



Species abundance distributions, statistical mechanics and the priors of MaxEnt



M.G. Bowler

Department of Physics, University of Oxford, Keble Road, Oxford OX1 3RH, UK

ARTICLE INFO

Article history:

Received 14 May 2013

Available online 19 December 2013

Keywords:

Species abundance
Statistical mechanics
Maxent
Priors

ABSTRACT

The methods of Maximum Entropy have been deployed for some years to address the problem of species abundance distributions. In this approach, it is important to identify the correct weighting factors, or *priors*, to be applied before maximising the entropy function subject to constraints. The forms of such priors depend not only on the exact problem but can also depend on the way it is set up; priors are determined by the underlying dynamics of the complex system under consideration. The problem is one of statistical mechanics and it is the properties of the system that yield the correct MaxEnt *priors*, appropriate to the way the problem is framed. Here I calculate, in several different ways, the species abundance distribution resulting when individuals in a community are born and die independently. In the usual formulation the prior distribution for the number of species over the number of individuals is $1/n$; the problem can be reformulated in terms of the distribution of individuals over species classes, with a uniform prior. Results are obtained using master equations for the dynamics and separately through the combinatoric methods of elementary statistical mechanics; the MaxEnt *priors* then emerge *a posteriori*. The first object is to establish the log series species abundance distribution as the outcome of *per capita* guild dynamics. The second is to clarify the true nature and origin of *priors* in the language of MaxEnt. Finally, I consider how it may come about that the distribution is similar to log series in the event that filled niches dominate species abundance. For the general ecologist, there are two messages. First, that species abundance distributions are determined largely by population sorting through fractional processes (resulting in the $1/n$ factor) and secondly that useful information is likely to be found only in departures from the log series. For the MaxEnt practitioner, the message is that the *prior* with respect to which the entropy is to be maximised is determined by the nature of the problem and the way in which it is formulated.

© 2013 Elsevier Inc. All rights reserved.

1. Introduction

A community of individuals from many species, sufficiently similar so as to constitute a guild, can contain thousands of individuals and perhaps hundreds of species. Such a community is a dynamical entity; individuals die and give birth and species can go extinct (at least locally), immigrate from a wider world or even (rarely) emerge by speciation. These properties positively invite the application of the techniques of statistical mechanics with the object of understanding emergent properties of the complex systems. The lowest level emergent property is the distribution of a specified number of individuals over the separate species in the guild—the species abundance distribution. The general characteristics of species abundance distributions are that there are few species represented by a large number of individuals and many species represented by a small number of individuals. The distributions are neither exponential nor any power law, rather they

can mostly be categorised by the number of species represented by n individuals, s_n , being distributed in n somewhere between a log series and a log normal distribution. This property has been explained within statistical mechanics (known to ecologists as MaxEnt, but they are not quite the same) rather successfully (Pueyo, 2006; Banavar et al., 2010; Bowler and Kelly, 2010, 2012) as a result of *per capita* birth and death rates, but this is not the only approach that generates this description. (Hubbell's (2001) neutral theory is not unrelated, in that birth and death rates are there *per capita* and that a guild contains a fixed number of individuals.) The alternative can loosely be described as resource partitioning, or stick breaking (MacArthur, 1960; Sugihara, 1980; Tokeshi, 1990, 1996), to which the methods of statistical mechanics can also be applied. The two approaches are not necessarily exclusive and this paper is concerned with both—crudely labelled as neutrality versus niches. There is then the question of what, if anything, we can learn from species abundance distributions (McGill et al., 2007) if very general principles underlie their form (my answer is 'very little'; the measure is too coarse grained).

E-mail address: m.bowler1@physics.ox.ac.uk.

I start by inviting the reader to travel back in time half a century, to the paper of MacArthur (1960) and his growth rate equation

$$r_i(t) = \frac{1}{N_i(t)} \frac{dN_i(t)}{dt}$$

formally integrated to Eq. (1) of that paper:

$$\ln N_i(t) = \ln N_i(0) + \int_0^t r_i(t) dt. \quad (\text{MacArthur equation (1)})$$

In the above equations the abundance of species i is N_i and r_i is the *per capita* time dependent growth rate of the abundance. The formal solution, written in this way, suggested two limits of potential ecological interest. First, if the integral is negligible, the value of the abundance stays at its initial value. This could happen if the dynamics were such that for every individual that dies, a new individual of the same species replaces it. The ecological interpretation is that each species inhabits and always fills its own niche. For a natural community, initial values $N_i(0)$ have nonetheless to be established by natural processes. Secondly, the integral might be such that over a long period the solution loses all memory of the initial conditions. For my purposes the suggestion that the second term on the right hand side of MacArthur equation (1) dominates and that the *per capita* (or fractional) growth rate undergoes a random walk with each time step is the most interesting. The reason is that if the r undergo random walks (due to environmental fluctuations, the sense of MacArthur, 1960) the integral on the right hand side of MacArthur equation (1) is normally distributed and the abundance is consequently log normally distributed. (Margalef, 1958 noted this and remarked that it might be relevant to Preston's (1948) observations of approximately log normal species abundance distributions; he did not encourage his readers to take this too seriously.) A log normal distribution only results if it is the *per capita* growth rates (as opposed to the rates $dN_i(t)/dt$) that execute a random walk; this seems natural for species where individuals are roughly identical and die or give rise to a new individual independently. More generally, I argue that the dominant feature of species abundance distributions depends on dynamics in which population changes are fractional—this is the key, but independent *per capita* birth and death rates are not the only possibility.

MacArthur suggested that log normal distributions would be attributable to opportunistic species and that the first term on the right hand side of MacArthur equation (1) would apply to a community in some static equilibrium. However, equilibrium can be dynamic and complex dynamical systems can forget initial conditions. It might also be that the two pieces on the right hand side should be thought of as results of processes operating on different time scales. The second term would operate on a time scale characteristic of the birth and death rates of individuals and the initial conditions appropriate to MacArthur equation (1) might have been set in place by slower processes. The first term might be associated with some fairly well defined niche and that niche could have a capacity that also varies with time, but on a much longer timescale.

In the next section I apply statistical mechanics to the case where the species abundance distribution is dominated by the effects of the birth and death of individuals. For this to be true, either there is no niche structure of any significance or the effects of niches must average out in the species abundance distribution (Pueyo et al., 2007; Bowler and Kelly, 2010, 2012). This suggests that niches must be flexible, changing perhaps on a time scale similar to that of the individuals that inhabit them. If birth and death rates are *per capita* (cf. MacArthur equation (1) above), statistical mechanics yields an underlying structure for species abundance distributions that is log series. This was demonstrated in Bowler and Kelly (2010, 2012) (see also Kendall, 1948a,b, Volkov

et al., 2003, Pueyo, 2006, Banavar et al., 2010); here I obtain those same results by a variety of different methods. I hope that this will remove any lingering doubts about statistical mechanics yielding, in this context, log series, and also prove enlightening on the somewhat vexed question of *priors*. This investigation of how to construct the correct function for maximisation also justifies a fundamental assumption in Harte's MaxEnt Theory of Ecology (Harte, 2011). In Section 3, I consider how rigid niches might be treated and what could go wrong with the birth and death rate argument. The biological principle underlying species abundance distributions, that they are largely determined by fractional sorting processes, is discussed in the final section.

