in their predictions.

We have developed a compromise between these two extremes, which is complex enough to describe real reef data, but simple enough that we can estimate parameters for a specific reef from a time series. In previous work, we fitted this model to a long-term data set from Heron Island, Australia, using maximum likelihood methods. To evaluate predictions from this model, we need estimates of the uncertainty in our parameters. Here, we obtain such estimates using Bayesian Metropolis-Coupled Markov Chain Monte Carlo. We do this for versions of the model in which corals are aggregated into a single state variable (the three-state model), and in which corals are separated into four state variables (the six-state model), in order to determine the appropriate level of aggregation. We also estimate the posterior distribution of predicted trajectories in each case.

In both cases, the fitted trajectories were close to the observed data, but we had doubts about the biological plausibility of some parameter estimates. We suggest that informative prior distributions incorporating expert knowledge may resolve this problem. In the six-state model, the posterior distribution of state frequencies after 40 years contained two divergent community types, one dominated by free space and soft corals, and one dominated by acroporid, pocilloporid, and massive corals. The three-state model predicts only a single community type. We conclude that the three-state model hides too much biological heterogeneity, but we need more data if we are to obtain reliable predictions from the six-state model. It is likely that there will be similarly large, but currently unevaluated, uncertainty in the predictions of other coral reef models, many of which are much more complex and harder to fit to real data.

*Corresponding author
Email address: m.spencer@liverpool.ac.uk (Matthew Spencer)

Key words: Coral reefs, Time series, Markov Chain Monte Carlo, Bayesian statistics, Community dynamics, Parameter uncertainty

1. Introduction

Coral reefs are one of several major ecosystem types for which models of the dynamics of sessile organisms are important (others include forests and rocky shores). Many reefs are thought to be at risk of large, perhaps irreversible, changes in composition, due to a combination of overfishing, disease and storm damage (Hughes, 1994). This has made the development of models that can make reliable long-term predictions of reef dynamics an important goal.

Many different kinds of models for coral reef ecosystems have been developed. At one extreme are generic, highly-simplified models whose value is in helping us to a qualitative understanding of the possible behaviour of the system (e.g. Mumby et al., 2007). At first sight, these models are too simple to describe particular real systems, and as a result, they are rarely fitted to data. At the other extreme are detailed simulation models (e.g. McClanahan, 1995; Langmead and Sheppard, 2004; Mumby et al., 2006). These models are mainly useful for numerical experiments. They cannot usually be fitted to data from a single system because they have too many parameters, or no rigorous way of comparing model output to data. Between these two extremes, it may be possible to produce models that can be fitted to real data, while at the same time having enough mechanistic detail to be able to make predictions. Early efforts in this direction focussed on linear Markov models (Tanner et al., 1994, 1996), which have also seen widespread use in other systems dominated by sessile organisms (e.g. Usher, 1979; Wootton, 2001c; Hill et al., 2004). More recently, attempts to add biological realism to this approach have resulted in nonlinear models whose mechanistic framework is no more complicated than that of the models used for qualitative understanding (Spencer and Tanner, 2008). Many of these studies have made use of the Heron Island data set. This long-term study of coral reef dynamics at Heron Island, Great Barrier Reef, Australia, provides a unique time series against which to test such models. The data consist of observations on the species present at large numbers of fixed points in space, at a series of irregularly-spaced times spanning almost thirty years. During this time, there have been dramatic changes to the composition of the reef (Connell et al., 1997).

In a previous study (Spencer and Tanner, 2008), we fitted a number of linear and nonlinear models to the Heron Island data using maximum likelihood methods. The best description of the data was a nonlinear Lotka-Volterra competition model, in which the intensities of transitions between states (such as species groups and free space) depend on the local abundance of the destination state. Linear Markov models, in which the intensities of transitions do not depend on local abundances, gave fitted trajectories that were much less like the observed data. The dynamics of this reef may therefore be dominated by local processes such as colony growth and interaction, rather than external processes such as dispersal.

For our model to be useful, we need to know how uncertain we are about parameter estimates, and how this uncertainty propagates through to the output of the model. In principle, this is straightforward in a maximum likelihood framework. However, for technical reasons (section 2.5) we were not able to estimate parameter uncertainty in our previous study. Here, we use Bayesian methods to estimate the joint posterior distribution of the parameters of interest around the estimates found by likelihood maximization methods. We first describe, in section 2.6, a basic Metropolis-Hastings algorithm which in principle allows us to sample from the required posterior distribution. However, this basic algorithm is impractically slow. The likelihood surface is multimodal, and the MCMC sampler only rarely moves from one mode to another. To solve the problem, we go on in section 2.7 to describe a Metropolis-Coupled Markov Chain Monte Carlo (MCMCMC: Geyer, 1991; Gilks and Roberts, 1996) method. MCMCMC runs a number of Markov chains in parallel, one of which (the cold chain) has as its stationary distribution the required posterior distribution of our parameters, while the others (heated chains) have similar but flatter distributions, making it easier to move between modes. Swapping between chains allows the cold chain to benefit from the improved mixing of the heated chains. This method allows us to search effectively around a number of starting points, which is particularly useful when the log likelihood is multimodal. The combination of initializing chains at modes found by a deterministic method, and designing the chains to jump between modes, has been suggested previously (Gelman, 1996, pp. 135-136).

We then use the joint posterior distribution of our parameters to estimate the induced posterior on the dynamics of the reef system. We also explore the effects of changing the level of aggregation of state variables. Models in which organisms are aggregated into too few state variables will be biased because of heterogeneity within these aggregated states. On the other hand, models with more state variables require more parameters, which increases the uncertainty in our estimates. In order to determine the right compromise between bias and uncertainty, we fit versions of our model containing three states (coral, algae, and free space) and six states (four different kinds of coral, algae, and free space). In earlier work, we used a six-state model. The three-state version requires fewer parameters, which is beneficial for estimation, but at the expense of aggregating several kinds of corals with different attributes into a single state. Finally, we discuss the implications of parameter uncertainty and model complexity for our ability to make predictions about the future dynamics of ecosystems.

2. Methods

2.1. Model structure

We use a model of community dynamics described in detail in Spencer and Tanner (2008). Briefly, we assume:

1. There is a fixed, finite number s of possible states for a point in space, of which one is the empty space state e and others are groups of species (e.g. acroporid corals, algae).

2. Conditional on its current state, the future state of a point is independent of its past states.
3. The rate of transitions of a point in state j to a non-empty state i depends on the abundance ($0 \le x_i \le 1$, dimensionless, $\sum_{i=1}^{s} x_i = 1$) of state i. These transitions represent colonists produced by state i successfully occupying space held by state j.
4. The rate of transitions from a non-empty state j to empty space e is independent of state abundances. These transitions represent death of organisms.
5. The interaction coefficients $a_{ij} \ge 0$ (dimensions T^{-1}, $i \ne j$) that determine transition rates from state j to state i do not vary over time.
6. The system is of infinite extent, and local spatial interactions are unimportant.

Given these assumptions, we derived the ordinary differential equation model

$$\frac{dx_i}{dt} = \begin{cases} -\left(a_{ei} + \sum_{j \ne e,i} a_{ji} x_j\right) x_i + \sum_{j \ne i} a_{ij} x_i x_j, & i \ne e \\ -\sum_{j \ne e} a_{je} x_j x_e + \sum_{j \ne i} a_{ej} x_j, & i = e \end{cases} \tag{1}$$

where the summations are over the s possible states. This model is analogous to a Lotka-Volterra competition model. Other attempts to approximate the dynamics of coral reefs using differential equations have resulted in similar models (e.g. Mumby et al., 2007).

The parameters of this model are the $s(s-1)$ interaction coefficients a_{ij} (for each state j, there are $s-1$ other states to which a transition may occur), and the s initial state probabilities $p_j(0)$, $j = 1 \ldots s$ (the probability that a point is in state j at the first observation time t_0). Thus in total there are s^2 parameters (although only $s-1$ of the initial state probabilities are independent, because they must sum to 1). Table 1 summarizes the parameters, state variables, and other symbols used in the model.

Linear Markov models, also with s^2 parameters, are often used for sessile organisms such as corals, mussels, and trees (Usher, 1979; Tanner et al., 1994; Wootton, 2001a; Hill et al., 2004), but unlike the nonlinear model described here, they do not allow density-dependent interactions. Our nonlinear model is a much better fit than a linear Markov model (Spencer and Tanner, 2008) to the Heron Island data. It is not clear whether this is a general pattern, but there is evidence for density-dependent interactions among sessile organisms from a number of other studies (Tanner et al., 2009).

