

Chapter 1. Introduction

“Truth proceedeth more readily from error than from confusion.”

Sir Francis Bacon

Whatever one might mean by the term “biodiversity,” the main thrust of all inquiries conducted under that name is the desire to say something meaningful about species of living things in their natural habitat. In particular, the population sizes of various species, both in relative and absolute terms, become the central focus. The most comprehensive studies address whole communities, namely all the species within a particular group and living within a specific area. One might wish to know how many species of butterflies there are in a particular forest, for example, along with their relative or absolute abundances. Or one might wish to study a community of aquatic plants in a tropical embayment with the same general aim in mind. Literally thousands of such studies have been conducted for more than 100 years, in every clime, habitat, and living kingdom. In all cases -- or nearly all cases -- the investigator has had recourse to only one principal mode of investigation, namely to take samples of the community in question.

This book is about such communities and the exact relationship of their (unknown) abundances to abundances that appear when samples are taken. Short of a Faustian magic mirror that shows us a whole community, along with all its living components, we are stuck with samples. As indicated in the description of the frontispiece, the (so-called) J-curve is ubiquitous in samples of natural communities and there is a clear consensus to that effect, reached in the last decade. (McGill et al. 2007). The random sample of a non-natural community such as a large vegetable garden is likely to have a rather different distribution of abundances.

When we examine any of the histograms that appear in this book, we shall not know which species have which abundances. The theory proposed here to account for this ubiquitous shape has an associated statistical manifestation called the logistic-J distribution (Dewdney 1998, 2003). Like other proposals to account for this shape (see the next section), the logistic-J theory is oblivious to the actual species that have a given abundance; only the number of species having that abundance is relevant.

1.1 The community and the sample

Figure 1.1 illustrates the relationship between a natural community and a sample of it. Each little square represents a species, with its position on the horizontal axis representing its abundance. It will be a general feature of the histograms shown in this book that the numerical label associated with each abundance appears at the right hand side of the corresponding interval on the abundance axis. This is a technical device that harmonizes discrete and continuous versions of the logistic-J distribution. Some theoretical abundances, as well as some empirical ones are fractional, with values such as 3.2, 3.7, and 3.8, for example, all gathered into the category 4.0. In general, the k th category would embrace all abundances lying in the half-open interval, $(k-1, k]$

It will be a common experience for readers of this book to encounter histograms like the ones in Figure 1.1 where, due to horizontal space limitations, not all species can be shown. In cases where it matters, the missing abundances are always given. In cases where it doesn't, the reader may imagine them as being present, nevertheless. Indeed, logistic-J theory indicates just how far out on the abundance axis such species might be expected to occur.

The histograms shown in Figure 1.1 are imaginary, but reflect what one might call realistic variability. Those with experience of actual sample histograms, may well accept the sample data as typical, but who has seen the histogram of an entire community? The community histogram is simply an indication of what a community might look like under logistic-J theory. The horrifying sight presented in the sample of so many species seemingly about to go extinct vanishes at the sight of the community with no such peak. It must be remembered that, in general, the abundance of a species in a sample is usually much lower than its abundance in the community from which the sample was taken. For example, a species of abundance 1 in the sample histogram might well have abundance 27 in the community. The sampling process has “mapped” it from a position of relatively high abundance to one of relatively low abundance.

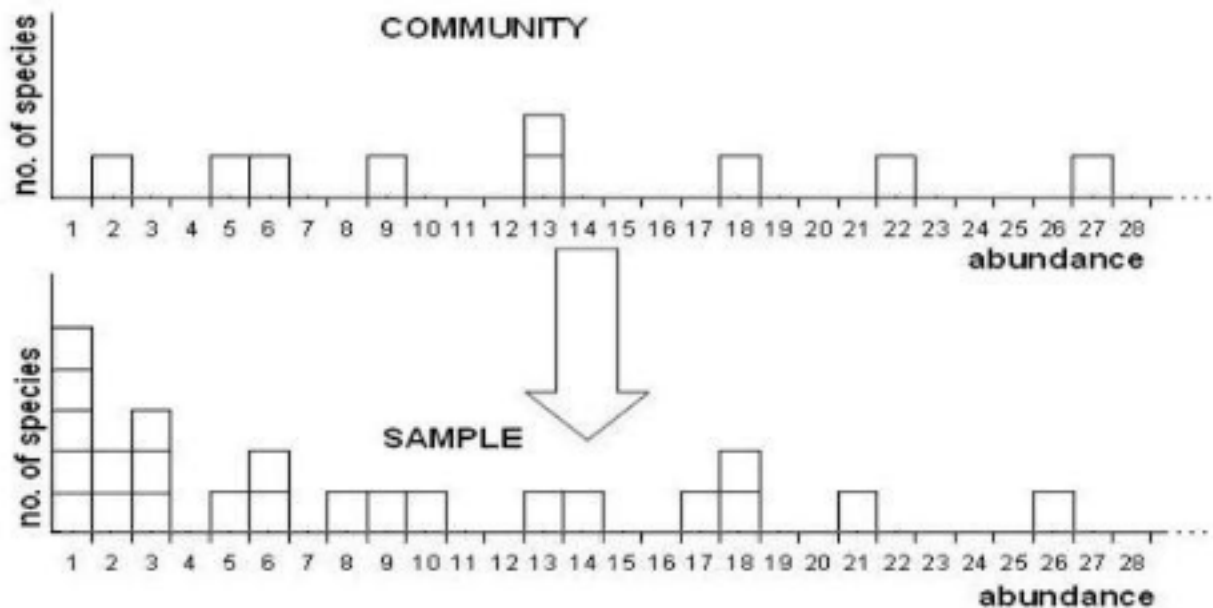


Figure 1.1 The relationship between abundances in a community and those of a sample

