UC Merced
Frontiers of Biogeography

Title
Fitting and comparing competing models of the species abundance distribution: assessment and prospect

Permalink
https://escholarship.org/uc/item/3gz504j3

Journal
Frontiers of Biogeography, 6(2)

Authors
Matthews, Thomas J
Whittaker, Robert J

Publication Date
2014

DOI
10.21425/F5FBG20607

Copyright Information

Fitting and comparing competing models of the species abundance distribution: assessment and prospect

Thomas J. Matthews[1,2,*] and Robert J. Whittaker[1,3]

[1]Conservation Biogeography and Macroecology Group, School of Geography and the Environment, University of Oxford, South Parks Road, Oxford, OX1 3QY, UK

[2]Azorean Biodiversity Group (ABG, CITA-A) and Portuguese Platform for Enhancing Ecological Research & Sustainability (PEERS), Departamento de Ciências Agrárias, University of the Azores, Rua Capitão João d´Ávila, Pico da Urze, 9700-042, Angra do Heroísmo, Portugal

[3]Center for Macroecology, Evolution and Climate, Department of Biology, University of Copenhagen, Universitetsparken 15, DK-2100 Copenhagen Ø, Denmark

*thomas.matthews@ouce.ox.ac.uk

Abstract. A species abundance distribution (SAD) characterises patterns in the commonness and rarity of all species within an ecological community. As such, the SAD provides the theoretical foundation for a number of other biogeographical and macroecological patterns, such as the species–area relationship, as well as being an interesting pattern in its own right. While there has been resurgence in the study of SADs in the last decade, less focus has been placed on methodology in SAD research, and few attempts have been made to synthesise the vast array of methods which have been employed in SAD model evaluation. As such, our review has two aims. First, we provide a general overview of SADs, including descriptions of the commonly used distributions, plotting methods and issues with evaluating SAD models. Second, we review a number of recent advances in SAD model fitting and comparison. We conclude by providing a list of recommendations for fitting and evaluating SAD models. We argue that it is time for SAD studies to move away from many of the traditional methods available for fitting and evaluating models, such as sole reliance on the visual examination of plots, and embrace statistically rigorous techniques. In particular, we recommend the use of both goodness-of-fit tests and model-comparison analyses because each provides unique information which one can use to draw inferences.

Keywords. Bayesian statistics, information theory, macroecology, maximum entropy, maximum likelihood, rank–abundance plot, species abundance distribution

Introduction

A species abundance distribution (SAD) describes the abundance of all species recorded within an ecological community, assemblage or sample. The study of SADs has a long history in biogeography and community ecology (e.g. Preston 1948, Tokeshi 1993, McGill et al. 2007), and the observation that ecological systems contain few very abundant species and relatively more rare species (the so-called hollow curve) is often labelled as one of the few universal laws in this subject space (McGill et al. 2007, Ulrich et al. 2010). SADs also provide the theoretical foundation for exploration of other biogeographical and macroecological patterns, such as the species–area relationship ('SAR', e.g. Preston 1948, Hubbell 2001, Whittaker & Fernández-Palacios 2007). The SAD has broad utility in applied ecology and biogeography and a number of applications are described in Table 1.

A plethora of SAD models have been described in the literature; indeed, a review by McGill et al. (2007) lists 27 different models. A summary of the models referenced in our review is given in Table 2. These models can largely be classified based on the mechanisms underlying the model, and traditionally models have been dichotomised as either statistical or biological (i.e. non-niche-oriented and niche-oriented). The distinction between statistical and biological models, as Magurran (2004, p.28) states, has become 'blurred,' and many distributions originally classified as statistical, such as the logseries and log-

Table 1. Applications of the species abundance distribution (SAD) in ecology and biogeography.

Broad application	Description	References
Disturbance ecology	The SAD has been used as ecological indicator to measure the effects of disturbance (e.g. pollution, land-use change) on biotic communities.	Gray & Mirza (1979)
Hotspot selection	Rank–abundance plots have been shown to aid in the selection of regional biodiversity hotspots for conservation purposes.	Dunstan et al. (2012)
Gradient analysis	The SAD has been used to analyse changes in community structure along ecological gradients, e.g. gradients of vegetation succession.	Bazzaz (1975)
Extrapolation	A number of studies have explored the scaling properties of SADs, including the prediction of the SAD of a large area from the SADs of its smaller constituent areas (upscaling) and vice versa (downscaling).	Hubbell (2001), Zillio & He (2010)
Community structure analysis	The SAD has been used to infer information relating to how ecological communities are structured. For example, the deconstruction of multimodal SADs has illustrated how non-native species can inflate the number of rare species in ecological samples.	Matthews et al. (2014)

normal, have since been interpreted biologically. For example, niche-based theories concerning resource allocation have predicted logseries and truncated lognormal distributions (Table 2, Sugihara 1980, Connolly et al. 2005), while demographic models based on ecological drift (such as Hubbell's interpretation of neutral theory, Box 1) predict logseries SADs at the metacommunity scale. As an aside, several authors (e.g. Tokeshi 1993, Magurran 2004) further classify SAD models into deterministic and stochastic models and stress the importance of knowing where a model falls in this dichotomy, in terms of fitting. In practice, the former include many of the statistical models, such as the lognormal and logseries, and make the assumption that individuals are distributed between species in a deterministic manner; that is, frequency values cannot vary. Stochastic models, which include the entire set of niche-oriented biological models with the exception of

Box 1. Neutral theory and the species abundance distribution

The Unified Neutral Theory of Biodiversity and Biogeography (Hubbell 2001) is a biogeographical theory which emphasises the importance of stochastic processes in ecological community structure. In this context, neutrality means that, for trophically similar species, all species are equivalent in terms of birth, death, speciation and dispersal rates, when measured on a per-capita basis. The theory is based on the assumption that differences between species are not important in determining ecological community properties, such as relative abundances. As such, it rejects two concepts that have formed the basis of traditional ecological research and niche theory, namely that species are ecologically and functionally different, and that environmental context is important. A number of different neutral models have been put forward under the 'neutral theory' banner and these can be largely dichotomised into spatially implicit neutral models (i.e. those in which spatial processes are only implied, e.g. Hubbell's classic two tier metacommunity–local community model) and spatially explicit neutral models (i.e. models with a multi-scale structure; see Rosindell et al. 2008).

Neutral models are capable of predicting many different macroecological and biogeographical patterns, including the species–area relationship and the species abundance distribution (SAD). However, the majority of empirical tests of neutral theory have focused on the SAD, usually testing the spatially implicit model of Hubbell (2001; for a review see Matthews and Whittaker 2014). Hubbell's model predicts a logseries distribution at the metacommunity scale and a zero-sum multinomial distribution (ZSM) at the local scale (see Table 1). The ZSM has three parameters which in theory can be measured empirically, including the number of individuals in the local community and a parameter specifying the degree of immigration from the metacommunity. Early tests of the ZSM were hampered by the lack of a full analytical solution for the distribution, meaning the model had to be fit, and the goodness of fit estimated, using simulation techniques. However, more recent advancements (e.g. Etienne 2005) have derived a full analytical solution, meaning that information-theoretic approaches, amongst others, can now be used to compare the fit of the ZSM with other candidate models (e.g. see Figure 1). These tests have shown that, despite its simple assumptions, the ZSM (and the SADs of other neutral models) often provides a very good fit to empirical data, but not necessarily a superior fit to certain competitor models.

Table 2. Descriptions of the various species abundance distribution (SAD) models referenced within the review. For each model, a description, selection of relevant references, and the model family according to McGill et al. (2007) are given.