2. Dynamics dominated by individual birth and death

2.1. Statistical mechanics and priors

In this section I apply the methods of statistical physics to establish, in a number of different ways, that the species abundance distributions of ecological guilds in which the dynamics is dominated by *per capita* birth and death rates are to a first approximation log series. It is otherwise technical and as such unlikely to be of much interest to the general ecologist. Section 2.4 through 2.6 are, as far as I know, new applications of the ideas of the microcanonical ensemble. Nonetheless it is intended to be both pedagogical and polemical, being concerned with the correct ways to approach the problem of relative species abundance through the techniques of statistical mechanics. That part of the physical sciences deals with very real entities (like atoms) and constraints (like conservation of energy) and while the techniques have much in common with those of maximum entropy (MaxEnt) as used in information theory, statistical mechanics is not just MaxEnt. It seems unfortunate that ecologists have mostly encountered these techniques through probability, Bayesian analysis and information theory, because the correct formulation is not always recognisable as maximisation of the Shannon entropy (see 2.6). In addition, the role of the (information theoretic) *prior* can all too easily seem arbitrary or mysterious—much effort has been dissipated on the theological debate concerning the true nature of the *maximally uninformative prior*. In statistical mechanics, analogues of priors have always a straightforward interpretation and are determined by physics or, here, biology. A complex system (be it a box full of gas or a forest full of trees) has a large number of accessible states which have the same gross properties; the more microstates available the more probable it is to find those gross properties. The correct way to calculate the most probable configuration depends on the properties of the complex system and that dependence translates into the appropriate *prior* if the problem is phrased in terms of MaxEnt.

In this section the underlying biological principle is taken to be that species move from one abundance class to another as a result of individuals dying and giving rise to new mature individuals independently (*per capita* birth and death rates). I formulate the problem in two ways; first by considering the most probable arrangement of species in abundance classes directly (as the problem is usually formulated using MaxEnt) and secondly by considering the distribution of the number of individuals over classes defined by all species with n individuals. In the first case a non-uniform *prior* is required; in the second the prior is uniform, yet the problem is the same. Both approaches are addressed through dynamical equations (the master equation approach of Volkov et al., 2003, 2005) and separately through the combinatoric methods of elementary statistical mechanics. I explicitly assume that, for calculating species abundance distributions, species in a guild can be treated on average as having the same properties; differences average out. The ecological significance of this extremely important aspect has been discussed at length in Pueyo et al. (2007) and Bowler and Kelly (2012).

2.2. Distributing species over population classes

We suppose that we have a set of species such that there are s_n species each containing n individuals. The community has a fixed number of species S and a fixed number of individuals N , thus providing constraints to which the s_n are subject as they fluctuate in response to the birth and death of individuals:

$$\sum_n s_n = S \quad (1a)$$

$$\sum_n ns_n = N. \quad (1b)$$

In the usual approach it is supposed that before applying these constraints there exists an *a priori* probability for each species to have a population n , $P_\pi(n)$. The probability of finding n_1 individuals in species 1, n_2 in species 2 and so on is then proportional to the continued product $\prod P_\pi(n_i)$, every species being counted individually. If instead of counting species individually we count the number in the class with population n , a particular choice of s_n species with population n contributes a factor $[P_\pi(n)]^{s_n}$. There are many ways of picking a specified set of all species $\{s_n\}$ from a grand total S and the total number of ways is calculated simply in combinatorics (the multinomial coefficient). The weight of a specified set $\{s_n\}$ is thus given by

$$W = \frac{S! \prod [P_\pi(n)]^{s_n}}{\prod s_n!}. \quad (2)$$

In (2) above, the continued product on the bottom line is of the numbers $s_n!$ over all populations n and similarly the continued product on the top line is of the numbers $[P_\pi(n)]^{s_n}$.

The most probable configuration is found by maximising $W(s_n)$ with respect to all s_n . It is convenient to maximise the logarithm of W and the constraints (1) are incorporated using the method of undetermined multipliers. Thus the function of each s_n to be maximised is

$$-\ln(s_n!) + s_n \ln P_\pi(n) + \mu s_n + \lambda ns_n \quad (3a)$$

where the constraints are introduced through the multipliers λ and μ .

On applying the familiar Stirling approximation this becomes

$$-s_n \ln s_n + s_n + s_n \ln P_\pi(n) + \mu s_n + \lambda ns_n. \quad (3b)$$

Eq. (2) was composed using combinatorics and the underlying supposition is that the most probable complexion is that with the greatest number of equivalent states, with the possibility of a weight factor. If we rephrase (3b) in terms of the probability of encountering s_n , $P(n) = s_n/S$, the equivalent function to be maximised with respect to all $P(n)$ is

$$-P(n) \ln P(n) + P(n) + P(n) \ln P_\pi(n) + \mu P(n) + \lambda nP(n).$$

The sum over all n of the first two terms is the information theory entropy (Jaynes, 2003, p. 350 ff) to be maximised relative to $P_\pi(n)$ (third term) subject to constraints (the last two terms). Hence MaxEnt for Maximising (information theory) Entropy. These two formulations are identical; compare Pueyo et al. (2007). However, starting from combinatorics as with Eq. (2) it is clear that the weight factor $P_\pi(n)$ must depend on the problem in question and be determined by physical or biological considerations. In approaching the problem from information theory, the information theoretic prior $P_\pi(n)$ comes in through Bayes' theorem and drags with it associations likely to be misleading.

Differentiating (3b) with respect to s_n and equating to zero, the solution is

$$s_n(n) = P_\pi(n) \exp(\mu + \lambda n). \quad (4)$$

The multipliers μ and λ are determined from the constraints (1); λ must be negative for an exponential distribution decaying

with n . If there are no constraints applied, then μ and λ are zero and then $s_n(n) = P_\pi(n)$; the solution prior to application of constraints. (Maximising (3) is equivalent to maximising the entropy in information theory relative to a prior P_π .)

