

# Diversity analysis: a fresh approach

## Part 1: Overview of diversity and estimation theory

### Chapter 1. Introduction to diversity theory

#### *The many faces of diversity*

In any science it is important to find meaningful measures that describe the systems under study. These measures will be the raw material for theories, and will allow meaningful comparisons between systems. Much of the success of a science depends on the choice of basic measures; poor choices can cause a field to stagnate, while insightful choices can make it easy to find regularities in nature. In this regard diversity analysis in ecology has always been a polemical field, with many differing opinions about how to best characterize biological diversity. There have been so many arguments about how to measure diversity that some authors (for example Hulbert 1971) have given up on the matter and concluded that there is no such thing as the concept of "diversity" per se. Indeed, the variety of diversity indices in the literature is bewildering. When these different diversity indices are applied to a system with  $S$  equally common species, some give  $S$ , some give  $\log S$ , some give  $1-1/S$ , etc. The mathematical behaviors of these indices are so different that there are bound to be problems when they are treated as if they were interchangeable in formulas or methods that use "diversity" as a variable. A method that might give good results with one index is likely to give meaningless results with some other indices. In a complex science like biology, the fuzziness of the diversity concept may not be a bad thing; different meanings of the term "diversity" may be useful in different situations. Nevertheless, it is important to be aware of the fuzziness, else it leads only to confusion.

Imagine a group of engineers who need to calculate the volumes of spherical objects for a construction project. Some of these engineers use the diameter of an object as their "index of volume", while others like to use the surface area as their index of volume, and still others use the actual volume. The engineers call each of their indices by the same word, "volume". After all, the engineers argue, each of the indices is somehow related to volume and it is easy to convert each to any of the others, so why not call them by the same word? The engineers that used diameter as their "volume" preferred it because it was less sensitive to measurement error than the indices that relied on the square or cube of the diameter. The partisans of the other indices also had reasons for their preferences. Is it likely that this group of engineers would be able to design a working plane or ship as long as they insisted on calling all these different quantities by the same word, and using them interchangeably in formulas requiring "volume"?

In ecology today we are like those imaginary engineers. We have created a confusing and messy field called diversity analysis, but the field is not really messy or confusing at all. We have made it so by confounding many distinct aspects of complexity and serving them up interchangeably under the single word, "diversity". All of these aspects are clearly related to diversity, but which of them, if any, match our intuitive concept of "diversity"? If we are not careful to keep these different aspects of compositional complexity separate, and if we use them interchangeably in general formulas requiring diversity, we are bound to end up with strange results.

For example, the Gini-Simpson index  $1-\sum p_i^2$  is often recommended as a diversity index (Magurran 2004, Lande 1996) and equated with "diversity". It measures the probability that two sampled individuals belong to different species, but this aspect of compositional complexity does not seem

to coincide with a biologically meaningful concept of diversity. Imagine a community with a million equally common species. Its Gini-Simpson index is 0.999999. A plague attacks this community and eliminates 99.99% of the species, leaving only a hundred species untouched. Ecologists and conservation biologists would consider this a huge drop in diversity, both in terms of the variety of interactions experienced by a constituent organism before and after the plague, and in terms of the conservation value of the pre- versus post-plague communities. The tools of diversity analysis should be able to clearly indicate the magnitude of this drop if they are to be useful in less dramatic situations. Yet the post-plague community's Gini-Simpson index is 0.99, only 1% less than the Gini-Simpson index of the pre-plague community. An ecologist equating the Gini-Simpson index with "diversity" would therefore conclude that the plague which killed almost all the species did not have a big effect on the community's "diversity"!

Many other often-used indices of diversity are also poor matches for a biologically meaningful definition of diversity. The popular Shannon entropy  $-\sum p_i \ln p_i$  (often called the Shannon-Wiener or Shannon-Weaver index in biology) exhibits the same problem as the Gini-Simpson index, though to a lesser degree. If we start with a continent that has 30 million equally common species, and a meteor impact wipes out half the species, this index drops by only 4%. An ecologist following standard textbook methods would therefore infer that the continent's "diversity" has hardly been affected by this meteor which removed half of its species.