SAD model	Family	Description	Relevant references
Gambin	Purely statistical	Gambin is a stochastic model which combines the gamma distribution with a binomial sampling method. Gambin has a single free parameter (α), which characterises the distribution shape.	Ugland et al. (2007)
Lognormal	Purely statistical	The lognormal represents a situation in which the logarithms of the different species' abundances follow a Gaussian distribution and as such it characterises a community with relatively few very abundant or very rare species; whereby the observed frequencies increase to a modal frequency and then decrease.	Preston (1948), May (1975)
Truncated lognormal	Purely statistical	As for the lognormal, but the distribution is truncated at the left hand side in accordance with Preston's (1948) concept of the veil line: the idea that rarer species are often missed because of under-sampling and thus that the left hand side of the lognormal distribution is not revealed. This truncation methodology has since been criticised as inexact and other authors have advocated techniques such as Poisson sampling and hypergeometric sampling, as more accurate alternatives (e.g. Dewdney 1998).	Preston (1948), Dewdney (1998)
Logseries	Purely statistical	The logseries (logarithmic series) results from the Poisson sampling of a gamma distribution after a certain relevant limit is taken, and conditional presence is considered. That is, it gives the conditional probability of attaining a certain abundance level given that the species is present. The number of species at different levels of abundance predicted by the logseries is given by: $S_n = (\alpha * x^n)/n$, where α and x are constants, and S_n is the number of species with n individuals. This distribution is characterised by a skewed J-curve with a modal value of 1 in arithmetic space. There is no fixed variance assumption and an increase in sample size leads to an increased variation in log abundance.	Fisher et al. (1943)
Poisson lognormal (PLN)	Purely statistical	An issue with the continuous lognormal is that it allows fractional abundances and does not have an associated sampling theory. One workaround in this regard is to incorporate Poisson sampling variability and fit the Poisson lognormal distribution: the Poisson sampling of individuals from a standard lognormal distribution.	Bulmer (1974)
Niche partitioning models	Niche partitioning	Models in which it is assumed that the total niche space of a community, often termed 'the stick', is sequentially divided in some manner between the species that comprise this community. Species' abundances are then related to the proportion of the stick acquired by each species. The development of such models has a long history, starting with Motomura's (1932, see Numata et al. 1950, for a biological interpretation) geometric series, and includes MacArthur's (1957, 1960) three niche apportionment models. Of these, the broken-stick model is of particular interest in that it is one of the only SAD models to be rejected by its developer. More recently, Tokeshi (e.g. 1993, 1999) proposed several niche partitioning models.	Motomura (1932, see Numata et al. 1950), MacArthur (1957, 1960), Sugihara (1980), Tokeshi (1993, 1999)
Zero-sum multinomial distribution (ZSM)	Population dynamics	The SAD predicted for the local community in Hubbell's (2001) spatially implicit neutral model. The ZSM has three parameters: the number of individuals in the local community (J), the probability that a dead individual is replaced by a randomly selected immigrant from the metacommunity (m), and the 'fundamental biodiversity number' (θ). The zero-sum assumption is simply that when an individual dies it is immediately replaced by another individual, i.e. resources are fully saturated at all times.	Hubbell (2001), Matthews & Whittaker (2014)

the geometric series (Table 2), incorporate a degree of randomness and assume that abundance will vary between communities structured through the same set of rules—thus allowing these models to better incorporate random variation such as may arise from sampling (Tokeshi 1993, Magurran 2004). The characterisation of a particular model as deterministic or stochastic has important implications when testing the fit of the model with real ecological data.

The easiest way to interpret SADs is to plot them. This produces graphical output which can be used to determine the form of the SAD, and as a basic means of evaluation of the goodness of fit (see Table 3). Two main methods are employed to plot SADs. First, a histogram (frequency distribution) of the species' abundances can be constructed (Figure 1a). The number of individuals (horizontal axis) can be left untransformed or may be logarithmically transformed before being 'binned'. The binning of data into abundance octaves is widespread in the SAD literature (e.g. Preston 1948, Volkov et al. 2003; see Figure 1a). Data can be binned in various ways; however, all result in the loss of information. Based on MacArthur (1957), Whittaker (e.g. 1965) developed an alternative plotting method (variously termed 'rank–abundance Plots,' 'Whittaker plots,' and 'dominance–diversity graphs'; see Figure 1b) to circumvent the issue of losing information when binning. Rank–abundance plots are constructed through plotting abundance (untransformed or log transformed) against rank order, where rank one corresponds to the species with the highest abundance, rank two corresponds to the species with the second highest abundance, and so on. As a method to determine the best SAD model for any given data set, rank–abundance plots are not perhaps as intuitive as frequency distributions, but they are beneficial in highlighting differences
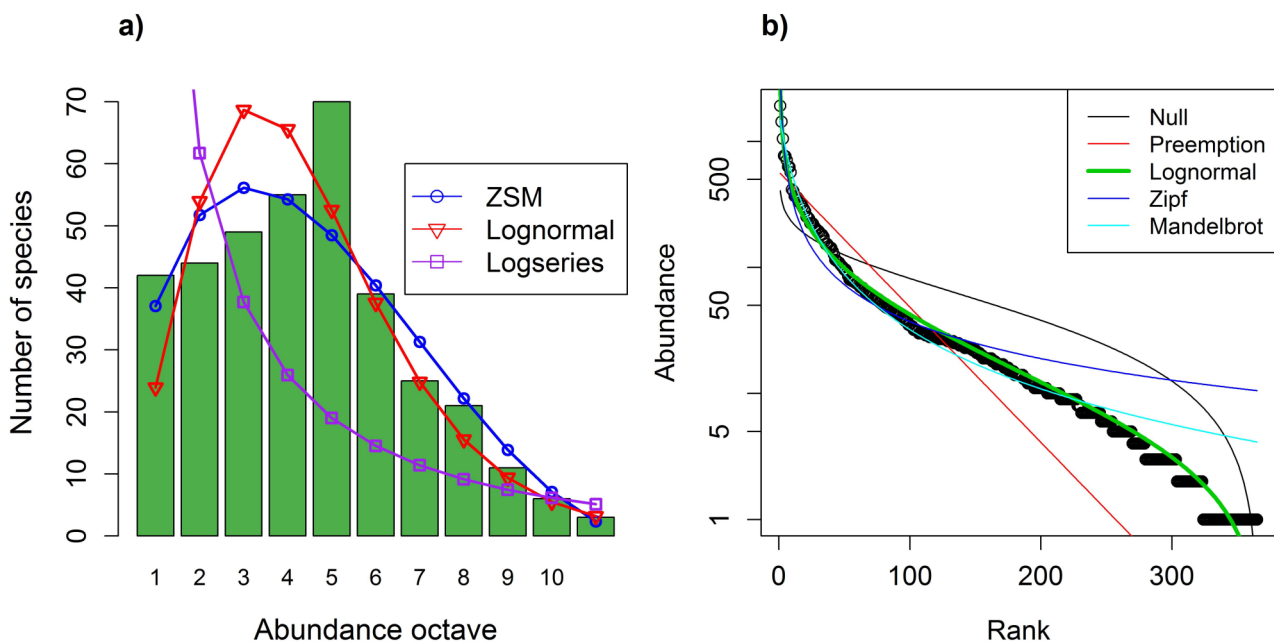

Figure 1. Exemplar SAD model fits illustrating the two major plotting methods: a) histograms and b) rank–abundance plots. In (a) the lognormal distribution (red line and triangles), logseries distribution (purple line and squares), and zero-sum multinomial distribution of Hubbell's (2001) spatially implicit neutral model (blue line and circles) have been fitted to binned simulated data (green bars, 365 species and 22945 individuals) derived through the random sampling of individuals from a simulated lognormal metacommunity, and fitted using maximum-likelihood methods. The predicted value of the logseries model for the first octave has been omitted for clarity. In (b) abundance is ordered in decreasing magnitude and plotted against the corresponding rank in this order. The broken-stick model ('Null', black line), geometric series ('Pre-emption', red line), log normal (green line), Zipf (dark blue line) and Mandelbrot (light blue line) distributions have been fitted to simulated data (open circles). Both plots were constructed in R: part (a) is based on a figure in Matthews & Whittaker (2014) and part (b) was constructed using the vegan R package (Oksanen et al. 2013).