Although the focus of this work is ultimately the community under investigation, we have only samples to work with. As a result, most of the histograms presented in this book belong to either hypothetical or actual samples. As illustrated in the frontispiece, they are generally characterized by a high initial peak that is followed by a rapid descent in bar height, whether smooth or ragged. The descent becomes more gradual with ever higher abundances, there being ever fewer species at these abundances. The theoretical curves that represent these histograms have the same general

shape, although they are smoother looking. As for communities, the histograms and curves that are associated with them tend to have rather low initial peaks.

Many biologists find the high initial peaks that occur in most samples rather alarming, as if many species were about to go extinct. However, as the example above illustrates, samples of ever greater intensity, all taken from the same community, would show a trend toward ever lower initial peaks. In any case, from the point of view of logistic-J theory the initial peaks are far from “anomalous” (Coddington et al., 2009). They are entirely natural.

The sections in this introductory chapter include a brief review of other proposals for abundance distributions, along with an explanation of goodness of fit tests and their use in both negative and positive mode. An introduction to the logistic-J distribution begins with the probability density function (pdf), an explanation of the two associated parameters and a derivation of a formula for the mean μ . The chapter ends with a brief history of my early work in this area, providing another doorway to understanding.

1.2 The need for an appropriate theory

It is an unusual development in any science that a plethora of formulae should result from attempts to capture a natural phenomenon in mathematical terms. In the normal course of research any formula found wanting would be discarded and a new formula developed. Since the early 1940s about a dozen formulas for species abundances (in samples) have been proposed and none of them discarded. Here are the principal ones discussed by Magurran (1988). She also mentions another three proposals, for a total of seven

1. The geometric series (May 1975)
2. The log-series distribution (Fisher, Corbett & Williams 1943)
3. The lognormal distribution (Preston 1948), (See Sections 3.6 and 3.6.1)
4. The broken stick model (MacArthur 1957)

Among the other proposals discussed by Magurran, special mention should be made of the dynamical model of Hughes (1986), the only proposal prior to the 1990s to be tested against an appropriately large number of samples from the field. Unfortunately, Hughes was never able to find a closed formula for the J-shaped curves that his data produced.

It is difficult to characterize succinctly, without committing the sin of oversimplification, how ecologists interested in theory have worked in the past. However, it is fair to say that the various proposals just indicated have been compared with various sets of field data over the years with the aim of discovering which proposals fit the field data best. This seems like a perfectly logical procedure, but it contains a very dangerous pitfall in the form of an unstated (and apparently unrealized) assumption. From the point of view of statistical theory, the outcome of this method is (when one thinks about it) predictable, with each distribution fitting best with certain data sets

and not fitting others as well. For example, four of the proposals listed above are described by Magurran, in her research summary, as follows. Based on the rank-abundance representation of field data (see Section 4.2), the geometric series best fitted (apparently by visual inspection) plant data from a sub-alpine forest, the log-series and the lognormal best fitted a histogram of plant data from a deciduous forest, and the broken stick model best fit a survey of nesting birds in a deciduous forest.

The methodology of comparing a handful of field histograms with various theoretical distributions has persisted into the new millennium. For example, Hubbell (2001) draws support for various manifestations of his zero-sum multinomial distribution from seven studies of various biotic groups, from trees to birds and, not surprisingly, finds a good (visual) fit within a milieu that contains previously proposed distributions implicitly. Connolly et al (2005) compared four different distributions with survey data involving corals and fish in Pacific Ocean reefs, finding that some distributions fit the data from some reefs better than others, where other distributions did better.

As more and more data sets are compared with the limited number of distribution proposals under active consideration, it is not surprising that each camp builds up followers, so to speak. A more recent summary of the situation provided by Diserud and Engen (2000) points out that “A large number of data sets from ecological communities correspond well to the model in which the abundances are lognormally distributed . . .”, later pointing out in the same breath that “. . . lognormal models may often provide a rather bad fit to observed data.”

As if it were a tacit admission that something is wrong, there has been a long tradition in theoretical ecology of attempting to unify somehow the various distribution proposals into one. It is not difficult to find mathematical generalizations of the proposals mentioned above. But was it the unsatisfactory results of such unification projects that has spawned a cottage industry of attempts? Beginning with Sugihara (1980) and culminating with (Diserud and Engen 2000), then (Hubbell 2001), the unification process has served to keep alive the idea that each proposal has its niche, so to speak. More recently, a group of 18 authors (McGill, Etienne, et al. 2007) have proposed “moving beyond single prediction theories to integration within an ecological framework.” The article demonstrates clearly a continuing confusion about the roles of species abundance distributions in an ecological context. It makes little sense to “move beyond” the present situation until it is corrected.