2.2. Data

The data to which we apply the model are from a long-term study of the coral reef at Heron Island, Great Barrier Reef, Queensland, Australia (Connell et al., 1997, 2004). Photographs of fixed $1\mathrm{m}^2$ quadrats at the Protected Crest site were taken at 17 unequally-spaced times between 1963 and 1989. We label the observation times $t_0, t_1, \ldots, t_N$. The organisms present in at least 1249 points with fixed spatial locations were recorded from each photograph (Tanner et al., 1994). Here, we analyze data in which the identities of organisms were aggregated into either $s = 6$ (acroporid corals, soft corals, algae,

massive corals, pocilloporid corals, free space) or $s = 3$ (corals, algae, free space) states. At time t_0, we record the initial numbers of points $n_j(0)$ in each of the states $j = 1, 2, \ldots, s$. For each adjacent pair of observation times t_{m-1}, t_m we record the number $n_{ij}(m, m-1)$ of points that were in state j at time t_{m-1} and state i at time t_m.

2.3. Transition probabilities

In order to fit our model to data, we need to be able to compute each of the transition probabilities $p_{ij}(m, m-1)$ (the probability that a point in state j at some specified time t_{m-1} is in state i at some later specified time t_m). Using Equation 1, we can write down a corresponding system of differential equations satisfied by $p_{ij}(m, m-1)$ which may then be solved numerically for any given values of the model parameters a_{ij}, $p_j(0)$. The details are given in Appendix A2.1 of Spencer and Tanner (2008). Note that because Equation 1 is written in continuous time, the interval between t_{m-1} and t_m can be any non-negative real number, rather than a fixed time step.

2.4. Log likelihood calculation

We have data on the states of a number of fixed points in space at times $t_0 \ldots t_N$, which are not necessarily equally spaced. We assume that these points are independent and have dynamics that can be described by the same interaction coefficients. The assumption of independence is unlikely to be true in practice, because nearby colonies are more likely to overgrow a point in space than distant colonies. Spatially explicit models exist for marine sessile organisms (e.g. Wilson et al., 1996; Burrows and Hawkins, 1998; Wootton, 2001b; Robles and Desharnais, 2002; Langmead and Sheppard, 2004; Dunstan and Johnson, 2005, 2006), but as yet there has been little progress on fitting them to data. In contrast, it is straightforward to calculate the log likelihood l of the data for the model in Equation 1, given the assumptions above:

$$l = \sum_j n_j(0) \log p_j(0) + \sum_{m=1}^{N} \sum_{i,j} n_{ij}(m, m-1) \log p_{ij}(m, m-1) \tag{2}$$

(Spencer and Tanner, 2008). The summations of i and j are over the s possible states.

2.5. Maximum likelihood estimation

Maximum likelihood estimation of the initial state probabilities and interaction coefficients for the model specified by Equation 1 is simple in principle, but difficult in practice. The log-likelihood surface is multimodal and has long, steep-sided ridges, making it difficult to find the global optimum. In addition, the optimization algorithms we used in Spencer and Tanner (2008) often terminate at points where the Hessian of the negative log likelihood is not positive definite, and therefore its inverse does not give a good estimate of the covariance matrix for the estimated parameters. However, the methods described in Catchpole and Morgan (1997) and Spencer and Tanner (2008, Appendix A3) show that the model is

locally identifiable given suitable data. The difficulties we experienced are therefore due to the relatively small number of observation times, and some parameters apparently lacking internal optima.

A simple Markov Chain Monte Carlo (MCMC) algorithm also fails to perform well, because in the absence of gradient information it becomes trapped in local optima that are much worse than those found by the gradient-based algorithms used previously. We therefore use the maximum likelihood method described in Spencer and Tanner (2008) to find good parameter estimates from which to initialize an MCMCMC algorithm, which we will then use to estimate the uncertainty in our parameters. We performed 200 replicate optimizations for the 3-state model, and 400 for the 6-state model, started from random initial estimates. The number of replicates was a compromise between having a reasonable number of starting points for MCMCMC chains, and the time needed for deterministic searches.

2.6. Basic Metropolis-Hastings MCMC algorithm

In general terms, Bayesian inference involves making statements about a set of parameters $\boldsymbol{\theta}$, given relevant data x and a probability distribution $\pi(\boldsymbol{\theta})$ (the prior) that represents our prior information or beliefs about $\boldsymbol{\theta}$. In our case, $\boldsymbol{\theta}$ is a vector containing all the parameters we want to estimate (the a_{ij} and the $p_j(0)$). The information in the data is contained in the likelihood function $L(x|\boldsymbol{\theta})$, which we have described above. Using Bayes' Theorem, the distribution of $\boldsymbol{\theta}$ conditional on the data can be written as

$$\pi(\boldsymbol{\theta}|x) = \frac{\pi(\boldsymbol{\theta})L(x|\boldsymbol{\theta})}{\int \pi(\boldsymbol{\theta})L(x|\boldsymbol{\theta})\, d\boldsymbol{\theta}}.$$

Here, $\pi(\boldsymbol{\theta}|x)$ is known as the posterior distribution of $\boldsymbol{\theta}$, and it contains all the information we have about the parameters $\boldsymbol{\theta}$ after analyzing the data. We can then use this posterior distribution to obtain the posterior distributions of functions of the parameters (such as, for our model, the posterior distribution of trajectories of Equation 1). Chapter 1 of Gelman et al. (2003) explains these ideas in more detail. In most cases, it is difficult to calculate the posterior distribution directly. MCMC allows us to sample from the posterior distribution, by simulating a Markov chain whose stationary distribution is the required posterior. At each iteration of the Markov chain, we propose new parameter values, which are accepted or rejected using a probabilistic rule which gives the correct stationary distribution.

Because we have little prior information about parameters, we use vague priors (priors which are designed to play a minimal role in the posterior distribution: Gelman et al., 2003, p. 61). For the interaction coefficients a_{ij}, with $0 \leq a_{ij} < \infty$ for $i, j \in \{1, 2, \ldots, s\}, i \neq j$, we use independent exponential priors each having mean 1×10^4. For the initial state proportions, we use a uniform prior on the set of feasible values

$$\{(p_1(0), p_2(0), \ldots, p_s(0)) : p_1(0), p_2(0), \ldots, p_s(0) \geq 0,\ p_1(0) + p_2(0) + \cdots + p_s(0) = 1\}.$$

To simplify the updating procedure for the initial state proportions, we define new parameters $A_1, A_2, \ldots, A_s$

such that $0 \leq A_i < \infty$ for $i = 1, 2, \ldots, s$, and set

$$p_i(0) \quad = \quad \frac{A_i}{A_1 + A_2 + \cdots + A_s} \text{ for } i = 1, 2, \ldots, s.$$

This transformation allows us to update the parameters A_i independently of one another while retaining the constraint $p_1(0) + p_2(0) + \cdots + p_s(0) = 1$. To ensure a uniform prior distribution for the initial state proportions, we take independent exponential priors for $A_1, A_2, \ldots, A_s$, each having mean 1×10^4 (Gelman et al., 2003, p. 582).

We use a proposal in which a single parameter is updated at each iteration. We choose the parameter to be updated uniformly at random from the full set of s^2 parameters (s transformed initial state probability parameters A_i and $s(s-1)$ interaction coefficients a_{ij}). If an interaction coefficient a_{ij} is chosen, the proposed value is

$$\log a_{ij}^*(k+1) = \log a_{ij}(k) + w(k)$$

where $a_{ij}(k)$ is the interaction coefficient at iteration k and $w(k) \sim \mathcal{N}(\mu_{ij}, \sigma_{ij}^2)$ (normal with mean μ_{ij} and variance σ_{ij}^2).

If the parameter chosen for updating is one of the parameters A_i corresponding to the initial state probabilities, we use a similar proposal:

$$\log A_i^*(k+1) = \log A_i(k) + w(k)$$

where $w(k) \sim \mathcal{N}(0, \tau_i^2)$.

The proposal variances σ_{ij}^2, τ_i^2 were chosen separately for each parameter by trial and error, to give suitable acceptance probabilities. To get good mixing (i.e. reasonably rapid exploration of the posterior distribution of the parameters), we aimed for acceptance probabilities in the range (0.2, 0.6) in short trial runs. In most cases, we set $\mu_{ij} = 0$. However, for a few interaction coefficients, the estimates obtained by the initial likelihood maximization search were very close to zero, and the likelihood surface for these parameters was very flat unless we moved far from zero. As a result, the acceptance probabilities were very high (typically > 0.99) even with very large proposal variances, and these $\log a_{ij}$ spent most of their time in random walks through large negative values for which a_{ij} is not different from zero when represented on a computer. Since it is important to know whether values further from zero are plausible, we used biased proposals ($\mu_{ij} > 0$) for these parameters. This allowed us to explore the space of larger values more efficiently, and reduced the acceptance probabilities to the target range. We checked with short preliminary runs that the marginal posterior distributions were similar whether or not we used biased proposals. The only exception was that for parameters that were very close to zero, the posterior means tended to be even smaller when an unbiased proposal was used (e.g. mean 7×10^{-9}, standard deviation 5×10^{-7} with an unbiased proposal, mean 5×10^{-5}, standard deviation 6×10^{-4} with a biased proposal, for a_{31} in the three-state model). However, for these parameters, the 95% credible intervals were always relatively wide, and the qualitative result that they were very close to zero was unaltered.

