

Biodiversity

Its measurement and metaphysics

D.M. Roche

A thesis submitted for the degree of Master of Science

University of Sydney

January 2001

Contents

Acknowledgements	ii
Abstract	iii
1 Biodiversity: introduction to a problem	1
2 Concepts of biodiversity	9
2.1 Scientific versus political concepts of biodiversity	9
2.2 Biological triage	14
2.3 Conservation value	18
2.3.1 Intrinsic value	20
2.3.2 Instrumental value	22
2.3.3 Inherent value	29
3 Measuring biodiversity	33
3.1 A tale of three taxonomies	34
3.1.1 Evolutionary taxonomy	35
3.1.2 Phenetics	36
3.1.3 Cladistics	38
3.2 Taxonomic biodiversity	41
3.2.1 Species richness	42
3.2.2 Higher taxon richness	43
3.2.3 Genetic diversity	44
3.2.4 Phylogenetic diversity measures	47
3.3 Raw biodiversity	57
4 Complexity	65
4.1 An intuitive complexity concept	66
4.2 Information content as a measure of complexity	72
4.3 Alternative complexity measures	80
4.4 The metaphysics of complexity	84
4.5 Effective complexity	86
5 Biocomplexity	91
5.1 Complexity as a measure of biodiversity	91
5.2 The uniqueness of biocomplexity	94
5.3 Defending biocomplexity	99
5.4 Biocomplexity—a summary	101
6 Applied biocomplexity	105
6.1 Estimating biocomplexity	105
6.2 Measures of biodiversity—a retrospective	113
6.2.1 Biocomplexity, species richness and higher taxa	113
6.2.2 Genetic biocomplexity	115
6.2.3 Biocomplexity and phylogenetic biodiversity	117
7 Conclusion	121
7.1 A rational argument for biocomplexity	121
7.2 An impassioned plea for biocomplexity	123
8 Bibliography	127

Acknowledgements

I would like to offer my sincerest thanks to Paul Griffiths for his extensive contributions to this project; in particular for his helpful guidance, patient supervision and strong coffee.

I would also like to thank James Maclaurin, Daniel Faith, Karola Stotz, John Collier and Jason Grossman for their helpful comments, and Professor Sahotra Sarkar for his contribution to a joint seminar on defining and measuring biodiversity.

Finally, to my wife Kathryn Phillips, I offer my most heartfelt thanks for, in her words, putting up with me.

Abstract

Biodiversity is a concept that plays a key role in both scientific theories such as the species-area law and conservation politics. Currently, however, little agreement exists on how biodiversity should be defined, let alone measured. This has led to suggestions that biodiversity is not a metaphysically robust concept, with major implications for its usefulness in formulating scientific theories and making conservation decisions.

A general discussion of biodiversity is presented, highlighting its application both in scientific and conservation contexts, its relationship with environmental ethics, and existing approaches to its measurement. To overcome the limitations of existing biodiversity concepts, a new concept of *biocomplexity* is proposed. This concept equates the biodiversity of any biological system with its *effective complexity*. Biocomplexity is shown to be the only feasible measure of biodiversity that captures the essential features desired of a general biodiversity concept. In particular, it is a well-defined, measurable and strongly intrinsic property of any biological system. Finally, the practical application of biocomplexity is discussed.

1 Biodiversity: introduction to a problem

All animals are equal but some are cetaceans.

— Neils Einarsson

As a rejoinder to the tradition of natural theology that seeks evidence of God's benevolence in the natural world, biologists are fond of a famous epigram attributed to J.B.S. Haldane. According to G. Evelyn Hutchinson,

Haldane ... found himself in the company of a group of theologians. On being asked what one could conclude as to the nature of the creator from a study of his creation, Haldane is said to have answered, "An inordinate fondness for beetles." (Hutchinson, 1959)

A nice line, but why did Haldane single out beetles? Recent estimates of the total number of formally named species stand at 1.4–1.6 million¹ (Wilson, 1986; Stork, 1996). Of this total, more than half are insects, and nearly half of this total are beetles. Hence beetles represent almost 25% of all named species, a total of about 400,000 (Hammond, 1992). This compares with the modest totals of about 4,000 mammal and 9,000 bird species (Wilson, 1986). Moreover, very few mammal or bird species remain to be discovered, whereas unnamed beetle species probably outnumber named species by at least an order of magnitude. Estimates of total species numbers are embarrassingly imprecise, ranging from 1.5 million to 300 million (Stork, 1996), with 30 million being a commonly quoted figure for tropical forest arthropods (Erwin, 1982). Given that beetles comprise about 40% of all arthropods, there may be as many as 12 million species of beetle on the planet. If, according to

¹ The estimated number of names defined by taxonomists is 1.8–2.0 million; however, the level of synonymy among insects is estimated at about 20% (Stork, 1996).

Haldane, we should one day meet the almighty face-to-face, he is much more likely to resemble a beetle than the Archbishop of Canterbury.

From a conservation perspective, this extraordinary abundance of beetle species in comparison to our favourite furred and feathered creatures raises an obvious but interesting question. Are some species more worthy of conservation than others? If we take this question seriously, the answer is an obvious ‘yes’. There are perhaps as many as 12 million species in the clade Coleoptera and only a few dozen in the clade Cetacea, and yet, despite the vastly greater extinction rate of the former, we seem far less concerned about the plight of beetles than of whales. Of course, whales are big, intelligent and spectacular, all of which contribute to our concern for their welfare, whereas beetles are small, stupid and non-descript. Moreover, despite the vast number of beetle species, most of which have never been named, described or even observed, all of them are just that—beetles. To most of us, one beetle species is not very different from any other. Even a well-trained entomologist may have difficulty distinguishing closely related species. Why should we care too much about losing one beetle species to extinction when there are 12 million others?

There is a well-recognised problem with this argument. It is what Thomas Lovejoy has called the ‘incrementalist problem’ (Lovejoy, 1986). If we accept the argument, then loss of one unknown beetle species is acceptable, both from an ethical and practical perspective. After all, the vast majority of all species that have ever lived are now extinct, and the evolution of all extant species has relied to some degree on these past extinctions (May, 1994). By the same reasoning, the loss of two species is also acceptable. At what point, however, is further species loss unacceptable? As Norman Myers puts it,

When do species losses shift from being marginal to becoming significant, serious, critical, crucial, and catastrophic? (Myers, 1991, p. 20)

Paul and Anne Ehrlich used a now famous metaphor to describe this problem. Imagine being on an aeroplane and seeing someone slowly removing rivets from a wing. The loss of one rivet is not important; the wing will stay on and the plane will fly regardless. Even with the loss of two, three or a handful of rivets, the plane will continue to fly. At some point, however, too few rivets will remain to hold the wing in place and the plane will fall from the sky. To preserve the wing intact, the Ehrlichs argue, we must fight to save each rivet (Ehrlich & Ehrlich, 1981).

The incrementalist problem has forced environmental campaigners to shift their focus away from efforts to preserve individual species towards the protection of whole ecosystems. The single-species approach to conservation works well for large, easily identified and charismatic species such as giant pandas, but it becomes hopelessly impractical for protecting organisms that provide a host of essential biological functions such as producing oxygen, purifying water and recycling soil nutrients. To protect these organisms required the invention of a new concept with its own neologism—biodiversity.

The term *biodiversity* was coined in 1985 by W.G. Rosen as a contraction of the phrase *biological diversity*. As Rosen later noted wryly, all he had to do was “take the ‘logical’ out of ‘biological’” (Takacs, 1996, p. 37). The exact meaning of biodiversity is somewhat nebulous. Etymologically, the word refers simply to the diversity of the biological world. The world has many different kinds of things, and biodiversity captures some notion of both the number of different kinds of things and the degree of difference between them. Many biologists define biodiversity as something like “the variety of life forms, the ecological roles they perform and the genetic diversity they contain” (Wilcox, 1984, p. 640). Biodiversity therefore captures more than just the number of species in the world; it also includes the rich variety of forms they represent, their complex interactions and the many ecological functions they perform.

By uniting the world's millions of small, unloved creatures under the banner of *biodiversity*, conservation biologists had suddenly created a new political weapon. Instead of arguing for the preservation of individual species, with the practical and logical weaknesses that such a position entails, they could now argue for the preservation of biodiversity—an all-embracing concept that includes all species from whales to beetles.

Since its invention, the concept of biodiversity has grown enormously in usage, both within the scientific and lay communities. As a keyword in *Biological Abstracts*, *biodiversity* does not appear at all in 1987, while *biological diversity* appears only three times. By 1992, *biodiversity* appears 87 times and *biological diversity* 36 times. Since 1997, *biodiversity* alone has appeared over 1000 times per annum.

The continuing existence of the biodiversity concept relies on a perception that biodiversity is an objectively specifiable property of biological systems that is worth preserving. While *biological diversity* was never intended to delineate anything precise, *biodiversity* includes a degree of scientific respectability, and not without some justification. After all, it is given a scientific as well as a political role in the reformulation of the species-area law of ecology.² But biodiversity has never been a clearly defined term. Just what is biodiversity? What is its metaphysical status? Is it an intrinsic or anthropocentric property? Does it have one or many forms? And why should we protect it? These questions are not without importance. Without some firm foundation, the concept of biodiversity may prove unable to sustain the weight imposed by both its scientific and political roles. How can we develop a science of ecology if our fundamental concepts are suspect? And why try to preserve something

² In its simplest form, the species-area law states that larger areas harbour proportionally more species than smaller ones. It dates back at least as far as the early 19th century work of Alexander von Humboldt (Rosenzweig, 1999).

called *biodiversity* if that term is based on the subjective preferences of some small groups of scientists?

My intention is to resolve these issues. I address two primary questions, namely:

1. What is the metaphysical status of biodiversity? More specifically, can we construe the concept of biodiversity in such a way as to make it an intrinsic³ property?
2. Can we measure biodiversity, and if so, how?

These two questions are far from independent. Developing a more metaphysically secure definition of biodiversity may suggest how we should go about measuring it, whereas developing a well-defined and independently-specifiable measurement procedure suggests that biodiversity may be an intrinsic property of any biological system.