Table 3. A summary of different methods (including some not otherwise discussed in the review) available for fitting, determining the goodness of fit, and comparing species abundance distribution models. For each method, a description and associated relevant references are given. Where appropriate, example tests or metrics are also given. Interested readers are directed to Magurran (2004) and Connolly & Dornelas (2011) for more information.

Method	Description	Example metrics	References
Fitting			
Least squares	Parameter values are selected which minimise the squared differences between the observed and predicted abundances	-	Gray et al. (2005)
Maximum likelihood (species abundance)	In maximum likelihood approaches the observed dataset is classed as a fixed observation, and a search is conducted for model parameter values that produce the most probable characterisation of the data, given the model	-	Matthews et al. (2014)
Maximum likelihood (rank–abundance)	As above, but for rank–abundance plots	-	Foster & Dunstan (2010)
Independent parameter estimation	Certain models (e.g. Hubbell's neutral model, Box 1) have parameters that can be estimated independently	-	Wootton (2005)
Goodness of fit			
Visual inspection	A subjective measure of fit in which the fit of models is determined by eye	-	Fisher et al. (1943)
Asymptotic goodness-of-fit tests	A set of tests which compare a given fit statistic with that of a null distribution	Pearson's χ^2	Matthews et al. (2014)
Linear regression	Only suitable for the geometric series and broken-stick model, because these have linear form in rank–abundance plots	R^2	Fattorini (2005)
Monte-Carlo test	Method to fit and evaluate a specific set of niche-partioning models. Preferred to Tokeshi's test, the leading alternative (see text)	-	Cassey & King (2001)
Parametric bootstrapping	A statistical resampling procedure whereby the properties of a metric are estimated through repeated sampling from an approximating distribution (see text)	$\hat{c}$	Connolly et al. (2009)
Model comparison			
Simple metrics and visual inspection	Comparison of models by eye, or by goodness-of-fit metrics and R^2 values	R^2, χ^2	Gray & Mirza (1979)
Akaike's information criterion	A measure of the relative quality of a model, given a set of data. It provides a relative estimate of lost information resulting from the use of a model as a representation of the process that produced the observed data. A version of the metric corrected for small sample size (AIC_c) is often used	AIC, AIC_c	Burnham & Anderson (2002)
Bayesian methods	A mathematical approach which incorporates prior beliefs/ insights and uses probability distributions to determine the level of uncertainty surrounding parameter values (see text)	Bayes factor	Etienne & Olff (2004)
Deviance information criterion	Derived in a Bayesian context, DIC is a model selection statistic based on deviance: the quality of fit of a given model. It is calculated using the log-likelihood of a given model, and comparing it to the log-likelihood of a hypothetical model providing a perfect fit	DIC	Mac Nally (2007)
Alternative approach			
Maximum entropy	Defines the most likely probability distribution as that which maximises entropy subject to particular information constraints (see text)	-	Harte et al. (2008)

in evenness between datasets, as well as not losing information in binning. Other SAD plotting methods exist, such as the empirical cumulative distribution function (e.g. McGill 2011), but are rarely used in practice.

There are three stages to fitting and evaluating SAD models: fitting the model and estimating the parameters, determining the goodness of fit of the model, and comparing the fit of the model with that of other models (Connolly & Dornelas 2011). Because SADs are often used to test different biogeographical theories, it is essential that rigorous statistical procedures are employed at each stage of model evaluation. A comprehensive review of the many different methods that have been employed at each of these three stages (see Table 3) is beyond the scope of this paper. Rather, we provide a synthesis of recent advances representing promising avenues for future SAD research. A number of general SAD reviews have been written (e.g. Pielou 1975, McGill et al. 2007), with a small number focusing specifically on methodology (e.g. Connolly & Dornelas 2011). However, our review differs from these in that we combine a general overview of SADs with a review of methodology. Furthermore, for each of the methodological advancements discussed, we provide recommendations and examples from our own experience of fitting and evaluating SAD models in biogeographical studies. Our review also differs from a number of others (e.g. Connolly & Dornelas 2011) in that we discuss themes such as maximum entropy and biological SAD models. Biological SAD models are particularly interesting from a biogeographical perspective, but are often overlooked.

The remainder of the paper is structured as follows. First, we summarise a number of issues in the fitting and evaluation of SAD models, using a simulated dataset in illustration, in order to contextualise recent methodological advances. We then evaluate novel advances in each of the aforementioned three stages of SAD model evaluation (fitting, assessing goodness of fit, comparing models). Finally, we focus on fitting and evaluating biological SAD models; these models are considered separately since they require different approaches to those used for statistical models.

The importance of considering methodology in the application of SADs

The choice of methods employed to fit and evaluate SAD models is an important consideration for any SAD study and can significantly affect the outcome of a particular analysis. For example, the selection of a particular SAD model as the 'best-fitting' model from a set of candidate models is often used as support for a particular theory, or as evidence for the importance of a particular biogeographical process. The early debate regarding Hubbell's (2001) spatially implicit neutral model was largely centred on the fit of the model's predicted SAD (the zero-sum multinomial, ZSM, Table 2) and acceptance of the model often came down to whether the ZSM provided a better fit to empirical data than other commonly used models (Box 1). However, this comparative approach has been criticised on a number of counts. This is largely due to the fact that: (a) the use of traditional goodness-of-fit tests (such as Pearson's χ^2, Table 3) generally reveals that most models fit the data well, (b) as such, small differences in factors such as sample size can result in changes to the model selected as 'best', (c) the model selected as 'best' often varies with the goodness-of-fit test used, (d) a number of commonly used tests are inappropriate (McGill 2003, Williamson & Gaston 2005), (e) multiple processes can produce the same SAD shape (McGill et al. 2007), and (f) the type of plotting method used (i.e. rank–abundance plots or histograms) can determine which model provides the best fit to a given dataset (Ulrich et al. 2010). With regard to the neutral model debate, the performance of the ZSM relative to other SAD models has been found to rely on the method used to fit the ZSM, the goodness-of-fit test used to compare models, and the plotting method used (e.g. Etienne & Olff 2004, Etienne 2005, Gray et al. 2006).

The use of the SAD in disturbance ecology provides an example of the importance of using appropriate methods. The SAD is often used to measure the impact of disturbance on community structure (e.g. Gray & Mirza 1979, Table 1). In this

context, the lognormal is theorised to characterise undisturbed communities (May 1975), while deviation from log-normality to more logseries-like shapes is expected in disturbed communities. Whether the lognormal provides a better fit than the logseries is thus an important part of this approach. However, different tests have been found to give conflicting results with regard to the fit of a given model (above, McGill 2003). To highlight this issue, we simulated an ecological community which followed a lognormal SAD (with left skew, as commonly observed in nature, 1000 individuals) and compared the fit of the logseries, Poisson lognormal (PLN) and truncated Poisson lognormal distributions to the data using a variety of metrics: χ^2 ($\log_2$ binning), χ^2 (arithmetic binning), AIC, a Kolmogorov–Smirnov test and visual inspection of the models plotted in rank–abundance form (Table 3). The PLN was the 'best-fitting' model according to AIC, χ^2 ($\log_2$ binning) and visual inspection, while the K–S test and χ^2 (arithmetic binning) identified the logseries as the 'best' model. Repeating this analysis for communities simulated to follow different SADs (see next paragraph) revealed that the five methods never universally selected a single model as the 'best'. Thus, the choice of metric can determine the choice of best-fitting model and it is essential that workers are aware of the most appropriate methods, several of which we outline below, in addition to recognising that method selection may influence any inferences drawn.