Some biologists are not concerned that these indices give difficult-to-interpret results. These biologists think that the actual values of the indices are unimportant, as long as they can be used to calculate the statistical significance of the drop in diversity after the plague or meteor or whatever. On this view, which is common in the diversity literature, the statistical significance of the result is the final product of a study. A moment's reflection, however, shows that this is not reasonable. The statistical significance of a change in the diversity index has little to do with the actual magnitude or biological significance of the change, which is the really important scientific question. To help make this clear, consider the example of tossing a coin  $N$  times to see if it is biased. A good measure of the *statistical significance* of the coin's bias is the  $p$ -value obtained (using the binomial distribution and the null hypothesis that the mean proportion of heads is 0.5). However, if  $N$  is large enough, even the most miniscule deviation from a mean number of heads of 0.5, say 0.50001, can result in a highly significant  $p$ -value, say 0.0001. This highly significant  $p$ -value proves that the coin is biased but says nothing about the size or practical importance of the bias. A measure that reveals the actual magnitude of the bias is the excess proportion of heads obtained. If the mean number of heads is 0.50001, the excess proportion of heads is 0.00001, a bias which is unimportant for most purposes, but if heads come up 75% of the time, the bias cannot be ignored even in a short trial. Ecologists are often guilty of stopping after calculating the  $p$ -value, when the more interesting question is the real size of the effect, the equivalent of the excess proportion of heads in the coin toss. However, in order to judge the actual magnitude of a change in diversity, we need measures that behave more intuitively than the indices discussed above. The aspects of compositional complexity measured by these indices are not good matches for what most biologists mean by "diversity", and do not help us judge the real magnitude of an effect.

#### *True diversity of order $q$*

Our imaginary engineers of the preceding section would do well to distinguish indices of volume from the volume itself. Diameter is an index of volume in the sense that it increases when volume increases, but it is not itself the volume. If the engineers think carefully about how volume behaves, they would note that a true volume should increase by a factor of  $2^3 = 8$  when all length dimensions are doubled. The engineers could use this test to see if any particular expression is a true volume. They should then all agree to convert any of their other indices to true volumes before comparing them or using them in formulas that require volumes (for example a formula that related volume to weight--this would give disastrous results if some of our engineers insisted on inserting their favorite indices, like diameter or surface area, into it in place of the actual volume). This would clear up their confusion and lead to planes that flew and ships that didn't sink.

Biologists should do the same with diversity. We should try to identify some measures of diversity which behave reasonably when applied to situations like the plague and meteor examples of the preceding section, and we should then try to pinpoint the reason why they behave reasonably. This will help us make our concept of diversity more precise, and give us a way of distinguishing diversity from the other things that covary with it. Several diversity indices do in fact behave reasonably. One example of a well-behaved and intuitive index is species richness (Chapter \_\_). If we apply it to our meteor example (30 million equally common species pre-impact, reduced to 15 million after the impact) we find the meteor has reduced this index by 50%, a biologically reasonable answer. Another example of a reasonably behaved index was discovered by Robert MacArthur, one of the first ecologists to introduce information theory to biology. He noticed that the exponential of Shannon entropy (Chapter \_\_) more closely matches ecologists' intuitive concept of diversity than the raw Shannon entropy (MacArthur 1965). For the meteor example the exponential of Shannon entropy gives the same reasonable answer as species richness, a drop of 50% after impact. Yet another index which behaves reasonably is the inverse Simpson concentration,  $1/(\sum p_i^2)$  (Chapter \_\_), giving the same answer as species richness and the exponential of Shannon entropy, a post-impact drop of 50%. These three indices also agree on a reasonable answer in the plague example above: if a million equally common species are reduced to a hundred equally common species by a plague, the value of each of these indices drops to 1/10000 (a hundred divided by a million) of its pre-plague value.