In real problems of the distribution of real objects over some quantity of interest, the weighting factor here denoted by P_π is determined by the nature of the problem and on how it is framed. In the statistical mechanics of gases, it is desired to calculate the distribution of atoms over accessible energy states, subject to the constraints of conservation of number and of conservation of energy. The dynamics of atomic collisions are such that all states with a unique set of quantum numbers are equally weighted *a priori*. Atoms are distributed over these unique states according to energy and the prior probability is uniform. However, if instead of asking how many atoms are to be found in a unique quantum state of energy E we were to ask how many atoms have energy E , the weight (or prior) would be the number of unique quantum states having that single energy E . In application of these notions to the problem of relative species abundance, the question is how the biological nature of the problem biases the number of species with occupation number n prior to the application of constraints on the total number of species and the total number of individuals in the guild. The biological processes of birth and death can affect profoundly the chance of finding a particular species in the class labelled by n , if birth and death rates are (at least approximately) *per capita*. Thus a species with at one moment n individuals will move into the class with $n - 1$ individuals at a rate proportional to n ; in any small interval of time all individuals have the same chance of dying and so the more there are the faster a species loses a single one. The same argument applies to individuals giving birth. Simply considering that birth and death rates are *per capita* suggests that for this problem of species abundance the appropriate weighting factor (or prior) is $P_\pi(n) = 1/n$, but verbal arguments cannot be trusted until the numbers work out. One way is through the master equation of Volkov et al. (2003, 2005), in which birth and death rates appear explicitly.

2.3. The master equation for species abundance

The gain or loss of one individual in one species in the class of n individuals in each species reduces the number of species in that class, s_n , by one. In the class there are ns_n individuals each with birth and death rates b, d . The rate at which s_n is depleted is therefore $(b + d)ns_n$. The number of species s_n is augmented by a death in class $n + 1$ or by a birth in class $n - 1$. The rate of change of s_n is therefore given by

$$\frac{ds_n}{dt} = -(b + d)ns_n + b(n - 1)s_{n-1} + d(n + 1)s_{n+1}. \quad (5)$$

For the equilibrium solution of this set of equations the right hand side of (5) is zero and it follows that

$$s_{n+1} = \frac{b}{d} \frac{n}{n + 1} s_n;$$

iterating this equation we have

$$s_n = s_0 \frac{1}{n} \left(\frac{b}{d} \right)^n$$

where s_0 is a normalisation constant. The second term on the right is an exponential in n ; thus the master equation with *per capita* birth and death rates gives the solution (4) with $P_\pi(n) = 1/n$. This is the log series distribution and it is a good approximation to species abundance distributions found in nature, departures occurring at the rare species end. Thus the necessary *prior* in a MaxEnt treatment of species abundance distributions can be understood in terms of the biological processes of birth and death.

So far the discussion has differed little from that to be found in Bowler and Kelly (2012); the material above is included for comparison with a different approach, in which individuals are distributed over species classes.

2.4. Distributing individuals over species classes

Rather than distribute numbers of species s_n over classes of occupancy, suppose instead we distribute individuals over classes labelled by n , but such that all individuals that are members of the species in class n are members of the new classes. The numbers of individuals in the new classes are $H_n = ns_n$. If the distribution of s_n is log series, the distribution of H_n is exponential. If the prior distribution for s_n is $1/n$, then the prior for H_n is uniform. The details of a complete treatment are less obvious but are illuminating; we start by working out a master equation for the quantities H_n , the total number of individuals that are members of a species with population n . Each individual can die or give rise to a new individual. Either event removes its species from the class n and thereby reduces H_n by n individuals. The death of a single individual in class $n + 1$ adds a species to class n and hence adds n individuals to H_n . A birth in class $n - 1$ similarly augments the population H_n by n individuals. H_n changes in units of n as species drop in or out of this class—this will be relevant when we consider combinatorics later. Thus we have the master equation

$$\frac{dH_n}{dt} = -n(b + d)H_n + nbH_{n-1} + ndH_{n+1}. \quad (6)$$

(Eq. (6) can also be obtained by multiplying (5) by n and making the general replacement $ms_m \rightarrow H_m$, but the above derivation is useful.)

The equilibrium solution of (6) is

$$H_{n+1} = \frac{b}{d}H_n;$$

iterating this we find $H_n = H_0 \left(\frac{b}{d}\right)^n$, where H_0 is a normalisation constant. The master equation approach to relative species abundance (for *per capita* birth and death rates) yields an exponential distribution and hence a uniform prior for H_n .

The relevant priors are obtained from the master equations by setting b equal to d . The paper of Pueyo et al. (2007) made no appeal to these dynamical equations (but see Pueyo, 2006); it was postulated that the correct prior be obtained by requiring (spatial) scale invariance of the prior distribution. This was implemented through the minimal step of changing volume so as to remove an arbitrary single individual; invariance under this operation in fact yields the content of (5) for b equal to d ; that is, a prior for s_n of $1/n$ (Bowler and Kelly, 2012). The same invariance principle yields equations equivalent to (6) for H_n ; that is, a uniform prior. Regardless of whether the underlying fundamental principle is scale invariance or *per capita* rates, priors depend on how the problem is formulated, but the solution to the species abundance problem does not.

2.5. Combinatorics for individuals

As in Section 2.4 above, rather than sorting species into classes defined by the number of individuals in a species, sort the individuals in the community into classes defined by the number of individuals in a species. From the master equation, we expect the distribution of H_n to be exponential and it might be expected that a straightforward application of the most elementary microcanonical treatment in statistical mechanics, or equivalently maximising the information theory entropy, would yield this result. In fact, a simple application of MaxEnt is not applicable to this formulation and the problem is much more interesting and revealing. Using combinatorics (equivalent here

to MaxEnt) there is a very well known way of generating an exponential distribution. It is not applicable to the real problem of species abundance distributions governed by *per capita* birth and death rates, but it is instructive to analyse this treatment in detail, showing why it does not apply here.

We first suppose we have N individuals and sort them into classes such that class 1 contains N_1 individuals, class 2 N_2 and so on. The first step is to calculate the number of ways of distributing N to achieve a particular set of N_i . When filling class 1 with N_1 individuals, there are N ways of choosing the first, $N - 1$ ways of choosing the second ... to $N - N_1 + 1$ ways of choosing the last. The total number of ways is conveniently represented by the ratio of two factorial functions, $N!/(N - N_1)!$. We then calculate the number of ways of choosing N_2 individuals out of the remaining $N - N_1$, which is $(N - N_1)!/(N - N_1 - N_2)!$. Continuing this process until we run out of options and multiplying the successive factors, the number of ways of achieving a particular set N_1, N_2, \dots is simply $N!$ and does not depend on the occupation numbers of the individual classes. However, if the order in which the N_1 individuals are chosen is of no significance, $N!$ is a gross overestimate of the number of significantly different ways of achieving our given set. The number of ways of ordering N_1 objects is simply $N_1!$ and if all are equivalent the number of significantly different ways of choosing N_1, N_2, \dots is the multinomial coefficient

$$M = \frac{N!}{N_1!N_2! \dots} = \frac{N!}{\prod N_i!} \quad (7)$$

where the continued product is taken over the index i .