The methods used to fit and evaluate SAD models are also important considerations when the aim of model fitting is parameter estimation—for instance, if the parameters of a particular model are to be used in a comparative analysis. The gambin distribution (Table 2) has been used in such a way. Gambin has a single shape parameter (α) which can be used as an ecological metric: low α values (around 1) indicate logseries-like SAD shapes, while higher values (e.g. 4) indicate lognormal-like SAD shapes (Ugland et al. 2007). As initially presented, the gambin model (e.g. Ugland et al. 2007) estimated α through a simulation-based approach. However, subsequent work (Matthews et al. 2014) derived the distribution's

likelihood function and thus α can now be estimated using maximum-likelihood (ML) methods, which is preferable to simulations (below and Table 3; see Box 1 for a similar example involving neutral theory and fitting the ZSM). We estimated α using both methods for five simulated communities of approximately 1000 individuals: a lognormal community, and a logseries, a broken-stick and two lognormal communities with varying degrees of left skew. All communities were simulated in R (R Development Core Team 2013). When using the simulation methodology, the estimated α values were 4.04, 1.03, 11.8, 0.01 and 2.71 for the five communities, respectively. When α was estimated using ML, the corresponding values were 4.46, 1.46, 12.01, 0.84 and 3.13. While these differences may not appear large in absolute terms, the deviation is considerable. For instance, the difference of 0.83 between the two values of α for one of the skewed lognormal communities is similar to differences in α observed between a number of disturbed and undisturbed Azorean arthropod communities (Matthews et al. 2014).

In sum, the choice of method used to fit and compare SAD models is an important consideration when applying the SAD to biogeographical problems. As such, several authors have argued for an updated SAD research agenda: one which incorporates improved tests for determining the performance of individual models (McGill et al. 2007). As we emphasize herein, many of these improved methods already exist and it is possible to circumvent a number of the aforementioned issues; it is just a case of increasing the uptake of the available 'statistical machinery' (McGill 2003). For example, the use of information-theoretic approaches to compare SAD models is more rigorous than simply visually comparing the fit of models, or comparing metrics such as χ^2 (Table 3). In order to encourage the use of appropriate methodology in SAD research, a prescription of best practice is provided at the end of this paper.

Fitting statistical SAD models

Traditionally, SAD models have largely been fitted to empirical data using one of two approaches:

least-squares approaches and ML methods. In least-squares approaches, the aim is to select the parameter values that minimise the squared differences between the observed and predicted species abundance values. This is generally conducted using rank–abundance plots so that the observed and predicted abundances of species at each rank are compared. However, least-squares approaches to fitting SAD models are generally inappropriate because the statistical assumptions of least-squares methods, such as the statistical independence of data points, are rarely met with SAD data (Connolly & Dornelas 2011). As such, we advise that least-squares approaches to fitting SAD models are avoided where possible. Mainly because of the issue of violated assumptions, most SAD models are now fitted using ML methods (e.g. Matthews et al. 2014). In ML approaches the observed dataset is classed as a fixed observation and a search is conducted for model parameter values that produce the most probable characterisation of the data, given the model. As such, the likelihood (L) of observing the parameters, given the data, is:

$$L(T \mid X) = P(X \mid T),$$

where T is a vector of model parameters and X is the observed data (e.g. the abundance of each species in a sample). Generally, ML methods are used with species' abundance values, but they have also been applied to rank–abundance data (e.g. Foster & Dunstan 2010). One issue with ML methods is that they generally use optimisation algorithms to find the parameters that maximise the likelihood. While such algorithms greatly speed up the process, they can fail to provide the optimal values when multiple peaks exist in the likelihood profile; that is, the algorithms may converge on sub-optimal parameter values (Matthews & Whittaker 2014). One workaround for this is to start the optimisation algorithms from multiple starting points (e.g. Matthews et al. 2014), and we recommend ML methods as the preferred general method for fitting most SAD models.

The choice of fitting method is not simply statistical nit picking: it can result in very different parameter estimates for the same model. Our gambin analysis (above) provides an effective example of this. By way of a further example we simulated another skewed lognormal community (1000 individuals) and fitted a lognormal model to the data using four different methods: 1) log_2-transformed rank–abundance fitted using a least-squares approach; 2) log_2-transformed rank–abundance fitted using a ML approach; 3) the model was fitted to the untransformed data using ML; and 4) binned data fitted by a ML approach, using Preston's method of binning (Preston 1948). The mean and standard deviation parameters of the distribution estimated using the various methods were 1.54 $\pm$1.12, 0.85 $\pm$0.27, 1.53 $\pm$1.08 and 2.55 $\pm$1.8, respectively. Again, this shows that the choice of method for fitting SAD models (e.g. least squares or ML) can have important implications.

Goodness-of-fit tests

Most SAD studies have determined the goodness of fit of a particular SAD model through either visual inspection of SAD plots (see 'Introduction'), or the use of asymptotic goodness-of-fit statistics, such as Pearson's χ^2 (Table 3). Both methods have significant drawbacks. Visual inspection methods are inherently subjective, and there is no need to rely upon them given modern advances in computing power (which is not to say that visual inspection should be abandoned as part of a checking process). Asymptotic tests, while evidently less subjective than visual methods, face similar issues to those associated with the use of least-squares approaches to fitting SAD models: the various test assumptions are often violated when using species' abundance data (see Connolly & Dornelas 2011).

Parametric Bootstrapping

A goodness-of-fit technique which potentially circumvents some of the aforementioned issues involves parametric bootstrapping (Connolly et al. 2005, 2009), which estimates parameters through repeated sampling from an approximating distribution (Varian 2005). Thus, the method allows a measure of accuracy to be applied to sample estimates (Efron & Tibshirani 1998, Varian 2005). In the context of SADs, Connolly et al. (2009) use

parametric bootstrapping to estimate a goodness of fit statistic, termed $\hat{c}$, based on model deviance from the best model for the data (Connolly et al. 2009). Multiple values of $\hat{c}$ can be calculated to produce a bootstrap distribution, which can indicate whether the data agree with the assumptions of the model (Connolly et al. 2009). However, parametric bootstrapping is not limited to this one metric ($\hat{c}$), and can be used in conjunction with a variety of different test statistics (e.g. Volkov et al. 2003). Parametric bootstrapping is an informative and relatively simple method that should gain in popularity in SAD research. We recommend it for evaluating the fit of a SAD model.

Model comparison

A number of different methods have been used to compare SAD models (Table 3), but recent studies have largely focused on two overlapping approaches: information-theoretic model selection criteria (generally Akaike's information criterion, AIC) and Bayesian methods.

