Species abundance dynamics under neutral assumptions: a Bayesian approach to the controversy

Crispin M. Mutshinda[1]*, Robert B. O'Hara[1] and Ian P. Woiwod[2]

[1]*Department of Mathematics and Statistics, PO Box 68 (Gustaf Hällströmin katu 2b), FIN-00014, University of Helsinki, Helsinki, Finland; and* [2]*Plant and Invertebrate Ecology Division, Rothamsted Research, Harpenden, Herts. AL5 2JQ, UK*

Summary

1. Hubbell's 'Unified Neutral Theory of Biodiversity and Biogeography' (UNTB) has generated much controversy about both the realism of its assumptions and how well it describes the species abundance dynamics in real communities.

2. We fit a discrete-time version of Hubbell's neutral model to long-term macro-moth (*Lepidoptera*) community data from the Rothamsted Insect Survey (RIS) light-traps network in the United Kingdom.

3. We relax the assumption of constant community size and use a hierarchical Bayesian approach to show that the model does not fit the data well as it would need parameter values that are impossible.

4. This is because the ecological communities fluctuate more than expected under neutrality.

5. The model, as presented here, can be extended to include environmental stochasticity, density-dependence, or changes in population sizes that are correlated between different species.

Key-words: ecological drift, hierarchical Bayesian framework, stochasticity, zero-sum multinomial

Introduction

Understanding the dynamics of biodiversity remains a great challenge for ecological science. Hubbell (2001) developed the unified neutral theory of biodiversity and biogeography (UNTB), a controversial theory of how biodiversity originates and is maintained. Neutral models of community dynamics are based on the assumption of ecological equivalence between individuals and species, so that any variation in traits between trophically similar species has no impact on their overall abundance or their speciation rates (e.g. Gaston & Chown 2005; McGill, Maurer & Weiser 2006). There are many implementations (models) of neutrality, but the UNTB on which this study is focused remains the best-known.

The UNTB applies to communities of trophically similar species occurring in sympatry. It aims to explain the species richness and relative species abundance on different spatio-temporal scales: the 'local' or 'ecological' community is defined as a group of species that actually or potentially compete in a local area for the same or similar resources (Hubbell 2001), while the much larger regional species pool or 'metacommunity' can be viewed as a regional collection of local communities, the arena where speciation occurs and the pool from which local communities are colonized (Magurran 2005). Under the UNTB, the dynamics of local communities are governed by birth and death events, and species turnover due to local extinction and immigration

events taking place on an ecological time-scale, whereas changes in composition of the regional species pool are ascribed to global extinction and random speciation occurring on a much larger (evolutionary) time-scale. The theory is neutral in the sense that all ecological properties (per capita birth, death, immigration and speciation rate) are considered identical for all individuals regardless of their species' identities. Hence, changes in species abundances result only from demographic stochasticity. Moreover, communities are treated as having a fixed number of individuals (zero-sum assumption; Hubbell 2001, p. 54). Consequently, no new individual can be added to an ecological community until a vacancy has been created, essentially by a death. When the ecological community is dynamically coupled to a regional species pool, a new recruit to the community will be either an immigrant from the regional species pool with probability *m*, henceforth referred to as the 'immigration rate', or an offspring of one of the species present in the community at the previous time with probability $(1 - m)$. The dynamics of the model community evolves then as a random walk constrained by the fact that all species abundances must sum to a constant, but absolute species abundances increase and decrease stochastically. Hubbell (2001) calls this particular version of demographic stochasticity 'ecological drift' by analogy with genetic drift (Kimura 1983), and calls the resulting species abundance distribution (SAD) the 'zero-sum multinomial' (ZSM) distribution.

Criticisms of Hubbell's theory have concerned both its assumptions and the resulting SAD. For example, Maurer & McGill (2004) gave two ways in which species may be

**Correspondence author. E-mail: crispin.mutshindamwanza@helsinki.fi

asymmetric. First, they may differ in their competitive ability for limiting resources so that some species can limit the population sizes of others by limiting their access to the resources or reducing the abundance of resources. Second, they may be differently adapted to environmental conditions so that they react differently to environmental heterogeneity. Engen *et al.* (2002) analysed a community of tropical butterflies and found that the large component of variance was due to ecological heterogeneity among species, strongly violating the assumption of neutral community structure. Poulin (2004) opposed the zero-sum assumption and gave the example of helminth communities characterized by huge inequalities in body sizes among species in the same community, defying the assumption of saturation and the zero-sum dynamics. However, Etienne, Alonso & Mckane (2007) have argued that the zero-sum assumption is not critical.

Moreover, the performance of the ZSM has been the centre of a debate. McGill (2003) argued that the log-normal distribution provides a better fit than the ZSM for the tropical tree community on Barro Colorado Island in Panama; Volkov *et al.* (2003) claimed the opposite; Etienne & Olff (2004) found again that the log-normal distribution performs slightly better, and later that the ZSM performs either as well as or slightly better than the Poisson log-normal (Etienne & Olff 2005). In contrast, Williamson & Gaston (2005) maintain that neither distribution is a suitable SAD.