If the probability of finding this configuration depends on this number of different ways of achieving it, the most probable or equilibrium configuration is found by maximising M with respect to all N_i , subject to constraints. In practice the logarithm is maximised, so we find the extremum of

$$-\sum_i \ln(N_i!) \cong -\sum_i (N_i \ln N_i - N_i). \quad (8)$$

Maximising this function can be equivalent to maximising the Shannon entropy—MaxEnt (Section 2.2). If N_i is a function of n_i , constraints on $\sum_i N_i$ and on $\sum_i n_i N_i$ are incorporated by maximising

$$-N_i \ln N_i + N_i + \mu N_i + \lambda n_i N_i.$$

The result is an exponential

$$N_i = \exp(\mu + \lambda n_i). \quad (9)$$

This invites identification of the H_n with the N_i above, in order to achieve an exponential distribution. However, there are two things wrong with such an application of this result, and they are related. First, we have taken no account of the species structure within the H_n ; in obtaining (7), N_i could have been the number of dried peas in the i th urn. Secondly, one of the constraints above on N_i , leading to (9), does not apply to H_n . The constraint on the total number of individuals $\sum_n ns_n = \sum_n H_n$ is correct, but the other constraint is on the total number of species; $\sum_n s_n = \sum_n H_n/n$. There is no constraint on the first moment $\sum_n nH_n$. If

$$-H_n \ln H_n + H_n + \mu \frac{H_n}{n} + \lambda H_n$$

is maximised with respect to H_n , the result is

$$H_n = \exp\left(\frac{\mu}{n} + \lambda\right)$$

and this cannot even be normalised. The combinatorics leading to (9) are formulated in a way inconsistent with the nature of the problem and the constraint on the number of species.

The mistake was that the number of biologically equivalent ways of ordering the number of individuals H_n is not, in this problem, $H_n!$. Suppose n is equal to 2; the total number of individuals in species with two members is H_2 . As soon as we choose one individual from among H_2 we have also chosen a second individual which is in the same species as our first choice. There are H_2 ways of choosing the first individual; the two different orderings of the two individuals in the pair are included in the H_2 choices. After this choice there are $H_2 - 2$ individuals left from which to select the next. This choice also entails another individual—we are selecting in pairs because of the structure of the class of all individuals in species with $n = 2$ members. The number of equivalent ways of ordering the H_2 individuals in this class is thus $H_2(H_2 - 2)(H_2 - 4) \dots$ (These notions are not quite standard microcanonical statistical mechanics. It may help some readers to consider, say, three individuals in the H_3 class as linked by a string. The linkage represents the structure of the class; pulling on any one of the three removes them all and there are three choices, thus three ‘orderings’.)

In the same way, the number of biologically equivalent ways of ordering the number H_n is not $H_n!$; rather we have to make the replacement

$$H_n! \Rightarrow H_n(H_n - n)(H_n - 2n) \dots = H_n!^{(n)}. \quad (10)$$

The factorial has been replaced by the multifactorial $H_n!^{(n)}$. Thus the maximum number of significant ways of achieving a given configuration, subject to the constraints on the number of individuals and the number of species, is obtained by maximising

$$-\ln(H_n!^{(n)}) + \mu \frac{H_n}{n} + \lambda H_n$$

and this does not look like maximising the Shannon entropy, subject to constraints.

The following identity holds for the multifactorial (it is easily demonstrated by extracting n from each factor in (10))

$$H_n!^{(n)} = \left(\frac{H_n}{n}\right)! n^{\frac{H_n}{n}}. \quad (11)$$

(The right hand side of (11) is just the expression $s_n! n^{s_n}$ encountered in (2) with a prior (for s_n) of $1/n$. The substitution (10) has, in this alternative formulation, done the job of the prior in (2).)

We now have to maximise

$$-\left(\frac{H_n}{n} \ln \frac{H_n}{n} - \frac{H_n}{n}\right) - \frac{H_n}{n} \ln n + \mu \frac{H_n}{n} + \lambda H_n$$

(where we have used the Stirling approximation). Differentiating with respect to H_n , equating to zero and multiplying through by n we have that

$$\ln H_n = \mu + \lambda n \quad (12)$$

so that the distribution of H_n is exponential in n and if no constraints were applied the distribution would be uniform. The appropriate prior for $H_n = ns_n$ is uniform.

The number of ways of ordering equivalently the H_n individuals is also the number of equivalent ways of extracting them all from the class defined by n . The extraction goes in groups of n . The change in the number H_n as a result of birth or death of single individuals, as expressed in the master equation (6), also occurred in groups of n , as species left or entered class n . The same underlying principles govern both the master equation and combinatoric approaches to the problem framed in terms of H_n .

2.6. Individuals and the species abundance distribution directly

In Section 2.5 N individuals were sorted into classes containing all individuals that are members of species with n individuals.

Instead, sort the N individuals into bins defined by species, so that n_i is the population of species i . The nature of the problem is now expressed in the number of equivalent orderings of the n_i individuals in species i . This is a question inherently different from asking how many equivalent orderings there are of n_i atoms in a quantum state of energy E_i , where the class is defined by E_i and the physics is the extraction of n_i as a function of E_i . There, the number of different equivalent ways of emptying class i is certainly $n_i!$. In the species abundance problem the number of species s_n in class n , each having n members, is reduced by 1 on the removal from species i of any one of its n members, n_i being equal to n if species i is a member of species class n . Thus for the species abundance distribution, the number of equivalent ways of removing the n_i individuals from species i is given by n_i and not by $n_i!$; this is equivalent to the replacement (10) encountered in the previous section. For the case of a single species $H_n = n$ and so (10) becomes $n! \Rightarrow n!^{(n)} = n$. Again, the number of equivalent orderings is defined by the structure of the problem. Thus the number of significantly different configurations $\{n_i\}$ is given by

$$\frac{N!}{\prod n_i}. \quad (13)$$

There are two aspects to these combinatorics. The first is that from the definition of the classes removal of one individual removes all. The second is that any of the individuals can be picked with equal probability—per capita rates.

It should be noted here that (13) is valid if all individuals within a species are equivalent. If this were not a valid approximation and each individual is distinct in some important way, there would not be n_i equivalent ways of removing n_i individuals from the class; each way (or ordering) would be unique and were that the case the denominator of (13) would be replaced by unity. This individual equivalence is of course implicit in the earlier Sections 2.3–2.5. We shall encounter a case where individual differences are important in species abundance distributions in Section 2.7; a different application is Kelly et al. (2011).