The manner in which these applications were determined illustrates an underlying weakness in the understanding of both the wide range of variation in the shapes of sample data and the reasons for that variation. The weakness has resulted in no little soul-searching, as when the 14-author paper, Doak, Estes et al. (2008) declare that, in view of science being full of surprises, “. . . it is not so surprising, so to speak, that we frequently face outcomes of experiments and observations that leave us scratching our heads, wondering how we could have been so wrong in our expectations.” From the point of view advanced here, it isn't surprising at all.

Two main assumptions appear to support the traditional approach to the assessment of proposals. It isn't clear from my reading of the literature that either assumption has ever been explicitly stated in some journal article, but most would agree that they would be necessary assumptions for the method just described to work:

- a) the shape of the sample histogram, however its abundances are plotted, reflects in a specific way the shape of the community.
- b) the shape of a community histogram is stable in the sense that two samples that are well separated in time or space will tend to reflect that (type of) community.

The first assumption is an article of faith with field biologists. If the samples they so painstakingly gathered in the wild turned out to have no such relation with the community, there could only be despair. In fact, the distribution underlying the sample will reflect the distribution underlying the community in a very direct manner. The underlying distributions are the same, on average, but for a change in parameter values. (Dewdney 2002) Thus, whatever one means by the (deliberately vague) word "reflect", the first assumption is completely solid.

A reasonable abstract description of any "community" can be framed in the context of the upper histogram in Figure 1.1. As I have pointed out, the histogram of a community would have a rather strung-out appearance, in comparison with the histogram of a sample of it. Over time, the various populations that make up the community will all fluctuate, some increasing, some decreasing, with occasional changes of direction, seemingly at random, among all of the populations. In terms of the histogram itself, one may visualize these fluctuations as a sort of vibration when the film is speeded up, so to speak. Species at the high abundance end vibrate at a higher frequency because their populations are much larger. At the low abundance end, in corresponding fashion, the species move in more sedate fashion, slowing rather quickly as they approach extirpation.

The reasons for the apparent randomness in the motion of these species are typically myriad. Suffice it to say that the overall shape of the histogram, however one described it, would also change. Would it change so much as to defy any attempt to describe it as "geometric", "log-series", "lognormal", or "broken stick"? The answer is simple: It would change its overall shape enough, over time, to defy any proposal, including the centrepiece of this book, the logistic-J distribution. I cannot account for the reasoning behind Assumption b, but it may have something to do with the notion that large sample size guarantees statistical significance. In other words, the large size of the field samples that have been used to promote one distribution over another were themselves thought to guarantee that the overall results were somehow determinative, all of that at a guess.

According to the first assumption, as a community distribution changes its shape over time, so would the samples taken of it. Studies that would attach significance to such changes (Magurran

2007) fall into essentially the same trap. Changes in the shape of a community distribution will happen willy-nilly and the new variations of the J-curve, as revealed (or not) by the sample distribution have little actual significance. Logistic-J theory tells us that living communities in general have no particular shape beyond lying within the class of perturbations claimed by that theory as having, collectively, a logistic-J distribution.

The sub-alpine community of forest plants that supported the geometric distribution today, may well betray it utterly 50 years later, favouring the log-series distribution instead, perhaps. With this view of natural communities in mind, none of the conclusions about proposed distributions reached by this method can be accepted as anything but coincidences. Made 30 years earlier on corresponding data, the order of the four distributions might well be scrambled. In view of this unfortunate methodology, one could substitute any four distributions one liked -- made up for the occasion -- then rewind the historical tape and watch very much the same papers appear, favouring one distribution or another or perhaps attempting to unify them! To put the same point more precisely, it is certainly possible to submit all 125 datasets in the metastudy reported in Chapter 7 to a simple test. It would be a bit labour-intensive, but one could compare each of the datasets with each of the distributions that have appeared in the literature to date, including the logistic-J distribution as well. It is, after all, an elementary observation that every proposal will have some datasets fit that distribution better than any other. Some will have more than others, of course, and I think it likely that the logistic-J distribution would have the most.

In the remainder of this chapter and in the chapter to come, I will argue that there can be only one universal statistical descriptor of abundances in natural communities.

As for other problems with the method just described, one might also point out that none of the goodness-of-fit tests that are standard in other fields were ever employed in the assessment of these shapes in sample histograms. Indeed, a search of the relevant literature by a paid professional library researcher (Galsworthy 2004) turned up no occurrences of the phrase “Chi square” or “kolmogoroff-Smirnov”, these being names of the principle tests of this kind. Such tests take two histograms and compare them, producing a numerical measure of the degree of the similarity. In a field that yearns for quantitative treatments, why should anyone neglect a perfectly good numerical measure of similarity between theoretical and empirical histograms? The fact that such measures of similarity have not been used may explain why there is so little awareness among theorists of the enormous variation to be expected from a single source distribution. The chi square curve (See Figure 1.2) has a long tail to tell, so to speak, of the many field histograms that will fit very poorly without failing to arise from the distribution in question.