Arguments about whether a function fits an empirical relative abundance slightly better than another are unlikely to advance the field (Harte 2003; McGill *et al.* 2007). In reality, all these disputes may be beside the point: most of the previous tests of the neutral theory have been limited to examining the consistency of the empirical abundance-frequency distributions of local communities with the equilibrium predictions of the neutral theory. However, as mentioned by Wootton (2005), such tests are weak because the observed and predicted data are ranked by abundance, guaranteeing that both functions decline monotonically. On the other hand, because the abundance curves assume equilibrium, the differences in the results might be merely caused by environmental changes. Indeed, over the last 30 years or so, ecologists have been moving away from the equilibrium ideas (e.g. Wallington, Hobbs & Moore 2005). The UNTB is a dynamical model (Alonso, Etienne & McKane 2006), so it can be examined without assuming equilibrium as we proceed here by using long-term time series, with the zero-sum assumption relaxed. This is one of the first times that the neutral model is tested on temporal data (but see Etienne *et al.* 2007). Most of the previous studies have either looked at data from a single sample or temporally pooled several samples, for example, Alonso & McKane (2004) used some early Rothamsted moth data sets published by C. B. Williams in the 1940s (Fisher, Corbet & Williams 1943) by pooling samples over a 4-year period.

We proceed by fitting a discrete-time version of the UNTB to long-term moth community time series with the zero-sum assumption relaxed; we then ask whether it makes ecologically sensible predictions, that is, if the parameter estimates are realistic. In particular, we ask whether the changes in species abundances can be explained by the neutral model with a realistic community size.

Ecological processes are often observed incompletely with large and unknown amounts of measurement error or process uncertainty. Understanding and eventually predicting such processes requires a modelling framework which is capable of accommodating these uncertainties. The Bayesian approach to statistical inference turns out to be well-suited to this purpose (e.g. Anderson 1998); in particular the refined hierarchical Bayesian framework (e.g. Gelman *et al.* 2003) adopted in this study. We are unaware of studies that have used the hierarchical Bayesian framework to examine the neutral theory. Etienne & Olff (2004, 2005) used a Bayesian approach to estimate the SAD, but their models did not have a hierarchical structure and did not model the community dynamics, only the abundance distribution.

Given that a significant number of ecologists are still unclear about the rationale for using such an approach, we devote the rest of this introductory section to a brief review of the principles of the Bayesian analysis and the hierarchical Bayesian modelling; we refer readers interested in more details to the appropriate literature such as Gelman *et al.* (2003).

Bayesian analysis and the hierarchical Bayesian framework

BAYESIAN ANALYSIS

Bayesian analysis is an alternative approach to statistical inference that is increasingly used to evaluate ecological models and hypotheses (Ellison 2004). A Bayesian analysis starts with the formulation of a probabilistic model, $p(y \mid \theta)$, intended to describe the distribution of the data, conditionally on the unknown parameter(s) of interest, $\theta \in \Theta$. A prior distribution, $p(\theta)$, is subsequently formulated to convey the analyst's state of knowledge about the plausible parameter values before observing the data. As data become available, the prior distribution is updated to the posterior distribution, $p(\theta \mid y)$, via the Bayes' theorem:

$$p(\theta \mid y) = \frac{p(\theta, y)}{p(y)} = \frac{p(y \mid \theta)\, p(\theta)}{\int_{\Theta} p(y \mid \theta)\, p(\theta) d(\theta)} \propto p(\theta) p(y \mid \theta). \quad \text{eqn 1}$$

The lack of relevant prior information about plausible parameter values leads to the use of so-called non-informative or 'vague' priors such as uniform distributions on large compact regions or centred normal distributions with large variance.

The posterior distribution is the target of the Bayesian inference as it conveys all necessary information about the parameter(s) of interest. Bayesian conclusions are essentially made in terms of probabilistic statements about plausible parameter values or outcomes of future observations.

THE HIERARCHICAL BAYESIAN FRAMEWORK

A hierarchical Bayesian (HB) model is a multilevel Bayesian model. This means parameters in the likelihood depend on

other parameters not mentioned herein, which themselves require priors that may depend on new parameters; the process coming to an end when no new parameters are introduced. The HB modelling provides a flexible way of representing complex phenomena through a series of simple structures. Berlinier (1996) and Wikle & Hooten (2006) delineate three basic stages in the HB modelling of hidden processes namely, the data or sampling model, the process model and the parameter model which are intended to be conditionally linked in a hierarchical structure usually consisting of the sampling model at the lowest level.

The sampling, the process and the parameter models

The sampling model specifies the distribution of the data, *y*, conditionally on the state, *x*, of the process of interest. It takes into account, for example, the fact that only a proportion of statistical units are sampled. We refer to the parameters involved in this stage as 'data parameters' hereafter denoted by β. The capture probability is an example of such a parameter. The process model describes the actual growth process. The parameters involved in this stage are referred to as 'process parameters' and denoted by α. The demographic rates: the birth rate, the death or the immigration rates are examples of process parameters. The parameter model consists of prior elicitation for the lately introduced parameter for both the data model and the process model.

Posterior estimation and Markov chain Monte Carlo (MCMC) method

The target of inference is $p(x, \theta \mid y)$, the joint distribution of the process and the vector parameter $\theta = (\alpha, \beta)$ conditionally on the observed data, which derives from the Bayes' theorem as

$$p(x, \theta \mid y) \propto p(y \mid x, \theta)p(x \mid \theta)p(\theta). \qquad \text{eqn 2}$$

The derivation of the posterior distribution via eqn 2 is usually difficult in practice due to the normalizing constant which typically involves a high-dimensional integral with no analytic solution and thus requires a resort to numerical methods of approximation such as MCMC (e.g. Gilks, Richardson & Spiegelhalter 1996). MCMC methods make it possible to sample from distributions with complex algebraic forms, so subsequent inferences can be based on posterior summaries of the quantities of interest or functions of them calculated from the samples.