The species abundance distribution gives the number of species with population n and is not concerned with the populations n_i of specific species i . There are s_n species with n members, so transforming (13) from classifying species by taxonomy to classifying species in terms of their numbers of individuals, the combinatoric (13) becomes

$$\frac{N!}{\prod n^{s_n}}. \quad (13a)$$

In this investigation of species abundance distributions, which species have population n is immaterial. The combinatoric (13a) was obtained from a specific assignment of populations to species and so the number of ways of achieving a set $\{s_n\}$ is obtained by adding together all possible versions of (13a); that is, multiplying by the number of ways of obtaining a particular set of numbers $\{s_n\}$ from a total number of species S . This is given correctly by the multinomial coefficient

$$\frac{S!}{s_1!s_2! \dots} = \frac{S!}{\prod s_n!}.$$

Thus the most probable configuration for the species abundance distribution is obtained by maximising with respect to s_n , subject to constraints, the function

$$W = \frac{S!}{\prod s_n! n^{s_n}}$$

and this is just Eq. (2) with a prior of $1/n$. Thus the combinatoric approach to the distribution of individuals over species, properly implemented, generates *a posteriori* the $1/n$ prior for the distribution s_n . The removal of any one of the n individuals in a member

of species class n removes that species and there are n equivalent ways of doing it—as in the master equation (5). This treatment embodies the same biology as the master equation, but the MaxEnt prior emerges from the combinatoric treatment, rather than requiring a transplant from the master equation. It might seem that this treatment is so general that nothing more needs to be said, but this is not completely true. This approach through statistical mechanics (or maximum entropy) calculates that distribution corresponding to the maximum number of ways of achieving it and this is not necessarily the right distribution. It is likely to be right if a dynamical system is continually exploring all configurations, but there is a sense in which log series species abundance distributions could be circumvented by niche structures. This is discussed in Section 3.

2.7. Relevance to Harte's METE

The discussion in 2.6 above is relevant to an important assumption in Harte's model, called by him the Maximum Entropy Theory of Ecology (Harte, 2011). He applies the MaxEnt formalism to produce a probability distribution that yields not only the species abundance distribution but also deals with the distribution of metabolic rate—the distribution of power consumed. He defines what he calls the ecological structure function $R(n, \varepsilon)$. This is the joint probability that a species selected from the community has population n and the probability density that an individual selected from the population of that species has power consumption ε . The form of this function is obtained by maximising $-\sum_n \int R(n, \varepsilon) \ln R(n, \varepsilon) d\varepsilon$ subject to constraints. The first is a normalisation condition, the second a constraint on the number of individuals in the whole guild. These are just the constraints usually applied in obtaining a species abundance distribution this way. The third is a constraint on the power consumption of the whole guild. These constraints are implemented in the usual way, using Lagrange multipliers, and the form of the function that emerges is

$$R(n, \varepsilon) \propto \exp(\lambda_1 n + \lambda_2 n \varepsilon)$$

normalised by summing over n and integrating over ε . If the dependence on ε is integrated out, the species abundance distribution results and is given by

$$R(n) \propto \exp(\lambda_1 n)/n$$

that is, a log series distribution. In the above equations the parameters λ_1, λ_2 are Lagrange multipliers, the values of which are determined by the constraints. The assumed MaxEnt prior was chosen to be uniform and no constraint was applied that would introduce an exponential in ε alone. (Had such a constraint been applied, the denominator of the species abundance distribution would have been rather $n + c$, of interest because with this modification to the log series the function can be very close to a log normal Volkov et al., 2005.)

In the light of Section 2.6, it is natural to ask whether in fact a prior of form $1/n$ should have been applied in maximising the MaxEnt function in METE; if so, the species abundance distribution would be an exponential divided by n^2 . Harte states explicitly that “On the grounds of lacking a reason to do otherwise, we invoke no non-uniform prior distribution” (Harte, 2011, p. 143). There is in fact no conflict between the absence of a $1/n$ prior in the maximisation of Harte's entropy function and our results, for we can apply the treatment of Section 2.6 to Harte's model and find (*a posteriori*) a uniform prior. In 2.6 above, we were sorting N individuals into species so that species 1 contained N_1 individuals and so on. The total number of ways of ordering N is just $N!$ and the number of equivalent ways of removing a species with n individuals was n (despite the fact that the number of ways of

ordering n objects is $n!$). Once any one individual has been removed the species is removed and along with it the other individuals. This is the origin of the MaxEnt $1/n$ prior in Section 2.6.

In Harte's model, the individuals in a species are not equivalent; they are distinguished by their metabolic rates ε . In terms of Section 2.6, individuals are sorted according not only to species but also according to metabolic rate. There is just one way of removing an individual, having metabolic rate ε , from a species with n members, not n ways. Thus the prior emerging for Harte's ecological structure function, applying the reasoning of Section 2.6, is the number 1 (or more generally, uniform). Another way of seeing this result is to construct a master equation for the functions $R(n, \varepsilon)$, along the lines of Section 2.3 or as formulated by Volkov et al. (2003). The rate at which the probability $R(n, \varepsilon)$ is depleted by (for example) individual death is simply given by the rate at which an individual in n with metabolic rate ε dies, a term $d(\varepsilon)R(n, \varepsilon)$. If it is assumed with Harte that ε enters the ecological structure function only in the product $n\varepsilon$, then the master equation so constructed yields just Harte's function. Thus thinking in terms of the machinery underlying the ecological structure function has removed an uncertainty in the appropriate form and, I suggest, provided the reason Harte's choice was appropriate. In Harte's model, again the MaxEnt prior results and again the species abundance distribution is log series.

The content of this Section 2 is indeed mostly technical. It reinforces the conclusion that where individual *per capita* birth and death rates dominate community dynamics, the species abundance distribution will be approximately log series. In Section 3 I consider the opposite extreme, where distribution of niche space is dominant.

3. The roles of niches, specific and flexible

3.1. Niches and broken sticks

The previous section demonstrated in various ways that if the dynamics of a community (I think of forests for historical reasons) are dominated by the (uncorrelated) birth and death of individuals, then to a first approximation species abundance distributions will be log series. It is not necessarily the case that an important role for individual rates conflicts with niche structure; a vivid discussion can be found in the first paragraph of the section ‘Random assortment model’ in Tokeshi (1990) and Bowler and Kelly (2012) argued that such niche structure would need to be flexible enough for per capita birth and death rates to work. This is not quite the way in which the resource partitioning approach to species abundance distributions has been envisaged. MacArthur (1960) devised his ‘broken stick’ model to deal with species with non-overlapping niches and the underlying idea (MacArthur, 1957) is that the stick represents some general resource partitioned among species. The lengths of the pieces of the broken stick give the resource available to the corresponding species and hence the numbers of individuals representing species. This basic idea also underlies the developments of Sugihara (1980) and Tokeshi (1990). The units of resource generated by breaking of the stick (sequential in the last two treatments) are regarded as being fixed in time and the niches thereby defined are envisaged as full and staying full. It seems implicit in models of resource or niche apportionment that an individual, on dying, is immediately replaced by another of the same species, in contrast to the picture of Section 2. The species abundance distribution would be controlled by distribution of specific resources, perhaps as modelled in MacArthur (1957, 1960), Sugihara (1980) and Tokeshi (1990, 1996). The sequential niche apportioning models seem far removed from biological machinery and could be regarded merely as algorithms for implementing the random walk of MacArthur equation (1), but they have the

attractive feature that the trees generated by sequential stick breaking look just like ecological dendrograms on the one hand (Sugihara et al., 2003) and phylogenetic trees on the other. Speciation would correspond to further sequential breaks, but extinction seems not to have been considered.