The effort to fit individual field histograms to specific theoretical distributions reminds one of the story of the child who was found one day with three buckets into which he was busily sorting pennies. The buckets were labeled “unlucky” “normal” and “lucky”. He would take a penny from a pile beside him and flip it 10 times, counting the number of heads, then placing the penny in the “lucky” bucket if it came up heads more than seven times. If it came up heads fewer

than four times, he would deposit it in the "unlucky" bucket. The remainder ended up in the bucket labeled "normal". The specific biotic source of a sample no more inheres in its field histogram than the element of luck inheres in a penny.

Given the difficulties just described, it may be suspected that the field of theoretical ecology is in the grip of an unacknowledged crisis. The failure to appreciate the statistical behaviour of abundances in communities has become a nursery of untested proposals. The crisis is only made worse by the adoption by some ecologists of a social constructivist ethic, as enunciated by Hilborn and Mangel in *The Ecological Detective* (1997). In this milieu, "there is no correct model." Ironically, the statement was largely true up to the time of that book's publication. The notion that all "models" are more or less acceptable is nevertheless unhelpful. Indeed, I have just reviewed a prime example, in which each "model" is seen as having a part to play in the description of natural communities. At this point, the quote from Sir Francis Bacon that begins this chapter comes into play. If, in a given field of inquiry, it is not possible to be in error, then it is not possible for that field to be a science. After all, that's exactly what the quote means.

1.3. A positive test using multi-sample data

As normally applied, goodness-of-fit tests are used to reject hypotheses. As such, the rejection, when supported by the test, applies to each sample tested. Non-rejection of a particular fit does not amount to "acceptance", as such, since the tests are not symmetrical in this respect. Non-rejection, when based on a single sample or just a few, may be taken as evidence in favour of the null hypothesis, but as "evidence" it is weak and cannot be used, by itself, to establish anything. Since such tests are used here in affirmative rather than rejective mode, a great many samples, rather than just a few, are needed to establish the presence of an underlying distribution, in any case.

This last point is important enough to expand upon so that the reasons for the multi-test requirement are made plain. As normally used in curve-fitting applications, a test such as the chi square compiles the chi square statistic or "score" as follows.

$$\chi^2 \text{ score} = \sum (t_i - e_i)^2 / t_i$$

The index i ranges through abundance classes 1, 2, 3, etc., and the variables t_i and e_i represent the number of species having abundance i in the theoretical and empirical distributions, respectively. The greater the difference between the theoretical prediction (t_i) and the empirical datum (e_i), the greater is the contribution to the statistic above. Thus higher chi square scores reflect a rather poor fit to the data, while lower scores represents better fits. The same thing is true of the Kolmogoroff-Smirnov test.

Used in the normal way (Hays & Winkler 1971) a goodness-of-fit test will determine the

likelihood that a given set of empirical data *does not have* a particular theoretical distribution. In such a case a chi square score, for example, can be used to reject the distribution if it exceeds a threshold or critical value, as given in a standard chi square table. (See Chapter 4.) In such a case, the null hypothesis (that the empirical data follow a particular theoretical distribution) is rejected. In normal parlance the hypothesis is said to be “accepted” if the score is at or below the appropriate critical value. As already pointed out, such terminology does not actually mean that the data in question have the distribution under test. After all, there are infinitely many curves, each with a different mathematical formula, that would fit the data as well or better than the distribution under test. How, then, could a particular form be confirmed as the underlying distribution?




Figure 1.2. Distribution of chi square scores: bar height = number of scores

To use a goodness of fit test in confirmatory mode, one must have a great many datasets (say 100) and one must perform a goodness-of-fit test on all of them, compiling the scores themselves into a new distribution that may be compared directly with the chi square distribution itself, as in Figure 1.2. The figure shows the chi square distribution (smooth curve) with 125 test scores superimposed upon it. The height of bar over a given numerical category represents the number of chi square scores that fell within the category. Thus four scores with values in the range (4.00, 5.00] happened to fall in category 5.

The theoretical chi square curve clearly indicates the expectation that some samples will have very poor (i.e., large) scores, while others have very good ones. Indeed the actual scores thus achieved fulfill this prediction, or nearly so. It is a key observation of chi square theory that if the null hypothesis is true for a great many such scores, their average value will equal the number of degrees of freedom of the tests themselves. If all the tests were conducted at 10 degrees of freedom, for example, the average expected score would be 10.0 or very close to it. Since the chi

square envelope is based on 10 degrees of freedom, while the actual average score of 10.42 was somewhat greater than 10, the bars representing scores of the field data appear to be collectively shifted to the right. In any event, it is not possible for results to be shifted to the left by any significant amount, since this would violate Pearson's theorem. The overall fits were optimal in this sense. If other proposed distributions have test scores that are far enough away from the optimal mean to be statistically separable from it, the logistic-J distribution is clearly preferable (Dewdney 2000).