My answers to both questions are affirmative. Biodiversity, properly interpreted, is an intrinsic property that can be measured. To demonstrate this, I formulate biodiversity in terms of biological complexity. Complexity, or at least a particular type of complexity, is an intrinsic and quantifiable property, and hence so is biodiversity construed as biological complexity. Moreover, I make an additional and stronger claim, which is that biological complexity is the only sensible interpretation of biodiversity that makes it a strongly intrinsic property.

³ There is much debate among philosophers over the exact meaning of *intrinsic*, which I do not wish to rehearse here. An intrinsic property of an object is generally considered to be one that does not depend on any relationship between that object and something else. It is therefore non-relational. In particular, it is non-anthropocentric. I include an additional criterion in my definition; that of non-arbitrariness. Under this definition, the mass of an object is an intrinsic property of that object, whereas its property of 'having mass between 0.31 and 0.35 kg' is not.

One other important question concerning biodiversity is, why should we care? Why protect biodiversity? I do not attempt to answer this question here. Foundational questions in ethics are notoriously intractable, and environmental ethics is no exception to this rule. However, if we are to have some hope of determining why it is important to protect biodiversity, we must at least begin with a good understanding of what we are trying to protect. So my project at least offers a more solid starting point for ethical enquiry.

In the following chapter, I outline what we want from a biodiversity concept, both from scientific and conservation perspectives. Whereas addressing the scientific aspect of this question is easy, the conservation aspect is not. To do so requires a venture into environmental ethics to look at the types of value we might ascribe to the natural world and how they relate to biodiversity. Although I make no grand claims for solving problems of environmental ethics, this discussion will, I hope, illustrate the interface between environmental ethics and my project. In Chapter 3, I turn to the issue of measuring biodiversity, including a survey of existing approaches to measuring biodiversity and a discussion of one attempt to resolve its metaphysical status. An understanding of such prior art is, of course, essential for placing my contribution in context. In Chapter 4, I discuss the concept of complexity, along with its measurement and metaphysics. This discussion is a key input for the following chapters, where I develop and discuss my own brand of biodiversity based on complexity. I label this concept, predictably enough, *biocomplexity*. In Chapter 5, I flesh out this concept of biocomplexity in more detail, discussing its metaphysics, demonstrating that it is a uniquely intrinsic property, and defending it from some potential criticisms. Finally in Chapter 6, I discuss the application of biocomplexity. In doing so, I revisit some of the existing measures of biodiversity discussed in Chapter 3, and illustrate their merits as estimator measures of biocomplexity.

Before I begin, there are two points I should make clear. The first is that I have adopted what to a crude approximation can be described as a scientific realist stance. This is not the place to fight out a realist-antirealist debate, and I will make no attempt to do so. I will simply assume that there are truths to be discovered and that it is the aim of science to discover those truths. The second point is that my central thesis is not normative. As I have already stated, my main goal is to understand what biodiversity is, not why we ought to save it. Although I sometimes suggest some answers to the latter question, none are intended to answer it definitively.

2 Concepts of biodiversity

This capacity of a thing to reveal itself in unexpected ways in the future, I attribute to the fact that the thing observed is an aspect of reality, possessing a significance that is not exhausted by our conception of any single aspect of it. To trust that a thing we know is real is, in this sense, to feel that it has the independence and power for manifesting itself in yet unthought of ways in the future.

— Michael Polanyi

Beauty in things exists in the mind which contemplates them.

— David Hume

2.1 Scientific versus political concepts of biodiversity

What do we want from a concept such as biodiversity? As I have outlined in the Introduction, there are at least two distinct applications of the biodiversity concept. One is scientific, and the other political. A scientific concept of biodiversity is one that plays some useful role in scientific theories, just as the concept of *mass* plays a useful role in Newtonian mechanics. Moreover, if biodiversity were something we could measure, perhaps we could discover or at least refine our understanding of the relationships between some attributes of ecosystems and their biodiversity; relationships such as the species-area law. A political concept of biodiversity is one that can be used for conservation purposes. Under this formulation, biodiversity is, or at least is an indicator of, something of value that we wish to preserve.

The key attributes we want from any scientific concept are that it is:

1. well-defined
2. based on some natural properties of the world
3. useful in some scientific theory.

The first criterion ensures we know what we are talking about, the second ensures long term stability in our attempts to agree on how to understand and manipulate the concept, while the third ensures that the concept actually does some scientific work. *Mass*, for example, meets all three criteria. It is well-defined (the amount of matter in a body⁴), is a natural property of physical objects that we can measure with precision, and is useful in several scientific theories such as Newtonian mechanics. *Fear* as a concept in psychology also meets the above criteria, even though, unlike *mass*, it cannot be measured precisely.

Species richness, defined simply as the number of species in a given area or population, at least partially matches all three criteria. Although there is as yet no agreement on what constitutes a species,⁵ there is at least some hope that such agreement can eventually be reached. The first and second criteria would then be at least partially satisfied. The third criterion is more problematic. Limited versions of the *species richness* concept do play useful roles in scientific theories such as the species-area law. For example, there are often weak but definite mathematical relationships between the number of mammal species on each of a group of islands and the areas of each of those islands.

Mathematical relationships involving a more general notion of *species richness* tend to be less precise. This is partly explained by the fact that different species are not directly commensurate. The set of factors that unites red kangaroos together as a single species is not identical to the set of factors that unites *E. coli* bacteria. We can

⁴ Matter is sometimes defined as that which has mass. As with many quantities in physics, there is a degree of circularity in attempts to define mass. In Newtonian physics, we can define (inertial) mass as the constant of proportionality between force and acceleration. The fact that there is such a constant, and that it obeys superposition, strengthens our conviction in the existence of mass as a measurable quantity. This two-way exchange between theory and experiment ensures the circularity in the definition of mass is not vicious.

combine counts of kangaroo and bacteria species together, but there is little reason to think that the resulting measure will be highly predictive because the underlying units are not, at the level of interest, the same sorts of things. Measuring biodiversity by simply counting species is like measuring the productivity of Queensland fruit-growers by counting individual pieces of fruit. At one level, grapes and watermelons are the same sort of thing (pieces of fruit), but they are not commensurate as units of productivity. Nevertheless, we might still expect the fruit-count measure to go up and down in response to good and bad years, just as we might expect to see striking but imperfect correlations between species richness and land area. Both concepts are at least somewhat useful.

As well as the species-area law, biodiversity as a scientific concept also arises in several questions posed by ecology. For example, are biodiverse ecosystems more or less stable than simple ecosystems?⁶ Does biodiversity reflect sustainability? Does biodiversity reflect the evolutionary time elapsed without major disturbance, or does it reflect the frequency of major disturbance in ecological or evolutionary history? (Harper & Hawksworth, 1994). There is a recognition in such debates that *biodiversity* expresses something more sophisticated than mere species richness. It includes diversity at several levels, from genes to species to ecosystems. Scientific biodiversity recognises that species are not commensurate—that some species count for more diversity than others. Just as fruit-count is an indicator of a more meaningful measure of productivity such as fruit biomass, species richness is at most an indicator of some more meaningful but yet-to-be-determined measure of biodiversity.

If a more meaningful scientific definition of biodiversity than species richness is possible, then our failure yet to devise one simply reflects our present ignorance. Just

⁵ See, for example, Sterelny and Griffiths (1999, chapter 9) for a recent review.

as the concept of *weight* is a slightly vague but useful approximation to the more precise concept of *mass*, *species richness* is potentially an indicator of some deeper concept of biodiversity. Ecology is, after all, an imprecise science. Generalities are often swamped by particulars, blurring and distorting mathematical relationships. Organisms are not as interchangeable with one another as are electrons. Different concepts of biodiversity may also be required for different applications. The biodiversity that underlies the species-area law may not be the same as the biodiversity that determines ecosystem stability.

Most existing attempts to better define and measure biodiversity do not address the question of whether or not biodiversity is a scientifically useful concept. Their primary motivation is to develop a concept of biodiversity that is useful for making conservation decisions. This is an approach I intend to emulate. Nevertheless, the concept of biodiversity I will develop may prove scientifically useful. It satisfies criteria 1 and 2 above and also overcomes the species commensurability problem. In addressing predominantly philosophical questions, this is as much as I can hope to do. Whether the concept does prove to be scientifically useful is an empirical question. It is not mine to answer.

As a concept in conservation biology, biodiversity is used as a tool for measuring some sort of ‘conservation value’. If, for example, a government wished to divide an area of land currently used for logging, and conserve one half by creating a new national park, a rational strategy would be to conserve that portion of land with the highest conservation value. In recent years, this conservation value has become more frequently measured in terms of biodiversity. So if, all other known factors being equal, one part of the area in question contains a relatively large number of endan-

⁶ See, for example, Naeem & Li (1997).

gered mammal species, it would make sense to include that area in the new national park. In this case, a larger number of mammal species is seen as more valuable than a smaller number of mammal species. Mammals may also be an indicator of diversity among species from other groups not surveyed, in which case selecting the area with more mammals automatically increases the probability of selecting the area with more species from these other groups. This sort of application of the biodiversity concept in conservation efforts is now quite common, despite the lack of agreement on what biodiversity is.

As discussed in the Introduction, there is a strong political motivation for questioning the metaphysical status of biodiversity. If biodiversity is an anthropocentric concept—one defined only in terms of what some group of individuals wishes to protect—the motivations of those defining biodiversity and calling for its protection may be called into question. Of course, there are plenty of reasons to protect various aspects of the natural world. The pollination of food crops by wild bees is one simple example. As a workable concept, however, an anthropocentric definition of biodiversity poses problem. A naïve account defining biodiversity simply as ‘that which we wish to protect’ is vacuous. It tells us nothing we do not already know. Such a concept of biodiversity is simply a banner under which we can list some gerrymandered collection of things. We could develop a more sophisticated anthropocentric concept of biodiversity in terms of the responses of ideal agents acting under certain circumstances with certain information. However, even such a sophisticated concept of biodiversity is unwieldy and brings us little closer in practical terms to defining what we mean by biodiversity.