Hill (1973) was the first to realize that these indices give intuitive, reasonable answers because they all possess something he called the "doubling" property. He noted that the aspects of complexity which most closely match an ecologist's or conservation biologist's intuitive concept of diversity are those aspects which double when two identically-distributed but distinct communities (no shared species) are added together in equal proportions. Put another way, the diversity of any community should double if every species is divided into two equal groups (say males and females) and each group is considered to be a distinct species (Hill 1973). This doubling property is biologically reasonable, since many biologically meaningful aspects of complexity double when communities are doubled in this way. For example, the variety of interactions an organism could experience in each encounter with another organism doubles in the doubled community. Likewise, if all species are treated as equally important, the conservation value of the doubled community is twice that of the first. More important, measures which possess Hill's doubling property are guaranteed to behave reasonably when combined into ratios, because such measures scale linearly when a community is doubled in the sense just described. If Community A has a true diversity four times as great as Community B, then we would need to combine four versions of Community B (each identically distributed but each with distinct species) to make a community with the same diversity as Community A. This linear yardstick in terms of community composition makes it easy to interpret diversity measures which have this property. If one of these indices doubles, it is equivalent to adding another community of the same size and species frequency distribution. It is easy to visualize what this means and easy to decide whether the change is biologically significant or not.

Ratios are used throughout diversity analysis, but unless the diversity measures forming these ratios have Hill's doubling property, they will not behave as expected. For example, a common method of judging the compositional similarity of two communities is to examine the ratio of the alpha diversity (average diversity of the individual communities) divided by the gamma diversity (total diversity of the pooled communities). If the communities were identical, this ratio would give unity. If the communities were completely distinct, ecologists have assumed that this ratio would be small. However, this assumption only holds if the diversity measure used in calculating the ratio scales linearly with respect to community composition. Let's examine what happens when we try to use various diversity measures in this ratio of alpha over gamma, to calculate the similarity of two equally large, identically distributed but completely dissimilar communities (that is, the two communities share no species). For definiteness let's take two communities with 100 equally common species each; remember, there are no species in common between the two communities:

First let's use diversity measures that have the doubling property, like species richness. For species

richness the ratio is 1/2 when two equally large communities are completely distinct, because the species richness of the pooled communities (200 species) is twice the species richness of the average single community (100 species). For species richness, then, this ratio is an easily interpretable similarity measure, ranging from a well-defined minimum of 1/2 indicating complete dissimilarity, to a well-defined maximum of unity indicating complete similarity. For the communities in this example any other measure that has the doubling property, such as the exponential of Shannon entropy or the inverse Simpson concentration, will also give a ratio of 1/2 if the two communities are completely dissimilar, because the denominator of the similarity ratio is the diversity of a community formed by composing two equally large, equally diverse, identically-distributed but distinct communities. For any measure that has the doubling property, this diversity is twice the diversity of the individual communities.

Now do the same with the Gini-Simpson index, which does not have the doubling property. The average Gini-Simpson index of the two communities is 0.99, and the Gini-Simpson index of the pooled communities is 0.995. Their ratio is 0.995, very close to 1.00, the value that would indicate complete similarity. Yet the communities are completely *dissimilar!* Indeed, for the Gini-Simpson index, a similarity ratio of 0.99 could arise from two almost identical communities, two completely distinct communities, or anything in between. The ratios of any other indices that lack Hill's doubling property give the same kind of misleading results when applied to diverse communities; this happens because most traditional indices do not increase linearly if two equally-large communities are pooled; the denominator in the ratio will therefore not be twice the numerator.

Diversity measures that possess the doubling property deserve to be distinguished from indices that measure other aspects of compositional complexity. As shown in the examples above, their intuitive and biologically reasonable behavior sets them apart from other indices. The doubling property is the biological equivalent of our imaginary engineers' test for a true volume measure. We will call a diversity measure which possesses the doubling property, and which is normalized so that it gives unity when applied to a community consisting of a single species, a "true diversity". Other diversity measures are just indices of diversity, not themselves diversity, in the same way that diameter and surface area are indices of volume but not really volume. By making this distinction, we avoid many of the confusions that have long plagued diversity analysis; true diversities provide a simple, powerful, and intuitive framework for a fresh approach to diversity analysis.