The histogram in Figure 1.2 represents the actual scores obtained in the metastudy which forms the empirical foundation for the claim of universality for the logistic-J distribution. Among the other proposals, the distribution closest to the logistic-J distribution in general shape is the log-series distribution. When put through the the same series of tests against empirical data, the log-series had a significantly higher score of 13.56, overall. It's chi square "envelope" would overlap the one shown in Figure 1.2, but it would be displaced several categories to the right. Other distributions proposed for the role of species/abundance descriptors, particularly the lognormal distribution, resemble the logistic-J far less and one does not need to subject them to the same test since their scores would be far too high to be in the running. To do as well as the logistic-J distribution, a proposed alternative must resemble it very closely indeed. I will return to the metastudy in Chapter 7, providing all the detail necessary for its evaluation as a research method.

It is an interesting historical fact that the British biologist C. B. Williams thought that the "hollow curves" he was seeing in lepidopteran light-trap data collectively resembled hyperbolae. (Williams 1964). He appears to have been right. His statistical colleague, R. A. Fisher, talked Williams out of the hyperbola, declaring that it was unsuitable for use as a statistical distribution, since it had a nonfinite area under it. Fisher was also right, but the idea of a displaced hyperbola truncated by its axes obviously never occurred to him. Given that the subtleties of randomly fluctuating populations would not become apparent until computers were widely available, Fisher can hardly be blamed for this. (See next section.) In any event, it was not Fisher who developed goodness-of-fit tests, but his arch-rival (statistically speaking), Carl Pearson. Fisher undoubtedly knew about goodness-of-fit tests, but seems to have avoided them for some reason. Ironically, in spite of Williams' insight, it may have been here that the pattern of inappropriate testing in theoretical ecology first developed.

The foregoing provides some background to the development of the logistic-J distribution that may be useful in clarifying its role in species-abundance studies. In the next section I will provide a mathematical description of the logistic-J distribution. In subsequent chapters I will show how to use the distribution in the field and to make reliable statistical inferences about communities from their samples.

1.4. The logistic-J distribution: probability density function

I would claim that the logistic-J distribution occurs universally in all natural communities (and in

all samples of them) in a collective sense. The J-prefix, as I explained earlier, indicates the ubiquitous shape of samples, while the word “logistic” indicates the presence of a finite limit to the abundances to be found in any particular community. This limit is implied by the equilibrium theorem. (See the next section.)

The logistic-J distribution is essentially the standard hyperbola, as shown in Figure 1.3, translated in the horizontal direction by the amount ϵ (epsilon) and in the vertical direction by δ (delta). If we use the variables x' and y' for the primary axis system, the formula for the hyperbola is

$$y' = 1/x'$$

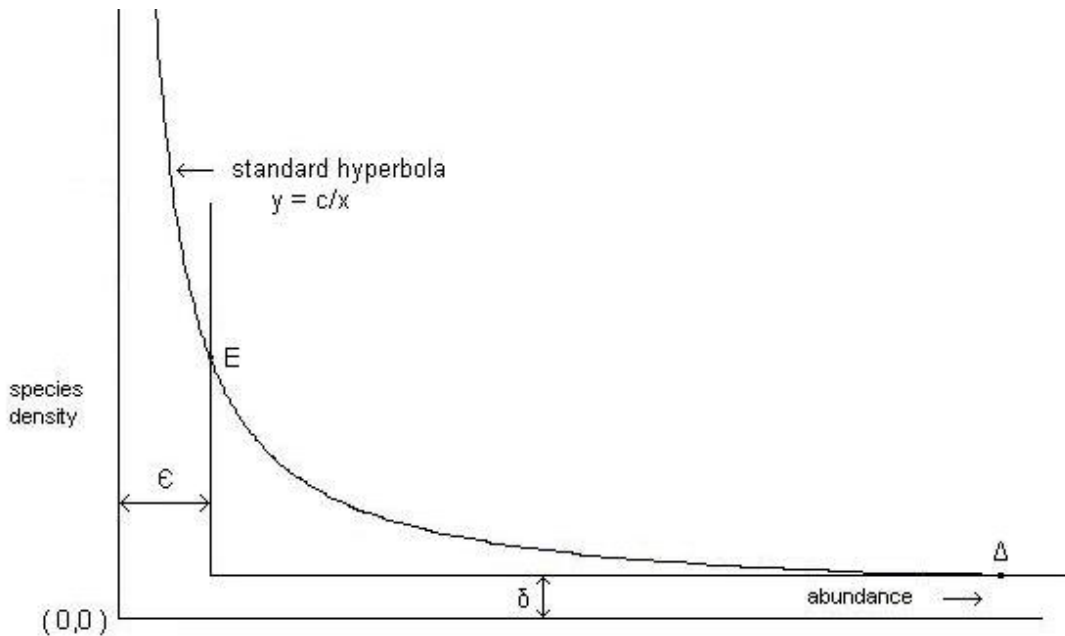


Figure 1.3. The hyperbola as the basis of the logistic-J distribution

The variables x and y for the displaced system obey the following equalities,

$$x = x' - \epsilon$$

$$y = y' - \delta,$$

and the formula for the probability density function (pdf) becomes

$$f(x) = c/(x + \epsilon - \delta); \quad 0 \leq x \leq \Delta - \epsilon,$$

$$= 0; \quad x \geq \Delta - \epsilon$$

where $\Delta = 1/\delta$ and c is a normalizing constant (a function of ϵ and δ) that gives the pdf an area of unity, as required by all probability density functions. This “density” represents the probability of finding a species with an abundance represented by an infinitesimal interval on the abundance axis. The constant c has the following formula:

$$c = (\ln((\Delta/\epsilon) - 1 + \epsilon/\Delta))^{-1}$$

In most cases, one may neglect the term ϵ/Δ , which tends to be very small.