Materials and methods

DATA

We use three long-term (> 25 years) time series of macro moths (*Lepidoptera*) from the Rothamsted Insect Survey's (RIS) large-scale light-trap network in the United Kingdom (Woiwod & Harrington 1994). Two are from a small piece of woodland on the Rothamsted farm in Hertfordshire (Geescroft I and II), and the third in
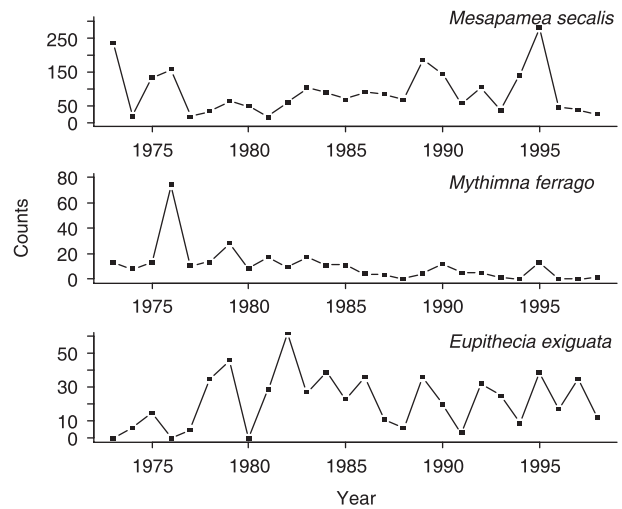
Fig. 1. Time evolution of the observed abundances of three species arbitrarily selected from the Geescoft 2 data set. Top: common rustic (*Mesapamea secalis*), middle: the clay (*Mythimna farrago*), bottom: mottled pug (*Epithecia exiguata*). Similar trajectories are present in the other data sets.

Table 1. Annual sample characteristics for the three full data sets

Data set	Characteristics	Number of species	Number of individuals
Geescroft I	2·5 percentile	154·95	3179·87
	Mean	180·05	5226·7
	97·5 percentile	202·4	10167·78
Geescroft II	2·5 percentile	124·37	1788·25
	Mean	159·69	3428·23
	97·5 percentile	183·75	5833·12
Tregaron	2·5 percentile	225·4	9772·6
	Mean	245·82	12715·93
	97·5 percentile	262·65	18543·12

woodland on the edge of a large National Nature Reserve in rural mid-Wales (Tregaron). Each data set consists of annual species totals. The numbers of species and years for the different data sets are: Geescroft I (352, 40); Geescroft II (319, 26); Tregaron (371, 28). Overall, the observed species' abundances proved to vary sensibly in consecutive years, with many species being not observed in some years. The sample characteristics for the three full data sets are summarized in Table 1, whereas three typical species' trajectories are displayed in Fig. 1.

Because light-traps rely on the physiological and behavioural response to light which may vary between species and families, we also carried out the same analyses on a single species-rich family alone, the geometrids (*Geometridae*), a moth family whose members are known to respond in a similar way to light and are sampled particularly efficiently by the RIS light-traps (Taylor & French 1974). The numbers of geometrid species in the three data sets are, respectively, 135, 127 and 135.

MODEL SPECIFICATION

Under the UNTB, the dynamics of an arbitrary species are governed by a generalized birth and death process (including speciation,

emigration and immigration) (Volkov *et al.* 2003). The UNTB assumes that individuals die and reproduce continuously, but this is unrealistic for many species (such as most of the moths considered here) which have an annual life cycle. This suggests using a model with discrete generations. We build a discrete-time neutral model of (local) community dynamics which is identical to Hubbell's UNTB in all other aspects except in one detail: we relax the assumption of a constant community size.

The UNTB makes the underlying assumption that all individuals are equivalent, and hence each has the same probability of producing offspring. For a discrete-time model, the process of reproduction can be viewed as one of sampling individuals at random (with replacement) from the current population to provide offspring for the next generation. If we condition on the new population size, the number of individuals of a species is binomially distributed with probability of success equal to the proportion of individuals of the focal species in the current population (the whole community is then drawn from a multinomial distribution). Crucially for us here, the variance in the number of offspring is related to the total population size, so that the more variation we see in the numbers of a species in a neutral community over time, the smaller the community size. This is because, when population size increases, the variance of the binomial distribution increases but the coefficient of variation decreases. Hence, the variation in the proportion of the focal species in the community decreases. This is exactly analogous to genetic drift, where the variation in allele frequencies can be used to estimate the effective population size (e.g. Waples 1989). In the same way, the variation in species frequencies can be used to estimate the 'effective' community size. We develop this argument more formally below.

We wish to emphasize that because we are considering the dynamics over short time-scales (evolutionarily speaking), we can reasonably assume that speciation (which is a rare event) does not occur, and restrict our attention to the ecological dynamics of the species.

The process model

Throughout, $N_{i,t}$ denotes the number of individuals of species i in the community at time t. So, $J_t = \Sigma_i N_{i,t}$ and $C_{i,t} = N_{i,t} \times (J_t)^{-1}$ denote, respectively, the community size and the relative abundance of species i in the community at time t. We designate by P_i the relative abundance of species i in the metacommunity, which is considered constant over an ecological time-scale where the local community dynamics are examined, and m_t denotes the immigration rate (i.e. the proportion of immigrants in the local community) at time t. For consistency with the data, we associate the value $t = 1$ with the first year of sampling.

As a consequence of the neutrality assumption, $N_{i,t}$ is entirely determined by $C_{i,t-1}$ and P_i through the reproduction and immigration processes, respectively, allowing for the drift in species abundances. More specifically, we assume that the expected number of individuals of species i in the community at time t is

$$\lambda_{i,t} = J_{t-1}[(1 - m_t) \times C_{i,t-1} + m_t \times P_i]. \qquad \text{eqn 3}$$

This means, on average, a proportion m_t of recruits to the community at time t are immigrants, and a fraction P_i of these are species i. The expression for $\lambda_{i,t}$ can alternatively be written as

$$\lambda_{i,t} = [(1 - m_t) \times N_{i,t-1} + m_t \times JP_{i,t}] \qquad \text{eqn 4}$$

where $JP_{i,t} = J_{t-1} \times P_i$ is real and non-negative.