If biodiversity is to be useful as a tool in conservation, it must not only correspond to something worthy of protection; we must also be able to agree, more or less, what it is. And we should ideally be able to measure it, at least in relative terms. In other words, if we have two biologically rich areas, A and B, and we want to decide which

one to give greater priority in terms of conservation using biodiversity as the criterion, we should be able to say either A has more biodiversity than B, B has more biodiversity than A, or A and B have the same amount of biodiversity. An intrinsic concept of biodiversity offers a possible solution to such problems.

In the remainder of this Chapter, I outline the sort of criteria that can be used for determining whether something is worth protecting, and hence that might form part of this conservation-based biodiversity concept. Before doing so, however, I want to discuss a useful tool for analysing comparative measures of conservation value.

2.2 Biological triage

Triage is a procedure for determining the order in which wounded soldiers are given medical treatment based on their condition. Priorities are assigned to each soldier in order to maximise the probable number of survivors, given the constraints on the rate at which medical treatment can be given. This strategy means giving high priority to those soldiers who will probably die unless treated urgently, moderate priority to those who will probably survive without immediate treatment and low priority to those who will probably die even with urgent treatment. Triage is similarly applied in the casualty wards of major hospitals, although more with a view to minimising some notion of the net total of patient suffering rather than simply the number of casualties.

The analogous concept of *biological triage* is useful for thinking through difficult conservation issues.⁷ Given the rapid rate at which species are becoming extinct and the limited resources available with which to help conserve them, how should such

⁷ Further discussions of biological triage can be found in Mann & Plummer (1995), Takacs (1996) and Maclaurin (1998*b*).

limited resources best be spent? As with conventional triage, arriving at the best possible outcome requires that resources be focussed on the more urgent cases, with the possible exception of some 'lost causes'. However, we should also focus on maximising conservation value. For example, if species A and B are equally likely to become extinct and we can only afford to spend money on protecting one of them, it makes sense to pick the one we value most highly. Most of us value giant pandas more highly than smallpox virus, so we sensibly focus more effort on protecting the former.

Biological triage is a useful tool for working out what we value and by how much (at least in relative terms) by asking what we are willing to sacrifice to protect it. If we were given the choice of protecting humpback whales or 2.5 million Amazon rainforest beetle species, and chose the former, then that tells us humpback whales have at least as much conservation value as those 2.5 million beetles. Of course, not everyone would agree about the number of beetle species that are equivalent in conservation value to the single species of humpback whale, but in most cases we can at least derive some rough estimate. I suspect nearly everyone would at least value humpback whales more than one beetle species.

The concept of biological triage is not without controversy, even though many biologists reluctantly accept the need for devising some means of prioritising conservation efforts. As Robert May puts it,

As more and more species face extinction in the wild over the next few decades, how do we go about making choices for the ineluctably limited number of places on the ark? (May, 1990)

Many other biologists, however, reject triage outright. Takacs (1996) provides an illuminating survey of various biologists opinions on the matter. Edward Wilson

asserts, “I think we ought to save it all, as a basic principle” (p. 60), while Reed Noss unequivocally states

I reject triage ... It’s ethically pernicious to me ... I think that the argument that there isn’t enough money, there aren’t enough resources in general to protect all species and some of them we’re just going to have to let go, is disingenuous. (p. 60)

Donald Falk is even more adamant.

You cannot persuade me that we are in a triage situation with respect to the resources to save life on earth. I do not accept it. I think it’s a completely fallacious argument. Just as medicine tries to save every patient they can, that has absolutely got to be our mission. If we fail, we fail. (p. 60)

Such views are difficult to reconcile with the present disparity between extinction rates and conservation efforts. Rejecting even the concept of triage does not answer difficult ethical questions; it simply avoids them. The unease of medical professionals when making decisions that preserve the lives of some over the lives of others does not prevent them from carrying out medical triage. Such decisions are simply forced upon them by the impossibility of saving everyone. Regardless of whether or not we think biological triage is an appropriate conservation policy, the concept of biological triage is a useful tool for thinking through difficult conservation issues. Even if we lived in an ideal world where we could protect everything, supposing that we are forced to make a choice about what to protect helps us understand what we value and how we value it. If we want to protect the biological world, it must have some value. And, if forced to make a choice, we prefer to protect some elements of it over others, then those elements must have more value than others.

Other biologists reject triage for practical rather than conceptual reasons. David Woodruff, for example, states

Pragmatically, we cannot save each and every species. If you force me—so then I think you have to rank or prioritize your species. I'd rather not use the word *triage*, because that means that you consciously have everything set out. You go past each stretcher and you put one of three colors on it. We don't know where the stretchers are, let alone who's on the stretchers. So instead, my attitude is to prioritize from what you do know. And then pick the key species through—by working, by concentrating on those key species, you will be able to save ecosystems and larger units. (Takacs, 1996, pp. 60–61)

In a similar vein, Lovejoy finds triage

... inadequate to the situation ... Conceptually, it's an easy way out. To actually do it intelligently would be a bitch. (Takacs, 1996, pp. 61–62)

There are good reasons for taking this practical objection to triage more seriously. Woodruff's argument is simply that triage could only work in a world where more information is available than in ours. Of course we would prefer to prioritise species and protect those with the highest priority first, but the act of trying to prioritise species distracts us from the main task at hand, which is to protect whatever we can. At the end of the day, most of our conservation decisions will be based on far more practical considerations than whether species A has more conservation value than species B. Moreover, a whole-ecosystem approach to conservation is likely to save far more species than a piecemeal species-by-species approach.

This objection to triage is well taken. However, it is not my task to discover the most effective means of achieving conservation aims. My goal is to determine what biodiversity is and how we can measure it. If I can answer that question, and show that biodiversity is intrinsic and measurable, then I can at least provide some insight into how we might define and measure biodiversity, and in doing so perhaps contribute to the goal of understanding what it is that we should be protecting, regardless of whether measuring biodiversity is part of the most efficacious procedure for doing so. Biological triage is one tool for assisting me in this effort.

2.3 Conservation value

If we want to protect the natural world, then it must have some value. And if, as the concept of biological triage shows us, we want to protect some aspects of it more than others, then that value can, at least in a crude and relative sense, be quantified. What is the nature of this value and how can we measure it? Moreover, what is the relationship between conservation value and biodiversity? Does biodiversity correspond closely with one or more types of value, or is biodiversity some independent property of which some types of natural value are manifestations? In the discussion that follows, I examine some types of natural value to assess whether any of these correspond closely with some sensible concept of biodiversity.

Questions about the value of natural objects suggest a role for environmental ethics. This relatively new discipline arose in the 1970s to address ethical questions applied to the natural world. Since its foundation, one of the central concerns of environmental ethics has been to determine the nature of environmental value. More specifically, can nature possess intrinsic value, value that exists independently of human valuers? Richard Routley devised a now famous thought experiment to highlight this central issue, involving a lone survivor of some human apocalypse who sets about a redwood grove with an axe (Routley, 1973).

Much of the early debate concerning this question has focussed on a two-way split between intrinsic value and non-intrinsic value. Whereas intrinsic value exists independently of human valuers, non-intrinsic value does not; it is inherently anthropocentric. In more recent years, the so-called Callicott-Norton debate has

injected some much-needed clarity into the dispute by suggesting the following three-way split:⁸

1. **Intrinsic value**, which can exist in the absence of any conscious valuers.
2. **Instrumental value**, which in some way benefits a valuer.
3. **Inherent value**, which requires but does not benefit a valuer.

Sentient creatures are considered by many people to have intrinsic value. They are valuable in themselves, regardless of what we think of them. My bicycle has instrumental value, but no intrinsic value. It has value only because it does something for me which I find useful. A painting is something that has inherent (as well as instrumental) value. The value of a painting is not intrinsic. It depends for its existence on a valuer. But neither is the value of a painting exclusively instrumental. Although a painting provides aesthetic pleasure (an instrumental value), we can also value a painting for itself (but not *in* itself). If the Uffizi gallery in Florence were burnt to the ground, most of us would mourn such a loss, even if we had never visited and never intended to visit the gallery. We therefore consider the contents of the Uffizi to be valuable, even though they provide us with no benefit. This value is clearly not instrumental (it does not confer a benefit on the user) but neither is it intrinsic, because it depends on valuers for its existence. In other words, it is inherent value.

⁸ Callicott (1996) is a response predominantly to Norton (1991). This three-way split does not correspond exactly with the views of either. In particular, Callicott makes no distinction between *intrinsic* and *inherent* value, and appears to use the term *intrinsic* when referring to what Norton calls *inherent*. While Norton makes a three-way split, his boundaries and terminology do not fully correspond with those used here, which are adapted from McQuillan (1998).

2.3.1 Intrinsic value

Whether the non-human world has intrinsic value is somewhat controversial. Each of us lives a life that depends for its continued existence on the death of numerous organisms to provide us with a steady supply of food, clothing and various materials. We clearly consider such organisms to be of lesser value than a single human life. However, there are very sensible arguments as to why we should treat certain sentient creatures such as dogs, chimpanzees and dolphins with a respect akin to our respect for fellow humans. If humans have intrinsic value, then, according to these arguments, so must such sentient creatures (although perhaps in some lesser quantity).

I am very sympathetic to the idea that humans have intrinsic value and that, by extension, some sentient creatures also have intrinsic value. However, I will not attempt to defend this position, because doing so would not assist me in clarifying the relationship between environmental ethics and biodiversity. Even if we accept that certain sentient creatures have intrinsic value, it is difficult to see how this can be used to develop a rational conservation strategy. Protecting Australia's arid-zone ecosystems relies more on culling large numbers of feral rabbits, cats and foxes than it does on preserving individual bilbies, parrots and geckos. Any intrinsic value of individual organisms, even highly sentient ones, is swept aside in a bid to prevent whole species of endangered animals from becoming extinct.