Thus the logistic-J is a two-parameter distribution. The inverses of ϵ and δ , namely E and Δ , represent the species peak and the logistic limit, respectively. The first parameter E (when multiplied by Rc , the number of species and the constant c) bounds the number of species in the lowest abundance category, as well as all others. The second parameter Δ bounds the abundance of the largest population. Both bounds operate in a statistical or average sense.

Early in the research program that resulted in this book, a second form of the logistic-J distribution called the “special” logistic-J was tested alongside the “general” logistic-J (an early name for the distribution treated above). It was identical to the general distribution, except for being defined only over the interval $[\epsilon, \infty]$, instead of $[0, \infty]$ (Dewdney 2000). It took the somewhat simpler form,

$$f(x) = c'(1/x - \delta),$$

with the variable x in place of $(x+\epsilon)$ in the denominator of the fractional part of the pdf. The unitizing constant c' is different from c , of course. Although it resulted in much the same average scores as the more general function, it did not have the same flexibility as the general form when dealing with communities, rather than samples.

In discussing particular logistic-J distributions, the notation $LJ[\epsilon, \Delta] \times R$ will be used to refer to the distribution with parameters ϵ and Δ and having R species (whether in a community or a sample). In the context of data on hand -- samples of a community -- we can estimate the two parameters ϵ and Δ by finding a best fit for the data, as described in Section 4.3. The parameter values that emerge from the best fit will be defined as the values pertaining to the sample. The parameter Δ can be estimated independently of this process by simply taking Δ as the maximum abundance. However, the parameter ϵ has a more indirect relationship with the other extreme point of the distribution, namely the number of species in the minimum abundance category. This relationship will be explored in the next chapter.

1.5. Emergence of the logistic-J distribution

The logistic-J distribution first emerged with certainty from a computer simulation written in the mid-1990s to simulate a mutually predatory community of protists. As a longtime sampler of local waters, I had plenty of microbial data on hand. As a longtime teacher of stochastic simulation and statistics, I had the tools and background necessary to the project.

When it first occurred to me to write such a program, I had already been observing the J-curve in plots of my abundance data. I had expected something more like a normal curve, with a concentration of species about some mean abundance. Although I was well aware that no microbial community could possibly consist of mutually predatory species of, say, ciliates, it seemed at the time that if any such “community” could produce a unimodal, normal distribution of abundances, it would be this mutually predatory one.

A simulation clock governed the flow of events in the system; at each tick of the clock two individual “organisms” would be selected from the total community population. Without having to represent the organisms at all, the simulation would simply ensure that the species of the first individual had its population incremented by 1 (the species-token moving, in consequence, one unit to the right along the abundance axis). Meanwhile the species of the second individual would have its population decremented, with its species shifting one position to the left. In such a system, the total number of individuals would be preserved by the “predation” operation.

The original purpose of the simulation was to demonstrate that such a community would achieve a balance of populations, with all species fluctuating about the same mean population size. This did not happen. Indeed, a J-shaped curve emerged every time the simulation was run. Each simulation run would begin with a single spike, with all populations occupying the same abundance category. Over time, the spike would flatten into a skewed, bell-shaped curve which itself flattened out while species piled up at the low end of the abundance scale or fled randomly to the high end. In most of these mini-experiments, I kept the average abundance low in order to encourage a more definite shape. Accordingly, I kept the extirpation switch in the “off” position in order to allow the shape to build.

During the equilibrium phase of the simulation, histograms with a low mean abundance bore a close similarity to the field samples that I had already begun to collect from biosurveys of other kinds of communities. Indeed, they were (visually at least) indistinguishable from them. Yet only a small fraction of the field histograms I had been examining involved predation of any kind. It took a considerable time to realize that the simulation was not about predation, with one organism dying while another reproduced in consequence. It was about the birth/death processes itself, in particular the equality of probabilities of births versus deaths. After all, there was no actual predation in the dynamical system. Whenever one species moved to the right, some other species would move to the left. It was a clear case of a specific model turning out to be rather general, after all.

An initial result that appeared to confirm the presence of a hyperbolic function relied on the

analysis of a slightly more general system in which a random individual was selected from the set of all individuals and either duplicated or deleted, both events taking place with equal probability. (See Section 2.3 for the Equilibrium Theorem)

After a year of studying this system, I obtained a slightly more exact result: the equilibrium solution of the system was indeed a hyperbola, but translated downward by a small amount that I decided to call “delta.” The formula was

$$k(1/x - \delta),$$

where k is a constant (Dewdney 1998a). (See Chapter 6 for the Logistic Theorem.) New simulation programs were written to explore the idea of equiprobable birth/death processes under varying conditions on the balance of probabilities governing them. The probability of birth, for example, could exceed that of death for a time, but it would eventually move in the other direction -- all at random. These more general versions of the dynamical system produced the same J-shaped curves as the earlier system. Indeed, they were indistinguishable from the histograms of field surveys that I had already been to collecting. I could hardly be blamed for suspecting that the J-curve was a universal phenomenon.