The neutrality assumption entails no selective difference between species. So, the fraction $(1 - m_t)$ of individuals of species i derived locally will have the same expected relative frequency as in the previous generation (i.e. $C_{i,t-1}$). Eqn 3 is in essence identical to Hubbell's (UNTB) model for local community dynamics, except that the zero-sum assumption is relaxed here. Random drift is introduced into the model by thinking of $\lambda_{i,t}$ as a birth rate, and then the actual numbers will follow a Poisson distribution. That is

$$N_{i,1} \sim \text{Poisson}(\lambda_{i,t}). \qquad \text{eqn 5}$$

From the properties of the Poisson distribution, we know that $E(N_{i,t}) = \lambda_{i,t}$ and $\text{Var}(N_{i,t}) = \lambda_{i,t}$. Further, if we condition on the total community size, J_t, then the vector of species abundances, will have a multinomial distribution with the expected proportion for species i at time t being $E(C_{i,t}) = \lambda_{i,t} \times (\Sigma_j \lambda_{j,t})^{-1}$ (e.g. Agresti 1990, p. 38; Gelman *et al.* 2003, p. 431). The variance of the proportion of species i in the community is therefore $(C_{i,t} \times (1 - C_{i,t}))/J_t$, so the fluctuations in $C_{i,t}$ decrease with increasing community size J_t. This formalises the verbal argument given above.

The initial abundances of all species are also unobserved, and as such, need to be estimated from the data (e.g. Buckland *et al.* 2004; Clark & Bjørnstad 2004). In the Bayesian paradigm, this involves specifying priors on them. We assume that

$$N_{i,1} \sim \text{Poisson}(\lambda_{i,1}) \ (\lambda_{i,1} > 0). \qquad \text{eqn 6}$$

Immigration tends to stabilize communities around the metacommunity relative abundance. If we assume equilibrium, then we would expect to have $C_{i,t} = P_i$. It follows from eqn 3 that the expected abundance of species i in the community at equilibrium is $J_{t-1} \times P_i$ as suggested by Hubbell (2001, p. 90) where $J_t = J, \forall t$.

It should be stressed that we are concerned here exclusively with the local community dynamics. We do not therefore make any assumption about the form of the distribution of the metacommunity. Instead, we allow the model to estimate the distribution. As with the total community size, restricting the model is not necessary for making our main point, so we choose to allow the model to be more flexible. This also means that our main results will not be due to these secondary assumptions.

The sampling model

The observed abundance of species i at time t ($t \geq 2$), $y_{i,t}$, (i.e. the number of individuals of species i trapped at time t) can be modelled assuming

$$y_{i,t} \sim \text{Poisson}(N_{i,t} \times q_t) \qquad \text{eqn 7}$$

where $q_t > 0$ is a parameter, henceforth referred to as 'sampling rates' at time t, and whose interpretation is as follows: at any time, q estimates the ratio of the observed community size to the (effective) size of a neutral community which corresponds to the observed level of variation. In case a community is completely observed, the expected sampling rate under neutrality is 1. If however the dynamics are not neutral, the excess of variation over the neutral expectations (i.e. a greater temporal fluctuation in numbers of each species than predicted from multinomial sampling) will tend to deflate the (effective) community size which would make the sampling rates to exceed 1, in virtue of the eminent inverse relationship between level of variation and effective population size. It is then clear that under neutrality, the sampling rates correspond to

the observation probabilities (probability of capture in this case). We note that the sampling rates at time points $t > 1$ are intrinsically identifiable (e.g. Haslett *et al.* 2006), owing to the additional information coming from the previous state conveyed by the underlying Markov structure.

As can be seen from the model specification, $N_{i,t}$ and q_t are parameters, and m_t and $JP_{i,t}$ are hyper-parameters, m_t and $JP_{i,t}$ being at a higher level in the hierarchy than $N_{i,t}$ and q_t. A full Bayesian specification of the model requires priors to be explicitly specified on all independent parameters.

Prior elicitation

The model was fitted with relatively vague and independent priors on the fitted parameters (m_t, q_t, $\lambda_{i,1}$, $JP_{i,t}$). We used Uni(0, 0·5) on m_t which allows the abundance of immigrants up to half of the entire community size, and Exp(0·01)on $\lambda_{i,1}$ which is positive and typically 'flat' far from zero, to allow for large initial values. A non-informative approach to expressing ignorance about the relative species abundances of the regional species pool is to consider that all species are equally abundant in there. Under the equilibrium perspective, this would suggest setting all $JP_{i,t}$ at the average species abundance in the (local) community. We placed the non-informative Uni(5, 1000) on all $JP_{i,t}$ and assigned Gam(3, 6) to q_t. Indeed the expected value and the variance of Gam(3, 6) match those of Uni(0, 1) whose support corresponds to the expected range of this parameter under neutrality.

MODEL FITTING

The required posteriors were simulated numerically by MCMC through OPENBUGS (Thomas *et al.* 2006), an interactive windows version of the BUGS software (Spiegelhalter *et al.* 2003) which is available at <http://mathstat.helsinki.fi/openbugs/> and a copy of the BUGS code is available from the authors. The quantities of interest were the actual species abundances, $N_{i,t}$, the sampling rates, q, and the immigration rate m. The community sizes at different time points were obtained by summing the abundances of all species in the community.