There is no question but that fairly rigid niches do exist and define rather different biomes. They are distinguished not just by the biota they contain, but also by significant abiotic differences, such as aridity or elevation or geography. Forest merges into savanna which merges into desert; on a smaller scale cactus gives way to oak which gives way to pine when ascending a sky island. Rainfall distinguishes the tropical dry forest on the Pacific coast of Mexico from moist forest on the Caribbean—they have few species of trees in common. Even such super-niches vary, on the timescale of thousands of years, and ecological guilds fit into super-niches. Niche structure within an ecological guild will be on a smaller geographic scale and will probably be less clearly distinguished by abiotic differences. It seems implausible that niches within guilds are rigid; they are likely to be affected by biotic factors and vary on a shorter timescale than super-niches. Thus niche structure in highly speciose communities (such as tropical forests where very many species can be found to some extent interdigitated) cannot be expected to correspond to niches both rigid and full.

Suppose, however, that groups of species or even individual species really do have strictly defined niches with fixed limits on capacity. To what extent would the results of assuming dominance by *per capita* birth and death rates of individuals be affected? First, no species could exceed the capacity of its niche, which would make the largest abundances not merely very unlikely but impossible. If the *per capita* birth and death rates were close, births and deaths uncorrelated, the individual niches would have populations drifting according to the master equation (5), up to the niche capacity, and the $1/n$ log series would be approximated. This is essentially the case studied by Chisholm and Pacala (2010). The results are potentially different if the (forest) dynamics are such that these niches are full all the time. Then the species abundance distribution becomes the distribution of niche space and individual birth and death rates become irrelevant, because death and birth are so tightly correlated that the master equation (which assumes no correlations between the individual steps) has no validity. It is also possible to modify birth and death rates in the master equation to concentrate the population close to the niche limit; similarly, if there is a strict upper limit N on the capacity of a niche, the log series obtained from the micro-canonical ensemble can concentrate very close to N if the λ parameter in (4) is large and positive. The species abundance distribution would then be dominated by the distribution of niche capacities; the highly N -skewed log series within the niches would be details subsumed in the overall picture. It suffices to consider the extreme case: every individual that dies instantly replaced by a new individual. If niches stay full, the question is how to determine the distribution of niche space, or niche apportionment.

Return to MacArthur (1960). He commences discussion of the first term in MacArthur equation (1) by supposing a fixed number of individuals and a fixed number of species and below his Eq. (3) refers to ‘niches non-overlapping’. He envisages the numbers of individuals controlled by the niche space available for the appropriate species, and his famous ‘broken stick’ model contains the assumption that resource is partitioned and allocated at random to a fixed number of species. This is not a dynamic model, but the most probable distribution is given by the methods of the microcanonical ensemble (as in Section 2 of this paper) with a uniform ‘prior’. The result is an exponential distribution of the number of species with the number of individuals. It is a simple way of implementing the assumed constraints of a fixed number of species and a fixed number of individuals, but does not seem

to have had any deeper ecological reasoning behind it and it does not work very well for real communities; they have distributions somewhere between a log series and a log normal. This could result if niche apportionment was dominated by a term like the second in MacArthur equation (1) rather than the first, albeit on a much longer timescale than for individuals.

Bulmer (1974) pointed out that log normal distributions result from breaking MacArthur’s stick sequentially, provided that the law of breakage is independent of the size of the piece to be broken, a form of scale invariance. At each step an extant piece is divided into two fractions according to the same splitting function and the probability of a piece being selected for splitting is independent of its length. He suggested that the corresponding hierarchical classification could correspond to the phylogenetic tree of a group of species or to their ecological dendrogram. Sugihara (1980) devised breaking rules successful in generating the canonical log normal of Preston (1962), and Tokeshi (1990, 1996) pointed out that MacArthur’s original model can be reformulated as sequential (but not scale invariant) breaking and investigated this and other variants. In fact, the distribution generated by scale invariant sequential breaking is left skewed log normal, because stick breaking imposes an upper limit but no lower limit on fragment size (Nee et al., 1991). The standard deviation and skewness parameters depend on the rules for splitting (fractionally) the pieces of stick at each step. The end product of sequential stick breaking thus corresponds to dominance of the random walk on the right hand side of MacArthur equation (1), discussed in Section 1, interpreted in terms of niche space. Stick fragments do not recombine and well defined niches are not going to maintain a fixed capacity; the sequentially broken stick must represent some suitable average over more complex behaviour. In both approaches the essential feature is random fractional variations and is not directly linked to any fundamental biological principle. Attempts to devise precise recipes for breaking the stick with a view to fitting various different species abundance distributions are purely heuristic; there seems to have been no biological machinery underlying such rules proposed. (There is a charming fable illustrating how multiple processes might generate an ecological dendrogram in a community of forest dwelling birds, to be found in Sugihara et al., 2003.) If niches are both filled and rigid, any not implausible hierarchical structure is likely to result in something between log series and log normal; in terms of fundamental biology there has not been much advance since the sequential fractional partitioning of Bulmer (1974). However, resource distributions might be determined through a mechanism more obviously related to MacArthur equation (1) and not conceived in terms of sequential stick breaking; the niches would be filled but not rigid, varying in capacity with varying environmental conditions.

3.2. A niche mechanism for scale invariance?

Suppose that niche spaces are continuous variables and vary with characteristic times that may be longer than for the birth and death of individuals. On the crudest scale niche capacity is likely to vary through such abiotic mechanisms as changes in long term aridity or mean temperature or prevailing winds. It might be that if a large niche space is so affected the change of niche space is large and if a small space is affected, small; that is, effects are fractional as in MacArthur equation (1), scale invariant in the sense of Pueyo et al. (2007). This is also a model for populations determined by multiplicative processes. If the growth rate equation for a niche space ξ were of the form of MacArthur equation (1), with $\ln \xi$ dominated by a random walk, then the distribution of specific resources would be roughly log normal. In terms of a MaxEnt treatment of niche space, imposition of scale invariance

on the prior, as implemented by Pueyo et al. (2007), results in a distribution of form $1/\xi$. The change in probability $\Psi(\xi)$ for any change in the resource allocation ξ is proportional to ξ ; this is equivalent to *per capita* birth and death rates for individuals and yields the same results, even though the niches are always full.