Denoting by $\boldsymbol{\theta}$ the full set of s^2 parameters, the Metropolis-Hastings acceptance probability is computed as

$$\alpha(k) = \min\left(1, \frac{\pi\left(\boldsymbol{\theta}^*_{k+1} \mid x\right) q(\boldsymbol{\theta}_k|\boldsymbol{\theta}^*_{k+1})}{\pi\left(\boldsymbol{\theta}_k \mid x\right) q(\boldsymbol{\theta}^*_{k+1}|\boldsymbol{\theta}_k)}\right) \tag{3}$$

(Gilks et al., 1996, p. 7), where x denotes the full set of data, $\pi(\boldsymbol{\theta} \mid x)$ the posterior density of $\boldsymbol{\theta}$, $\boldsymbol{\theta}_k$ the parameter values at iteration k, and $\boldsymbol{\theta}^*_{k+1}$ the proposed values at iteration $k+1$. The density of proposing $\boldsymbol{\theta}_k$ from parameter $\boldsymbol{\theta}^*_{k+1}$ is $q(\boldsymbol{\theta}_k|\boldsymbol{\theta}^*_{k+1})$. Denoting the likelihood function for our model as $L(x \mid \boldsymbol{\theta})$, and the joint prior density of our parameters as $\pi(\boldsymbol{\theta})$, then (3) can be written as

$$\alpha(k) = \min\left(1, \frac{\pi\left(\boldsymbol{\theta}^*_{k+1}\right) L\left(x \mid \boldsymbol{\theta}^*_{k+1}\right) q(\boldsymbol{\theta}_k|\boldsymbol{\theta}^*_{k+1})}{\pi\left(\boldsymbol{\theta}_k\right) L\left(x \mid \boldsymbol{\theta}_k\right) q(\boldsymbol{\theta}^*_{k+1}|\boldsymbol{\theta}_k)}\right).$$

Parameters are updated by setting

$$\boldsymbol{\theta}_{k+1} \quad = \quad \begin{cases} \boldsymbol{\theta}^*_{k+1} & \text{with probability } \alpha(k), \\ \boldsymbol{\theta}_k & \text{otherwise.} \end{cases}$$

2.7. MCMCMC chain swapping algorithm

Unfortunately, the basic MCMC algorithm outlined in the previous section is impractically slow to converge to the desired stationary distribution, due to both multimodality and the presence of long, steep-sided ridges in the log likelihood surface. Metropolis Coupled Markov Chain Monte Carlo (MCMCMC) improves mixing by running j chains in parallel (Geyer, 1991; Gilks and Roberts, 1996). One chain, known as the cold chain, has the correct posterior distribution $\pi(\boldsymbol{\theta} \mid x)$ as its stationary distribution. The other chains have flatter stationary distributions $\pi_i(\boldsymbol{\theta} \mid x) \propto \pi(\boldsymbol{\theta} \mid x)^{\beta_i}$, with a heat parameter $0 < \beta_i < 1$. This helps the heated chains to move between modes. At each iteration k, each chain is updated using a standard Metropolis-Hastings algorithm. We then pick the indices a, b of two chains for which to propose a swap from independent discrete uniform (1,j) distributions. We accept the swap and exchange the states $\boldsymbol{\theta}^{(a)}_k$, $\boldsymbol{\theta}^{(b)}_k$ of the two chains with probability

$$R = \min\left(1, \frac{\pi_a(\boldsymbol{\theta}^{(b)}_k \mid x)\pi_b(\boldsymbol{\theta}^{(a)}_k \mid x)}{\pi_a(\boldsymbol{\theta}^{(a)}_k \mid x)\pi_b(\boldsymbol{\theta}^{(b)}_k \mid x)}\right).$$

As a result, the cold chain will sometimes exchange states with one of the heated chains, which are more likely to move between modes. We did not exclude swaps where $a = b$, where the chain exchanges with itself, which are always accepted. We used incremental heating, $\beta_i = 1/(1 + \Delta T(i-1))$. We chose ΔT so that chain swaps were accepted 20-60% of the time in short preliminary runs, which should give a sufficient amount of mixing (Altekar et al., 2004).

2.8. Starting points and run length for MCMCMC

We sorted the parameter sets found by likelihood maximization search in decreasing order of posterior density. We want to use all the solutions that we expect to make a non-negligible contribution to the

posterior as starting points for MCMCMC chains. Let ρ_{k1} be the ratio of posterior densities in the kth and the best solutions. We added the best j solutions to our set of chain starting points, where j was chosen to be the largest integer such that $\rho_{j1}/\sum_{k\leq j}\rho_{k1} \geq \epsilon$, where ϵ is some small value. For the 3-state model, where there were many solutions with similar log likelihoods, we chose $\epsilon = 1 \times 10^{-6}$, which gave $j = 19$ chains. For the 6-state model (where the solutions had relatively large differences in log likelihood), we chose $\epsilon = 1 \times 10^{-7}$, which gave $j = 3$ chains. We initialized chains with decreasing β_i with the selected sets of parameters in decreasing order of posterior density.

We ran the MCMCMC algorithm for long enough to establish that there were no obvious trends in any parameters, or in the log likelihood. Due to the slow mixing of this chain, we would not expect to achieve convergence to the stationary distribution from arbitrary starting points in a reasonable length of time. Because we started close to modes of the log likelihood surface, we would not expect parameter estimates to change much. Nevertheless, we discarded all the iterations from an initial burn-in period, before sampling parameter values (every 10th iteration) and predicted trajectories (every 100th iteration) from the posterior distribution. We calculated predicted trajectories over a 40-year period (up to 1 January 2003).

Figure 1 summarizes the modelling process.

2.9. Implementation

We implemented the MCMCMC algorithm in Matlab Release 2008a (The Mathworks, Inc., Natick, MA). We used the `ode15s` stiff differential equation solver in log likelihood calculations. The code is available from `http://www.liv.ac.uk/~matts/coral_MCMCMC.html`. Details of the starting points, proposal distributions, and acceptance probabilities are given in the Supplementary Data.

3. Results

3.1. Three-state model

For the three-state model, the marginal posterior distributions of interaction coefficients were not strongly bimodal (Figure 2; the means and standard deviations of the posterior distributions for all parameters are given in the Supplementary Data, Table 1). The most striking result was that the coefficients a_{21} (transition from coral to algae: Figure 2d), a_{23} (free space to algae: Figure 2f) and a_{32} (algae to free space: Figure 2h) were all very large, with modes at more than 60 years^{-1} (throughout, we report statistics about posterior distributions based on the burn-in times and sampling frequencies given in the figure legends). Algae rapidly colonize areas of free space or dead coral, and algal blooms have been observed at the study site shortly after cyclones that caused substantial coral mortality (J.H. Connell, personal communication). Some algae can overgrow live coral, although overgrowth ability varies among algal species (Jompa and McCook, 2003) and the evidence that algal overgrowth is a cause, rather than a consequence, of coral mortality is fairly limited (McCook et al., 2001). Algae

are also subject to high grazing pressure (e.g. Fox and Bellwood, 2008), leading to rapid turnover. The high interaction coefficients between algae and other states are qualitatively consistent with these observations. However, due to their short-lived nature, algal blooms were not captured in our data. Small, short-lived algal species were not recorded in our surveys, because although they were present, they could not be reliably distinguished in photographs. The dominant algal species recorded in our data is *Chlorodesmis fastigiata* (77% of all algae recorded). Patches of *C. fastigiata* tend to persist long-term (J.E. Tanner, personal observation), although they may undergo cycles of die-back and regrowth over periods of months (Jompa and McCook, 2003). In contrast, the interaction coefficients in the three-state model suggest turnover on a weekly time scale. It is possible that these high estimates of transitions in and out of algal states are partly due to the sampling method used. Fronds of *C. fastigiata* have only a small attachment. At high tide, the fronds will float erect. At low tide, when samples were taken, they lie flat and may cover live coral or free space. A point where this occurred would be recorded as algae. Since the fronds may lie in different directions at different low tides, sample points could switch frequently between being recorded as algae and coral or free space. Thus it is not clear whether the estimated transition rates in and out of algae are due to sampling artefacts, insufficient data, or model misspecification.

Although much smaller than the coefficient for overgrowth of coral by algae, the posterior distribution of the coefficient for overgrowth of algae by coral (a_{12}) has substantial mass far from zero, with a mean of 2.0 years^{-1} (Figure 2b). This is at least qualitatively plausible. Corals are able to overgrow algae in some circumstances (McCook et al., 2001), including regeneration of lesions overgrown by algae, on a timescale of several weeks (Meesters and Bak, 1993).