We ran 300 000 iterations of three chains, discarding the first 100 000 as burn-in and thinning the remainder to one in every tenth observation. We carried out an extensive sensitivity analysis of the results to the prior inputs by substantially varying the range of the prior inputs, but the results remained broadly robust to these changes.

Results

The results for the full macro-moth data sets and for the geometrids alone are shown in Figs 2 and 3, respectively. As might be expected, the two Geescroft data sets show a similar pattern, with a downwards temporal trend in abundance. In contrast, there does not appear to be any trend at Tregaron (see solid lines in Fig. 2a–c).

In all three data sets, the community sizes undergo substantial fluctuations (Fig. 2a–c), suggesting that community size is not constant.

The model-predicted community sizes are far too low to be realistic (Fig. 2a–c) as they are much lower than the actual sample sizes (i.e. the number of individuals caught). Indeed,

the estimated sampling rates (Fig. 2d–f) suggest that, if the neutral model is true, the communities should contain much fewer individuals than were actually caught. The dashed horizontal lines in panels d–f of Figs 2 and 3 are drawn at height 1, which corresponds to the expected sampling rate for a completely observed community under neutrality as discussed above.

The immigration rates were estimated to be very low (Fig. 2g–i). In the light of the results above, this is not surprising, as a high immigration rate would tend to stabilise the dynamics, as expected relative frequencies of immigrants are constant over time.

The results are similar for the analysis using only the geometrid species: a similar level of excessive temporal variation was observed in all three populations (Fig. 3).

To verify that the biologically unintuitive results are not due to any feature of our model, we fitted the model to 1000 time series simulated using the neutral model with and without immigration. In all cases, the estimated sampling rates were found to be sensible (always below 1 and close to the true values). Figure 4 displays the sampling rates obtained by fitting the model to a data set simulating a 10-species neutral community dynamics over 30 renewal periods with the parameters q_t and m_t set to 0·5 and 0·2, respectively. The model was fitted using Gam(3, 6) priors on q_t. We ran 200 000 iterations of three chains, discarding the first 50 000 iterations as burn-in and thinning the remainder to each tenth observation. As can be seen from Fig. 4, the posterior means of q_t are broadly contained within the range 0·4 to 0·6 with an overall mean at 0·5 which corresponds to the true value.

Discussion

Our results suggest much more variation in the moth time series than the assumed neutral model can explain. Indeed as discussed above, if the community dynamics were conform to the neutral theory/model, the sampling rates q_t would lie between zero and one, and should correspond to the probabilities of capture. Keeping in mind the inverse relationship between community size and magnitude of variation in relative species abundances emphasized in the Materials and methods section, it is clear that an excess of variation over the neutral expectations will slim down the (idealised) community size which may even drop below the observed values, resulting in estimated sampling rates which are higher than 1, as found here.

As already mentioned, the sampling rates correspond to probabilities of capture provided the actual community dynamics are neutral. In case a neutral community is completely observed, we expect the sampling rates to be 1. Given that we are dealing herein with data on only partially observed species, neutral dynamics would necessarily yield sampling rates between 0 and 1. Because the model-predicted sampling rates are far beyond 1, we can logically conclude that the assumed neutral model and/or demographic stochasticity alone cannot explain the large fluctuations in the three moth communities.
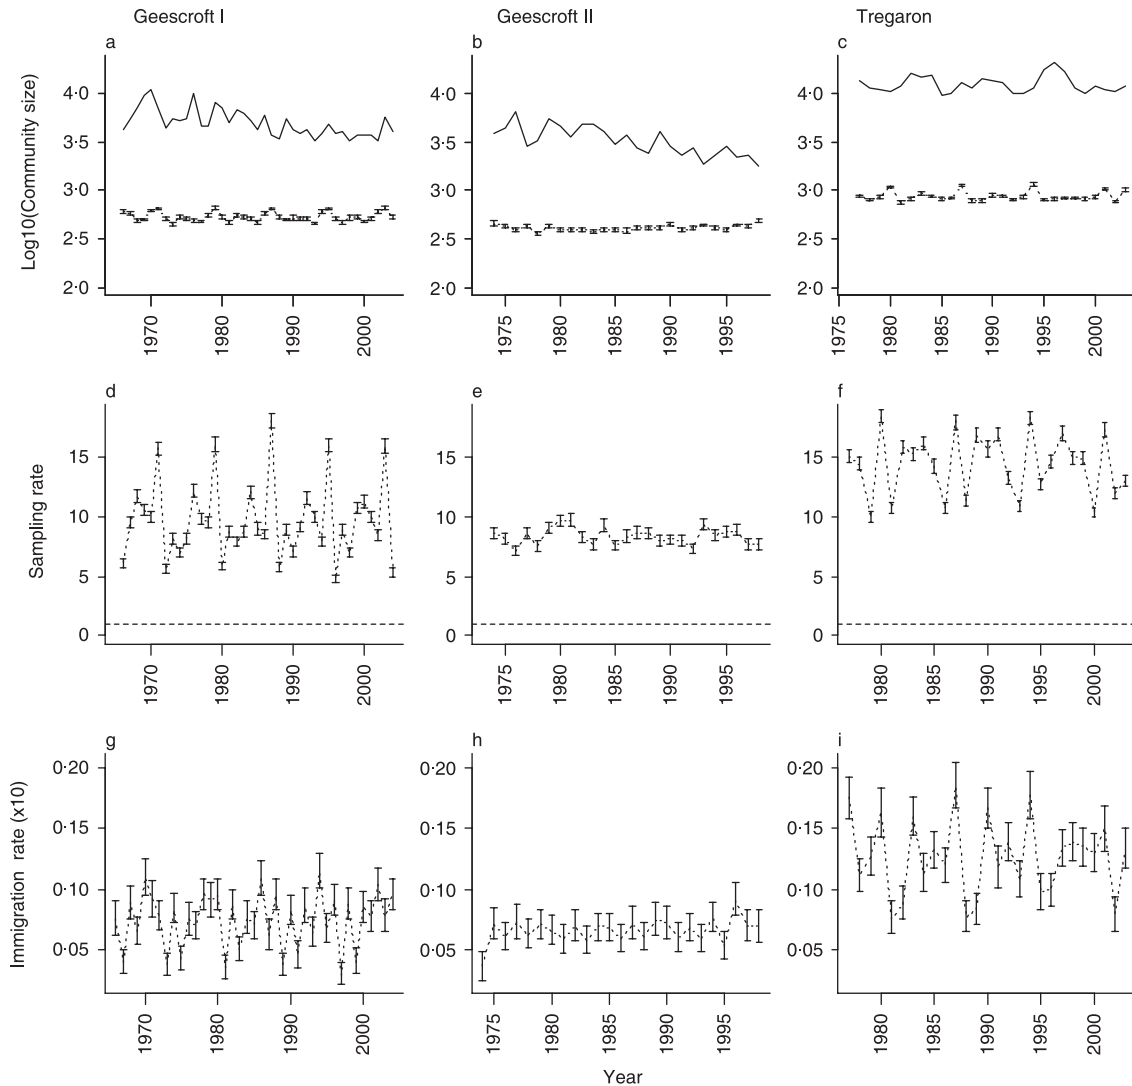
Fig. 2. Posterior medians and 95% credibility sets for the estimated community sizes and the corresponding observed values (solid lines) (a–c); posterior medians and 95% credibility sets for the estimated sampling rates (d–f) and for 10× immigration rates (g–i) for the full macro-moth data sets. Note that the dashed horizontal lines in panels d–f are drawn at height 1.