If we are to argue for the protection of a species on the basis of intrinsic value, then we require that the species possess some intrinsic value that is more than just the sum of the intrinsic value of its individual members (O'Neil, 1997). If we captured

the last remaining kakapos⁹ and distributed them individually to zoos around the world, we would extinguish their species as surely as if we had shot each one dead. Yet each individual kakapo might live its remaining days in splendour. Without devaluing each individual kakapo, we can devalue the kakapo species.

Many people, myself included, consider intrinsic value to be rooted in sentience. Humans, and presumably dogs, chimpanzees and dolphins, are sentient and hence have intrinsic value. Species, however, are not sentient. They might consist of sentient individuals, but they are not in themselves a sentient being. Moreover, most species, such as Wollemi pines, for example, do not even consist of sentient individuals. Hence it is difficult to see how a species could possess the same sort of intrinsic value that a human or chimpanzee possesses.

I reject the idea that a species, or indeed any natural collective, could have intrinsic value beyond the intrinsic value of individual sentient creatures within that collective. Again, I will not attempt to defend this position any further, since doing so would not contribute to my central thesis. However, adopting such a position has some significant consequences for my concept of biodiversity. Biodiversity is not a value measure of individual organisms. A concept of biodiversity that is useful for making conservation decisions must incorporate natural collectives such as species and ecosystems. If I accept that the value of these natural collectives is non-intrinsic, then I am forced to accept that the value of biodiversity is at least partially non-intrinsic.

By adopting this position, I might be justified in relaxing the metaphysical requirements of a conservation-based biodiversity concept. If the value of biodiversity is non-intrinsic, then biodiversity itself may also be a non-intrinsic property of

⁹ A highly endangered flightless parrot that survives on a few small New Zealand islands.

biological systems. In other words, biodiversity need only reflect whatever anthropocentric values we wish it to include. However, the concept of biodiversity I will develop is metaphysically much stronger than this. It is an intrinsic property of any biological system. This more robust formulation of biodiversity offers several advantages. First, it corresponds with our intuitions. It seems that there should be some independently specifiable property of biological systems that at least roughly correlates with the common notion of biodiversity. Secondly, it makes the concept of biodiversity meaningful. As I have already said, a concept of biodiversity defined only as “that which we wish to protect” tells us nothing we do not already know. Thirdly, an intrinsic definition of biodiversity offers hope that we can agree on what biodiversity is and how we should measure it. If biodiversity were not some well-defined intrinsic feature of the world, attempts to define and measure it are prone to disagreement. By defining biodiversity in terms of some intrinsic property, we can avoid such arguments. Fourthly, such a concept of biodiversity has more political weight than an exclusively anthropocentric concept. Those calling for the protection of some socially-constructed concept of biodiversity would forever be liable to accusations of reinventing the concept to suit their particular political interests. Finally, a metaphysically strong form of biodiversity can accommodate those who do not wish to reject the intrinsic value of natural collectives. Indeed, for those who maintain that natural collectives have intrinsic value, and that biodiversity should correspond in some way to this value, a realist concept of biodiversity is a necessary prerequisite.

2.3.2 Instrumental value

The natural world clearly has instrumental value. It provides us with a host of economic benefits, in the form of foods, building materials, medicines, recreational activities and aesthetic pleasure. Wheat, for example, is a cash crop with significant economic value, as well as a foodstuff with direct utility. Trout provide a recreational

pursuit and echidnas aesthetic pleasure. Non-human instrumental value undoubtedly exists.

If there were a clear link between biodiversity and one or more types of instrumental value, then quantifying biodiversity might simply require us to quantify the types of instrumental value we wish to protect. Some types of instrumental value are, at least in principle, very easy to quantify. A good approximation of the annual commodity value of wheat could be made by summing the total monetary value of wheat traded over a year. This value could then be easily compared with the commodity value of, say, tobacco, to determine which is more worthy of conservation. This is, however, a rather bizarre way of assessing biodiversity. Biodiversity is not simply about commodity value, but about diversity. Wheat has lots of commodity value, but few would suggest it constitutes significantly more biodiversity than any other species of grass. It is just a particularly valuable aspect of biodiversity.

Another type of instrumental value that we could try to quantify is what we might loosely call aesthetic pleasure. We want to protect giant pandas, humpback whales and eastern bilbies because we think they are beautiful, lovable, interesting, awe-inspiring or in some other way endearing. Exactly how we would quantify this form of value is far from clear. Is a naked mole rat more endearing than a yellow-footed rock wallaby? Not only does the answer depend on who you ask; for many individuals there simply is no meaningful answer. One possible solution is to measure the aesthetic value of any species by counting the total monetary value that the public would be willing to pay to prevent its extinction, or else that it would be willing to accept in compensation for its loss (Randall, 1986). For many people, there is something fundamentally hollow about this reduction of aesthetic value into a purely economic form. However, there is a more fundamental objection to counting biodiversity in terms of aesthetic value. The whole point of biodiversity as a political concept is that it transcends the species-by-species approach to conservation. By

focusing our attention exclusively on *charismatic megafauna*, such as giant pandas, highland gorillas and Sumatran tigers, we risk losing the millions of far less charismatic species. Most of these species have never been observed, let alone appreciated for their aesthetic qualities. Focussing exclusively on aesthetic value largely misses the point of why biodiversity is such a potentially useful concept.

One of the central arguments for the protection of biodiversity is that it plays some function in maintaining healthy ecosystems. In the words of Paul Ehrlich,

Other organisms have provided humanity with the very basis of civilization in the form of crops, domestic animals, a wide variety of industrial products, and many important medicines. Nonetheless, the most important anthropocentric reason for preserving diversity is the role that microorganisms, plants, and animals play in providing free ecosystem services, without which society in its present form could not persist. (Ehrlich, 1986)

According to Ehrlich and many other biologists, these unglamorous species of microorganisms, plants and animals—what Edward Wilson calls “the little things that run the world” (Takacs, 1996, p. 57)—provide a host of functions such as producing oxygen, purifying water, maintaining healthy soils and controlling pests. These functions are essential both to human agriculture and the maintenance of habitats on which humans depend. To maintain these essential functions, we must maintain biodiversity. Losing one spider species may not have much effect on aphid numbers, for example, since other spider species are available to fulfil the same function, but losing lots of spider species may result in insufficient aphid predation to keep their numbers in check.

We could feasibly develop a concept of biodiversity based solely on a valuation of ecological functions—*ecological-function biodiversity*. Much of the value associated with ecological function biodiversity is what we might call *indirect* value. An indirect value depends on some other direct value for its existence. If a commercially valuable

timber species, for example, can only survive in the presence of some symbiotic fungus, then that fungus has indirect value. If we lose the fungus, we lose the valuable timber species.

There are several problems with ecological-function biodiversity, both on practical and conceptual levels. Its main practical limitation is that we have no idea what the full range of essential ecological functions is, let alone which organisms provide these functions. Such practical difficulties, however, are not my main concern. They need not bar us from developing a robust concept of biodiversity, regardless of how impractical it would be to apply such a concept. The main conceptual difficulty with ecological-function biodiversity is that it does not capture much of what we want a biodiversity concept to capture. The charismatic megafauna on which many conservationists focus, for example, provide virtually no essential ecological functions. Such species often occupy high trophic levels, and their numbers are, in most cases, simply too low to have much ecological impact. We, and the biological systems on which we depend, would comfortably get by without giant pandas, spotted owls and white rhinos. Furthermore, there are many whole ecosystems on which humans clearly do not depend. It is simply not true that destroying an isolated coral reef, say, would necessarily have a major impact on ecological functions essential to human life. Such an ecological system is largely autonomous from systems of human food production, which would therefore be largely unaffected.

Both charismatic megafauna and isolated ecosystems should, it seems, form an integral part of any sensible biodiversity concept. Defining biodiversity solely in terms of ecological functions is therefore problematic. The maintenance of ecological functions might be a reason for protecting biodiversity, but it is not the only reason. Biodiversity, it seems, is a property too fundamental to be encompassed by ecological functions alone.

Consider another popular argument for protecting biodiversity. According to Edward Wilson,

Biological diversity must be treated more seriously as a global resource, to be indexed, used, and above all, preserved. (Wilson, 1986, p. 3)

This view sees biodiversity not as just a provider of ecological functions, but as a *resource*—something valuable that we should protect now so we can exploit its usefulness both now and in the future. The value associated with this ‘future usefulness’ aspect of biodiversity is what we might call ‘option value’. An option value derives from the values that something *might* possess in future. Maclaurin calls this the value associated with ‘hedging our bets’ (Maclaurin, 1998*b*, p. 174). A wild strain of wheat, for example, may prove to be uniquely resistant to some wheat virus that will emerge some time in the future. If we can exploit the beneficial properties of this strain, it will have enormous commercial value. At present, there is some non-zero probability that this will happen.¹⁰ The present option value of the wheat strain can be thought of as the product of this probability and the value the wheat strain would have if the future scenario were to come true. Like indirect value, option value is auxiliary; it depends on other values for its existence.

A jar full of old screws, nails and other assorted bits of hardware is an example of something with option value. Chances are, most of the items inside the jar will eventually be lost, discarded, rust away or otherwise end their lives without performing any useful function. The reason we keep them is that some of them might come in handy. One day we might just need that single 10 mm M3 Philips head screw, so we keep it just in case this situation arises. If we were somehow required to discard

¹⁰ There are numerous historical examples of wild strains or close relatives of domesticated crops being used to improve their commercial counterparts. Ittis (1986), for example, discusses cases involving tomatoes and maize.

half the jar's contents, we would be forced into carrying out a form of hardware triage. How should we go about deciding which items to keep? One sensible rule would be to discard those items that have very little chance of ever being useful. A bent nail might fall into this category. A second rule might be to keep those items that would be of great value should they eventually be required, even if the probability of this happening is relatively small. The irreplaceable spare screw for fastening the arms on your sunglasses might fall into this category. A third rule would be to cull predominantly from those items of which there are many identical or similar examples. If there are 27 one-inch flat head nails and you suspect you will only ever need three at the most, then you can safely throw away the other 24. This last rule is telling us to preserve the most diverse set of items, where diversity is measured with respect to some set of functionally useful attributes.¹¹

Carrying out triage with the constraint of maximising biological option value would mean choosing that set of biological elements that maximises the sum total of all the probabilistically-weighted future values of all the potentially useful resources of the biological world. Since in most cases we do not know which elements of the biological world are potentially useful resources, it makes sense to select a diversity of elements to increase our chances of picking the right ones. In other words, we should hedge our bets by picking a biodiverse set. Maximising option value means maximising biodiversity.