It is also striking that the interaction coefficient for transitions from coral to free space (a_{31}, Figure 2g) is very small and right-skewed, with posterior mean 5.49×10^{-5} and 95% credible interval [4.40×10^{-238},1.48×10^{-6}]. This does not match the apparently substantial mortality of corals in this system, due to both desiccation and cyclones (J. E. Tanner, personal observation). Peak coral cover during the study period was 68% in 1969, and minimum coral cover 19% in 1989 (Connell et al., 2004). Our model reproduces much of this variation in coral cover (Figure 3a). However, the dominant pathway of coral loss in the model is transitions from coral to short-lived algae, rather than from coral to free space. This is the most important disagreement between the posterior rate estimates from the three-state model and biological knowledge.

The posterior distribution of trajectories is similar to the observed temporal pattern for coral (Figure 3a) and free space (Figure 3c) up to 1989. The agreement is less good for algae (Figure 3b), but this is based on few observations, because algae were at low abundance throughout the observation period. The uncertainty in the posterior distribution of trajectories remains small throughout the simulation period, with all sampled parameter sets resulting in the prediction that by 2003, there will be less than

5% coral (Figure 3d), less than 1% algae (Figure 3e), and more than 93% free space (Figure 3f).

3.2. Six-state model

The behaviour of the six-state model is rather different. For many parameters, the posterior distribution contains a small but distinct secondary mode (Figure 4; the means and standard deviations of the posterior distributions are given in the Supplementary Data, Table 2). The largest coefficients are a_{25} (pocilloporid corals to soft corals: Figure 4k, which has a very large value for the secondary mode), a_{31} (acroporid corals to algae: Figure 4m), a_{34} (massive corals to algae: Figure 4p), a_{35} (pocilloporid corals to algae: Figure 4q), a_{36} (free space to algae: Figure 4r), a_{52} (soft corals to pocilloporid corals: Figure 4z), and a_{63} (algae to free space: Figure 4ag). As in the three-state model, many of the large rates involve transitions to and from algae, which may not be biologically plausible. However, the mortality rates of corals are not negligible (Figure 4ae, af, ah, and ai). This appears more consistent than the three-state model with what we know about the biology of the system.

The parameter estimates are different from those we obtained for the six-state model in our previous study (Spencer and Tanner, 2008). In that work, we relied on maximum likelihood with only 10 starting points, and the best-fitting parameter estimates had a log likelihood of -1.6208×10^4. Here, the best starting point for MCMC had a substantially better log likelihood (-1.6099×10^4), and our previous estimate would make only a negligible contribution to the posterior distribution. Although we showed that our previous methods were reasonably good at estimating the parameters of simulated data, our new approach is likely to be much better (although also much slower), especially on real data.

The six-state model has the potential to capture some of the major differences in life histories between different groups of corals, rather than assuming that all corals (including hard and soft corals) are homogeneous. Corals display a broad range of life histories (e.g. Hughes et al., 1992; Hall and Hughes, 1996). Here, we divide them into four somewhat homogeneous groups based on taxonomic and morphological grounds. Recruitment rates tend to be high for acroporids (especially when asexual fragmentation is taken into account), intermediate for pocilloporids, and low for massives (e.g. Highsmith, 1982; Wallace, 1985; Wallace et al., 1986). Growth rates show the same pattern (e.g. Buddemeier and Kinzie, 1976; Guzmán and Cortés, 1989; Babcock, 1991; Tanner, 1997). Conversely, massive corals have low mortality rates and are resistant to mechanical disturbance, acroporids generally have higher mortality and can be susceptible to disturbance, while pocilloporids are generally short-lived and vulnerable to disturbance (e.g. Harriott, 1985; Marshall, 2000; Baird and Marshall, 2002). Soft corals are in a different order from the other corals, and have a mixture of life history strategies. Only one species, *Zoanthus vietnamensis*, became abundant in the late 1980s, when the site became drier. Little is known about its biology, but it is apparently a good competitor and tolerant of dry conditions (Connell et al., 2004).

In our model, growth and recruitment are represented by transitions from free space (state 6) to corals. The observations above suggest that the corresponding interaction coefficients should be ordered

$a_{16} > a_{56} > a_{46}$: growth and recruitment highest in acroporids (state 1), intermediate in pocilloporids (state 5), and lowest in massive corals (state 4). In fact, they are ordered $a_{56} > a_{46} > a_{16}$ (Supplementary Data, Table 2). Similarly, mortality is represented by transitions from corals to free space. The observations above suggest that the corresponding interaction coefficients should be ordered $a_{65} > a_{61} > a_{64}$: mortality highest in pocilloporids, intermediate in acroporids, and lowest in massive corals. In fact, they are ordered $a_{61} > a_{64} > a_{65}$ (Supplementary Data, Table 2). Thus, although our model allows differences in life histories between types of coral, the estimated parameters do not reflect what we know about these life histories. Possible explanations include the model being misspecified in some way, or the small amount of data for some states such as pocilloporid corals, which were rare throughout the study period (Figure 5e).

The posterior distribution of trajectories from the six-state model (Figure 5) is tightly constrained by the data up until 1989 (the year of the last observation used in this study). From then onwards, uncertainty about the predictions of the model rapidly increases. The posterior distribution of trajectories has three distinct modes, which are coloured in Figure 5 by the amount of free space $x_6(40)$ predicted in 2003, at the end of the 40-year run ($x_6(40) < 0.5$: black; $0.5 \leq x_6(40) < 0.75$: blue, $0.75 \leq x_6(40)$: red). 85% of trajectories fall into the blue cluster, which is dominated by free space (mean 68%, Figure 5l) and soft corals (mean 22%, Figure 5h). All other states are rare in the blue cluster (means: 4% pocilloporid corals, Figure 5k; 3% acroporid corals, Figure 5g; 1% algae, Figure 5i; and 1% massive corals, Figure 5j). The red cluster (4% of trajectories), while quantitatively distinct, has similar qualitative behaviour (means: 79% free space, Figure 5l; 15% soft corals, Figure 5h; 3% algae, Figure 5i; 2% pocilloporid corals, Figure 5k; 0.5% acroporid corals, Figure 5g; 0.3% massive corals, Figure 5j). These two clusters also show similar qualitative behaviour to the three-state model (Figure 3). However, the black cluster (11% of trajectories) has very different behaviour, dominated by acroporid corals (mean 50%, Figure 5g), pocilloporid corals (mean 33%, Figure 5k), and massive corals (mean 10%, Figure 5j), with mean 7% free space (Figure 5l) and almost no soft corals (mean 0.1%, Figure 5h) or algae (mean 4×10^{-11}%, Figure 5l). We checked by simulation that the differences between clusters persisted over much longer time scales, although there were small, sustained oscillations in trajectories from all clusters (results not shown).

The presence of three modes in the posterior distribution of trajectories is a result of there being three starting points, close to modes on the log likelihood surface, with large differences in log likelihood (the short vertical bars at the top of each panel in Figure 4, and Supplementary Data, Table 2). This contrasts with the three-state model, in which the differences in log likelihood among the 19 starting points were much smaller (the short vertical bars at the top of each panel in Figure 2, and Supplementary Data, Table 1).

4. Discussion

Many kinds of uncertainty affect our ability to predict the natural world (Regan et al., 2002). Here, we have dealt with one of them, parameter uncertainty, in a quantitative way. The effect of parameter uncertainty on our model is similar to that described in van Nes and Scheffer (2003). In their simulation model of competition between two aquatic plant species, moderate uncertainty in parameter values resulted in large uncertainty in the biomasses of the two species. This was because there were two attractors in their model, one in which both species were present, and another in which only one of the two species persisted. Changes in interaction parameter values switched the model from one attractor to the other. The effects of parameter uncertainty are likely to be even stronger in cases such as the three-state coral model of Mumby et al. (2007), where there is more than one locally stable attractor for some sets of parameter values. In such cases, changes in initial conditions can cause switching from one attractor to another.

Our Bayesian analysis has given us much more information than we would have got from a typical maximum likelihood method. Inverting the Hessian of the negative log likelihood is (when possible) a quick and easy way to estimate standard errors for maximum likelihood estimates of parameters. However, in doing this we would have missed the multimodal nature of the posterior distribution of many parameters in the six-state model (Figure 4). Consequently, we would have lost all information on the small but important probability with which the six-state model predicts high rather than low abundance of acroporid corals in 2003 (Figure 5).

The models we have analyzed here are conceptually simple and based on data from a single long time series. Although we cannot be certain, it seems likely that the effects of parameter uncertainty will be of a similar order in other reef models, where it has not been quantified. We will consider three examples. The model described in Mumby et al. (2006) is spatially explicit, contains two different kinds of corals, three different classes of algae, grazing by fish and urchins, varying nutrient levels, and hurricanes. This model appears to have 22 parameters, which were estimated from diverse sources and whose uncertainty is unknown, in addition to the initial abundance of each of seven states. The complexity of this model leads us to expect that parameter uncertainty would have substantial effects. Langmead and Sheppard (2004) describe a spatially explicit model with ten species of corals, and 40 parameters estimated from diverse sources. They did a sensitivity analysis in which each parameter was altered by what seems to be an arbitrary but plausible amount. 20% of such changes resulted in a change in cover of more than 10%. Again, given the complexity of the model and the limitations of data, we expect the consequences of parameter uncertainty to be large. McClanahan (1995) developed an energy-based model of coral reef fisheries, including corals, algae, two kinds of herbivores, two kinds of carnivores, human fishing, and calcium carbonate accretion. The 27 parameters were estimated from diverse sources. A sensitivity analysis was carried out. It is not described in detail in the paper, but calcium carbonate accretion was

described as 'very sensitive' to changes in two processes, and fisheries yields as 'sensitive' to changes in another two processes.