What causes this excess of temporal variation over the neutral expectations? An obvious explanation is the effect of fluctuations in the environment. Indeed, the importance of this environmental stochasticity, which is known to affect the dynamics of populations regardless of their sizes has been recognized by population ecologists (e.g. Lande, Engen & Saether 2003; Chase 2005), so it is not surprising that we find it in community dynamics as well. The poor fit of the model may also be due to other model assumptions. For example, the assumption that the birth-death-immigration processes are linear and do not involve any kind of density-dependence, although density-dependence has been shown to occur in a larger set of similar RIS moth data (Woiwod & Hanski 1992).

Whilst the UNTB has provoked an upsurge of interest in community dynamics, its failure to include some major sources of fluctuation in population size limits its predictive ability. Nonetheless, it may continue to serve as a null model

in some contexts, in particular when functional hypotheses are being tested (Bell 2000; Etienne & Alonso 2005; Nee 2005; Alonso *et al.* 2006; McGill *et al.* 2006), and its value might be as a springboard for merging community and population biology.

The neutral model, as presented here, can be extended to include, for example, environmental stochasticity (e.g. Alonso, Etienne & McKane 2007; Benedetti-Cecchi 2007), density-dependent effects, or changes in population size that are correlated between different species. In practice, these models can be developed and fitted through the hierarchical modelling framework that is becoming popular amongst Bayesian statisticians.