It is easier to accept scale invariance, as a principle, when dealing with niche space or resource availability than when dealing with numbers of individuals directly. The notion that niche space responds to environmental conditions in a scale invariant way is entirely speculative, but it does suggest a mechanism for the scale invariance of Pueyo et al. (2007), there justified by supposing a heterogeneously uninformative prior. If such a mechanism governed filled niches, then scale invariance would be the underlying principle. For flexible niches, that scale invariance is an outcome of *per capita* birth and death rates (Bowler and Kelly, 2012). However it may be, these different mechanisms illustrate another curious feature to do with *priors*.

3.3. A polemic on minimally informative priors

If the problem of applying statistical mechanics to species abundance distributions is approached by first working with discrete, integer numbers of species and individuals, as in Pueyo et al. (2007) and Eq. (2), the *least informative prior* in the sense of MaxEnt is uniform, according to Jaynes (2003, p. 373). Pueyo et al. (2007) adopted the $1/n$ prior from scale invariance (arguing for a *minimally informative prior* and hence no information on scale) and made an analogy with the Jeffreys prior for a continuous variable x , $1/x$. For the continuous niche space variable ξ , exhibiting scaling as discussed above, the prior $1/\xi$ can be held to be *minimally informative*, Jaynes (2003). There is no obvious reason why the prior appropriate to a particular real problem should contain ‘minimum information’, but here apparently two different approaches to the problem of species abundance distributions yield two different ‘minimally informative’ priors for the same problem, depending on whether the discrete integer character of numbers of individuals is put in *a priori* or imposed *a posteriori*.

There is a real difficulty here if the notion of ‘minimally informative’ priors is to be taken seriously. If the biology of *per capita* rates is used to determine priors, then the uniform prior is appropriate to a formulation in terms of the discrete distribution H_n , yet if the species abundance distribution is expressed in terms of the discrete distribution s_n then the prior of Eq. (2) is $1/n$. The uniform prior is supposedly the *least informative* for a discrete distribution, but here two different priors describe the same problem and only one is ‘minimally informative’. The theory developed in Pueyo et al. (2007) rather postulates that the prior should contain no information on the (geographical) scale of the species abundance and hence that the prior should be invariant under changes of scale. This single application of a precisely defined principle of ‘minimal information’ yields both the uniform MaxEnt prior for H_n and the $1/n$ prior for s_n , Eq. (2).

The moral is that any real problem must be thought through. In applying statistical mechanics by maximising (subject to constraints) a weight function, the correct weight function for the problem must be worked out. If there is not enough information about the nature of the problem to determine the prior unambiguously, any assumed MaxEnt prior is a guess. A prior that is said to be *minimally informative* is not especially likely to be a good guess.

4. Discussion

This paper is an attempt to discuss fairly generally the techniques of applying statistical mechanics to the problem of species abundance distributions. I have considered the case where birth

and death of individuals dominates the dynamics of a community and the opposite extreme where the distribution of niche space determines species abundance. The two may not be distinguishable, particularly as the time scale for change of niche space approaches that of individual birth and death. Technically, my approach has been to find the weight function appropriate to the biological nature of the machinery underlying community dynamics, and maximise that function subject to constraints on the total number of individuals and the number of species. I found the weight function appropriate to individual dynamics in several different ways. I claim that there can be no doubt that under these dynamics the correct weight function incorporates the $1/n$ prior of MaxEnt. I emphasised that the recipes of what is called MaxEnt need to be approached very carefully and suspiciously; in one of my approaches the *prior* emerged *a posteriori* and in the section above are ways of obtaining two different ‘minimally informative’ priors for the same problem. They cannot both be correct. Section 2 formulated the species abundance problem in two ways that yielded the same answer, but have different priors.

If species abundance distributions instead reflect rigid and specific apportionments of resource, the correct approach in the application of statistical mechanics is less clear. We have no treatment to offer as simple and clear cut as that of Section 2.6. It is merely plausible that such niche abundance distributions (hence species abundance distributions) would be represented approximately by a left skewed log normal form. This can be argued on the basis of multiplicative processes producing the composition of the final community (the sequentially broken stick of Bulmer, 1974). A separate possibility, admittedly speculative, is that niche space is divided up dynamically and niches add or lose small pieces of niche space at rates proportional to the space already there. This speculative notion leads to a scale invariant prior for niche space and hence log series species abundance distributions; very minor modifications generate skewed log normals (Volkov et al., 2005; Pueyo et al., 2007). The general argument of Section 2.6 is then resurrected. (If, however, niches add or lose small pieces of niche space independently of the space already present, exponential species abundance distributions would be generated. Something of this kind seems to be going on in the very different ecological problem of the distribution of species over the number of sites at which they are alien Kelly et al., 2011.)

Thus I have reinforced and extended the conclusion of Bowler and Kelly (2012); that the natural minimal form of species abundance distributions is log series (see also Volkov et al., 2003, 2005; Pueyo, 2006; Pueyo et al., 2007). I envisage rigid and slowly changing niches as dominated by largely abiotic factors and likely to contain whole guilds rather than differentiating among individual species within a guild. Niche structure within a guild is likely to be on a smaller scale, change faster and contain a strong biotic component. Working down the phylogenetic tree (or ecological dendrogram) I expect niche structure to have an increasing biotic component and such niches to be increasingly less exclusive. If niches are flexible in this sense (as postulated in Bowler and Kelly, 2012; see also the data and analysis in Clark et al., 2010) then in the words of Koch (1974b) “the niche concept must be redefined in a broader way”. The species abundance distributions will become increasingly dominated by the dynamics of the birth and death of individuals and the effects of niche structure on species abundance distributions likely average out (Pueyo et al., 2007; Bowler and Kelly, 2012). The demonstration (Koch, 1974a,b; Armstrong and McGehee, 1980) that an unlimited number of species can in principle coexist on very few resources, if population densities of individual species fluctuate with time, is surely relevant. Such oscillations could be generated either

internally (biotically) or by externally imposed environmental fluctuations.