¹¹ Defining future functionally useful attributes can be problematic. For example, on the basis of present knowledge, we might discard all black nails because we predict that the otherwise identical silver ones will be functionally identical. We might eventually find, however, that the black nails match a particular piece of furniture better or that their blackness was due to a nitride coating that makes them more corrosion resistant and therefore more suitable for certain applications. Nevertheless, it seems we have no choice other than to make a best guess at what we suspect are the future functional attributes based on our limited existing knowledge.

Consider a simple example. Suppose we have a set of three endangered species and we can only afford to spend money on conserving two of them. With all other considerations being equal, we should focus our efforts on those two species that maximise option value. Suppose two of the species are wild tomatoes and the other is a wild potato. In the absence of any other knowledge, our only rational strategy would be to select the two most dissimilar species—presumably the potato and one of the tomatoes.

The main failing of option-value as way of defining or even understanding biodiversity is that, like *ecological-function biodiversity*, it is not robust enough to capture everything we want from a biodiversity concept. Option value may provide a strong reason for preserving biodiversity, but biodiversity does not in itself capture every aspect of option value. Recall that with the jar-of-hardware example, there were two rules other than maximising diversity. One was to discard obviously useless things and the other was to keep obviously useful things. Rather than rely solely on diversity as the criterion for selection, we should use whatever knowledge we have at our disposal to assist us in maximising option value. Similarly, in maximising biological option value, we should use present knowledge to help us assess which species will be useful in future. All domesticated crops, for example, are extremely useful now and will almost certainly remain so in future. They should therefore figure prominently in any option value measure. However, wheat or any other domesticated crop does not appear to represent significantly more biodiversity than a host of wild counterparts. Hence there is an incongruity between measuring option value and measuring biodiversity.

We might concede this point but argue that the lack of congruence between option value and biodiversity only applies to specific species such as domestic crops. Since we are vastly ignorant of the potential uses of most species, biodiversity remains the only sensible criterion for assessing the option value of such species. This is not an

unreasonable argument. However, it leaves us no closer to spelling out what biodiversity is. All we are left with is some sense of congruence between biodiversity and a limited form of option value without any criterion for independently determining either. If we are to formulate option value in terms of biodiversity, we need first to determine what biodiversity is. Alternatively, if we are to use option value to define biodiversity, we need to spell out ‘likely future usefulness’ in more concrete terms. In Chapter 3, I will discuss this issue in more detail when examining Maclaurin’s concept of ‘raw biodiversity’.

Whatever the concept of biodiversity is supposed to mean, its relationship with the many types of instrumental value is not straightforward. Biodiversity is fundamentally about *diversity*, not instrumental value. Lineages with high instrumental value such as domestic wheat do not constitute significantly more biodiversity than many other species with low instrumental value. Conversely, taxonomically and morphologically distinct species such as echidnas are often thought to represent more biodiversity than less distinct species such as bottlenose dolphins, but do not necessarily have more instrumental value. Whatever our concept of biodiversity might come to represent, there are many forms of instrumental value that the concept will be unable to encompass.

2.3.3 Inherent value

Formulating biodiversity in terms of intrinsic value proved difficult because there is no straightforward way of assigning intrinsic value to the sorts of units that biodiversity deals with, such as species. There is simply too much doubt over the very existence of such collective intrinsic value. The existence of collective instrumental value, on the other hand, is uncontroversial, but there is a lack of congruence between the various forms of instrumental value and the sort of thing we want a

biodiversity concept to capture. Inherent value offers a possible solution to this impasse.

All of the types of instrumental value discussed above can be used to generate persuasive arguments for protecting the biological world. Together, however, they fail to capture all of our reasons for wanting to protect biodiversity. Suppose we were able to strip the biological world of all its instrumental value. So we might have machines that performed various functions such as making food, creating oxygen and so on, as well as artificially satisfying all recreational and aesthetic desires. We would also have to satisfy all option value by, for example, curing all diseases. Having thus satisfied all instrumental value, would it matter if we destroyed all remaining rainforests, coral reefs and other zones of high biodiversity? Nearly all of us, I suspect, regardless of whether we think the biological world can have intrinsic value, would unhesitatingly answer “yes”. This shows us that the biological world has some value that is neither intrinsic nor instrumental; in other words inherent value.

We have a duty *regarding* but not necessarily *to* something with inherent value. A valued painting, for example, ought to be protected from destruction, not because we owe it to the painting, but because we owe it to those who value the painting. Similarly, we may have a duty to cull individual African elephants if that will help protect the species, not because we owe it to the elephants (or the species), but because we owe it to those who value African elephants, including ourselves if we are one of those people.

The inherent value of the biological world can be expressed simply as a love of nature—what Edward Wilson calls *biophilia*; an “innate affinity for the natural world” (Wilson, 1984). Whether or not we possess biophilia ourselves, the existence of it cannot be disputed. It is empirically observable. People love nature. The sort of value on which biophilia is based is explicitly non-intrinsic. Without a human valuer it

would not exist. However, neither is it instrumental. Inherent and instrumental value are quite different. Instrumental value can be reduced to basic human desires, such as the desire to be alive rather than dead, warm rather than cold, nourished rather than hungry. The desire to protect that which has instrumental value can always be given a rational justification in terms of these more basic human desires. However, we are at a loss when it comes to mounting a rational argument for preserving the natural world once stripped of instrumental value. The residual value—inherent value—seems to spring directly from a desire without any intermediate rational justification.

McQuillan uses the moral philosophy of David Hume to make some sense of this situation. For Hume,

Reason is and ought to be the slave of the passions and can never pretend to any other office than to serve and obey them. (Hume, 1978, p. 415)

In other words, passions such as desire are primary, while reason is secondary. Hence there is no justification for protecting nature's inherent value other than our desire to do so. McQuillan states this position quite plainly.

... in keeping with Hume's privileging of passion over reason, the source of nature's inherent ... value lies not in its instrumentality but in the passion that we have for it. (McQuillan, 1998, p. 322)

The sort of biodiversity concept commonly employed in environmental debates is, I think, deeply connected to the inherent value of nature. Citing ecosystem functions and cures for cancer as justifications for protecting biodiversity are weak rationalisations of a deeper passion. Hence inherent value is likely to be the best starting point for a conservation-based concept of biodiversity. However, this leaves me with some problems. I am trying to construct a metaphysically strong concept of biodiversity. I want biodiversity to be an intrinsic property; something that exists outside the human mind. Yet, the type of value that we want our biodiversity concept to best represent

is explicitly anthropocentric. It exists only in the human mind. Although there is no contradiction involved in our anthropocentric valuing of some non-anthropocentric property, inherent value leaves me no closer to specifying what an intrinsic concept of biodiversity looks like. Moreover, decomposing inherent value into simpler elements is problematic. Nature's inherent value, it seems, springs forth *in toto* from some collective passion.

Despite these problems, the type of biodiversity concept I have in mind is a good candidate for capturing nature's inherent value. It at least comes close to describing what gives rise to my impassioned plea for protecting the rich diversity of life. Later on I hope to provide some justification for linking the type of biodiversity concept I will develop and nature's inherent value. The justification is necessarily weak, however, because inherent value is not a quality that can be easily analysed. I can only hope that my desires correspond with yours.

3 Measuring biodiversity

When you can measure what you are speaking about, and express it in numbers, you know something about it: but when you cannot measure it, your knowledge is of a meagre and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science.

— Lord Kelvin

It is generally acknowledged that all organic beings have been formed on two great laws—Unity of Type and the Conditions of Existence. By unity of type is meant that fundamental agreement in structure, which we see in organic beings of the same class, and which is quite independent of their habitats of life. On my theory, unity of type is explained by unity of descent. The expression of conditions of existence, so often insisted upon by the illustrious Cuvier, is fully embraced by the principle of natural selection. For natural selection acts by either now adapting the varying parts of each being to its organic conditions of life; or by having adapted them in long-past periods of time

...

— Charles Darwin

The diversity of living forms is apparent to us all. Beetles and whales are vastly different creatures. Just as apparent is the unity of such forms; the extent to which organisms can be grouped by similarity. Virtually all multicellular organisms can be grouped by species, and these species can in turn be organised into hierarchical categories. Taxonomy is the science of classifying organisms. It is a science whose roots extend at least as far back as Aristotle and for which parallels can be found in the folk taxonomies of numerous cultures. At least since Darwin, the patterns of life recognised by taxonomy have been explained in terms of both common descent—Darwin’s Unity of Type—and adaptation to the environment—his Conditions of Existence. Interpreting the patterns of nature through the lens of these ‘two great laws’ has provided modern taxonomy with a powerful means of organising biological information. In recent years, taxonomy has in turn been used as a tool for carrying out what Robert May calls ‘the calculus of biodiversity’ (Takacs, 1996, p. 61).

The concept of actually measuring biodiversity in an apparently mathematically rigorous fashion has attracted some strong criticism. Daily and Ehrlich have described it as

... ‘crackpot rigour’ (detailed mathematical analyses of an intractable problem) or ‘suboptimization’ (doing in the very best way something that should not be done at all). (Daily & Ehrlich, 1992)

Ehrlich also describes it as “mental masturbation” (Takacs, 1996, p. 62). However, as with biological triage, these objections are largely aimed at practical rather than conceptual aspects of measuring biodiversity. No one doubts that attempting to measure biodiversity with any sort of mathematic rigour is an Herculean task. However, my main concern is not to address the practical aspects of measuring biodiversity, but simply to understand what biodiversity is, how we might measure it and what sort of property it might be. If I am able to demonstrate that biodiversity can, at least in principle, be measured in a mathematically rigorous fashion, then I have taken a large step towards understanding what biodiversity is and perhaps showing that it is an intrinsic property of biological systems. However, before I look at taxonomic approaches to measuring biodiversity, I need to consider the aims and methodologies of taxonomy itself.