In conclusion, although our analyses have shown that the neutral theory/model does not fit the Rothamsted moth data well, we feel that Hubbell's model is wrong in informative ways that will continue to stimulate the development of community ecology theory.
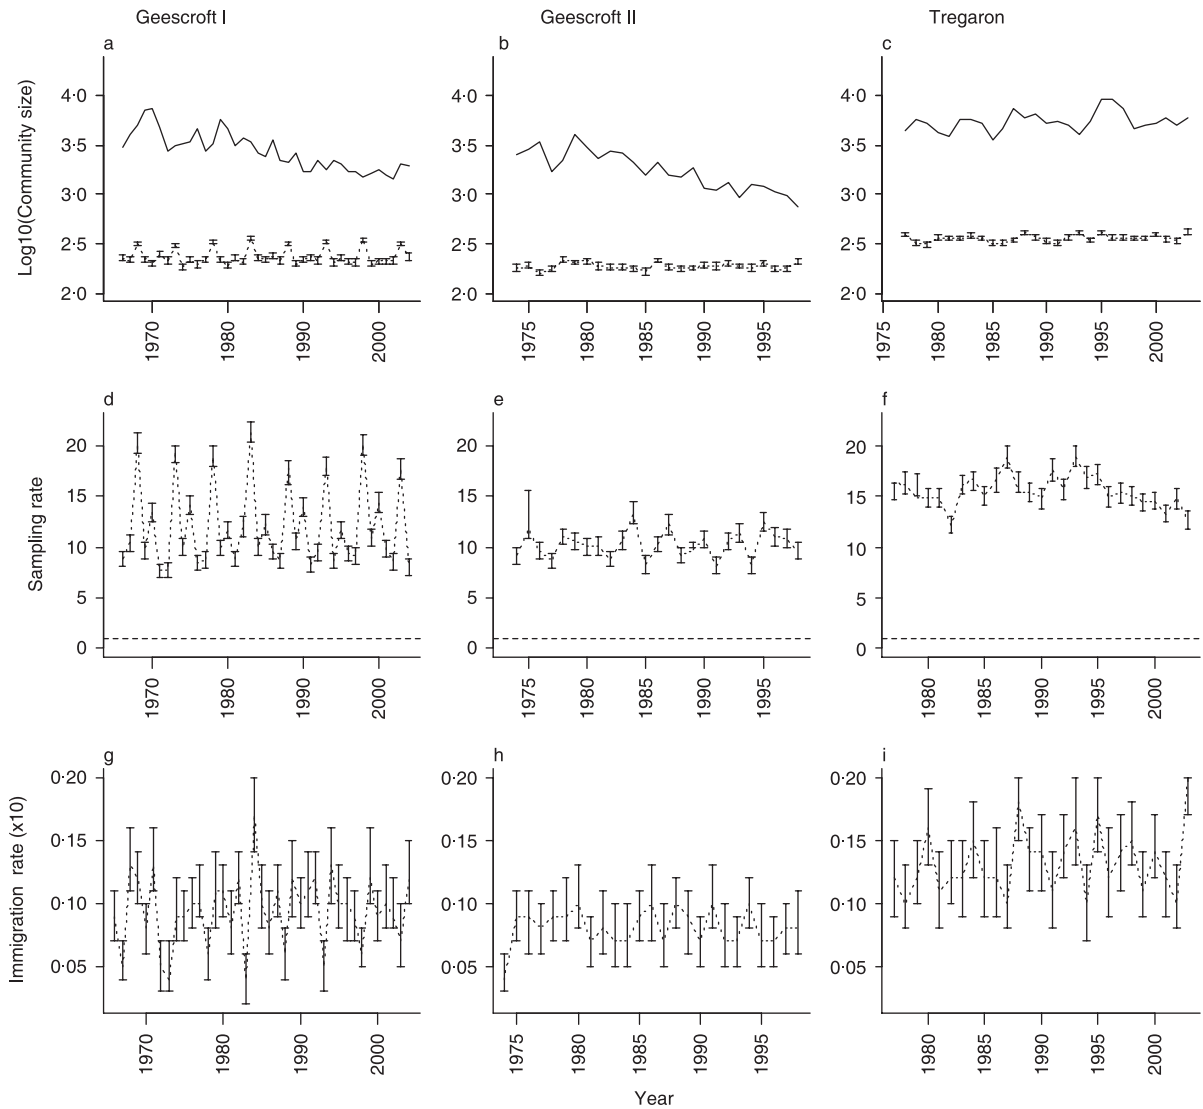
Fig. 3. Posterior medians and 95% credibility sets for the estimated community sizes and the corresponding observed values (solid lines) (a–c); posterior medians and 95% credibility sets for the estimated sampling rates (d–f) and for 10× immigration rates (g–i) for the geometrid species. The dashed horizontal lines in panels d–f are drawn at height 1.
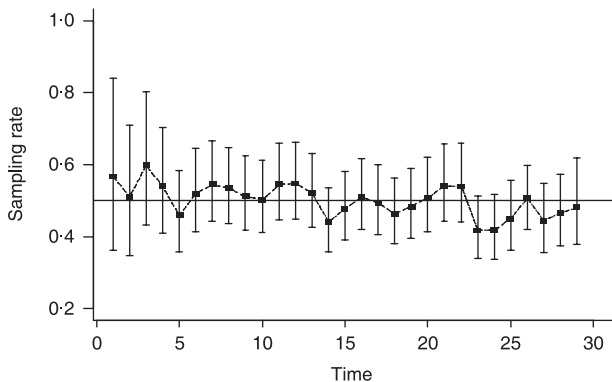
Fig. 4. Posterior means and 95% credibility sets of sampling rates for a 10-species neutral community simulated with the true value of the sampling rate uniquely set to 0·5 (at the level of horizontal line) over 30 renewal periods.

References

Agresti, A. (1990) *Categorical Data Analysis*. John Wiley & Sons, Inc, New York.

Alonso, D. & McKane, A.J. (2004) Sampling Hubbell's neutral theory of biodiversity. *Ecology Letters*, **7**, 901–910.

Alonso, D., Etienne, R.S. & McKane, A.J. (2006) The merits of the neutral theory. *Trends in Ecology and Evolution*, **21**, 451–457.

Alonso, D., Etienne, R.S. & McKane, A.J. (2007) Response to Benedetti-Cecchi: neutrality and environmental fluctuations. *Trends in Ecology and Evolution*, **22**, 232.

Anderson, J.L (1998) Embracing uncertainty: the interface of Bayesian statistics and cognitive psychology. *Conservation Ecology* **2**(1), 2. <http://www.consecol.org/vol2/iss1/art2/>.

Bell, G. (2000) Distribution of abundance in neutral communities. *American Naturalist*, **155**, 606–617.

Benedetti-Cecchi, L. (2007) Neutral theory and 1/f noise make similar predictions of assemblage dynamics. *Trends in Ecology and Evolution*, **22**, 231.

Berlinier, L.M. (1996) Hierarchical Bayesian time series models. *Maximum Entropy and Bayesian Methods* (eds K. Hanson & R. Silver), pp. 15–22. Kluwer Academy Publishers, Dordrecht, Netherlands.

Buckland, S.T, Newman, K.B., Thomas, L. & Koesters, N.B. (2004) State-space models for the dynamics of wild animal populations. *Ecological Modelling*, **171**, 157–175.

Chase, J.M. (2005) Towards a really unified theory for metacommunities. *Functional Ecology*, **19**, 182–186.

Clark, J.S. & Bjørnstad, O.M. (2004) Population time series: process variability, observation errors, missing values, lags, and hidden states. *Ecology*, **85** (11), 3140–3150.

Ellison, A.M. (2004) Bayesian inference in ecology. *Ecology Letters*, **7**, 509–520.