The way the uncertainty in our six-state model increases after the period for which we have data (Figure 5) suggests that even with observations at 17 times spread over 26 years, we cannot make meaningful predictions more than a very few years into the future. This is because the data do not constrain our estimated parameters tightly enough. Furthermore, in both the three-state and the six-state models, we had doubts about the biological plausibility of some parameter estimates, especially the low mortality of corals in the three state model and the very high coefficients for transitions to and from algae in both models. Other coral models such as those discussed above (McClanahan, 1995; Langmead and Sheppard, 2004; Mumby et al., 2006) made extensive use of literature data from diverse sources to estimate parameters. In principle, we could use such data and expert judgement to provide informative priors. This would help ensure that our parameter estimates are biologically plausible, and reduce the uncertainty in our posterior distributions. Turning expert knowledge into probability distributions is not trivial (Burgman, 2005, section 4.4), although guidelines exist (Garthwaite et al., 2005; O'Hagan, 2005). One difficulty is that expert knowledge about a particular system, such as the Heron Island reef, may be largely shaped by the same observations that provide our data. The result may be posterior distributions that are narrower than they should be.

We are not aware of other published models of coral reef communities that attempt to estimate the consequences of parameter uncertainty, although at least one other unpublished study has used MCMC for parameter estimation (Żychaluk et al., 2005). A search on the ISI Web of Science for 'coral parameter uncertainty' on 6 April 2009 gave three results, of which two are relevant. van Nes and Scheffer (2003) look at a two-species aquatic plant model, discussed above. Nguyen and de Kok (2007) describe a Monte Carlo uncertainty analysis of an integrated systems model for coastal zone management, which included the effects of processes such as blast fishing on living coral area. Both studies showed that parameter uncertainty is likely to have substantial effects on model predictions, but neither used probabilistic methods to estimate this uncertainty based on field data.

Model uncertainty (uncertainty arising from the choice of processes to include in the model and the choice of mathematical construct used to represent them: Regan et al., 2002) is recognized as being important in ecological models, but difficult to deal with. We have only considered model uncertainty in a qualitative way. First, among a set of candidate six-state models with different mathematical forms that have the potential to make predictions about the future, the six-state version of the model we used here is the only one with a non-negligible Akaike weight (Spencer and Tanner, 2008). Thus, of the model forms we have so far evaluated, we are justified in using only Equation 1. Second, we have considered two versions of this model, with three and six states respectively. These two versions of the model behaved in substantially different ways. Aggregating state variables reduces parameter uncertainty because there

are fewer parameters to estimate. This is done at the expense of increasing bias, because we may be aggregating states (for example the four different kinds of corals in the six-state model) that have very different dynamics. Our subjective impression is that the six-state model has too many parameters, given the amount of data we have (although it does not have any more parameters than other simple models for the same type of data and the same number of states, such as a linear Markov model). On the other hand, the three-state model is probably over-aggregated, given how different the dynamics of acroporid corals (Figure 5a), soft corals (Figure 5b) and massive corals (Figure 5d) appear. A quantitative treatment of this kind of model uncertainty would involve giving a probabilistic weight to each level of aggregation. This could in principle be done using reversible-jump MCMC (Green, 1995), treating transitions within aggregated states as missing data (Susko and Roger, 2007). Reversible-jump MCMC could also be used to deal with more general kinds of model uncertainty, encompassing models that do not fit into the class defined by Equation 1. However, this would involve specifying what those alternative models are. One obvious candidate is a model in which the algal grazing rate per unit area declines as algal cover increases (Mumby et al., 2007). It seems likely that considering more models would increase our overall estimate of uncertainty.

We have not attempted to address natural variation (in the sense of Regan et al., 2002). We have assumed that the parameters of the model remain constant over time. In reality, changes in the environment, including gradual uplift and storms (Connell et al., 2004) may alter parameters. Stochastic models in conservation biology often incorporate temporal variability in parameters. However, the data requirements are demanding. Even in the simplest single-species models, reliable estimates of the risk of population decline below a specified size may only be possible over time horizons of 10-20% of the length of the time series from which parameters were estimated (Fieberg and Ellner, 2000). The situation is likely to be worse for a nonlinear, multi-species model. Nevertheless, species do not exist in isolation, so ecological risk analyses based on multispecies models may have a place.

The absence of spatial effects in our model might partly explain why some of the estimated parameters are implausible. Cellular automata models are an obvious solution. The parameters of discrete-time cellular automata are easy to estimate when the model time steps are the intervals between observations (e.g. Balzter et al., 1998). However, our data are observed at irregular time intervals. We would therefore need either a discrete-time cellular automata model with a time step chosen so that all observation intervals are approximately integer multiples of it, together with a way of integrating over missing data between observations (O'Hara et al., 2002), or a continuous-time cellular automata model, together with some way of approximating its likelihood function.

5. Conclusions

We showed that Bayesian Markov Chain Monte Carlo methods can be used to quantify the effects of parameter uncertainty on predictions about the dynamics of coral reef models. Our approach could be extended to incorporate literature data and expert knowledge (which our results suggest are necessary to constrain parameter estimates), and to deal with temporal variability. Given the data available to us, predictions based on coral reef models are necessarily probabilistic, and it seems likely that the uncertainty in these predictions will be large. We think it important to focus on conceptually simple models in order to reduce this uncertainty.

Acknowledgments

This work was started as an EPSRC-funded undergraduate project by SM under the supervision of DC and MS. We are particularly grateful to Joe Connell for initiating and conducting the field study, and making the data available, and Terry Hughes for his role in data collection and processing. Chris Frid, Cheryl Knowland, Kamila Żychaluk, and two anonymous reviewers made many helpful suggestions. Author contributions: DC and MS planned the study. SM, MS and DC wrote the code. JET provided biological interpretation of the results. DC, JET and MS wrote the manuscript.

References

Altekar, G., Dwarkadas, S., Huelsenbeck, J. P., Ronquist, F., 2004. Parallel Metropolis coupled Markov chain Monte Carlo for Bayesian phylogenetic inference. Bioinformatics 20 (3), 407–415.

Babcock, R. C., 1991. Comparative demography of three species of scleractinian corals using age- and size-dependent classifications. Ecol. Monograph. 61 (3), 225–244.

Baird, A. H., Marshall, P. A., 2002. Mortality, growth and reproduction in scleractinian corals following bleaching on the Great Barrier Reef. Mar. Ecol. Prog. Ser. 237, 133–141.

Balzter, H., Braun, P. W., Köhler, W., 1998. Cellular automata models for vegetation dynamics. Ecol. Model. 107, 113–125.

Buddemeier, R. W., Kinzie, III, R. A., 1976. Coral growth. In: Barnes, H. (Ed.), Oceanography and Marine Biology: an annual review. Allen and Unwin, London, pp. 183–225.

Burgman, M., 2005. Risks and Decisions for Conservation and Environmental Management. Cambridge University Press, Cambridge.

Burrows, M. T., Hawkins, S. J., 1998. Modelling patch dynamics on rocky shores using deterministic cellular automata. Mar. Ecol. Prog. Ser. 167, 1–13.

Catchpole, E. A., Morgan, B. J. T., 1997. Detecting parameter redundancy. Biometrika 84 (1), 187–196.

Connell, J. H., Hughes, T. P., Wallace, C. C., 1997. A 30-year study of coral abundance, recruitment, and disturbance at several scales in space and time. Ecol. Monograph. 67 (4), 461–488.

Connell, J. H., Hughes, T. P., Wallace, C. C., Tanner, J. E., Harms, K. E., Kerr, A. M., 2004. A long-term study of competition and diversity of corals. Ecol. Monograph. 74 (2), 179–210.

Dunstan, P. K., Johnson, C. R., 2005. Predicting global dynamics from local interactions: individual-based models predict complex features of marine epibenthic communities. Ecol. Model. 186, 221–233.

Dunstan, P. K., Johnson, C. R., 2006. Linking richness, community variability, and invasion resistance with patch size. Ecology 87 (11), 2842–2850.

Fieberg, J., Ellner, S. P., 2000. When is it meaningful to estimate an extinction probability? Ecology 81 (7), 2040–2047.

Fox, R. J., Bellwood, D. R., 2008. Direct versus indirect methods of quantifying herbivore grazing impact on a coral reef. Mar. Biol. 154, 325–334.

Garthwaite, P. H., Kadane, J. B., O'Hagan, A., 2005. Statistical methods for eliciting probability distributions. J. Am. Stat. Assoc. 100, 680–700.