3.1 A tale of three taxonomies

In recent decades, the science of taxonomy has undergone significant change as competing approaches to the taxonomic classification of organisms have been advocated (Hull, 1988). Each approach can generally be classified under one of three labels—evolutionary taxonomy, phenetics and cladistics. These three approaches differ markedly, not just in the classifications they produce, but in their aims, methodologies and philosophical outlooks. These differences must therefore be examined in order to understand the biodiversity measures to which they give rise.

3.1.1 Evolutionary taxonomy

Modern evolutionary taxonomy is a system that combines elements of evolutionary thinking with aspects of the Linnean and other systems of classification that predate the Darwinian revolution. The traditional element of evolutionary taxonomy is the recognition that similarities between organisms allow them to be grouped and ranked to form a nested hierarchy of taxa. Thus a highly similar population of interbreeding organisms are grouped into a species (e.g. *Homo sapiens*), several highly similar species are grouped into a genus (e.g. *Homo*), several genera into a family (Hominidae) and so on through the ascending levels of order, class, phylum and kingdom. The notion of common ancestry central to modern Darwinian thinking later explained these patterns of similarity and diversity among organisms, thus allowing pre-Darwinian taxonomy to be reinterpreted within this Darwinian context without undergoing radical change.

The main feature of evolutionary taxonomy is its attempt to classify and rank organisms on the basis of two criteria:

1. phylogenetic branching (i.e. the relationships between species in respect of common ancestry), and
2. the degree and type of evolutionary change between branching points.

In other words, the evolutionary taxonomist attempts to group organisms on the basis of both ancestry and morphology.

In recent decades, evolutionary taxonomy has been heavily criticised and largely replaced by alternative systematic techniques (Hull, 1988). One of the central criticisms is that evolutionary taxonomy attempts to classify organisms on the basis of two incompatible criteria: phylogeny and morphology. Although these two criteria often produce congruent classifications, frequently they do not. Similarity often belies ancestry. For example, there is significant morphological (and ecological)

affinity between tuataras and lizards, despite their closest common ancestor predating the divergence of lizard and snake lineages. Deciding how to classify organisms in such cases can be unclear. Should tuataras be classified as primitive lizards or something completely different to either lizards or snakes? An evolutionary taxonomist is forced to decide such questions by using some arbitrary weighting of both morphological and phylogenetic factors. Because of this 'intuitive' approach to classification, the methods of evolutionary taxonomy are almost impossible to codify explicitly. Instead they are developed over many years of experience. Although different taxonomists working in the same area tend to reach general agreement in their classifications, problems arise when attempts are made to explicitly justifying these classifications, teach the methods used to derive them, or apply the methods used in one class or phylum to another.

3.1.2 Phenetics

In response to some of the problems posed by evolutionary taxonomy, an alternative classificatory scheme known as *phenetics*¹² was developed in the early 1960s, principally by Sokal and Sneath (Hull, 1988). Unlike the 'intuitive' approach used by evolutionary taxonomists, pheneticists use a standardised classificatory procedure based on quantitative techniques to measure the degree of 'overall similarity' among groups of organisms. The similarity measure employed is based on the presence or absence of numerous unweighted characters or character states.

By using a standardised quantitative technique, pheneticists hoped to remove subjectivity and ambiguity from taxonomy. Although some convergent or highly variable characters might obscure the observed pattern, reproducible results, it was

¹² Also referred to as *numerical phenetics* or *numerical taxonomy*.

claimed, should arise whenever sufficient characters are used for analysis. However, such hopes have proved unfounded, leading to the abandonment of phenetic techniques of classification (de Queiroz & Good, 1997; Hull, 1988). The calculated degree of similarity between any two organisms depends on which characters are measured. So, for example, cuttlefish and rabbits could be classified as similar on the basis that both have two eyes, can be kept as pets, and are found on restaurant menus in provincial France. Furthermore, significant morphological differences frequently exist between different members of a single polymorphic species.¹³ A purely phenetic classification would necessarily classify these organisms into separate taxa. The infinity of similarities and differences between any two organisms means that the use of more characters cannot, on its own, guarantee a stable classification. The hope that phenetic measures would converge on a single classificatory scheme by including more characters has proved false. Hence the objective ‘theory-free’ approach of phenetics has proved ill-founded, since some theory must be introduced to decide which characters are relevant and which are not (Hull, 1988).

Despite its problems, phenetics represents a positive step towards separating the two divergent aims of evolutionary taxonomy. In grouping organisms only by morphological similarity, a phenetic classification eschews any claims about ancestry.¹⁴ Furthermore, by demonstrating the limitations of evolutionary taxonomy and employing mathematical analysis to classification, phenetics proved highly influential on a third approach to classification, an approach that has come to be known as *cladistics*.

¹³ For example, caterpillars and butterflies, or even males and females of many species such as garfish (Hull, 1988).

¹⁴ An implicit hope of many pheneticists, however, was that a purely phenetic classification would nevertheless reveal ancestral patterns.

3.1.3 Cladistics

Like phenetics, cladistics also recognises the conflicting aims of evolutionary taxonomy in trying to capture both phylogenetic and morphological affinity. With cladistics, however, this conflict is resolved by focusing on the phylogenetic component. The central idea of cladistics is that systematics represents evolutionary history; the job of the taxonomist is to deduce the phylogenetic relationships among organisms, not their morphological similarities. Cladistics makes a further metaphysical claim that real taxonomic units are necessarily the *monophyletic* groups. A monophyletic group is one for which all and only the descendants of a single ancestral species are members.¹⁵ Thus the cetaceans are a monophyletic group, since all cetaceans share a common ancestor that was not the ancestor for any other living organism. In contrast, the monkeys are *paraphyletic*, since there was no common ancestor of all monkeys that was not also an ancestor of the apes, while the cyanobacteria are *polyphyletic*, since this grouping consists of an assemblage of organisms that evolved common traits independently.

As with phenetics, the methodology of cladistics is based on trait analysis. Only certain types of traits, however, are informative for constructing evolutionary history. A *unique trait* possessed by only one species gives no information about that species' relationship with other species. Similarly, a *primitive trait* that was inherited by all members of a particular group gives no information about the relationships within that group, although at a higher level in the tree, it may help infer the relationships linking the entire group with other groups. The only type of traits that are informative in cladistic analysis are *derived traits*; ones that vary within a group.

¹⁵ The monophyletic group consisting of an organism and all of its descendants is known as a 'clade'.

In cladistic analysis, species are linked on the basis of shared derived traits. If species A and B possess a certain trait, and species C does not, then A and B are likely to be more closely related to one another than either is to C. Of course, this may not be the case. C may be more closely related to B than either is to A, either because C lost the trait that the common ancestor of A, B and C all shared, or because A and B evolved the trait independently. These possibilities are illustrated in the ‘cladograms’ of Figure 1 below.




Figure 1: Cladistics infers phylogeny from character analysis. In the data, ‘1’ indicates the possession of a particular trait, whereas ‘0’ indicates its absence. In the absence of other information, the true phylogenetic relationships are more likely to correspond to (a) than (b), since (a) requires fewer character state changes; i.e., it is more parsimonious (after Sober, 1988).

In cladistic analysis, we decide between the two competing hypotheses of ancestry illustrated in Figure 1 by choosing the most ‘parsimonious’ cladogram; that is, the one that requires the fewest character state changes. The hypothesis represented by Figure 1a requires the assumption only that a single character state change occurred somewhere along the lineage that gave rise to A and B, whereas that represented by Figure 1b requires the assumption of at least two character state changes, with the two possible ways in which such changes could have occurred illustrated. Figure 1a is therefore more parsimonious than 1b. In practice, numerous character state changes

are possible, and hence the most parsimonious cladogram will not always reflect true ancestral relationships. Nevertheless, such a methodology reflects a 'best guess' inference of such relationships on the basis of available evidence. Moreover, the inclusion of additional data into the analysis can help resolve such discrepancies, resulting in convergence towards a single, stable classification (Patterson, 1977).

The key advantage of cladistics over both evolutionary taxonomy and phenetics is that its central aim is clearly defined. This aim is to ascertain knowledge of the 'one true tree of life'; the phylogenetic relationships among existing organisms. The existence of such a tree is uncontroversial. In contrast, the notion of 'biological similarity' that both evolutionary taxonomists and pheneticists try to measure has yet to be clearly defined. Cladistics avoids this problem by reformulating the role of biological similarity. Rather than using similarity to define taxonomic affinity, cladistics uses taxonomic affinity, defined phylogenetically, to explain similarity. In other words, cladistics has removed similarity from a classificatory role and made it part of taxonomy's explanatory agenda.

Cladistics is not without its methodological problems. However, the achievements of cladistics are more than just methodological. It has clearly established phylogenetics as the central aim of taxonomy, giving rise to what is now commonly referred to as 'phylogenetic systematics'. Once agreement has been established about the fundamental aim of taxonomy, many mathematical methods for estimating phylogeny other than parsimony analysis become permissible (even phenetics). Once derided as a higher form of stamp collecting, taxonomy has now secured a central position in the biological sciences. It can no longer be seen as simply collecting and classifying the gifts of a benevolent creator to impose an order upon nature's untidiness, but as a means for compiling detailed knowledge of earth's evolutionary history. Taxonomic data now play an important role in resolving important theoretical and conceptual issues in biology, such as in understanding the notion of evolutionary constraint

central to debates over ‘adaptationism’.¹⁶ Taxonomy is also pivotal to the biodiversity question, because it is taxonomy on which existing biodiversity measures are based.