Engen, S.B.J., Lande, R., Walla, T. & Devries, P.J. (2002) Analyzing spatial structure of communities using the two-dimensional poisson lognormal species abundance model. *American Naturalist*, **160**, 60–73.

Etienne, R.S. & Alonso, D. (2005) A dispersal-limited sampling theory for species and alleles. *Ecology Letters*, **8**, 1147–1156.

Etienne, R.S. & Olff, H. (2004) A novel genealogical approach to neutral biodiversity theory. *Ecology Letters*, **7**, 170–175.

Etienne, R.S. & Olff, H. (2005) Confronting different models of community structure to species-abundance data: a Bayesian model comparison. *Ecology Letters*, **8**, 493–504.

Etienne, R.S., Alonso, D. & McKane, A.J. (2007) The zero-sum assumption in the neutral biodiversity theory. *Journal of Theoretical Biology*, **248**, 522–536.

Fisher, R.A., Corbet, A.S. & Williams, C.B. (1943) The relation between the number of species and the number of individuals in a random sample of an animal population, Part 2. Results obtained by means of a light-trap at Rothamsted. *Journal of Animal Ecology*, **12**, 42–58.

Gaston, K.J. & Chown, S.L. (2005) Neutrality and niche. *Functional Ecology*, **19**, 1–6.

Gelman, A., Carlin, J.B., Stern, H.S. & Rubin, D.B. (2003) *Bayesian Data Analysis*, 2nd edn. Chapman and Hall, London, UK.

Gilks, W.R., Richardson, S. & Spiegelhalter, D.J. (1996) *Markov Chain Monte Carlo in Practice*. Chapman and Hall, London, UK.

Harte, J. (2003) Tail of death and resurrection. *Nature*, **424**, 1006–1007.

Haslett, S., Millar, R.B., Newman, K.B., Plummer, M. & Welsh, A.H. (2006) Invited discussion of the paper by Domijan, Jorgensen Reid. *Australian and New Zealand Journal of Statistics*, **48**, 393–413.

Hubbell, S.P. (2001) *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton University Press, Princeton.

Kimura, M. (1983) *The Neutral Theory of Molecular Evolution*. Cambridge University Press, New York.

Lande, R., Engen, S. & Saether, B.E. (2003) *Stochastic Population Dynamics in Ecology and Conservation*. Oxford University Press, New York.

Magurran, A.E. (2005) Species abundance distributions: pattern or process? *Functional Ecology*, **19**, 177–181.

Maurer, B.A. & McGill, B.J. (2004) Neutral and non-neutral macroecology. *Basic and Applied Ecology*, **5**, 413–422.

McGill, B.J. (2003) A test of the unified neutral theory. *Nature*, **422**, 881–885.

McGill, B.J., Etienne, R.S., Gray, J.S., Alonso, D., Anderson, M.J., Benecha, H.K., Dornelas, M., Enquist, B.J., Green, J.L., He, F., Hurlbert, A.H., Magurran, A.E., Marquet, P.A., Maurer, B.A., Ostling, A., Soykan, C.U., Ugland, K.I. & White, E.P. (2007) Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. *Ecology Letters*, **10**, 995–1015.

McGill, B.J., Maurer, B.A. & Weiser, M. (2006) Empirical evaluation of neutral theory. *Ecology*, **87** (6), 1411–1423.

Nee, S. (2005) The neutral theory of biodiversity: do the numbers add up? *Functional Ecology*, **19**: 173–176.

Poulin, R. (2004) Parasites and the neutral theory of biodiversity. *Ecography*, **27** (1), 119–123.

Spiegelhalter, D., Thomas, A., Best, N. & Lunn, D. (2003) WinBugs version 1·4 User manual. <http://www.mrc-bsu.cam.ac.uk/bugs/>.

Taylor, L.R. & French, R.A. (1974) Effects of light-trap design and illumination on samples of moths in an English woodland. *Bulletin of Entomological Research*, **63**, 583–594.

Thomas, A., O'Hara, R.B., Ligges, U. & Sturtz, S. (2006) Making BUGS Open. *R News*, **6**, 12–17.

Volkov, I., Banavar, J.R, Hubbell, S.P. & Martian, A. (2003) Neutral theory and relative species abundance in ecology. *Nature*, **424**, 1035–1037.

Wallington, T.J., Hobbs, R. & Moore, S.A. (2005) Implications of current ecological thinking for biodiversity conservation: a review of salient issues. *Ecology and Society*, **10**, 15.

Waples, R.S. (1989) A generalized approach for estimating effective population size from temporal changes in allele frequency. *Genetics*, **121**, 379–391.

Wikle, C.K. & Hooten, M.B. (2006) Hierarchical bayesian spatio-temporal models for population spread. *Applications of Computational Statistics in the Environmental Sciences: Hierarchical Bayes and MCMC Methods* (eds J. S. Clark & A. Gelfand), pp. 145–169. Oxford University Press, Oxford.

Williamson, M. & Gaston, K.J. (2005) The lognormal is not an appropriate null hypothesis for the species-abundance distribution. *Journal of Animal Ecology*, **74**, 409–422.

Woiwod, I.P. & Hanski, I. (1992) Patterns of density dependence in moths and aphids. *Journal of Animal Ecology*, **61**, 619–629.

Woiwod, I.P. & Harrington, R. (1994) Flying in the face of change: the Rothamsted Insect Survey. *Long-term Experiments in Agricultural and Ecological Sciences* (eds R.A. Leigh & A.E. Johnston), pp. 321–342. CAB International, Wallingford.

Wootton, J.T. (2005) Field parameterization and experimental test of the neutral model. *Nature*, **433**, 309–312, doi: 10.1038/nature03211.