We have mentioned three true diversities: species richness, the exponential of Shannon entropy, and the inverse Simpson concentration. In 1971 Shannon's concept of entropy was generalized by Renyi into a one-parameter family of "generalized entropies" (see Chapter 1). Just as MacArthur took the exponential of Shannon entropy to produce a more intuitive measure of diversity than the Shannon entropy itself, Hill (1973) took the exponential of Renyi's family of generalized entropies in order to obtain a whole family of diversity measures that possess the intuitive "doubling" property: Eq. 1.

$${}^qD(p_1, p_2, \dots) \equiv (\sum p_i^q)^{1/(1-q)}$$

(with  $q$  nonnegative). These functions are often called "Hill numbers" (e.g. Magurran 2004). Because they possess the doubling property and equal unity when there is only one species, they are all true diversities according to our test. In fact they possess a stronger additive property: if any  $N$  completely distinct communities, each with true diversity  $X$ , are added together in equal proportions, the true diversity of the pooled communities is  $N \cdot X$ . The species frequency distributions in the  $N$  communities do not have to be identical. This strong additive property ensures that these measures match our intuitions about diversity. These kinds of functions are not unique to biology; they play a fundamental role in thermodynamics, where they give the number of states available to a system of given entropy, and they are also used in economics, where they are called the "numbers equivalent" of a diversity index (Patil and Taillie 1982, Adelman 1969). Unfortunately Hill's profound paper introducing this idea to ecology was not sufficiently appreciated by biologists.

Hill's formalism unexpectedly reveals that species richness, the exponential of Shannon entropy, and the inverse Simpson concentration are all fundamentally the same measure, differing only in the parameter  $q$  in Eq. 1. When  $q = 0$ , Eq. 1 gives species richness, since raising the species

frequencies to the zeroth power gives unity for each species in the summation. When  $q = 2$  Eq. 1 gives inverse Simpson concentration. When  $q = 1$  Eq. 1 is undefined because of division by zero in the exponent  $1/(1-q)$ , but its limit as  $q$  approaches unity exists, and it equals  $\exp(H_{\text{Shannon}})$ . Thus biologists do not need to borrow Shannon entropy from information theory or physics; on the contrary, much of information theory follows from our analysis of diversity measures. In Chapter \_\_ we will show that the formalism used here not only generates standard information-theoretic measures but also their more recent generalizations, and sheds new light on some aspects of this generalized information theory.

Hill noted that all true diversities are based on the mean frequency of species in an ecosystem--they are just the reciprocal of mean species frequency. There are many ways of taking the mean of species frequencies, and the different ways produce the different members of the family of true diversities. For example, the simplest mean is the expectation value of the species frequencies:

Eq. 2. 
$$F = E(p_i) = \sum p_i \cdot p_i.$$

It is easy to see that this mean frequency drops by a factor of 2 when the ecosystem is doubled in Hill's sense, regardless of the frequency distribution of the original ecosystem:

Eq. 3. 
$$F(\text{doubled community}) = 2 \sum (p_i/2) \cdot (p_i/2) = F/2.$$

The reciprocal of this mean frequency, the inverse Simpson concentration, therefore doubles in the doubled community, and is a true diversity. It has units of "effective number of species" or "equivalent number of species", the number of equally common species needed to produce an ecosystem with the same mean frequency  $F$  as the original ecosystem. Hill noted that the other common true diversity measures -- species richness and the exponential of Shannon entropy -- correspond to the reciprocals of the  $p_i$ -weighted harmonic mean and  $p_i$ -weighted geometric mean of the species frequencies, respectively:

Eq. 3. 
$${}^0D = 1 / [p_1(1/p_1) + p_2(1/p_2) + p_3(1/p_3) + \dots]^{-1}$$

Eq. 4. 
$${}^1D = 1 / [(p_1)^{(p_1)} (p_2)^{(p_2)} (p_3)^{(p_3)} (p_4)^{(p_4)} \dots]$$

The precise meaning of these special averages is discussed in the chapters on individual diversity indices (Chapters \_\_).

More generally, Kolmogorov (1930) and Nagumo (1930) defined the generalized expectation value of a random variable as

Eq. 5. 
$$\text{Generalized mean frequency} = f^{-1}(E(f(p_i))).$$

By comparison with Eq. 1, we see that all of Hill's true diversities of order  $q$  are the reciprocals of the generalized expectation values of species frequency, using  $p^{q-1}$  as  $f(p)$  in Eq. 5. (The result for  $q = 1$ , Eq. 4, follows from Eq. 1 in the limit as  $q$  approaches unity.)