Information-theoretic model-selection criteria

Information-theoretic model-selection criteria are used to determine which of a set of candidate models most closely represents the 'true' model; in this way they allow researchers to choose between competing models. Their increase in prevalence is due to the realisation that, since models are only characterisations of nature, it may be more conducive to compare the fit of competing models, rather than simply trying to determine whether a model provides a statistically significant fit (i.e. goodness-of-fit tests; Burnham & Anderson 2002). In practice, the most effective protocol is to use both approaches in tandem. The last decade has seen a large increase in the use of information-theoretic approaches for comparing SAD models. For instance, the derivation of a full analytical solution to Hubbell's (2001) neutral zero-sum multinomial SAD model (ZSM, Etienne 2005) has enabled the use of information-theoretic approaches for comparing the ZSM with other SAD models, an advance which has resulted in more rigorous tests of the model (see Box 1).

Various information-theoretic model-selection statistics can be used to compare models. For instance, AIC (Table 3) is the most commonly used metric in recent SAD studies. The use of Akaike's information criterion corrected for small sample size (AIC_c) is often preferable as it is more applicable when sample size is small, and because AIC_c converges to AIC when the sample size is large (Burnham & Anderson 2002). Metrics such as AIC and AIC_c are advantageous because they include a penalty for the number of model parameters. A large number of parameters means a model will almost inevitably provide a good fit to the data, and thus a strong statistic for determining and comparing the fit of models should include a penalty for extra parameters (McGill 2003). In addition to calculating the metric values, it can be useful to normalize the model likelihoods such that they sum to 1, and to regard them as probabilities (Akaike weights). Akaike weights are interpreted as the probability that a given model is, in fact, the best model for the data (Burnham & Anderson 2002). One limitation with information-theoretic model-selection criteria is that they are only of use when comparing models: they provide a relative measure of goodness of fit rather than an absolute measure. For this reason information-theoretic metrics should be used in conjunction with some sort of goodness-of-fit metric. For example, if combined with likelihood-ratio tests (a likelihood-based test used to compare the fit of two nested models with varying numbers of parameters), the use of AIC has been labelled the most robust strength-of-fit test in macroecology (McGill 2003). A number of information-theoretic model-selection metrics have also been derived in a Bayesian context, and are discussed below. Information-theoretic approaches to model comparison have been discussed extensively elsewhere, and interested readers are directed to Burnham & Anderson (2002) as a useful starting point. Given their relative ease of calculation and the fact that they account for the number of parameters, we recommend that information-theoretic model-selection criteria should be used to compare SAD models in preference to metrics such as χ^2 and R^2.

A Bayesian approach

Several recent studies have used a Bayesian framework for parameter estimation (and thus Bayesian methods can also be listed under the 'model fitting' heading) and comparing SAD models (e.g. Etienne & Olff 2004, 2005). Bayesian analysis involves the use of probability distributions to determine the level of uncertainty surrounding parameter values. Prior knowledge of potential parameter values is incorporated into the process and forms the 'prior distribution' (priors). If no prior knowledge exists, non-informative priors (i.e. a prior distribution which has minimal impact on the posterior distribution) can be employed. The posterior distribution represents the 'result' in Bayesian analysis and comprises a probability distribution of parameter values, and thus can be seen as more informative than the single parameter value derived through more traditional methods (Bolker 2008). The posterior distribution is derived through:

$$P(A|B) = (P(B|A) \cdot P(A)) / c$$

where $P(A|B)$ is the posterior probability, $P(B|A)$ is equivalent to the model likelihood (the probability of observing B, given A), $P(A)$ is the prior distribution of parameter values, and c is a normalising constant which allows the posterior distribution to be a probability distribution. Numerous algorithms and statistical programs exist to compute Bayesian analyses, e.g. WinBugs. Model comparison can be conducted within a Bayesian framework using a criterion such as the posterior Bayes factor (e.g. Etienne & Olff 2005). The advantages of adopting a Bayesian approach are that the probability distribution of parameter values allows for greater inference, and the Bayes factor provides a competent, easy-to-interpret model comparison criterion which is largely parameter independent (Etienne & Olff 2005). Furthermore, a Bayesian view allows for decisions to be made on the plausibility of SAD model parameter estimates, a process which is an important component of model testing (Etienne et al. 2007). The downside is that most Bayesian approaches are analytically complex and they generally require a greater amount of expert knowledge than information-theoretic approaches. For example, many general and SAD-specific software packages fit particular SAD models and automatically return the log-likelihood value. However, this is not a reason to avoid Bayesian methods and they can be particularly useful in situations where the user is interested in deriving a probability distribution of the estimated values of a particular parameter, for example the immigration parameter of the ZSM (Box 1; e.g. Etienne & Olff 2004). Thus, depending on the aim of a study, we recommend the use of Bayesian methods in SAD studies in combination with other approaches, in order to provide additional insights regarding parameter inference and model comparison. However, when using Bayesian approaches it is important that users evaluate how their choice of prior distribution has influenced results.

In addition to the Bayes factor, the Bayesian information criterion (BIC) and deviance information criterion (DIC, Table 3) are model-selection metrics derived in a Bayesian context that enable comparison of competing SAD models. BIC is easier to calculate than the Bayes factor, and is superficially more similar to AIC. Often studies present model selection results using a number of different model selection metrics (e.g. Matthews et al. 2014). However, it is important to remember that AIC, BIC and DIC have different properties and assumptions (see Burnham & Anderson 2002), and workers should adopt a reasoned position on which is most appropriate for a given analysis.

Maximum Entropy

The maximum entropy principle has recently been applied to the field of SADs (e.g. Banavar et al. 2010, Harte 2011, White et al. 2012). We place it under a separate heading because it does not operate independently from a given SAD model and thus works differently from other approaches discussed. Maximum entropy is a tool developed in the physical sciences and works by defining the most likely probability distribution as that which maximises entropy, subject to particular information constraints; the process starts with what we already know about the system and then fills out what we don't know by maximising entropy

(Jaynes 2003, Harte et al. 2008, Frank 2011). The method has been applied in the field of statistical mechanics to investigate how spheres (e.g. atoms) are distributed into boxes (energy levels, Banavar et al. 2010). In SAD research, the spheres can represent individual organisms and the boxes can represent species. The theory behind the method is that multiple random factors that influence a particular pattern operate in different 'directions' and will ultimately cancel each other out in the aggregate (Frank 2009, 2011), leaving an entirely random pattern with the exception of the restricting constraints. Once the user has defined the constraints, the least biased distribution can be determined via a constrained optimisation procedure, using the method of Lagrange multipliers (Jaynes 2003). As such, it can be considered to minimise the information 'expressed in the final pattern' (Frank 2011). For instance, Frank (2009) provides the example of the binomial distribution: if all the available information is minimised to the knowledge that a given observation is a set number of Bernoulli trials with an average number of successes, the observation will approximate the binomial distribution.

Opinions on the use of maximum entropy in SAD research are divided. Studies have shown that logseries and lognormal distributions, as well as neutral-model SAD patterns (for summary of neutral SAD models see Box 1), can be generated using a maximum-entropy method (Banavar & Maritan 2007, Pueyo et al. 2007, Harte 2011, White et al. 2012). For example, using four constraints, including the average per-species metabolic flux and the average number of individuals per species, Harte (2011) has shown that the SAD predicted by maximum entropy follows a logseries distribution. However, a number of workers have questioned the applicability of maximum entropy in ecology (e.g. Haegeman & Loreau 2008, Haegeman & Etienne 2010, He 2010). Part of the problem with the use of maximum entropy in general is that, in order to obtain accurate results, the user must make a number of assumptions and select appropriate constraints; the user cannot identify missing constraints from the method itself. These choices are not trivial and have been shown to

affect the outcome of the process (Haegeman & Etienne 2010). Nonetheless, a recent application of the maximum entropy principle to over 15,000 ecological communities has shown that it is possible to characterise the majority of SAD structure when the only known information is the number of species and total number of individuals (White et al. 2012, see also Pueyo et al. 2007), thus largely circumventing the issue of constraint selection. The same study also proved that the parameter values of the logseries estimated by maximum entropy and by maximum-likelihood methods are equivalent (White et al. 2012; see their supplementary material). As such, we label maximum entropy as a method with considerable potential, but one that requires further research and evaluation. In particular, there is a need to determine the consequences for particular methods of species missed during sampling (He 2010), and a greater understanding is needed of how initial assumptions affect results (Haegeman & Etienne 2010).