3.2 Taxonomic biodiversity

In recent years, a relatively rich literature has been produced concerning the taxonomic measurement of biodiversity. There are several motivations for constructing a measure of biodiversity. One is to demonstrate that biodiversity is a robust concept; defining what biodiversity is through the development of a well-defined measurement procedure. A more practical motivation is the development of procedures to assist conservation decision-making, such as which regions of a forested area to protect from logging. This may involve developing not just a technique of biodiversity measurement, but also things like optimisation procedures for area-selection.¹⁷

Given the difficulties we have already encountered in trying to define a precise notion of biodiversity, it is perhaps unsurprising that there exists a variety of taxonomic biodiversity measures; there is simply not yet enough agreement about such fundamental questions as what biodiversity is. Hence many of the differences between measures arise because their target properties differ—they are trying to measure different things. Biodiversity measures also differ in many instances because they are based on different *indicators*.¹⁸ Measuring biodiversity is not just conceptually but also practically difficult, and so some other property is often measured that,

¹⁶ See, for example, Gould & Lewontin (1978), Dennett (1995) and Sterelny & Griffiths (1999).

¹⁷ A key notion here is *complementarity*. A region is more worthy of protection if it not only has high biodiversity; but also contains species not already contained in other protected regions.

¹⁸ Sarkar (1999) introduces the terms *surrogate* (or *true surrogate*) and *estimator surrogate* in this context. For Sarkar, an estimator surrogate is akin to what I call an indicator, whereas a true surrogate is a target property that supposedly represents ‘overall biodiversity’. Sarkar presupposes that ‘biodiversity’ is an ill-defined concept, not reducible to a single target property. I eschew this terminology because it is unnecessarily confusing for my purposes.

hopefully, is a good indicator of a more fundamental target property (*true biodiversity*). In the remainder of this section, I will give an overview of the various types of taxonomic biodiversity measure.

3.2.1 Species richness

In Chapter 1, we met the concept of ‘species richness’¹⁹, which is simply the number of species within a particular area or ecosystem. Although species richness is often used as a crude definition of biodiversity, most biologists think it is simply an indicator of some more fundamental property (Takacs, 1996). The main limitations of species richness as a biodiversity measure have already been discussed. Measuring the species richness of an assemblage accounts only for the number of species, and not how much those species might differ from one another. Beetle species are numerous and in many cases highly similar, whereas whale species are few and relatively disparate. Red wolves attract considerable conservation effort, despite their status as a mere subspecies. It is generally accepted that an ideal biodiversity measure should somehow account for these differences between species, and not just sum them with equal weighting.

Despite such limitations, there are many good reasons for retaining species richness as a biodiversity measure. Gaston (1996, pp.78–79) lists five.

1. Species richness can be correlated to many measures of ecological diversity and it functions better than most other potential indicators including species diversity.

¹⁹ ‘Species diversity’ is sometimes defined as a different measure that also accounts for the relative abundance of different species (Sarkar, 1999).

2. There is usually a positive correlation between species richness and higher taxon richness (see following section), allowing the latter to be used as a good indicator of species richness.
3. When species richness is relatively high, it is correlated with character richness.
4. Some parameters used to measure the complexity of community webs (including the number of edges and the length of directed paths) seem to be correlated with species richness. (This remains somewhat controversial.)
5. Relatively high species richness is correlated with increasing topographic diversity. This is only to be expected insofar as topographic diversity not only potentially allows more niches but also may encourage reproductive isolation and speciation.

Given these advantages, as well as its conceptual simplicity,²⁰ species richness remains the most prevalent measure of biodiversity.

3.2.2 Higher taxon richness

Higher taxa are the various groupings of species that form the hierarchical classifications of taxonomy; the genera, families, orders, classes, phyla and kingdoms inherited from the Linnean system. Measuring higher taxon richness is analogous to measuring species richness, but carried out at a higher level of the taxonomic hierarchy. There are two clear advantages to such an approach. One is that higher taxa are less numerous and easier to distinguish than individual species, making them an easier sampling unit in field surveys. A second is that such ‘coarse graining’ partly avoids the problem of counting each member of speciose taxa as a single unit of biodiversity, regardless of their degree of similarity (Williams & Gaston, 1994).

²⁰ Although a single species definition has proved elusive (Mallet, 1995; Sterelny & Griffiths, 1999), particular species are, at least outside the microbial world (Sogin & Hinkle, 1996), typically well-defined. As mentioned in Chapter 1, however, species from diverse clades are not necessarily commensurate.

A serious problem with using higher taxon richness as a measure of biodiversity is that higher taxa are not robustly defined. Cladistics enables us to organise species into a nested hierarchy of monophyletic groups, but it offers no advice on where to draw the boundaries between genera, families or any other higher taxonomic rank. From a cladistic perspective, these higher taxa make little sense. There is simply no well-defined notion of ‘evolutionary divergence’ for determining whether a particular monophyletic group of species constitutes a genus, a family or an order (Sterelny & Griffiths, 1999).

Despite this problem, there are nevertheless some conventions, albeit not always well-defined ones, which are used to delineate higher taxa. Such conventions are typically based on morphological disparity; just the sort of target property that we might want a biodiversity measure to capture. The suitability of higher taxon diversity is further supported by empirical evidence that it provides a good indicator for more detailed morphological measures based on ‘character richness’, as discussed below (Williams & Humphreys, 1996).

3.2.3 Genetic diversity

Biodiversity is commonly held to be manifest at three levels: genes, species and ecosystems. Genetic diversity is the lowest of these levels. Advocates of this type of biodiversity typically make either a weak or strong claim regarding the status of genetic diversity. The weak claim is that genetic diversity provides a simple and practical indicator of biodiversity. The strong claim is that genetic diversity *is* the fundamental unit of biodiversity.

Advocating the utility of genetic diversity measures is not without foundation. DNA is operationally useful. Base pairs are discrete and easily counted units that can be compared across widely different species. Moreover, there is evidence to support the use of genetic distance as an indicator of phenotypic diversity (Templeton, 1994;

O'Donnell *et al.*, 1994; Williams & Humphries, 1996) and hence for the substitution of genetic distance for phylogenetic distance in taxonomic distinctness measures (Crozier, 1992).

Nevertheless, there are substantial problems with using genetic diversity to measure phenotypic diversity. Primary among these is the lack of any simple mapping between DNA structures and phenotypic outcomes. In most organisms, most of the genome does not code for protein structures, whereas some portions of the genome may code for multiple traits. Small changes in DNA structure can lead to massive phenotypic change, while relatively large genetic changes such as reversals, translocations or allopolyploidy can lead to trivial phenotypic change. Determining the outcome of a genetic change in order to meaningfully define genetic distance requires some sort of 'reading back' from the phenotypic and other higher levels. These observations undermine the assumption that the probability of a character state change is linearly related to genetic distance. Hence genetic distance is at best an unreliable indicator of phylogenetic distance.

The second and stronger approach to using genetic diversity is to claim it as a more fundamental form of biodiversity of which phenotypic and ecosystem diversity are simply manifestations. Harper and Hawksworth (1994), for example, claim that:

Unlike higher taxa which may be based on characters which are not necessarily comparable, the DNA and RNA found in all living cells can provide a basis on which to make direct comparisons between diverse organisms. There is a sense in which the biodiversity of a community is expressed as the sum of the variety of genetic information coded within the genotypes of the inhabitants. A biodiversity calculus could be envisaged for which we ask of the various species (and individuals) in a community how many new base sequences they contribute to the genetic vocabulary of the whole. (p. 8)

In a similar vein, Mallet (1996) asserts that:

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. (p. 13)

This stronger thesis relies on the notion that phenotypes can in some sense be reduced to their genotypes. The points already made against genetic diversity as an operational measure undercut this argument. Without rehearsing in detail the debate over ‘genetic reductionism’²¹, an additional point taken from this debate, namely the parity thesis, weakens it further. It is well known that many factors other than the genome are necessary for embryonic development. Whatever can be said about the causal and informational roles of genes in development can also be said about epigenetic factors (Oyama, 1985; Griffiths & Gray, 1994). Cell membranes, methylation patterns, parental instruction and a host of other essential developmental resources are transmitted from parent to offspring in much the same manner as genetic resources. Genotypes do not define their phenotypes any more than do epigenetic factors, and hence genetic diversity is no more fundamental than phenotypic diversity.

Another reason we might reject genetic diversity as the fundamental unit of biodiversity is that genetic diversity is just not the sort of thing we care about. We might marvel at the blue feet of a blue-footed booby, but few of us would get too excited about the alleles that produce this blueness. Preserving a particular species may require us to maintain a certain degree of genetic diversity, but this is a property we need only satisfy, not one we want to maximise for its own sake. However, this argument does not rule out biodiversity being some fundamental property that we do

²¹ See, for example, Sterelny and Griffiths (1999) for a recent summary.

not value in its own right, but which is nonetheless something worth preserving because we value its consequences.

3.2.4 Phylogenetic diversity measures

Phylogenetic diversity measures are based on the observation that more distantly related species tend to be more distinct in terms of their morphological or other characters. They therefore use relatedness as a criterion for assessing what weighting each species should be given. Atkinson sums up the rationale behind this approach.

... given two threatened taxa, one a species not closely related to other living species, it seems reasonable to give priority to the taxonomically distinct form. (Atkinson, 1989)

There are two basic strategies for measuring phylogenetic diversity. One strategy, as exemplified by Vane-Wright *et al.* (1991), is to define a procedure for calculating the relative *weights* for each taxon within a set of taxa. The weight of a taxon is akin to the value of that taxon. If we were then to use this information to carry out triage, we would give highest priority to protecting those taxa with the greatest weights. The second approach, as exemplified by Faith (1992), is to define a procedure for calculating the weights, not of individual taxa, but of subsets of the main set of taxa. Hence if we were to use this information to carry out triage, we would first ask how many taxa we are able to protect, and then select the subset of this size with the greatest total weight. This approach recognises the conservation value of a taxon as being context-dependent; it varies depending on which other taxa are being considered for protection. In other words, it accounts for complementarity.