Gelman, A., 1996. Inference and monitoring convergence. In: Gilks, W. R., Richardson, S., Spiegelhalter, D. J. (Eds.), Markov Chain Monte Carlo in practice. Chapman and Hall/CRC, Boca Raton, pp. 131–143.

Gelman, A., Carlin, J. B., Stern, H. S., Rubin, D. B., 2003. Bayesian Data Analysis, 2nd Edition. Chapman and Hall/CRC, Boca Raton.

Geyer, C. J., 1991. Markov Chain Monte Carlo maximum likelihood. In: Keramidas, E. M. (Ed.), Computing science and statistics: proceedings of the 23rd symposium on the interface. Interface Foundation, Fairfax Station, pp. 156–163.

Gilks, W. R., Richardson, S., Spiegelhalter, D. J., 1996. Introducting Markov Chain Monte Carlo. In: Gilks, W. R., Richardson, S., Spiegelhalter, D. J. (Eds.), Markov Chain Monte Carlo in practice. Chapman and Hall/CRC, Boca Raton, pp. 1–19.

Gilks, W. R., Roberts, G. O., 1996. Strategies for improving MCMC. In: Gilks, W. R., Richardson, S., Spiegelhalter, D. J. (Eds.), Markov Chain Monte Carlo in practice. Chapman and Hall/CRC, Boca Raton, pp. 89–114.

511 Green, P. J., 1995. Reversible jump Markov Chain Monte Carlo computation and Bayesian model de-
512 termination. Biometrika 82 (4), 711–732.

513 Guzmán, H. M., Cortés, J., 1989. Growth rates of eight species of scleractinian corals in the Eastern
514 Pacific (Costa Rica). Bull. Mar. Sci. 44 (3), 1186–1194.

515 Hall, V. R., Hughes, T. P., 1996. Reproductive strategies of modular organisms: comparative studies of
516 reef-building corals. Ecology 77 (3), 950–963.

517 Harriott, V. J., 1985. Mortality rates of scleractinian corals before and during a mass bleaching event.
518 Mar. Ecol. Prog. Ser. 21, 81–88.

519 Highsmith, R. C., 1982. Reproduction by fragmentation in corals. Mar. Ecol. Prog. Ser. 7, 207–226.

520 Hill, M. F., Witman, J. D., Caswell, H., 2004. Markov chain analysis of succession in a rocky subtidal
521 community. Am. Nat. 164, E46–E61.

522 Hughes, T. P., 1994. Large-scale degradation of a Caribbean coral reef. Science 265, 1547–1551.

523 Hughes, T. P., Ayre, D., Connell, J. H., 1992. The evolutionary ecology of corals. Trends. Ecol. Evolut.
524 7 (9), 292–295.

525 Jompa, J., McCook, L. J., 2003. Coral-algal competition: macroalgae with different properties have
526 different effects on corals. Mar. Ecol. Prog. Ser. 258, 87–95.

527 Langmead, O., Sheppard, C., 2004. Coral reef community dynamics and disturbance: a simulation model.
528 Ecol. Model. 175, 271–290.

529 Marshall, P. A., 2000. Skeletal damage in reef corals: relating resistance to colony morphology. Mar.
530 Ecol. Prog. Ser. 200, 177–189.

531 McClanahan, T. R., 1995. A coral reef ecosystem-fisheries model: impacts of fishing intensity and catch
532 selection on reef structure and processes. Ecol. Model. 80, 1–19.

533 McCook, L. J., Jompa, J., Diaz-Pulido, G., 2001. Competition between corals and algae on coral reefs:
534 a review of evidence and mechanisms. Coral Reefs 19, 400–417.

535 Meesters, E. H., Bak, R. P. M., 1993. Effects of coral bleaching on tissue regeneration potential and
536 colony survival. Mar. Ecol. Prog. Ser. 96, 189–198.

537 Mumby, P. J., Hastings, A., Edwards, H. J., 2007. Thresholds and the resilience of Caribbean coral reefs.
538 Nature 450, 98–101.

539 Mumby, P. J., Hedley, J. D., Żychaluk, K., Harborne, A. R., Blackwell, P. G., 2006. Revisiting the
540 catastrophic die-off of the urchin *Diadema antillarum* on caribbean coral reefs: Fresh insights on
541 resilience from a simulation model. Ecol. Model. 196, 131–148.

542 Nguyen, T. G., de Kok, J. L., 2007. Systematic testing of an integrated systems model for coastal zone
543 management using sensitivity and uncertainty analyses. Environ. Modell. Softw. 22 (11), 1572–1587.

544 O'Hagan, A., 2005. Elicitation. Significance June 2005, 84–86.

545 O'Hara, R. B., Arjas, E., Toivonen, H., Hanski, I., 2002. Bayesian analysis of metapopulation data.
546 Ecology 83 (9), 2408–2415.

547 Regan, H. M., Colyvan, M., Burgman, M. A., 2002. A taxonomy and treatment of uncertainty for ecology
548 and conservation biology. Ecol. Appl. 12 (2), 618–628.

549 Robles, C., Desharnais, R., 2002. History and current development of a paradigm of predation in rocky
550 intertidal communities. Ecology 83 (6), 1521–1536.

551 Spencer, M., Tanner, J. E., 2008. Lotka-Volterra competition models for sessile organisms. Ecology 89,
552 1134–1143.

553 Susko, E., Roger, A. J., 2007. On reduced amino acid alphabets for phylogenetic inference. Mol. Biol.
554 Evol. 24 (9), 2139–2150.

555 Tanner, J. E., 1997. Interspecific competition reduces fitness in scleractinian corals. J. Exp. Mar. Biol.
556 Ecol. 214, 19–34.

557 Tanner, J. E., Hughes, T. P., Connell, J. H., 1994. Species coexistence, keystone species, and succession:
558 a sensitivity analysis. Ecology 75 (8), 2204–2219.

559 Tanner, J. E., Hughes, T. P., Connell, J. H., 1996. The role of history in community dynamics: a
560 modelling approach. Ecology 77 (1), 108–117.

561 Tanner, J. E., Hughes, T. P., Connell, J. H., 2009. Community-level density dependence: an example
562 from a shallow coral assemblage. Ecology 90, 506–516.

563 Usher, M. B., 1979. Markovian approaches to ecological succession. J. Anim. Ecol. 48, 413–426.

564 van Nes, E. H., Scheffer, M., 2003. Alternative attractors may boost uncertainty and sensitivity in
565 ecological models. Ecol. Model. 159, 117–124.

566 Wallace, C. C., 1985. Seasonal peaks and annual fluctuations in recruitment of juvenile scleractinian
567 corals. Mar. Ecol. Prog. Ser. 21, 289–298.

568 Wallace, C. C., Watt, A., Bull, G. D., 1986. Recruitment of juvenile corals onto coral tables preyed upon
569 by *Acanthaster planci*. Mar. Ecol. Prog. Ser. 32, 299–306.

570 Wilson, W. G., Nisbet, R. M., Ross, A. H., Robles, C., Desharnais, R. A., 1996. Abrupt population
571 changes along smooth environmental gradients. Bull. Math. Biol. 58 (5), 907–922.

572 Wootton, J. T., 2001a. Causes of species diversity differences: a comparative analysis of Markov models.
573 Ecol. Lett. 4 (1), 46–56.

574 Wootton, J. T., 2001b. Local interactions predict large-scale pattern in empirically derived cellular
575 automata. Nature 413, 841–844.

576 Wootton, J. T., 2001c. Prediction in complex communities: analysis of empirically derived Markov
577 models. Ecology 82 (2), 580–598.

578 Żychaluk, K., Blackwell, P. G., Mumby, P. J., Foster, N., May 2005. Modelling evolution of coral reefs.
579 Poster at *Workshop on recent advances in modelling spatio-temporal data*, Southampton.

Tables

Table 1: Parameters, state variables, and symbols used in the model

Symbol	Definition	Units
Parameters		
a_{ij}	Interaction coefficient between source state j and destination state i	years^{-1}
$p_j(0)$	Probability of state j at time t_0	
State variables		
$p_{ij}(m, m-1)$	Probability that a point in state j at time t_{m-1} will be in state i at time t_m	
x_i	Abundance of state i	
Symbols used to describe data		
e	Index of the empty space state	
$n_{ij}(m, m-1)$	Number of sampled points in state j at time t_{m-1} and state i at time t_m	
$n_j(0)$	Number of sampled points in state j at time t_0	
s	Number of states in the model	
$t_0, t_1, \ldots, t_N$	Times at which observations were made	years
Symbols used in model fitting		
A_i	Transformed initial state probability	
$\alpha(k)$	Acceptance probability at iteration k	
β_i	Heat parameter for chain i	
l	Log likelihood	
$L(x\|\boldsymbol{\theta})$	Likelihood of parameters $\boldsymbol{\theta}$	
μ_{ij}	Mean of proposal increment for $\log a_{ij}$	
$\pi(\boldsymbol{\theta})$	Prior density of parameters $\boldsymbol{\theta}$	
$\pi(\boldsymbol{\theta}\|x)$	Posterior density of $\boldsymbol{\theta}$ given data x	
$q(\boldsymbol{\theta}_k\|\boldsymbol{\theta}^*_{k+1})$	Density of proposing $\boldsymbol{\theta}_k$ from parameter $\boldsymbol{\theta}^*_{k+1}$	
R	Chain swap acceptance probability	
σ^2_{ij}	Variance of proposal increment for $\log a_{ij}$	
$\boldsymbol{\theta}$	Vector of model parameters (a_{ij} and $p_j(0)$)	
$\boldsymbol{\theta}_k$	Vector of model parameters at iteration k	
τ^2_i	Variance of proposal increment for $\log A_i$	
x	Data	

Figure captions

Figure 1. Conceptual diagram summarizing the way in which data, prior distributions, and likelihood calculation are used to obtain posterior distributions of parameters and trajectories for our coral model.