Fitting and evaluating biological SAD models

Biological models are largely niche-oriented models in which it is assumed that the total niche space of the community, often termed 'the stick', is sequentially divided in some manner between the species that comprise this community (Marquet et al. 2003, Magurran 2004). The analogy of 'the stick' was introduced by MacArthur (e.g. 1957, 1960), who proposed a number of different biological SAD models, including the broken-stick model (Table 2). MacArthur argued that if the focus was ecologically similar species competing in a community, a structured pattern of relative abundance results. According to MacArthur's model, niche partitioning can be estimated as a random partition of a common limiting resource (breaking the stick). The relative abundances of the species are proportional to the total fraction of the resources each species appropriates. However, this is just one particular model (i.e. set of rules for linking abundance with niche space), and many others have been proposed (Table 2). Within all biological models, species that then immigrate into the community must apportion what-
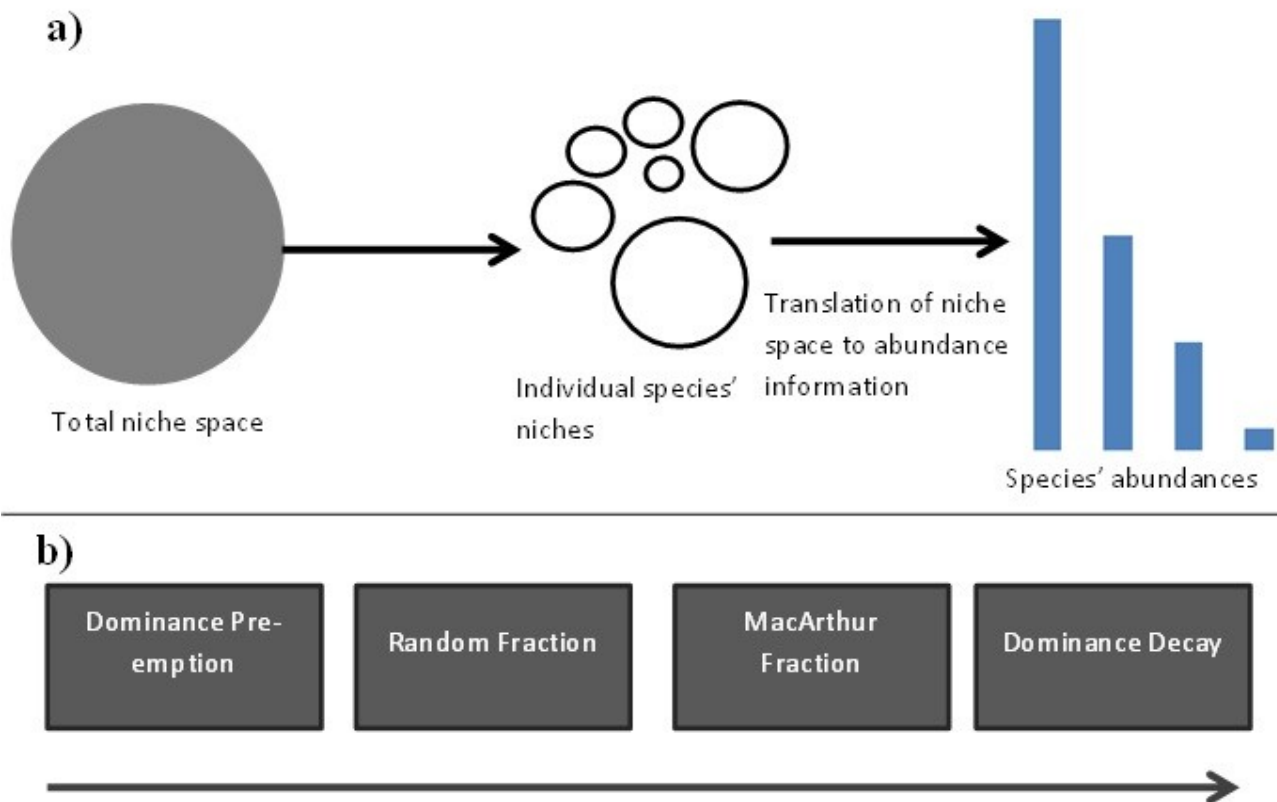
Figure 2. Niche-oriented species abundance distribution models, where (a) illustrates the conceptual basis of niche-oriented SAD models, and (b) sets out the spectrum of a selection of Tokeshi's sequential-breakage niche models. These are models in which it is assumed that the total niche space of the community, often termed 'the stick', is sequentially divided between the species that 'invade' the community according to a set of rules. Going from left to right the arrow represents the increase in evenness/uniformity of species' abundances in the community. The Dominance Pre-emption Model (Tokeshi 1990, 1993) predicts a fixed dominance hierarchy and can be seen as conceptually similar to the geometric series (Table 1). The first species, the superior competitor, pre-empts more than 50% of the total niche. The second species, the next best competitor, then apportions more than 50% of the remaining niche space, and so on. The inverse of this division is the Dominance-decay Model (Tokeshi 1990), whereby it is always the largest niche in the community that is divided; thus, the probability of a niche being apportioned is related to the niche size. The MacArthur Fraction Model (Tokeshi 1990, 1993) specifies a sequential process of niche subdivision in which total niche space is randomly divided in two, with one of these resulting niches further chosen (again, this is based on the idea that the niche space of the most abundant species is most likely to be invaded) and divided randomly. The Random fraction model is similar, except that it assumes that niches of varying size have the same probability of invasion. See Tokeshi (1990, 1993, 1999) for a full description of these and additional niche-oriented SAD models. Part (a) is adapted from Tokeshi (1993).

ever parts of the stick remain and break such sticks according to the rules of the particular model (see Figure 2a for an illustration of this process). In this way, biological models are mechanistic, providing ecological rationale to the way in which abundance is proportioned across species. This modus operandi assumes a one-to-one relationship between species' niches and species' abundances, a correspondence which is hard to test in practice, especially considering that a species has both a fundamental and (more restricted) realised niche (Hutchinson 1957, Sugi-hara 1980, Tokeshi 1993). Nonetheless, this assumption is backed up by certain studies (e.g. Marquet et al. 2003). Niche models are also largely zero-sum models in the sense that the number of individuals within the local community is generally fixed; an increase in the number of individuals of one species must equate to a decrease in the number of individuals in one or more other species. Several authors have advised removing the very rare species from a sample when calculating niche-based biological models (e.g. Tokeshi 1993, Magurran 2004, but see Fesl 2002).

This suggestion is valid in theory but requires careful consideration of what constitutes rarity at the level to be omitted.