In the new simulations, each individual “organism” had an approximately equal probability of dying or reproducing at each moment. Populations fluctuated randomly but ultimately, a few became larger, while many became smaller. The result, as in the case of the predation system, was distinctly counterintuitive and illustrates the dangers of armchair theorizing. At equilibrium, which never failed to develop, the histogram would be frozen and compared with a hyperbola (visually). The fit was often rather good, taking into account normal statistical fluctuations. Since the net effect of the births and deaths in this program was to preserve the total “biomass” of the system, the logistic limit could be applied and the same formula derived from the predation system applied here, as well.

At this point a general mechanism suggested itself. Called the stochastic species hypothesis, it postulated equal probabilities of birth and death over small periods of time, with corresponding changes in population over longer periods.

Following an intensive literature review of other species-abundance models (a disappointing experience), a massive study was launched. With the aid of graduate students in our Department of Biology, I began a random collection of biosurvey papers, ensuring that all the major classes of biota were covered as they went. Each of the resulting histograms was compared with a version of the logistic-J distribution that shared the same mean and height of initial peak; given values for these quantities, the values of the parameters ϵ and Δ are readily determined. (See Transfer Equations in Appendix A.1.) The resulting chi square scores were normalized to 10 degrees of freedom to make cross-comparisons possible. (See the scores histogram in Figure 1.1) The average score that emerged from the study was 10.4, very close to the optimum score of 10.0 and

well separated from the average score that emerged from parallel tests of the distribution that most closely resembled the logistic-J curve. The log-series distribution mentioned in the previous section scored over 13. It was unnecessary to test other distributions, since they resembled the logistic-J distribution even less. Only something with a very similar shape could hope to score as well (Dewdney 2003). In short, there is no “room” for another distribution in view of the optimality of the scores. The results of this metastudy amounted to what can only be called “strong support” of the stochastic species hypothesis and therefore of the logistic-J distribution that emerges from it as a universal natural phenomenon.

1.6. Fluctuating populations

The stochastic species hypothesis (see section 6.1.1 for a full description) might seem, at first sight, to justify a great throwing up of the hands: if it’s all random, what’s the point of trying to tease out ecological mechanisms to account for population changes? The answer is simply that the stochastic species hypothesis provides a framework of understanding within which such studies may be pursued. Populations fluctuate for a huge variety of reasons, each with its own mechanism. The net effect of these mechanisms is effectively random, even though one mechanism may dominate during one period of time, another later. Weather is the ultimate determinant of many of these mechanisms and the weather itself is known to be subject to the grand mechanism of chaos (Gleick 1985) which itself produces events that are effectively random.

Randomness in natural populations has long been suspected, but has faced an uphill struggle in the academic forum against mechanistic theories which can be classified into a continuum. The first view is that all species maintain more or less fixed populations (Marsh 1865), a common 19th century understanding of nature. Amazingly, this view persisted to some degree throughout most of the 20th Century. For example, in a book on population genetics Wallace (1981) declares more than 100 years later, “The second thing we can say about populations is that, despite temporal fluctuations, in the long run they tend to remain constant in size.”

This “common understanding” was replaced by the notion of predator-prey cycles, as predicted by the Lotke-Volterra equations (Leigh 1968). Backed by famous datasets such as the Hudson’s Bay trapping data, the theory gained wide acceptance and is still believed by many ecologists today, in spite of experimental evidence that such cycles don’t always develop, e.g. the Gause experiments (Gause 1934) as described in (Botkin 1990). However, a more modern view, that of “density regulated” populations, held that populations did, indeed, fluctuate randomly, but only within limits imposed by the density (relative abundance) of a species (Smith 1935). Despite the extensive literature on density-regulation that has developed since Smith’s time, this concept of population behavior is not universally accepted. Thanks in part to the popularity of chaotic population dynamics (May 1974) and in part to the failure of the population-regulation school to come to firm conclusions (Cappuccino and Price 1995), there would appear to be a growing suspicion that fluctuations in natural populations are indeed random, the same view pursued by

Hubbell (2001). An interesting example of accommodation between the two views addresses the concept of “density-vague” behaviour in which populations are bounded away only from extreme density and extirpation (Strong, 1986). Not surprisingly, the literature is sprinkled with negative results, as in the comparison of a great many populations with random walks that yielded very little difference. (den Boer, 1991).

About the problems of coming to grips with population fluctuations , Botkin (1990) writes:

“At the heart of the issues are ideas of stability, constancy, and balance, ideas intimately entwined in theories about nature. Perhaps one reason that the deficiencies of the theories were not examined or tested adequately by observation in the field -- out in nature -- was that ecologists were typically uncomfortable with theory and theoreticians. Doing science and creating theory were commonly distinguished as separate activities. Although theory was typically considered not to be necessary or important to the practicing ecologist, . . . theory played a dominant role in shaping the very character of inquiry and conclusions about populations and ecosystems (i.e., about nature). As Kenneth Watt wrote in 1962, ecologists had tended to believe that their science had lacked theory, while in fact it had ‘too much’ theory -- in the sense that the theory had been utilized and was influential even though it was not carefully connected to observations. ‘Field ecologists,’ those making measurements and observations in the forest and field, generally did not understand the mathematics of the logistic and of the Lotka-Volterra equations. But since physicists and mathematicians had the highest status among scientists, and since what physicists and mathematicians generally said was generally right, field ecologists tended to regard the logistic and the Lotka-Volterra equations as true. Lacking the understanding to analyse and thereby to criticize these equations, they accepted them on the basis of authority.”