Figure 2. Marginal posterior distributions of initial state proportions $p_j(0)$ (diagonal panels, dimensionless) and interaction coefficients a_{ij} (off-diagonal panels, years^{-1}) in the three-state model (states: 1=coral; 2=algae; 3=free space) for the Heron Island data. The short vertical lines at the top of each panel are the starting parameter values. Produced from an MCMCMC run with 19 chains, $\Delta T = 0.5$. Run for 112799 iterations (sampling every 10th iteration) after a burn-in of 1×10^5 iterations. All panels have a common vertical scale, all diagonal panels have a common horizontal scale $[0, 1]$, but the horizontal scales for off-diagonal panels vary.

Figure 3. Posterior distribution of three-state model trajectories over 40 years, starting from 1963. Circles are the observed proportions of each state in the Heron Island data (sampled 1963-1989), and lines are a sample of approximately 2000 trajectories from the posterior distribution of the model, for: (a) corals; (b) algae; and (c) free space. The vertical scale is different for each of these. Small vertical panels are histograms of the predicted proportion of: (d) corals; (e) algae; and (f) free space on 1 January 2003, with vertical scales that match panels a to c respectively. Produced from an MCMCMC run with 19 chains, $\Delta T = 0.5$. Run for 112799 iterations (sampling every 100th iteration) after a burn-in of 1×10^5 iterations.

Figure 4. Marginal posterior distributions of initial state proportions $p_j(0)$ (diagonal panels, dimensionless) and interaction coefficients a_{ij} (off-diagonal panels, years^{-1}) in the six-state model (states: 1=acroporid corals; 2=soft corals; 3=algae; 4=massive corals; 5=pocilloporid corals; 6=free space) for the Heron Island data. The short vertical lines at the top of each panel are the starting parameter values. Produced from an MCMCMC run with 3 chains, $\Delta T = 0.5$. Run for 505199 iterations (sampling every 10th iteration) after a burn-in of 4×10^5 iterations. All panels have a common vertical scale, all diagonal panels have a common horizontal scale $[0, 1]$, but the horizontal scales for off-diagonal panels vary.

Figure 5. Posterior distribution of six-state model trajectories over 40 years, starting from 1963. Circles are the observed proportions of each state in the Heron Island data, and lines are a sample of approximately 5000 trajectories from the posterior distribution of the model, for: (a) acroporid corals; (b) soft corals; (c) algae; (d) massive corals; (e) pocilloporid corals; and (f) free space. Lines are coloured by the proportion of free space $x_6(40)$ predicted after 40 years (1 January 2003): $x_6(40) < 0.5$ black; $0.5 \leq x_6(40) < 0.75$ blue; $0.75 \leq x_6(40)$ red. The vertical scale is different for each panel. Small vertical panels are histograms of the predicted proportion of: (g) acroporid corals; (h) soft corals; (i) algae; (j) massive corals; (k) pocilloporid corals; and (l) free space on 1 January 2003, with vertical scales that match panels a to f respectively. Produced from an MCMCMC run with 3 chains, $\Delta T = 0.5$. Run for 505199 iterations (sampling every 100th iteration) after a burn-in of 4×10^5 iterations.

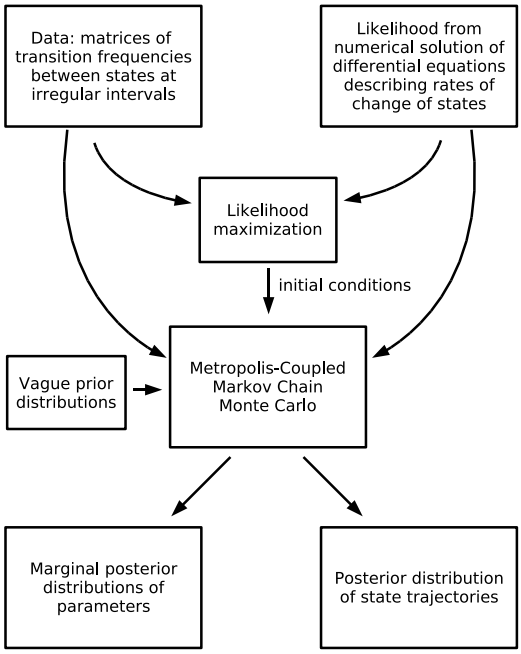
Data: matrices of transition frequencies between states at irregular intervals
Likelihood from numerical solution of differential equations describing rates of change of states
Likelihood maximization
initial conditions
Vague prior distributions
Metropolis-Coupled Markov Chain Monte Carlo
Marginal posterior distributions of parameters
Posterior distribution of state trajectories

Figure 1:

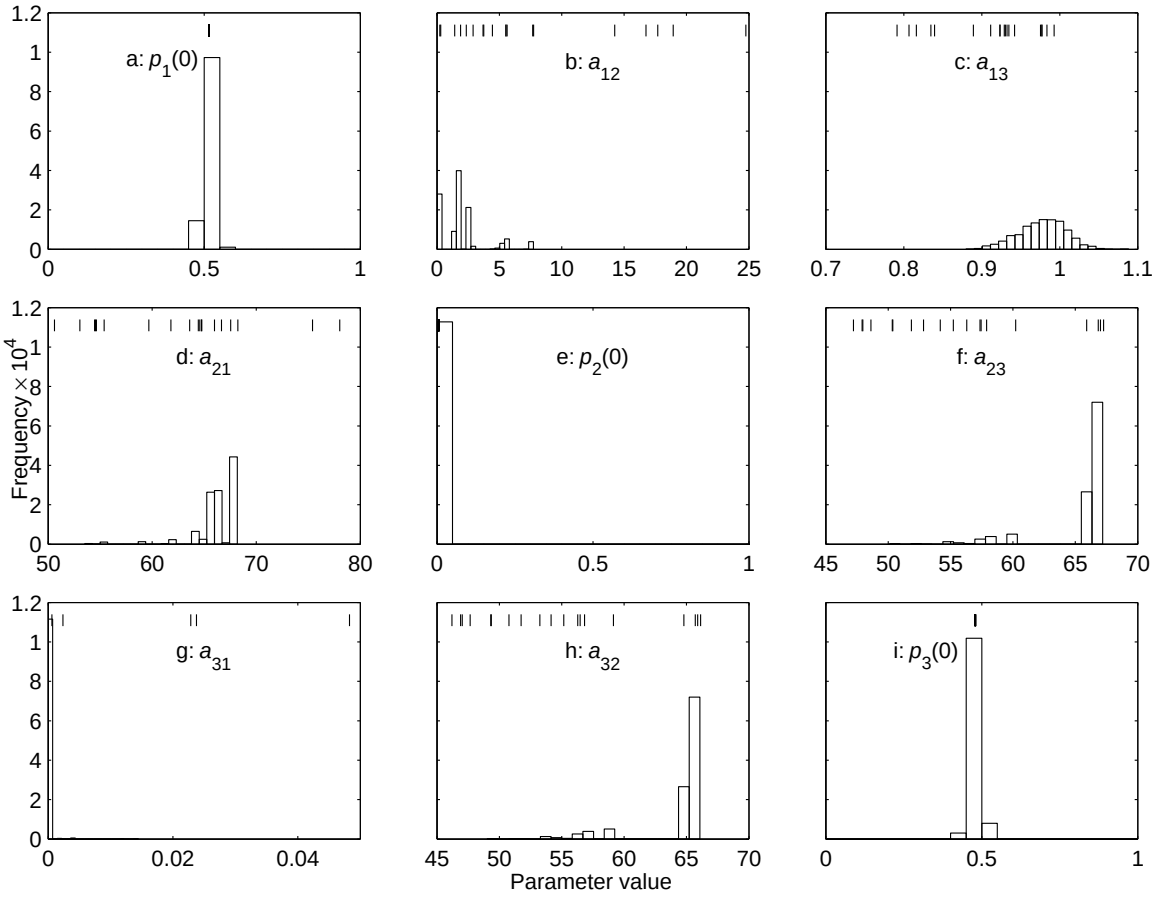
a: $p_1(0)$
b: a_{12}
c: a_{13}
d: a_{21}
e: $p_2(0)$
f: a_{23}
g: a_{31}
h: a_{32}
i: $p_3(0)$
Frequency $\times 10^4$
Parameter value