The extreme limit of niche averaging is Hubbell's neutrality. There is little doubt that this provides a good description of species abundance distributions as a result of averaging and that the success of the log series distribution also depends on averaging (Bowler and Kelly, 2012). Within guilds there exist correlations between populations of similar species (Sugihara et al., 2003, related to the structure of ecological dendrograms, Kelly et al., 2008, 2010, related to the structure of phylogenetic trees) that could well be indicative of niche structure and probably are. John et al. (2007) report correlations between the distribution of species and the distribution of soil nutrients at three different tropical forest sites. The correlations in Kelly et al. (2008), Kelly et al. (2010) and John et al. (2007) result in no signatures legible in the species abundance distributions of those communities. The question as to what extent species abundances are the products of population drift on the one hand and niche structure on the other is unanswered, but the default distribution is log series and the log series distribution is determined only by the number of individuals and the number of species. (The same is true of any two parameter distribution function, such as the canonical log normal.) If there is information in species abundance distributions it must be looked for in departures from the log series, although such departures may not give a clear signal of underlying processes—see Volkov et al. (2005) and the subsequent exchange in *Nature* 441 (4 May 2006). Questions about the contribution and nature of niches are not likely to be answered through the coarse grained species abundance distribution and since the essential input is the number of species and number of individuals, no model of species abundance distributions alone addresses the enormous question of species richness. With few exceptions, species abundance distributions tell us nothing about species and their interactions; here I differ from the conclusions of McGill et al. (2007). The $1/n$ factor that sets the scale for the log normal distribution and turns an exponential into a log series largely determines the dominant feature of species abundance distributions (the 'hollow curve' of McGill et al., 2007), generated by *per capita* birth and death rates if niches are flexible and hypothetically via fractional variations in niche space in the event of niches mostly full. Biological machinery is responsible for the $1/n$ prior of Eq. (2) and so for the 'hollow curve'. In the former case the machinery of *per capita* rates is familiar and not invented to explain species abundance distributions. Treatments involving niche space involve postulates devised to generate approximate log normals. The common key factor is that population sorting processes are fractional; this is the biological principle underlying species abundance distributions. Insofar as species abundance distributions contain further biological information, it must be looked for in departures from the underlying log series or, where appropriate (perhaps communities of fish or of plankton), in change of the distributions over time.

No model of species abundance distributions predicts the number of species to be found in a guild; given that the gross structure of species abundance distributions is understood, the real question is the origins of species richness.

Acknowledgment

I am happy to acknowledge very many relevant conversations with Colleen Kelly.

References

- Armstrong, R.A., McGehee, R., 1980. Competitive exclusion. *Am. Nat.* 115, 151–170.
- Banavar, J.R., Maritan, A., Volkov, I., 2010. Applications of the principle of maximum entropy; from physics to ecology. *J. Phys.: Condens. Matter* 22, 0631001.
- Bowler, M.G., Kelly, C.K., 2010. The general theory of species abundance distributions. *arXiv:1002.5008*.
- Bowler, M.G., Kelly, C.K., 2012. On the statistical mechanics of species abundance distributions. *Theor. Popul. Biol.* 82, 85–91.
- Bulmer, M.G., 1974. On fitting the Poisson log normal distribution to species-abundance data. *Biometrika* 30, 101–110.
- Chisholm, R.A., Pacala, S.W., 2010. Niche and neutral models predict asymptotically equivalent species abundance distributions in high-diversity ecological communities. *Proc. Natl. Acad. Sci. USA* 107 (36), 15821–15825.
- Clark, J.S., et al., 2010. High-dimensional coexistence based on individual variation: a synthesis. *Ecol. Monograph* 80, 569–608.
- Harte, J., 2011. *Maximum Entropy and Ecology*. Oxford University Press, Oxford.
- Hubbell, S.P., 2001. *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton University Press, Princeton, NJ.
- Jaynes, E.T., 2003. *Probability Theory: The Logic of Science*. Cambridge University Press, Cambridge.
- John, R., et al., 2007. Soil nutrients influence spatial distributions of tropical tree species. *Proc. Natl. Acad. Sci. USA* 104, 864–869.
- Kelly, C.K., et al., 2008. Phylogeny, niches and relative abundance in natural communities. *Ecology* 89 (4), 962–970.
- Kelly, C.K., et al., 2010. Fractional abundance and the ecology of community structure. *arXiv:1008.2527*.
- Kelly, C.K., et al., 2011. The statistical mechanics of community assembly and species distribution. *New Phytol.* 191, 819.
- Kendall, D.G., 1948a. (a) On some modes of population growth leading to R.A. Fisher's logarithmic series distribution. *Biometrika* 35, 6–15.
- Kendall, D.G., 1948b. (b) On the generalized "birth-and-death" process. *Ann. Math. Stat.* 19, 1–15.
- Koch, A.L., 1974a. Coexistence resulting from an alternation of density dependent and density independent growth. *J. Theoret. Biol.* 44, 373–386.
- Koch, A.L., 1974b. Competitive coexistence of two predators utilising the same prey under constant environmental conditions. *J. Theoret. Biol.* 44, 387–395.
- MacArthur, R.H., 1957. On the relative abundance of bird species. *Proc. Natl. Acad. Sci. USA* 43, 293–295.
- MacArthur, R.H., 1960. On the relative abundance of species. *Am. Nat.* 94, 25–36.
- Margalef, D.R., 1958. Information theory in ecology. *Gen. Syst.* 3, 36–71.
- McGill, B.J., et al., 2007. Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework. *Ecol. Lett.* 10, 995–1015.
- Nee, S., Harvey, P.H., May, R.M., 1991. Lifting the veil on abundance patterns. *Proc. R. Soc. Lond. Ser. B* 243, 161–163.
- Preston, F.W., 1948. The commonness, and rarity, of species. *Ecology* 29, 254–283.
- Preston, F.W., 1962. The canonical distribution of commonness and rarity. *Ecology* 43, 185–215.
- Pueyo, S., 2006. Diversity: between neutrality and structure. *Oikos* 112, 392–405.
- Pueyo, S., He, F., Zillio, T., 2007. The maximum entropy formalism and the idiosyncratic theory of biodiversity. *Ecol. Lett.* 10, 1017–1028.
- Sugihara, G., 1980. Minimal community structure: an explanation of species abundance patterns. *Am. Nat.* 116, 770–787.
- Sugihara, G., et al., 2003. Predicted correspondence between species abundances and dendrograms of niche similarities. *Proc. Natl. Acad. Sci. USA* 100, 5246–5251.
- Tokeshi, M., 1990. Niche apportionment or random assortment: species abundance patterns revisited. *J. Anim. Ecol.* 59, 1129–1146.
- Tokeshi, M., 1996. A new explanation of relative abundance patterns in species-rich assemblages. *Oikos* 75, 543–550.
- Volkov, I., et al., 2003. Neutral theory and relative species abundance in ecology. *Nature* 424, 1035–1037.
- Volkov, I., et al., 2005. Density dependence explains tree species abundance and diversity in tropical forests. *Nature* 438, 658–661.