Tests of niche models have produced mixed results (Tokeshi 1990, Fesl 2002, Mouillot & Wilson 2002). Fesl (2002) compared a number of niche-oriented models, including six of Tokeshi's (Table 2, Figure 2b), using data sourced from larval Chironomidae assemblages in the River Danube, Austria. Interestingly, Fesl's study found that none of the models provided a good fit to the data when all species were included. However, when the whole assemblage was deconstructed into functional subgroups it was found that the 'random fraction model' (RAM: one of Tokeshi's models in which species apportion available niche space in a random manner; Figure 2b) generally provided the best fit amongst functional groups. Strong evidence also existed for the 'random assortment model' (another of Tokeshi's models in which abundances vary independently of one another; see Tokeshi 1990, 1993). Based on the latter finding it was deduced that, in this system, species' abundances are largely independent of each other (Fesl 2002). The RAM was also found to provide the best fit to a fish parasite community, again indicating that species' abundances vary independently (Mouillot et al. 2003).

Part of the problem with determining the goodness of fit of such sequential niche apportionment models is that there is not an 'industry standard' goodness-of-fit test (Cassey & King 2001). Tests such as χ^2 are not valid in this context because niche models produce a set of distributions as opposed to a single fixed distribution (Pielou 1975, Bersier & Sugihara 1997). A simple approach is to use linear regression analysis to fit the geometric series and broken-stick models (Table 2) in rank–abundance plots, and as a means of comparing the goodness of fit of the two models (Fattorini 2005). This works because the geometric series has a linear form in rank–abundance plots in which the abundance data are log transformed, and the broken-stick model a linear form when the ranks are log transformed. However, this approach is limited to these two distributions. A Bayesian approach has also been used to compare the broken-stick model with various statistical models (Etienne & Olff 2005). Tokeshi (1990) developed a specific goodness-of-fit test to circumvent the issue of niche models producing sets of distributions, in which the mean of each rank of the theoretical distribution proposed by the different models is compared with the observed abundances of each rank within set confidence limits. However, as Bersier & Sugihara (1997) attest, Tokeshi's method, while intuitively appealing, has several drawbacks, most problematic being that it is sensitive to species richness. This problem has been addressed by a computationally intensive Monte Carlo simulation method proposed by Bersier & Sugihara (1997), and improved by Cassey & King (2001) and Mouillot et al. (2003). Given the aforementioned problems with Tokeshi's goodness-of-fit test, the Monte Carlo simulation method (cf. Cassey & King 2001) is preferable when fitting and evaluating niche SAD models. While often receiving less focus than methods to fit statistical SAD models, the development and advancement of tests that allow for the evaluation of biological models, and the comparison of niche-oriented models with statistical models, are key areas for future research on biological SAD models.

Conclusions and a prescription for best practice

The SAD is a key component of much biogeographical and ecological theory (McGill et al. 2007, Whittaker & Fernández-Palacios 2007) and while it has not received as much coverage as other macroecological patterns, such as the SAR, SAD research has been at the forefront of methodological progress in ecological biogeography. From the increases in computer power over the last two decades, coupled with recent advances in statistical theory, a number of novel methodological opportunities have arisen with the potential to significantly progress the field of SAD research. As a result of these methodological advancements, we feel it is time for SAD studies to move away from many of the traditional methods available for fitting and evaluating models, such as sole reliance on the visual examination of plots, and embrace many of these new approaches. Statistically

rigorous techniques should be applied at each of the three stages of SAD model evaluation and, in particular, we recommend the use of both goodness-of-fit tests and model-comparison analyses because each provides unique information which one can use to draw inferences. We are confident that the application and development of approaches synthesised in this review, in addition to others (Table 3), will result in the continued advancement of a vibrant and exciting field within ecological biogeography and macroecology. To encourage best practice, we conclude with a set of guidelines for use when comparing the fit of a set of SAD models to empirical data, while recognising that alternative approaches exist and others may be developed in the future:

1) Select a sensible set of candidate models for comparison, based on the aim(s) of the study. The selection of a priori (i.e. before step #2) models is an important first step that requires careful thought (Burnham & Anderson 2002).

2) Plot the SAD to visually determine whether the models may poorly fit any parts of the empirical distribution.

3) Where possible ML methods should be employed to fit the models and estimate the parameters. For complex models, optimisation algorithms should begin from multiple starting values.

4) Combine parametric bootstrapping with a simulation approach to generate a distribution of, for example, the $\hat{c}$ goodness-of-fit statistic (cf. Connolly et al. 2009) for each a priori model.

5) Compare models statistically, not simply by visual examination of plots. Information-theoretic approaches are both appropriate and easy to use. If the assumptions are met, multiple information criteria (e.g. AIC_c , BIC) should be compared to ensure robustness of outcomes. It is important to test the goodness of fit of models in addition to comparing models in this way because, for example, the use of AIC assumes that a model is an accurate representation of the data.

6) For biological SAD models, linear regression can be used to analyse the fit of the broken-

stick model and geometric series in rank–abundance plots, while Cassey & King's (2001) Monte Carlo simulation method is preferable to Tokeshi's (1990) test for other biological models.

Acknowledgements

References

Banavar, J.R. & Maritan, A. (2007) The maximum relative entropy principle. [online] URL: http://arxiv.org/abs/cond-mat/0703622 (accessed 03/04/2013)

Banavar, J.R., Maritan, A. & Volkov, I. (2010) Applications of the principle of maximum entropy: from physics to ecology. Journal of Physics: Condensed Matter, 22, 063101.

Bazzaz, F.A. (1975) Plant species diversity in old-field successional ecosystems in southern Illinois. Ecology, 56, 485–488.

Bersier, L.-F. & Sugihara, G. (1997) Species abundance patterns: the problem of testing stochastic models. Journal of Animal Ecology, 66, 769–774.

Bolker, B.M. (2008) Ecological models and data in R. Princeton University Press, New Jersey.

Bulmer, M.G. (1974) On fitting the Poisson lognormal distribution to species-abundance data. Biometrics, 30, 101–110.

Burnham, K.P. & Anderson, D.R. (2002) Model selection and multi-model inference: a practical information-theoretic approach, 2nd edn. Springer, New-York.

Cassey, P. & King, R.A.R. (2001) The problem of testing the goodness-of-fit of stochastic resource apportionment models. Environmetrics, 12, 691–698.

Connolly, S.R. & Dornelas, M. (2011) Fitting and empirical evaluation of models for species abundance distributions. In: Biological diversity: frontiers in measurement and assessment (ed. by A.E. Magurran and B.J. McGill), pp. 123–140. Oxford University Press, Oxford.

Connolly, S.R., Hughes, T.P., Bellwood, D.R. & Karlson, R.H. (2005) Community structure of corals and reef fishes at multiple scales. Science, 309, 1363–1365.

Connolly, S.R., Dornelas, M., Bellwood, D.R. & Hughes, T.P. (2009) Testing species abundance models: a new bootstrap approach applied to Indo-Pacific coral reefs. Ecology, 90, 3138–3149.

Dewdney, A.K. (1998) A general theory of the sampling process with applications to the "veil line". Theoretical Population Biology, 54, 294–302.

Dunstan, P.K., Bax, N.J., Foster, S.D., Williams, A. & Althaus, F.

(2012) Identifying hotspots for biodiversity management using rank–abundance distributions. Diversity and Distributions, 18, 22–32.

Efron, B. & Tibshirani, R.J. (1993) An introduction to the bootstrap. Chapman & Hall, New York.

Etienne, R.S. (2005) A new sampling formula for neutral biodiversity. Ecology Letters, 8, 253–260.

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., Apol, M.E.F., Olff, H. & Weissing, F.J. (2007) Modes of speciation and the neutral theory of biodiversity. Oikos, 116, 241–258.

Fattorini, S. (2005) A simple method to fit geometric series and broken stick models in community ecology and island biogeography. Acta Oecologica, 28, 199–205.

Fesl, C. (2002) Niche-oriented species-abundance models: different approaches of their application to larval Chironomid (Diptera) assemblages in a large river. Journal of Animal Ecology, 71, 1085–1094.