We follow Hill in calling the value of  $q$  in Eq. 1 the "order" of the diversity. The order of a true diversity determines how much weight the measure gives to each species' frequency in the sample. When  $q$  equals zero, the measure is completely insensitive to frequency (since  $p^0 = 1$  regardless of the value of  $p$ ). This is species richness, the true diversity of order zero. Values of  $q$  between zero and unity disproportionately emphasize rare species relative to their frequency in the community, while values of  $q$  greater than unity disproportionately emphasize dominant species. The crossing point, the value of  $q$  that weighs each species exactly by its frequency in the community, is  $q = 1$ . Because this measure with  $q = 1$ , the exponential of Shannon entropy, weighs species exactly according to their frequency in the community, it is the most appropriate general-purpose measure of diversity. It is also the only true diversity that can always be partitioned into independent components such as alpha and beta diversity. It will be discussed in more detail in Chapters \_\_ and \_\_.

Another way of thinking about the role of the order  $q$  is in terms of the relative importance of "richness" and "evenness" to the overall diversity. "Richness" refers to the simple count of the number of species in the community, without regard to their frequencies. "Evenness" refers to the equitability of the species frequencies in a community; a community in which all species are equally common is a perfectly even community, while a community dominated by just a few species is highly uneven. Both richness and evenness contribute to the compositional complexity of a community. The evenness of a community affects complexity because very rare species have little contact with the majority of the individuals in an ecosystem, and do not contribute much to the

variety of interactions in that ecosystem. If rare species become slightly more common, this would increase the average variety of interactions among individuals. However, if a few species became superabundant, this would reduce the average variety of interactions among individuals, since so many of the interactions would be between the same superabundant species. For a given total number of species, the point at which the variety of interactions is maximized is the completely even community, in which all species are equiprobable. The same argument applies if, instead of the variety of interactions, we had used the average uncertainty in the identity of a sampled species, or any other frequency-dependent aspect of compositional complexity. Indeed, this may be considered a defining characteristic of all increasing measures of compositional complexity.

For a community of  $S$  species, true diversity of any order  $q$  is maximized when that community is perfectly even (and the value of the true diversity is then just  $S$ ). If the perfectly even community with  $S$  species becomes uneven, the true diversities of low order  $q$  would not drop much (indeed the true diversity of order zero, species richness, would not drop at all), but the true diversities of high order  $q$  would drop significantly. This reflects their different emphases on the richness and evenness components of diversity. The true diversity of order zero is completely insensitive to species frequencies and only reflects the richness component of diversity. The true diversity of order 1 balances richness and evenness. True diversity of order 2 emphasizes evenness more than richness. As  $q$  approaches infinity, the true diversity becomes completely insensitive to the richness component.

A good way to show the relative contributions of evenness and richness to a community's diversity is through the community's "diversity profile", which is a graph of true diversity versus  $q$ . The graph starts on the left with the community's species richness (the true diversity of order zero). This reflects only the richness component of diversity. The graph's value at  $q = 1$  reflects a balanced combination of the richness and evenness components of diversity. The graph eventually drops down to the Berger-Parker index (the true diversity of order infinity) on the extreme right; this measure reflects only the evenness component of diversity. It is sufficient to graph the diversity profile from  $q = 0$  to  $q = 5$  or  $6$ ; beyond this, there is usually very little change in the graph. The shape of this graph conveys the amount of dominance in the community; a completely even community has a perfectly flat diversity profile with no drop at all, while a community with a high degree of dominance has a diversity profile that drops steeply as  $q$  increases. Diversity profiles are the most complete way to convey diversity information about a system. We will return to them in Chapter 10.

The quantity  $\sum p_i^q$ , which is at the heart of the definition of the true diversity of order  $q$ , plays a central role in diversity analysis and requires its own symbol. By analogy with Simpson concentration  $\lambda = \sum p_i^2$ , we define

$$\text{Eq. 6.} \quad {}^q\lambda = \sum p_i^q$$

In statistics and information theory this is sometimes called the  $q$ -th frequency moment. It is the basis not only for the true diversities of order  $q$  but for almost all other diversity indices used in the sciences. Even information-theoretic measures that appear not to contain  ${}^q\lambda$  are the limits of certain functions of  ${}^q\lambda$  as  $q$  approaches unity. All measures based on  ${}^q\lambda$  share common properties, which can be exploited to derive fundamental index-independent relations. We will refer to all diversity measures that are monotonic functions of  ${}^q\lambda$ , or limits of such functions as  $q$  approaches unity, as "standard diversity indices". These include species richness, Shannon entropy and its exponential, the Gini-Simpson index, the Simpson concentration and its reciprocal, the Hurlbert-Smith-Grassle index for  $m=2$ , Patil and Taillie's general diversity index, all Tsallis or HCDT entropies, all Renyi entropies, Varma entropy, Arimoto entropy, Sharma-Mittal entropy, and others. Since they are all based on  ${}^q\lambda$ , their sensitivity to rare or common species depends on the value of  $q$ , just as with true diversities, and they can be classified as diversity indices of order 0, order 1, order 2, etc, just like true diversities. For example, the Gini-Simpson index, inverse Simpson concentration, and second-degree Renyi entropy are all diversity indices of order 2. All diversity indices of a given order share the same behavior with respect to rare and common species.

If two communities share the same value of a given index, those communities are identical with respect to the aspect of compositional complexity measured by that index. Diversity indices thus

establish equivalence relations among communities; the set of all communities which share a given value of the index form an equivalence class. For example, all communities with a Gini-Simpson index of 0.98 are equivalent with respect to the aspect of complexity measured by this index (which is the probability that two sampled individuals belong to different species). Since the species frequencies only enter into the calculation of this index through  ${}^2\lambda$ , communities which share a particular value of the Gini-Simpson index must all share the same value of  ${}^2\lambda$ ; the equivalence classes determined by this index are really equivalence classes based on the values of  ${}^2\lambda$ . Likewise the inverse Simpson concentration creates equivalence classes, based in this case on shared values of  $1/{}^2\lambda$ . (which measures the expected value of the species frequencies). However, since the species frequencies only enter the calculation of the inverse Simpson concentration through  ${}^2\lambda$ , the equivalence classes formed by this index are also really equivalence classes determined by shared values of  ${}^2\lambda$ . The same is true for any other diversity measure of order 2, including true diversities; all standard diversity indices of order 2 have identical equivalence classes, based on the "level surfaces" of  ${}^2\lambda$ . Therefore any two communities that have identical true diversities of order 2 also have identical Gini-Simpson indices, inverse Simpson concentrations, second degree Renyi entropies, Hurlbert-Smith Grassle index with  $m = 2$ , and all other standard diversity indices of order 2. This means that even though each of these indices measures a different aspect of compositional complexity, these aspects are fundamentally related and are completely determined by the true diversity. All standard diversity indices of a given order produce the same "level surfaces" and differ only in the way they label these level surfaces. It is therefore reasonable to standardize on the labelling system that gives the most intuitive results, the true diversities; in doing so we are not ignoring these other aspects of complexity but rather converting them all to common units.

It is straightforward to find the transformation that converts a particular standard diversity index  $H^{(q)\lambda}$  to a true diversity  ${}^qD^{(q)\lambda}$ . As we have just explained, the equivalence classes of any standard diversity index of order  $q$  are the same as the equivalence classes of the true diversity of order  $q$ . Suppose we have an equivalence class consisting of all communities whose diversity index  $H^{(q)\lambda}$  is equal to  $X$ ; what is the true diversity of this set of communities? One member of that equivalence class is a community with all its species equally common. We know from the definition of true diversity that the true diversity of that community is just  $S$ , the number of equally common species it has. Therefore the true diversity of any community in that equivalence class is  $S$ . Finding  $S$ , the number of equally common species which give the value  $X$  of the diversity index  $H^{(q)\lambda}$ , is a matter of simple algebra. For example, suppose we want to transform the Gini-Simpson index  $1 - \sum p_i^2$  to a true diversity. If the value of the Gini-Simpson index is  $X$ , we need to find the value of  $S$  that gives a Gini-Simpson index of  $X$ . We first write an expression for the Gini-Simpson index applied to a community with  $S$  equally common species; this is  $1 - S \cdot (1/S)^2 = 1 - 1/S$ . Then we set this equal to  $X$ , and solve for  $S$ :

$$1 - \sum (1/S)^2 = X$$

$$1 - 1/S = X$$

$$1/S = 1 - X$$

Eq. 7.  $S = 1/(1 - X)$ .