Figure 2:

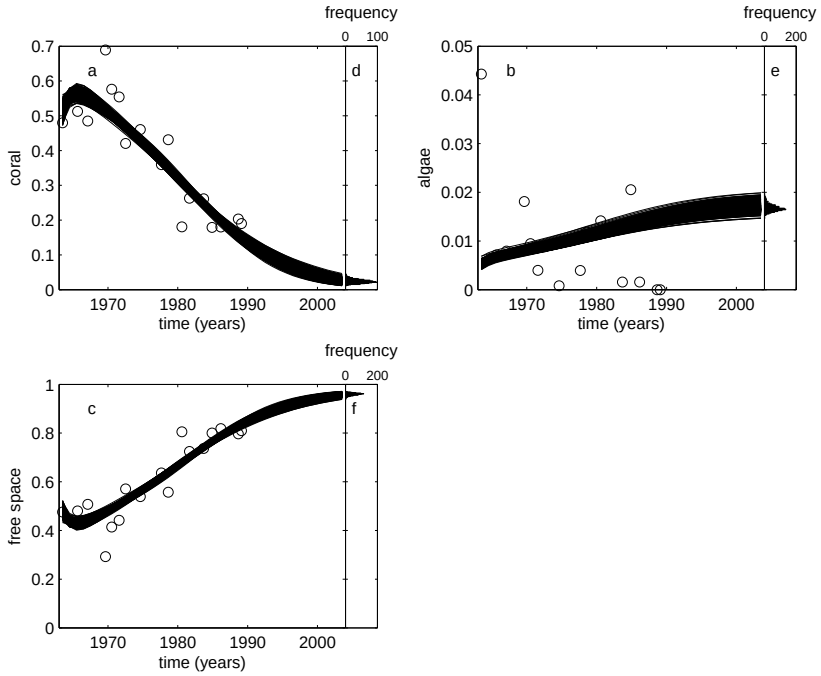
frequency
0 100
a
d
coral
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1970 1980 1990 2000
time (years)
frequency
0 200
b
e
algae
0.05
0.04
0.03
0.02
0.01
0
1970 1980 1990 2000
time (years)
frequency
0 200
c
f
free space
1
0.8
0.6
0.4
0.2
0
1970 1980 1990 2000
time (years)

Figure 3:

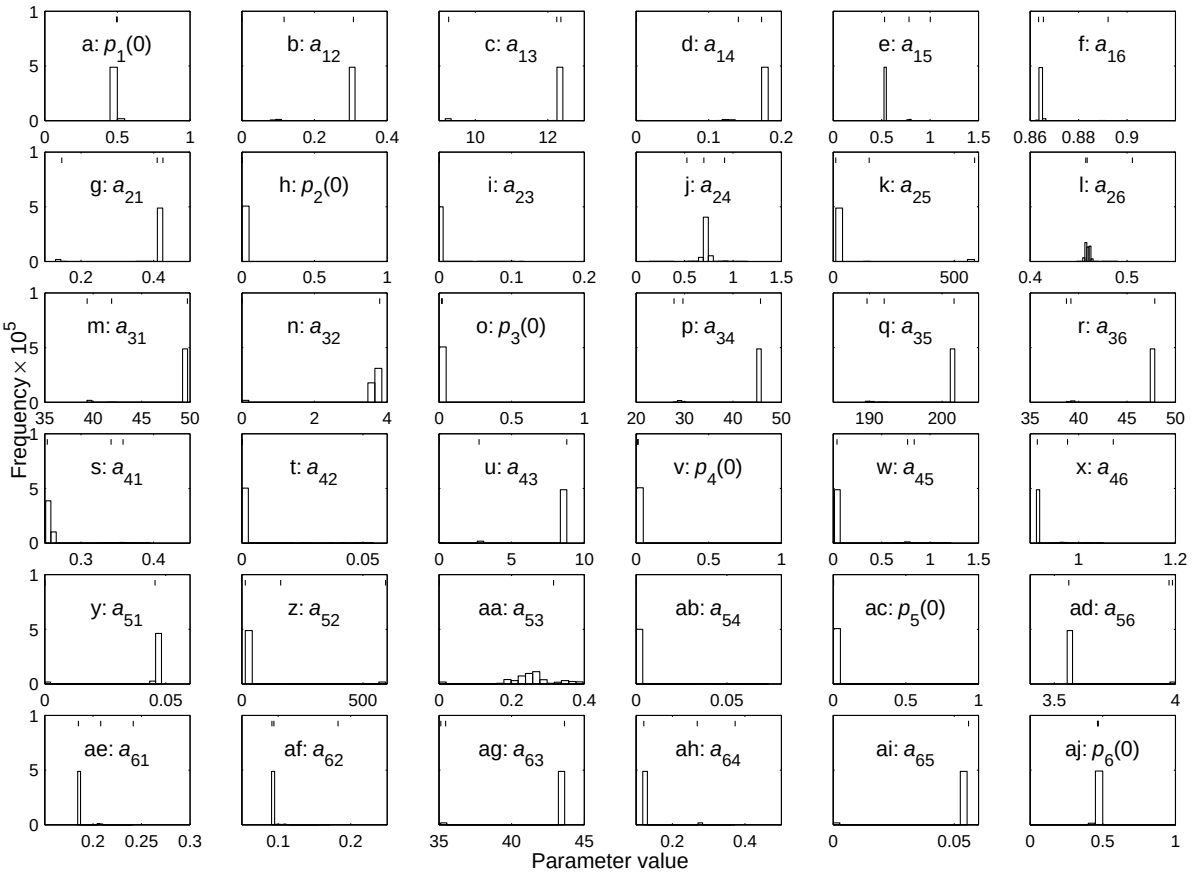
a: $p_1(0)$
b: a_{12}
c: a_{13}
d: a_{14}
e: a_{15}
f: a_{16}
g: a_{21}
h: $p_2(0)$
i: a_{23}
j: a_{24}
k: a_{25}
l: a_{26}
m: a_{31}
n: a_{32}
o: $p_3(0)$
p: a_{34}
q: a_{35}
r: a_{36}
s: a_{41}
t: a_{42}
u: a_{43}
v: $p_4(0)$
w: a_{45}
x: a_{46}
y: a_{51}
z: a_{52}
aa: a_{53}
ab: a_{54}
ac: $p_5(0)$
ad: a_{56}
ae: a_{61}
af: a_{62}
ag: a_{63}
ah: a_{64}
ai: a_{65}
aj: $p_6(0)$
Frequency × 10^5
Parameter value

Figure 4:

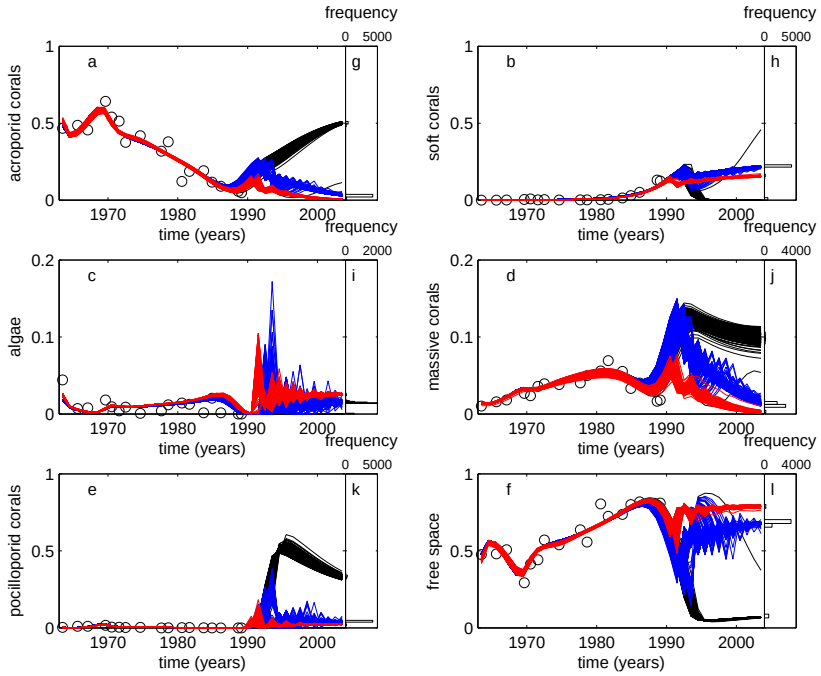
frequency
0 5000
a
g
acroporid corals
1
0.5
0
1970
1980
1990
2000
time (years)
b
h
soft corals
frequency
0 2000
c
i
algae
0.2
0.1
frequency
0 4000
d
j
massive corals
e
k
pocilloporid corals
f
l
free space

Figure 5: