7 *Conservation Biology: The Measurement Problem*

7.1 INTRODUCTION

In chapters 2–6, we focused on the explanatory and predictive significance of biodiversity properties, on their roles in driving important biological dynamics. We did not entirely neglect conservation issues, but we did not focus on biodiversity properties as targets of conservation policy. The example of the food web illustrates the important connection between biodiversity as cause and biodiversity as a policy target: biodiversity properties are targets of conservation policy because biodiversity properties, and changes in those properties, drive biological dynamics of fundamental importance. Identifying the causally salient features of systems identifies the sites in those systems at which interventions change outcomes (for an eloquent and detailed articulation of this view of causation, see Woodward 2003). Interventions can be deliberate human interventions, side effects of human activities, and (of course) disturbances that are entirely independent of us. Sunspots flare, volcanoes vent, faults shift, and soils erode independently of human action. So, for example, if individualist models of ecology of the kind we discussed in the last chapter are right, the policy implications are profound. On the one hand, individualism implies that ecological communities are predominantly modular, and hence the removal of one species is unlikely to have important consequences for most other populations. On the other hand, individualism also implies that these systems can be quite sensitive to perturbations in abiotic conditions; diverting water for irrigation, or allowing nutrient-filled runoff into a wetland might utterly transform it. So the causal and predictive considerations of the last few chapters are of great importance to conservation biology. These theoretical programs, when successful, identify levers of change in biological systems. But they cannot by themselves settle

policy issues; they cannot tell us the human costs on intervention, and neither can they tell us what outcomes to aim for, and which to avoid.

In this chapter and the next, conservation biology becomes our central focus. In this chapter, we focus on measurement issues. These are difficult and controversial for two reasons. The first replays the theme of this whole book: measurement requires us to identify the explanatorily salient dimensions of diversity, because there will always be some way of comparing (say) one wetland to another that will count the first as the more diverse, and another procedure that will reverse the result. The point is the same as that made about the phenetics movement in systematics, and has the same rationale: there is no theory-neutral notion of overall richness any more than there is a theory-neutral notion of overall similarity. The second reason is that measurement procedures must be tractable. We must be able to measure features of biological systems even given the constraints on time, of resources, and information imposed on conservation projects. These resource limits seriously constrain measurement. As a consequence, conservation biologists almost never measure directly the full range of phenomena that they take to constitute the biodiversity of a system. Rather, they sample that diversity, or rely on measurable signs that vary (they believe) with biodiversity itself. Samples and signs are biodiversity surrogates, and this chapter will mostly be concerned with the evaluation of such surrogates.

While biodiversity and its protection is fundamental to the goals of conservation biology and the policies that discipline has devised, consensus on the importance of biodiversity has not been matched by consensus on the technical problem of biodiversity measurement. The last two decades have seen a proliferation of biodiversity measurement strategies, but a paucity of theory aimed at evaluating and comparing them. This proliferation is widely recognized in research volumes such as *Biodiversity: Measurement and Estimation* (Harper and Hawksworth 1995) and *Biodiversity: A Biology of Numbers and Difference* (Gaston 1996), and in textbooks on biodiversity such as *Biodiversity* (Lévêque and Mounolou 2003) and *Biodiversity: An Introduction* (Gaston and Spicer 2004). These works give thorough inventories of current measurement techniques, but are much less forthcoming on how measurement strategies ought to be compared with one another or how the success of biodiversity measurement strategies in general ought to be evaluated.

Formal conservation policy is even less useful than the technical literature in articulating a measurable concept of biodiversity. The United Nations Convention on Biological Diversity defines biodiversity in Article 2:

> "Biological diversity" means the variability among living organisms from all sources including, inter alia, terrestrial, marine and other aquatic ecosystems and the ecological complexes of which they are part; this includes diversity within species, between species and of ecosystems.

These pieties treat "biodiversity" as a synonym for "all living things." Such a definition is of little use to conservation biologists trying to develop and evaluate methodologies for biodiversity measurement, and is of equally little use to conservation planning. Planning always involves choices, sacrificing one system to save another. So we begin this chapter by setting out a group of biodiversity measurement strategies. This is not a complete survey. We want instead to focus on how widely these strategies differ and on the considerations that are supposed to favor one rather than another. In the next chapter we move to a different set of issues: those involving costs and goals.

We begin our investigation of the place of biodiversity in conservation biology with a description of its use in current science, identifying the phenomena scientists actually measure when making judgments about diversity, and the phenomena they would measure if unconstrained by considerations of cost and effort. Once we turn to actual practice, we confront the problem of biodiversity surrogates noted above. We do not measure temperature by directly measuring the kinetic energy of particles and taking a mean. Instead, we use a substance (namely, mercury) with characteristics that are both highly sensitive to changes in temperature and that are easily measured. Analogously, it would be ideal to discover a sort of biodiversity thermometer. The strategy of using surrogates to detect biodiversity is the strategy of devising such biological thermometers, of identifying properties of biological systems that are reliable indicators of biodiversity properties. This strategy is almost universal in conservation biology, and many surrogates have been proposed. If conservation biologists are getting it right, these surrogates are reliable indicators of important characteristics of biological systems. Whether or not they are getting it right, these surrogates are reliable indicators of what conservation biologists take to be important about biological systems. So we now turn to a quick sketch of the most important surrogacy suggestions. As we shall see, there is a good deal of ambiguity about the status of these measured variables. Sometimes they are interpreted as signs of biodiversity, but not themselves as actual components of biodiversity. Counting family-level diversity in a system as a proxy for its morphological diversity exemplifies this approach. Sometimes they seem to be taken as representative samples, parts of

the whole that indicate the whole. The use of indicator taxa exemplifies this approach. Counting butterfly species in two forests gives a component of species richness in each forest, and also can be used as a sign of the overall species richness of the two areas. Sometimes the measured variables seem to be taken to be a measurement of biodiversity itself, as in some views of genetic diversity.

7.2 COUNTING TAXA

We begin with the simplest idea, one that has been central to chapters 2–6. Perhaps we should measure biodiversity just by counting taxa, for the most widely used strategy for the measurement of biodiversity is counting taxonomic groups and estimating their frequency. These strategies typically distinguish between estimating alpha and beta diversity. The alpha diversity of a particular habitat patch is its local taxon richness (usually species richness): the number of taxa found in the community, weighted by abundance. A system with one very numerous species and a few rare ones is less alpha diverse than one in which the species are equally abundant (see Box 7.1 for details). The beta diversity is a relational measure; it measures the additional richness this patch adds to the regional system, and the species added to the count through surveying this community. Beta diversity (and its relatives) is very important to conservation planning, because that planning typically involves the selection of an ensemble of sites to maximize the overall protection of biodiversity. The difference between one community and others already protected (or considered for protection) is often as important as the intrinsic richness of a community.

As we have just noted, information about species richness is often joined with information about abundance; measures that combine information in this way include the Shannon Wiener Diversity Index and Simpson's Index (see Box 7.1). The intuitive background to such measures is the thought that a sample of (say) 100 organisms representing 10 species is not very diverse if 85 of the organisms belonged to a single species. If this were a plant community (for example), the characteristics of the community would depend largely on the phenotype of the hyperabundant species. Of course, the idea that ecological processes are controlled by the phenotype of the numerically abundant species can be trumped by special features of the rare species: if the hyperabundant species is an annual wildflower, and the other nine species are all species of large tree, we might well make no such assumption. Phenotype matters, and we will soon consider ways of making its importance explicit.

BOX 7.1: Diversity Indices

Diversity indices supplement species richness. The number of species represented in a sample (s) is supplemented with information about the evenness with which individuals are distributed between the species present. Evenness information is often represented as p_i (the fraction of individuals belonging to the i^{th} species). Two common measures are:

Simpson's Index

$$D = \sum_{i=1}^{S} p_i^2$$

This is a measure of the probability that any two individuals in a sample will belong to the same species.

The Shannon Wiener Diversity Index

$$H' = -\sum_{i=1}^{S} p_i \ln p_i$$

This is a measure of the disorder of the sample (strictly the "entropy" as understood in mathematical information theory). On this measure, a highly diverse group is one with a great number of different types of individuals and roughly the same number of individuals of each type.

Counting species involves surveying (perhaps several times to account for seasonal variations) the organisms in a particular habitat, and sorting the specimens collected into species. One advantage of this strategy is that, for some taxa it is relatively simple. Because organisms of different species tend to be morphologically distinct, workers with limited training in taxonomy can roughly estimate the number of species in an area. Estimates of species numbers made by those without formal taxonomic training will be "rough" because they will be confounded by cryptic species (populations that do not interbreed despite a high degree of morphological similarity), radical sexual dimorphism (species in which males and females are so different as to appear to be members of different species), and radical morphological differences in successive life stages (common among invertebrates). Moreover, our ability to distinguish between species is much more reliable for some taxa (for example, vascular plants and vertebrates) than others (for example, fungi and protists) (Berlin 1992). So while there are practical advantages to species counting, there are practical disadvantages as well.

The vertebrates and vascular plants in a region can usually be identified fairly accurately, but the same is not true of invertebrates, fungi, and microbes, and these are important components of taxonomic richness. Abundance is difficult to estimate reliably, too. Hence conservation biologists often use proxy taxa, like bird diversity, as indicators of overall taxonomic diversity, and of changes in diversity.

Counting species is also theoretically well motivated. As we have argued in chapters 2–6, if there is a decent candidate for a good overall measure of biodiversity, a measure relevant to many of the theoretical and practical projects of the life sciences, it is based on the species richness of a biota. Despite the controversy over species definitions, there is widespread agreement that species are objective features of the biological world: species are the crucial units of evolution. Moreover, as we have noted already, there are natural ways of supplementing information about species richness. We can add abundance data. In chapters 2–6, in talking about species richness as an overall measure of biodiversity, we talked of information about the species and their genealogies. So we can add phylogenetic information, to represent the difference between a biota that represents a number of ancient clades, and a biota dominated by a large population of recently evolved close relatives. The small mammal fauna of Tasmania contrasts with that of North Queensland in this regard: both are diverse, but North Queensland has a large number of recently evolved true rodents, where Tasmania has more representatives of ancient marsupial lineages.

However, while in principle it is possible to supplement a species-richness-based account of biodiversity with phylogenetic information, in practice it is not obvious how to do this in a precise and tractable way. This problem is particularly pronounced in estimates of beta diversity. While we might plausibly estimate the total species count of a large region, it would be much more difficult to estimate a phylogenetically adjusted account of its species diversity. As we have remarked, almost all biologists share the judgment that different species represent different amounts of biodiversity. The two surviving species of tuatara (genus *Sphenodon*) are remarkable both morphologically (for the possession of a hidden third eye) and phylogenetically (as the last survivors of the order Rhynchocephalia (Sphenodontia), sister group to the snakes and lizards). Given this, many think that conserving a species of tuatara represents a much greater saving of biodiversity than, say, preserving a species of minnow.

The tuatara are such classic examples of "living fossils" that they make the intuition that species are not all equally unique very vivid. But we do not need such a vivid example to make the point, as is shown

by a thought experiment of Harper and Hawskworth (1995, 7). They suggest that we consider how much biodiversity is present in a series of hypothetical sites. Each site contains just two species. One is a species of *Ranunculus*, a genus of flowering plant within the buttercup family (Ranunculaceae), and the other is:

1. Another species of *Ranunculus* from the same section of the genus.
2. Another species of *Ranunculus* from a different section of the genus.
3. A species from a different genus in the same family (Ranunculaceae).
4. A species from a different family within the same order as the Ranunculaceae.
5. A species from a different family and in a different order (for example, a grass).
6. A rabbit.
7. A fungus of the genus *Agaricus*.
8. A protozoan of the genus *Amoeba*.
9. An archaebacterium.
10. A eubacterium of the genus *Pseudomonas*.

In some important sense of biodiversity (the thought goes), these samples are not equally biodiverse. As Robert May puts it:

> One of the basic conceptual issues in quantifying biological diversity is the extent to which a "species" does or does not represent the same unit of evolutionary currency for a bacterium, a protozoan, a mite, and a bird. (May 1995, 15)

Thought experiments like these have led ecologists to search for a measurement strategy that more accurately reflects the differences among organisms. We need some representation of species structure, not just the numbers of species present. Family-level diversity is sometimes suggested as a surrogate for this structure. So some taxon-counting measures of biodiversity count families instead of, or as well as, species. The family is a common choice because families are less subject to taxonomic revision than genera and they are more informative than more inclusive taxonomic levels such as orders and classes. This is one way we can, in practice, add information about the evolutionary history and morphological disparity to our measure of biodiversity. A biota that includes ten families of arthropod represents more evolutionary history and disparity than a biota that includes two. That said, we have already

seen the serious limits on the use of higher levels of the Linnaean system to capture biodiversity. There is no robust scientific theory that allows us to settle disputes about whether a particular group of taxa constitutes a family or not. This is not to say that we could pick any assemblage of species and call it a family (at the very least such groupings must be monophyletic). As a clade grows by speciation at the tips, the tree of species so formed gets larger and larger. Within any large tree, there will be many branches that we could pick out and name, but that science has chosen not to name. Perhaps in a rough-and-ready way, family-level diversity is a surrogate for phylogenetic diversity. But this will be at best a rough measure. Conservation biologists influenced by cladism have tried to do better.

7.3 Measuring Phylogenetic Diversity

The great theoretical strength of cladistics is that it does latch onto something real in the world: phylogenetic structure, the massively complex set of relationships that is the "genealogy" of species. There is nothing conventional or subjective about the claim that a bat and a bear really are closer phylogenetic relatives than are a bat and a bee. It's not surprising then, that many have sought to exploit this fact about nature in the measurement of biological diversity. Instead of relying on intuitive judgments of phenotype distinctiveness, one of the aims of those who want to measure biodiversity directly from cladistic principles has been to try to devise a measurement strategy that treats all speciation events as contributing equally to biodiversity. There is a wide range of strategies available, but the most widely used[1] measure of phylogenetic diversity is due to Daniel Faith (1994). However, see also Owens and Bennett (2000), Posadas et al. (2004), and Barker (2002).

One aim of the strategy is to pick out the group of species (from a larger group being studied) whose members are most distantly related to one another. To do so, Faith defines closeness of phylogenetic relationship in terms of the number of speciation events that separate a group of taxa. So, for example, two sister species are separated by one event. The direct offspring of those two sister species are separated by three, and so forth. So we might think of the basic strategy as tracing a line between taxa on the phylogenetic tree and counting the number of nodes (that is, speciation events) along that line. The other aim of this strategy is to try to capture the rate at which particular lineages evolve. One of the reasons why phylogeny is not a perfect predictor of phenotype is that species evolve at different rates. So if two sister species

experience very different selection pressures, then one may evolve much faster than the other and thus end up looking much less like the parent species than its sibling. Faith thinks that we ought to take this evolution between speciation events into account when measuring biodiversity. He proposes to plot the evolution of character states onto the phylogenetic tree. When we trace the line between taxa, we can count not just how many nodes we pass but also the number of character states that have evolved along the way.

To calculate Faith's phylogenetic diversity we must first construct a cladogram that includes feature information (information about character state changes that occur either at or between speciation events). An example of such a cladogram is given in figure 7.1. The idea behind phylogenetic diversity is that if, for example, we could save some but not all of the taxa shown in figure 7.1 then we would set about this task by looking for a "minimum spanning path." Assume that we only have funding sufficient to save four out of the ten taxa shown. We then find all the paths on the cladogram that connect four species and choose the path out of that group that includes the greatest number of speciation events as well as the greatest number of character state changes. Those four species are the ones we should save.

If this seems a bit abstract, analogy might help. Think of figure 7.1 as a road map. At the tip of each branch is a destination and each of the dots represent potholes. The minimum spanning path is just the bumpiest way of getting to a given number of destinations. The minimum spanning path for the tree in figure 7.1 is shown in figure 7.2. Despite acknowledging phenotypic difference, this is explicitly a cladistic theory. Faith argues (1994, 4) that the advantage of using phylogenetic diversity

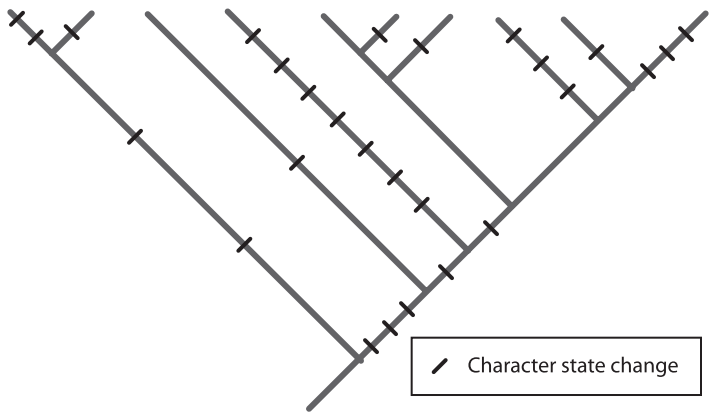

FIGURE 7.1. Cladogram with character state changes. This diagram depicts the ancestry of ten extant species.

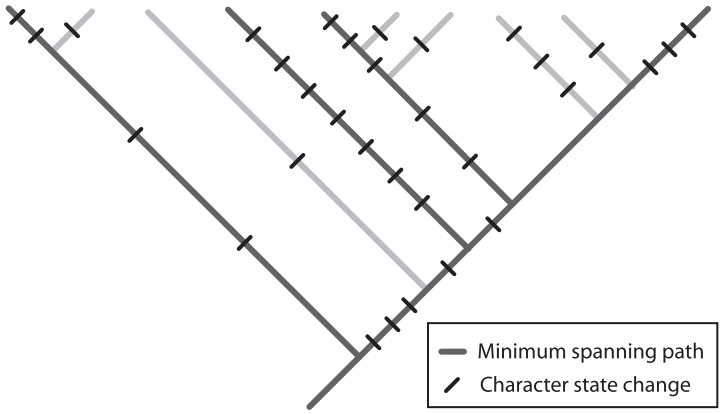

FIGURE 7.2. The minimum spanning path.

based on minimum spanning paths is that it will count traits that are structurally identical, but that result from evolutionary convergence, as different traits.

However, whatever the in-principle merits of Faith's proposal, in all but the simplest of cases it does not seem practicable. Most existing cladistic analyses do not contain the amount of information required for a measurement of phylogenetic diversity that includes comprehensive information about phenotypic difference. Moreover, as we noted in our early discussion of phenetics, the notion of *complete* information about phenotypic difference, is itself ill defined. So there will be difficult choices to make in deciding which information to include. Furthermore, cladistic systematics is increasingly dominated by cladograms derived from molecular data. Of course we can treat molecular change as character state change, but given that molecular difference does not covary cleanly with phenotypic difference, we cannot base our measure of phylogenetic diversity on both types of data. Faith's method might still be an ideal toward which we might work, but it would be vastly more labor intensive than species or family counting.

Moreover, and most importantly, despite the in-principle objectivity of the method, it is theoretically unmotivated. What exactly would distinguish a regional biota that was more Faith-diverse than one that was less Faith-diverse? Would it show more evolutionary flexibility on short or long time scales? Would it provide more resilient ecosystem services? Would it be more phenotypically disparate? If Faith-diversity is a measure of a causally important dimension of biological systems, we need an explicit case for that view. Equally, if Faith-diversity is a goal, a measure of some valuable feature of biological systems, that case must be explicit too (we will see a sketch of such a case for a measure similar

to that of Faith in the next chapter). Measurement strategies need to be explicitly linked to claims about value or claims about intervention points.

Phylogenetic diversity is a blend of phenotype and phylogeny, but it is not a satisfactory blend. It is committed to the view that every character state change is of equal importance in measuring biodiversity, and that is no more plausible than the idea that every species is of equal importance. Given the fundamental implausibility of this view one might expect to see purely phenotypic measures of biodiversity, and indeed such approaches have been advocated (for example, Roy and Foote 1997). However, they are not common methodological choices for reasons that we discussed in chapters 3 and 4; once we abstract phenotype differences from a phylogenetic context, we have lost the most objective way to choose the traits to measure and compare. So we shall suggest that one option worth considering is the use of local morphospaces to explore the fate of a clade in different regions. We could, for example, compare the phenotypic diversity of New World versus Old World monkeys or Australasian versus American parrots using such local morphospaces. The common history of the clade makes them phenotypically commensurable; we can use the same dimensions to plot their spread in a common morphospace. Theoretical morphology is an important tool for thinking about biodiversity differences, but only in combination with genealogical information about the history and relationships of species.

7.4 MEASURING GENETIC DIVERSITY

Genetic diversity is crucial to conservation biology. As we noted in 5.1, populations on the brink of extinction often have too little genetic diversity. Selection pressures that would simply delete unfortunate phenotypes from larger populations may well destroy small populations because they lack the variations that would allow them to respond successfully. Moreover, measuring genetic diversity certainly has methodological attractions. DNA sequences are relatively easily identified, and the differences between sequences are more discrete and therefore more countable than phenotypic characters. A new and important research effort aims at identifying DNA bar codes, short DNA sequences that show little within-species variation compared to their variation between species. There has been some success in identifying a characteristic class of such sequences of animals; the situation with other taxa seems less promising. If we can find such bar codes, they will be an important tool for taxonomy and hence conservation biology, revealing the presence of

sibling species, and enabling field workers to identify morphologically cryptic organisms. Many invertebrates have life cycles that involve stages that do not advertise their specific identity (Savolainen et al. 2005). Perhaps the most promising role for studies of genetic diversity is in understanding microbial diversity. Importantly, we can sample and amplify the DNA in a substrate, and thus get some information about both the variety and number of microorganisms present in the environment from which the substrate has been extracted. This technique has been used to estimate microbial diversity and community organization in environments as different as soils, human guts, and the open ocean (Falkowski and de Vargas 2004; Fierer and Jackson 2006; Gill et al. 2006). There are many uncertainties about these methods because the fragments of DNA that are amplified have to be assembled into putative organism genomes. Even so, measuring genetic diversity is a window onto an important aspect of biodiversity that is largely invisible to other methods for its assessment. These uses of DNA bar codes are uncontroversial. Much more controversial is the idea that DNA bar coding can largely replace traditional systematics. We agree that this more ambitious aim for DNA bar coding is wrongheaded; DNA bar codes need to be calibrated against an independently identified species phylogeny (Herbert and Gregory 2005; Smith 2005; Will et al. 2005). As always with a biodiversity surrogate, we can never just assume that there is a reliable relationship between the indicator property and the target property.

So there are good reasons to focus on measuring genetic diversity within biological systems. Genetic diversity is causally important (it is certainly part of the real diversity of biological systems) and it may covary well with other important aspects of diversity. Genetic similarity is certainly a reasonable predictor of important phenotypic similarity (Williams and Humphries 1996, 57). But there are also confused reasons; in particular, the idea that genetic diversity is fundamental and other dimensions of diversity are not. This confuses a surrogate for biodiversity with diversity itself.[2] For example, James Mallet argues:

> Biodiversity consists of the variety of morphology, behaviour, physiology, and biochemistry in living things. Underlying this phenotypic diversity is a diversity of genetic blueprints, nucleic acids that specify phenotypes and direct their development. (1996, 13)

It is certainly true, as we have noted, that the biochemical structure of genetic material provides us with quantifiable differences. But base pair similarity and difference is one thing; gene similarity and difference is another. Functioning genes are typically in the range of hundreds to

thousands of base pairs. Furthermore, some portions of our genomes appear to play no protein-coding role in the development of phenotype, though it is increasingly likely that much untranscribed DNA has a regulatory function. Given this, it is at least theoretically possible for two species to display a high degree of similarity with respect to base pairs without sharing many genes. Moreover, the relationship between genotype and phenotype is complex. We discussed some of those complexities in 5.2–5.4; another symptom of that complexity is the so-called C-value paradox, the fact that there is so little relationship between genome size and (apparent, intuitive) morphological complexity. The variation in genome size, and its lack of connection with phenotype complexity is really quite striking. Genome size varies by a factor of 200,000 in eukaryotes (Ryan Gregory 2001), and not because some eukaryotes are small and simple and others are huge and complex, as the following data (taken from Zimmer 2007) show:

GENOMES SIZE FROM SMALL TO LARGE
Nematode (*Caenorhabditis elegans*): 100 million bp (bp = base pairs)
Thale cress (*Arabidopsis thaliana*): 160 million bp
Fruit fly (*Drosophila melanogaster*): 180 million bp
Puffer fish (*Takifugu rubripes*): 400 million bp
Rice (*Oryza sativa*): 490 million bp
Human (*Homo sapiens*): 3.5 billion bp
Leopard frog (*Rana pipiens*): 6.5 billion bp
Onion (*Allium cepa*): 16.4 billion bp
Mountain grasshopper (*Podisma pedestris*): 16.5 billion bp
Tiger salamander (*Ambystoma tigrinum*): 31 billion bp
Easter lily (*Lilium longiflorum*): 34 billion bp
Marbled lungfish (*Protopterus aethiopicus*): 130 billion bp

Indeed, Ryan Gregory points out that the 200,000-fold range is found across single-celled eukaryote lineages; the genome of *Amoeba dubia* is more than 200,000 times larger than that of the microsporidium *Encephalitozoon cuniculi* (Ryan Gregory 2001, 66).

In the light of this complex relationship between genome and phenotype, it has increasingly been argued that it is misleading to think of the genome as a program that controls or organizes development (see, for example, Gerhart and Kirschner 1997; Oyama et al. 2001). While the genome does direct development, it doesn't do so alone. A host of behavioral, embryological, and environmental resources are required for the development of an individual, and changes in these factors can produce radical differences in the developed individual (for a comprehensive

survey of these phenomena, see Jablonka and Lamb 2005). The maintenance of stable and diverse global gene pools is an invaluable tool in the fight to achieve stable and diverse global ecosystems. Moreover, measuring genetic diversity gives us some insight into the otherwise hidden world of microbial diversity and community structure. Finally, there are genuine measurement advantages in focusing on gene diversity; it is an important diversity surrogate. That said, we see no reason in general to equate biodiversity in conservation biology with genetic diversity.

7.5 BIODIVERSITY SURROGATES

Biodiversity surrogates, in all probability, do not vary independently from one another. There is clearly an important correlation between, for example, species richness and family richness.[3] Nonetheless, the various measurement strategies rest on different foundations. Some tie biodiversity to speciation. Others tie it more closely to phylogenetic structure. Some include a morphological component. Others come close to tracking common intuitions about biological diversity. But measurement strategies in conservation biology have to be especially responsive to tractability issues; often conservation biologists measure what they can, with the expectation (or hope) that the facts that can be measured in the field track those believed to be of causal importance. It has long been recognized that conservation biology is a "crisis discipline" (Soulé 1985). Its raison d'être is to be found in overpopulation, intensive exploitation of environmental resources, habitat loss, and pollution. These factors lead to species loss and environmental degradation. Global conservation is a daunting task performed by too few people and with insufficient funds. These facts constrain methodology. Conservation biologists must therefore concentrate their efforts on "what is feasible, what is too crude to be useful, and what is unnecessarily detailed" (Fjeldså 2000).

Resource constraints sometimes bite very hard indeed, and hence there are simpler and cruder surrogates than species richness. As organisms tend to be specialized to niches in which they occur, as a rough regularity (since it ignores generalists), different niches will likely be filled by different organisms. The greater the difference in niche, the more the occupants will differ in their genetics, morphology, and behavior. As we noted in the last chapter in discussing the value of phenomenological communities as a guide to beta diversity, we can use features of environments as surrogates for the biodiversity that inhabits those environments. So, for example, environmental parameter diversity rests on the assumption that any available niche will be occupied by

at least one species (for a good discussion of this rather complex idea, see Sarkar 2002, 142–43). What it measures is diversity with respect to niches, but (if the basic assumptions are correct) what it detects is biodiversity.

There are even cruder measures: using satellite photography to estimate vegetation cover, and treating this as an index of biodiversity and biodiversity change. These measures are crude, but one of the main worries of those concerned with the conservation of biodiversity is the impracticality of strategies that involve the measurement of large numbers of properties of vast numbers of organisms. That is why we returned again and again in chapters 2–6 to the idea of phylogenetically enriched species information as a surrogate for biodiversity in general. It is a plausible compromise between what we would like and what we can do. Typically, here is information about species present in biological systems, and traditional taxonomy still encodes a lot of information about the genealogy of a species for all its subjectivity, failures to include stem species, and its use of paraphyletic groups (dinosaur, reptile). Thus a good flora and fauna (supplemented by some rough-and-ready abundance data) provides a sensible starting place in any study of biodiversity (where we are otherwise uncommitted to the nature of the diversity that is driving the system in question).

It is one thing to estimate the diversity of a system; it is another to be confident that the system continues to be as diverse. Even using surrogates, estimating diversity is often difficult and expensive, and yet systems are in a state of flux, and we can rarely assume that they are in equilibrium. Conservation biology badly needs surrogates for detecting change in previous baseline states. It is common to use proxy taxa to detect change. The idea here is to detect disturbance and estimate its severity by using change in abundance of some indicator taxon: a canary species whose loss or decline is a good indicator of general loss or decline. Thus an ideal indicator taxon is one that is very sensitive to habitat change, can easily be surveyed, and whose taxonomy and natural history are well known. Invertebrates make particularly good indicators as their short life spans mean that a change in breeding rates is easy to detect (Greenslade and Greenslade 1984).[4] But, clearly, even if there are indicator taxa in a habitat, they are difficult to identify with confidence, for (as with all surrogacy methods) the use of indicator taxa involves extrapolation from observed facts about the ecologies of known taxa in studied environments, to predictions about biodiversity in different environments under different conditions.

In 7.1 we noted that a good surrogate must be both practically usable in the field and a reliable indicator of its target property. In his

recent introduction to the philosophy of conservation biology, Sahotra Sarkar discusses surrogacy extensively as part of his defense of the idea of ranking places according to their *relative* biodiversity value (Sarkar 2005, chap. 6). Sarkar thinks it is neither necessary nor possible to give an explicit definition of *absolute* biodiversity. Instead, he suggests that biodiversity can be implicitly defined by a ranking procedure using surrogates, a procedure that takes into account both the objective biological richness of places we have identified as candidates for protection and the practical constraints on our abilities to measure and protect this richness. Sarkar accepts the idea that there is an element of choice in the selection of surrogates, but we think he understates the problem of evaluating surrogates. In our view, we can assess the adequacy of surrogates only by explicitly addressing the question: what aspects of biological richness do we wish to conserve, and why? Butterflies, for example, have prima facie advantages as surrogates because (as with birds) natural history enthusiasts have generated a good database about their abundance and distribution. Moreover, as adults, they are readily identifiable. Butterfly richness may be a true surrogate for species-level taxonomic richness. But by itself, that does not tell us that butterflies are a good surrogate for other aspects of biodiversity. It is true that conservation biology would not have to address this problem if we knew that the various kinds of biodiversity covaried well with one another, if phenotypic distinctiveness covaried well with ecological complexity, which covaried well with levels of endemism or with phylogenetic distinctiveness. If various versions of biodiversity covaried, a good surrogate for any form of diversity—for example, species richness—would be a decent surrogate for all the others. But we do not know that (for some initial reservations in the conservation context, see Andelman and Fagan 2000).

The role of surrogates and index species has added both complexity and confusion to the literature on biodiversity in conservation biology. As we complained in 7.1, it is often not clear whether the features of a system being measured are seen as direct measures of target properties or whether they are surrogates: measurable proxies for causally relevant properties. In some cases the situation is unambiguous. The recent and increasing use of satellite images to assess the extent of vegetation cover in making conservation assessments is clearly the use of a mere surrogate. This technique is chosen because the data are easily available, not because anyone thinks we are thereby directly measuring the biodiversity that matters (see, for example, Margules and Pressey 2000). In contrast, Faith's phylogenetic diversity is probably conceived as a measure of the target property itself. But in other cases, the profusion

of surrogates has led to much confusion, as our discussion of genetic diversity illustrates. Mallet tells us that "the diversity of life is fundamentally genetic" (1996, 13), whereas Williams and Humphries (1996) talk as if genetic diversity is better thought of as a surrogate for biodiversity, particularly in conservation settings.

Further, there is often little calibrating information about proxies and their reliability. This is no accident. They are used because it is difficult to get direct information about the causally relevant target properties of the system. That very fact makes proxies difficult to calibrate. For example, the coevolutionary interactions between butterflies and flowering plants probably make it safe to assume that areas rich in butterfly species are species rich. But it is not safe to make the converse assumption: that butterfly poor patches are species poor. So there are severe practical problems in calibration. But conservation biology faces theoretical problems in choosing target properties: we cannot choose what properties to conserve without an account of conservation aims. The literature is often not explicit (as we saw in discussing Faith diversity) on why particular target properties are chosen. To make further progress on this issue, we finally have to move beyond purely empirical issues about the driving properties of systems to claims about goals of conservation biology.