This  $S$  is the true diversity corresponding to a Gini-Simpson index of  $X$ . This same formula could also be obtained by writing both the Gini-Simpson index and the true diversity of order 2 in terms of  ${}^2\lambda$ : the Gini-Simpson index  $H^{(2)\lambda} = 1 - {}^2\lambda$  and  ${}^2D = ({}^2\lambda)^{-1}$ . The operations which transform  $1 - {}^2\lambda$  to  $({}^2\lambda)^{-1}$  create the transformation function: subtract 1 from  $H$ , change sign, and take the reciprocal.

Another example: Suppose we wish to convert a Shannon entropy,  $X$ , to true diversity. We find an equivalent community made up of  $S$  equally common species with the same Shannon entropy,  $X$ :

$$-\sum 1/S \ln 1/S = X$$

$$-S(1/S \ln 1/S) = X$$

$$\ln S = X$$

Eq. 8.  $S = \exp(X)$

This value of  $S$  is the true diversity corresponding to a Shannon entropy of  $X$ .

Since all standard diversity indices of a given order  $q$  are interconvertible in this way, all the

information contained in all possible standard diversity indices is encapsulated in the true diversities. This means that the diversity profile of a community (the graph of true diversity versus  $q$ ) contains all possible information about all aspects of complexity measured by any standard diversity index. Diversity profiles therefore play a fundamental role in our approach to diversity analysis. The diversity profile can be used to calculate any standard diversity index of order  $q$ , by reading the appropriate value of  ${}^qD$  from the graph and then transforming this true diversity to the desired diversity index using the inverse of the kinds of transformations discussed in the preceding paragraphs. For example, if we wanted to know the value of the Gini-Simpson index for an ecosystem, we could look up the value of  ${}^2D$  on the ecosystem's diversity profile graph (we must use  ${}^2D$  because the Gini-Simpson index is an index of order 2). The conversion of this to the Gini-Simpson index is the inverse of Eq. 7:

Eq. 9. 
$$\text{Gini-Simpson index} = 1 - 1/{}^2D.$$

Because true diversities convert all standard diversity indices into common units-- "effective number of species" -- this approach permits the meaningful comparison of different indices applied to the same ecosystem, and makes possible the unified treatment of all standard diversity indices. This unification provides the framework for most of the chapters that follow.

A diversity measure is of no use if it does not have an intuitive interpretation. One way to interpret any diversity measure is to find an easily visualized ecosystem which is equivalent to the ecosystem under study with respect to the given diversity measure. The most easily visualized community is one whose species are all equally abundant. Any biologist can imagine what a community of, say, 40 equally common species would be like, and any biologist can judge the difference between a community of 40 equally common species and a community of 10, or 60, equally common species. True diversities are precisely the number of species in these equivalent communities. As we have seen above, a biologist could also find the equivalent community for any other standard diversity index, as an aid in the interpretation of that index, and this number of equally common species in the equivalent community is given by converting the index to a true diversity (of the same order  $q$  as the index). The true diversity can be thought of as a shortcut for interpreting any standard diversity index in this way. True diversities transform standard diversity indices onto a linear, easy-to-visualize scale involving perfectly even communities, so that we can intuitively judge the importance of an effect.

It may sometimes be easier to judge the magnitude of a change in diversity by using the original community itself as the base of comparison. Suppose the original diversity is  $D_1$ , and the diversity after some treatment or event is  $D_2$ . One way to judge the magnitude of the difference between  $D_1$  and  $D_2$  is to ask how many distinct copies of the original community (each with the same frequency distribution but with completely different species) must be added together to reach a diversity of  $D_2$ ? If  $D_2$  is double  $D_1$ , we must add an entire community of the same size and with the same species frequency distribution (but with different species) as the original community. If  $D_2$  is triple  $D_1$ , we must add two distinct communities of the same size and species frequency distribution as the original. In terms of the true diversity, and the many other aspects of complexity which are monotonic transformations of it, the original ecosystem transformed in this way is equivalent to the second ecosystem.