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. Journal of Animal Ecology, 12, 42–58.

Foster, S.D. & Dunstan, P.K. (2010) The analysis of biodiversity using rank–abundance distributions. Biometrics, 66, 186–195.

Frank, S.A. (2009) The common patterns of nature. Journal of Evolutionary Biology, 22, 1563–1585.

Frank, S.A. (2011) Measurement scale in maximum entropy models of species abundance. Journal of Evolutionary Biology, 24, 485–496.

Gray, J.S. & Mirza, F.B. (1979) A possible method for the detection of pollution-induced disturbance on marine benthic communities. Marine Pollution Bulletin, 10, 142–146.

Gray, J.S., Bjørgesæter, A. & Ugland, K.I. (2005) The impact of rare species on natural assemblages. Journal of Animal Ecology, 74, 1131–1139.

Gray, J.S., Bjørgesæter, A. & Ugland, K.I. (2006) On plotting species abundance distributions. Journal of Animal Ecology, 75, 752–756.

Haegeman, B. & Etienne, R.S. (2010) Entropy maximization and the spatial distribution of species. The American Naturalist, 175, E74–E90.

Haegeman, B. & Loreau, M. (2008) Limitations of entropy maximization in ecology. Oikos, 117, 1700–1710.

Harte, J. (2011) Maximum entropy and ecology: a theory of abundance, distribution, and energetics. Oxford University Press.

Harte, J., Zillio, T., Conlisk, E. & Smith, A.B. (2008) Maximum entropy and the state-variable approach to macroecology. Ecology, 89, 2700–2711.

He, F. (2010) Maximum entropy, logistic regression, and species abundance. Oikos, 119, 578–582.

Hubbell, S.P. (2001) The unified neutral theory of biodiversity and biogeography. Princeton University Press, Princeton.

Hutchinson, G.E. (1953) The concept of pattern in ecology. Proceedings of the Academy of Natural Sciences of Philadelphia, 105, 1–12.

Jaynes, E.T. (2003) Probability theory: the logic of science. Cambridge University Press, New York.

MacArthur, R.H. (1957) On the relative abundance of bird species. Proceedings of the National Academy of Sciences of the USA, 43, 293–295.

MacArthur, R.H. (1960) On the relative abundance of species. The American Naturalist, 94, 25–36.

Mac Nally, R. (2007) Use of the abundance spectrum and relative abundance distributions to analyze assemblage change in massively altered landscapes. The American Naturalist, 170, 319–330.

Magurran, A.E. (2004) Measuring biological diversity. 2nd edn, Blackwell, Oxford.

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

Marquet, P.A., Keymer, J.A. & Cofre, H. (2003) Breaking the stick in space: of niche models, metacommunities and patterns in the relative abundance of species. In: Macroecology: concepts and consequences (ed. by T.M. Blackburn and K.J. Gaston), pp. 64–86. Cambridge University Press, Oxford.

Matthews, T.J. & Whittaker, R.J. (2014) Neutral theory and the species abundance distribution: recent developments and prospects for unifying niche and neutral perspectives. Ecology and Evolution, doi: 10.1002/ece3.1092

Matthews, T.J., Borges, P.A.V. & Whittaker, R.J. (2014) Multimodal species abundance distributions: a deconstruction approach reveals the processes behind the pattern. Oikos, 123, 533–544.

Matthews, T.J., Borregaard, M.K., Ugland, K.I., Borges, P.A.V., Rigal, F., Cardoso, P. & Whittaker, R.J. (2014) The gambin model provides a superior fit to species abundance distributions with a single free parameter: evidence, implementation and interpretation. Ecography, doi: 10.1111/ecog.00861

May, R.M. (1975) Patterns of species abundance and diversity. In: Ecology and evolution of communities (ed. by M.L. Cody and J.M. Diamond), pp. 81–120. Harvard University Press, Cambridge, MA.

McGill, B. (2003) Strong and weak tests of macroecological theory. Oikos, 102, 679–685.

McGill, B.J. (2011) Species abundance distributions. In: Biological diversity: frontiers in measurement and assessment (ed. by A.E. Magurran and B.J. McGill), pp. 105–122. Oxford University Press, Oxford.

McGill, B.J. et al. (2007) Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. Ecology Letters, 10, 995–1015.

Motomura, I. (1932) On the statistical treatment of communities. Zoological Magazine Tokyo, 44, 379–383.

Mouillot, D. & Wilson, J.B. (2002) Can we tell how a community was constructed? A comparison of five evenness indices for their ability to identify theoretical models of community construction. Theoretical Population Biology, 61, 141–151.

Mouillot, D., George-Nascimento, M. & Poulin, R. (2003) How parasites divide resources: a test of the niche appor-

tionment hypothesis. Journal of Animal Ecology, 72, 757–764.

Numata, M. (1950) The homogeneity of plant communities – studies on the structure of plant communities VI. (In Japanese). Botanical Magazine (Tokyo), 63, 747–748.

Oksanen, J.F., Blanchet, G., Kindt, R., Legendre, P., Minchin, P.R., O'Hara, R.B. & Simpson, G.L. (2013) Package 'vegan'. R package version 2.0-9.

Pielou, E.C. (1975) Ecological diversity. Wiley, New York.

Preston, F.W. (1948) The commonness, and rarity, of species. Ecology, 29, 254–283.

Pueyo, S., He, F. & Zillio, T. (2007) The maximum entropy formalism and the idiosyncratic theory of biodiversity. Ecology Letters, 10, 1017–1028.

R Development Core Team (2013) R: A language and environment for statistical computing. R Foundation for Statistical Computing.

Rosindell, J., Wong, Y. & Etienne, R.S. (2008) A coalescence approach to spatial neutral ecology. Ecological Informatics, 3, 259–271.

Sugihara, G. (1980) Minimal community structure: an explanation of species abundance patterns. The American Naturalist, 116, 770–787.

Tokeshi, M. (1990) Niche apportionment or random assortment: species abundance patterns revisited. Journal of Animal Ecology, 59, 1129–1146.

Tokeshi, M. (1993) Species abundance patterns and community structure. Advances in Ecological Research, 24, 112–186.

Tokeshi, M. (1999) Species coexistence: ecological and evolutionary perspectives. Blackwell Science, Oxford.

Ugland, K.I., Lambshead, P.J.D., McGill, B., Gray, J.S., O'Dea, N., Ladle, R.J. & Whittaker, R.J. (2007) Modelling dimensionality in species abundance distributions: description and evaluation of the Gambin model. Evolutionary Ecology Research, 9, 313–324.

Ulrich, W., Ollik, M. & Ugland, K.I. (2010) A meta-analysis of species–abundance distributions. Oikos, 119, 1149–1155.

Varian, H. (2005) Bootstrap tutorial. Mathematica Journal, 9, 768–775.

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

White, E.P., Thibault, K.M. & Xiao, X. (2012) Characterizing species abundance distributions across taxa and ecosystems using a simple maximum entropy model. Ecology, 93, 1772–1778.

Whittaker, R.H. (1965) Dominance and diversity in land plant communities. Science, 147, 250–260.

Whittaker, R.J. & Fernández-Palacios, J.M. (2007) Island biogeography: ecology, evolution, and conservation, 2nd edn. Oxford University Press, Oxford.

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

Wootton, J.T. (2005) Field parameterization and experimental test of the neutral theory of biodiversity. Nature, 433, 309–312.

Zillio, T. & He, F. (2010) Inferring species abundance distribution across spatial scales. Oikos, 119, 71–80.