The logistic-J distribution breaks this deadlock by implying that notions of stability are best applied at the community level and not at the level of individual populations. Species are constrained only at the high abundance end by logistic influences embedded in the myriad interactions among species within a community, between communities, and between species and the physical environment.

1.6.1 Perturbation theory

A key conceptual tool in the exploration of fluctuations in the populations of a community is the idea of a perturbation. The concept is implicit in the proof of Pearson’s theorem (MIT, 2006) that chi square test scores of random samples arising from a given source distribution must follow the chi square distribution. I have made it explicit here in order to play a role in the application of Pearson’s theorem

Let F be any finite, nonnegative, real-valued function defined over the integers 1, 2, 3, etc.. In this

context, the function F will be called the *source distribution*. Let $V(k)$ be a real-valued random variable having a normal distribution with mean 0 and standard deviation $\sigma(k)$. Then the random variable $F(k) + V(k)$ will be called a *perturbation* of the k th column. A *perturbation of the function* F will be the multivariate function consisting of perturbations of its values $F(k)$. In practice, a perturbation of F will consist of specific values for $V(k)$ selected by a random number generator, while the function F may be a theoretical distribution for a community or a sample. It may also be a histogram, whether real- or integer-valued, of a community or a sample. The standard deviations $\sigma(k)$ are defined on the basis of $F(k)$, the variance being proportional to

$$1 - k/N.$$

Perturbations enable one to give a specific meaning to variations within communities or samples of them. For example, I will claim as an hypothesis that the histograms of all natural communities (could we but know them) are perturbations of the logistic-J distribution in a collective sense, just as I shall claim that all samples of them are perturbations of the logistic-J distribution, albeit with different parameter values. A single random perturbation has no particular relevance or meaning in this context. In the case of multiple random perturbations of a particular source distribution F , however, a specific collective behaviour is expected: a goodness-of-fit test applied to the perturbations in question should have scores that reflect the test in relation to F . The scores for a chi square test should follow the chi square distribution optimally, with an average that supports the null hypothesis. This requirement will be made clearer in Chapter 7, where the metastudy mentioned in Section 1.3 above is explained in detail.

This hypothesis can be applied directly to the notion of fluctuating populations described in the foregoing section. Over time, the histogram of an entire community might undergo slow changes, with some functional values $F(k)$ increasing, others decreasing. Such changes would merely reflect the result of fluctuating populations, some increasing, others decreasing. Random motions along the abundance axis are reflected in random changes in the bar heights that represent the function F . When a species of abundance k adds or subtracts an individual, $F(k)$ decreases by unity, for example. Multiple population changes would therefore induce multiple changes in the histogram. The hypothesis implies that at any given moment the histogram of a community will be the perturbation of the corresponding logistic-J distribution. Over time, such perturbations, whether they change quickly or slowly, will follow the logistic-J distribution, the shape being not a property of individual perturbations, but of their collectivity. The more such histograms could be observed, the more they would come to resemble the logistic-J distribution and no other. As already mentioned, that is the thesis supported by the metastudy.

When one samples a community, there are two sources of variation in the histogram of the field sample. The first source lies in the community itself, with its already-perturbed version of the logistic-L distribution present in the community histogram. The second source lies in the vagaries of the sampling process itself, with samples taken closely together in time showing what are typically the variations inherent in the sampling process. If the samples are taken far apart in

time, however, both sources of variation would come into play and the samples would tend to resemble each other less. Such a claim hardly needs any support, since no one has trouble believing that both sources of variation exist. In the context of what I call exact ecology, however, these sources must be taken explicitly into account. A sample, one might say, is the perturbation of a perturbation.

1.7 The value of a uniform methodology

In an age when the natural environment is under threat from habitat loss, pollution, and climate change, it is more important than ever that ecologists develop and employ uniform procedures for assessing the condition of all the canaries in the mine, so to speak. Unless results can be compared directly, confusion is likely to result.

The lack of a universally accepted (theoretical) species abundance distribution hampers the field even more greatly. I have mentioned some of the principal proposals in this introduction, along with a few others. Attempts to “unify” these proposals, from Sugihara (1980) to Hubbell (2001) are not likely to be helpful since the presence of these distributions was never demonstrated in the first place and would be unlikely to pass a metatest, as described in Chapter 7.

A widely accepted, and well established distribution would be greatly preferable to the confusion of tongues that now characterizes the field. I would of course propose the logistic-J distribution for this role. However, should it turn out to be incorrect in the end, it can only be said that the replacement distribution will have to be mathematically very similar, if not a simple modification of it. In any event, to have the present proposal roundly rejected in a purely scientific manner would of course be a disappointment, but one balanced by the pleasure of seeing *at least one* theory so rejected!