We can think of true diversities of order  $q$  as a way of generalizing species richness to take species frequencies into account while preserving its intuitive doubling property. It is very important for diversity measures to take species frequencies into account, because almost all ecologically meaningful aspects of complexity depend on frequency, not mere presence-absence. If there are only one or two members of Species X in a large community, this species will not contribute much to its complexity, yet species richness gives this vagrant Species X as much weight as the more common species in the ecosystem. This may sometimes be justified in conservation biology (where the mere presence of a species in an area may be significant), but in ecology (where the functioning of ecosystems is important) it is not usually reasonable. Many ecologists use species richness anyway, even when frequency data is available, because of a mistrust of traditional frequency-based diversity indices. The examples above show that the mistrust may have been justified, but frequency data should not be thrown away if it is available. True diversities that take



species frequencies into account are almost as easy to interpret as species richness, and it is our hope that the theory presented here will change people's attitudes towards frequency-based diversity measures.

Apart from supplying easily interpreted practical tools for diversity analysis, true diversities also provide us with a powerful new mathematical tool for the derivation of general index-independent theorems about diversity. True diversities all have the same set of mathematical properties, and so it is possible to prove theorems which apply to all true diversities. Such a theorem can be converted into index-specific relationships by using the transformations discussed earlier. In addition, since theorems about true diversities apply (via suitable transformations) to the entire class of standard diversity indices, it becomes possible to make generalizations and impossibility proofs that apply to the entire class of such indices, including those not yet invented. These ideas will be used to solve the problem of defining alpha, beta, and gamma diversity in Chapter 2.

Many of the results we discuss in this book are related to results in information theory and its recent generalizations. This is because diversity is intimately connected with uncertainty, whose mathematical quantification was the great achievement of Shannon's theory of information. The usual biological concept of diversity is analogous to the uncertainty *per symbol* in information theory. This is the uncertainty in the identity of a symbol received over some channel, or the uncertainty in the species identity of a randomly sampled individual in an ecosystem. Information theory also deals with another kind of uncertainty, uncertainty *per unit time* or *per message*. An ecological analogue would be the total uncertainty in the identities of all the organisms found in a square or cubic meter of an ecosystem. This kind of diversity has only recently been proposed by Ricotta (2002) but has interesting applications. This subject will be discussed in Chapter 3, but it is important to keep in mind that most of this book (like virtually all of the literature on diversity) deals with the first kind of uncertainty or diversity: the diversity per sampled individual.

Likewise it is important to keep in mind that all of the theoretical relationships discussed in this chapter involve the true frequencies  $p_i$  of the species in a community. These are *not* the same as the species frequencies observed in a sample taken from the community. Unless the sample is truly enormous, some species in the community will be missed by the sample. This means that the sample frequencies will systematically overestimate the actual frequencies of the species in the community, and hence underestimate the diversity. It has been common for ecologists to calculate diversity indices using the actual sample frequencies, but much better estimates of their true values can be obtained by correcting for this inherent bias in the sample frequencies. The theoretical basis for this correction is set out in the next chapter.

Compositional complexity is important not only in biology and information theory but also in other sciences such as economics, and especially in the mother of all sciences of complexity, thermodynamics. Each of these sciences has advanced at its own rate, and each has travelled down a slightly different path, with very little communication between them. Each of these fields has something to teach the others. For example, recent advances in nonextensive thermodynamics have led to a deeper understanding of the mathematics of complexity, and some of these advances have implications for diversity analysis in ecology. These will be discussed in Chapters 4-6. Conversely, some of the advances in ecological diversity analysis presented here have interesting applications in information theory, nonextensive thermodynamics, and other sciences. Though this book is written in the language of biology and is aimed at ecologists, much of it applies equally well to these other fields (see Chapter 7), and we hope that it leads to some interdisciplinary cross-fertilization.