Maclaurin (1998*b*) introduces a nice analogy to illustrate the differences between these two strategies. Suppose we liken two followers of each of the above two strategies to stamp collectors. The first Maclaurin labels ‘the investor’, and the second he labels ‘the enthusiast’ (p. 144). The investor is interested only in the value

of her collection, whereas the enthusiast is interested only in its completeness. The difference between them is best exemplified when the two collectors are given the opportunity to buy a large number of rare but similar stamps. The investor will buy as many of them as she can afford. Since they are all rare, they will all increase the value of her collection. The enthusiast, on the other hand, will buy only a few. He cares only about the completeness of his collection. Once he has bought enough to satisfy completeness, the remaining stamps on offer are of less value to him. Applying this analogy to phylogenetic measures of biodiversity, the investor strategy would have us save both species of tuatara. They are both rare and both valuable. The enthusiast strategy, on the other hand, would have us save at least one species of tuatara. Once we have secured its future, saving the other species is of less importance.²²

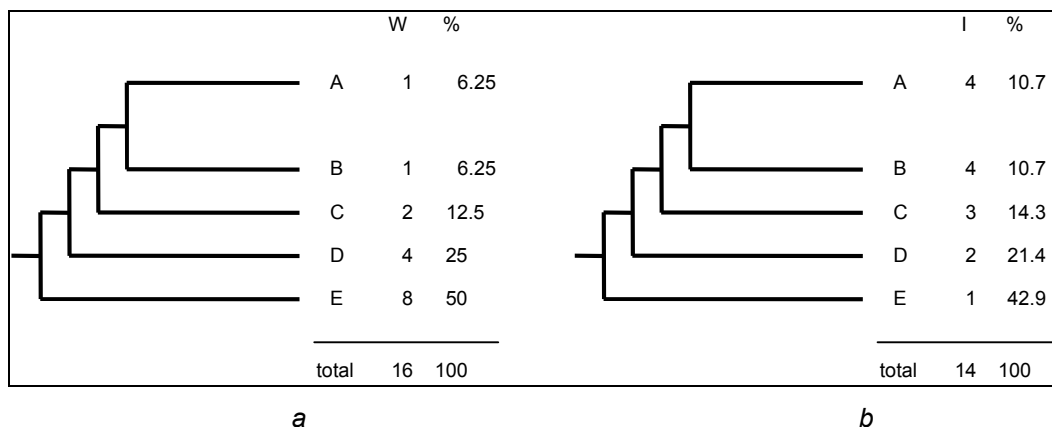


Figure 2: Phylogenetic biodiversity measures. (a) Equal weighting for sister groups. Column W lists the weighting applied to each species, which equals the aggregate weighting of its sister group. (b) Taxonomic distinctness. Column I lists the number of clades to which each species belongs, the normalised reciprocals of which are used to determine each weighting (after Vane-Wright *et al.*, 1991).

²² In practice, however, trying to protect both species of tuatara may be the best strategy for securing the future of at least one species.

Figure 2 above shows two phylogenetic diversity measures of the investor type, introduced by Vane-Wright *et al.* (1991). These are labelled ‘equal weighting for sister groups’ and ‘taxonomic distinctness’.

In cladistics, sister groups are two clades separated by a single speciation event. So in Figure 1*a*, the clade consisting solely of the species C is the sister group of the clade consisting of species A and B, while the clades formed by A and B are themselves sister groups. The phylogenetic diversity measure based on equal weighting for sister groups considers sister groups to be of equal importance. So if we could only afford to preserve two species from A, B and C in Figure 1*a*, we would preserve C and either A or B. Figure 2*a* illustrates this procedure for a more complicated cladogram. The two most closely related species are first given a weighting of one, giving a total weighting of two for the clade they comprise. The sister group to this clade is also given a weighting of two, and so on down the cladogram.

The principal objection to applying equal weighting to sister groups is that it seems to weight phylogenetically distinct species too heavily (Vane-Wright *et al.*, 1991; May, 1994; Maclaurin, 1998*b*). The coelacanth,²³ for example, is the sister group of a large clade that includes all terrestrial vertebrates. According to a biodiversity measure that applies equal weighting to sister groups, coelacanths are more valuable than every species of mammal, bird, lizard and frog combined. Vane-Wright *et al.* (1991) offer an alternative weighting scheme of ‘taxonomic distinctness’ that avoids this problem, as shown in Figure 2*b*. In this system, each species is assigned a number corresponding to the number of clades of which it is a member. These numbers express the ‘information content’ represented by each branch of the cladogram; in other words,

²³ A rare and highly unusual species of lobe-finned fish found near the Comoros Islands, between Africa and Madagascar, and recently discovered off the coast of Sulawesi in Indonesia.

the number of visible speciation events²⁴ between the ancestor of the entire group and the species in question. The weighting for each species is then the reciprocal of this number, normalised to some arbitrary total. This system has several advantages over the ‘equal weight for sister groups’ strategy. One is that species of equal taxonomic rank are given equal weighting. Another is that groups with many species have a higher aggregate weight than sister groups with fewer species, avoiding the problem of excessive weighting of phylogenetically distinct species.

The basic approach of enthusiast-strategy phylogenetic diversity measures is to take the cladogram for the set of taxa in question, apply a weight (length) to each branch according to some rationale regarding ‘evolutionary distance’, and then select the subset of taxa that maximises overall intervening branch length (Humphries *et al.*, 1995). Figures 3–5 illustrates this procedure. Three cladograms of ten taxa are shown, with branch lengths weighted using different criteria. The subsets of three taxa with maximum intervening branch length are highlighted.

The three weighting schemes typically used to determine evolutionary distance are:

Clock model: Branch lengths are weighted chronologically. Evolutionary change is assumed to occur at a constant rate along each lineage (Figure 3).

²⁴ Visible in the sense that they are represented by extant species.




Figure 3: Clock model of branch weighting (after Humphries *et al.*, 1995).

Sample model: Branch lengths are weighted by extrapolating from some sample of character state change data (typically genetic or morphological). Evolutionary change is assumed to occur at varying rates, as represented by the character state change data (Figure 4).

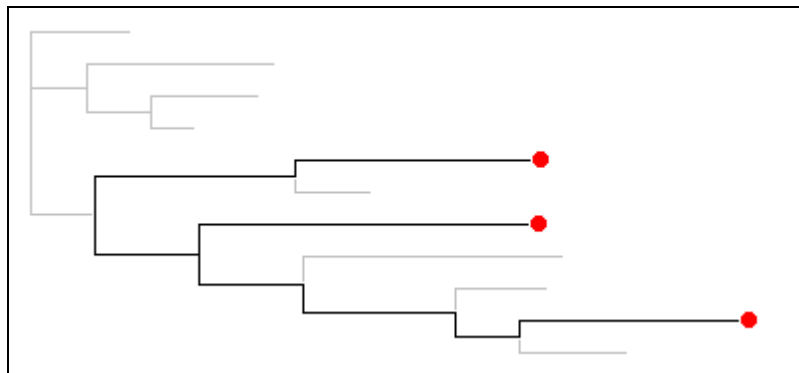


Figure 4: Sample model of branch weighting (after Humphries *et al.*, 1995).

Saltatory model: Branch lengths are weighted by the number of intervening speciation events. A fixed amount of evolutionary change is assumed to be associated with each speciation event (Figure 5).




Figure 5: Saltatory model of branch weighting (after Humphries *et al.*, 1995).

All three of the above weighting schemes have been developed to cope with the central problem of such phylogenetic diversity measures; the sheer quantity of data required to accurately determine branch length. However, none of the three weighting schemes is without its own difficulties. Clearly, rates of evolution have varied greatly over time. Such morphologically diverse species as whales and goats diverged from a common ancestor only a few tens of millions of years ago, while over similar time scales, ‘living fossils’ such as coelacanths and Wollemi pines have undergone almost no evolutionary change. Only the sample model has some hope of adequately accounting for this problem, although the quantity of data required is in many cases still prohibitive.

The sample model of measuring phylogenetic distance also goes under the name of ‘character richness’. Figure 6 provides a further illustration of this approach, based on the original method developed by Faith (1992). In this example, the character state change data are superimposed onto the branches of the cladogram where such state changes are inferred to have occurred. The subset of, say, four organisms with the greatest proportion of total character richness can then be found by choosing the spanning path connecting four organisms that intersects the most character state changes (as well as speciation events, in Faith’s example).

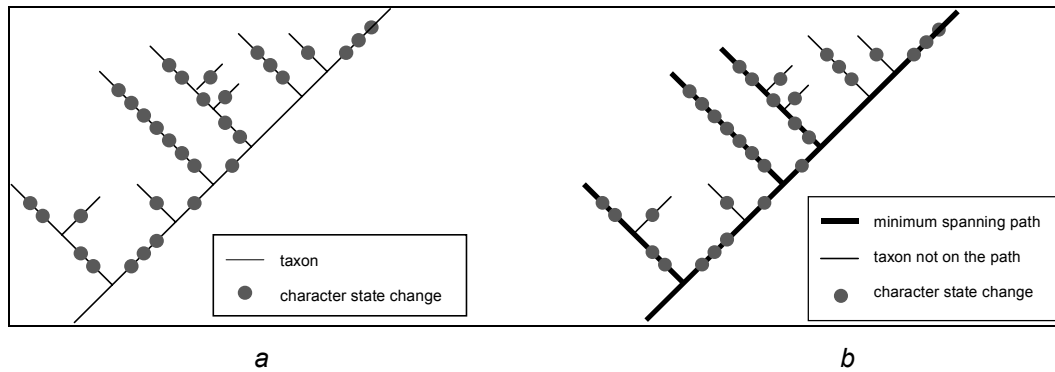


Figure 6: Another phylogenetic biodiversity measure based on character richness. (a) A cladogram on which has been superimposed the inferred character state changes. (b) The most diverse set of four species is that which includes the most character state changes and speciation events, as shown (after Faith, 1992).

Of the various approaches to measuring biodiversity examined so far, character richness appears to come closest to a sort of universal biodiversity measure that accounts for variation both within and between species. Measures based on species and higher-taxa may provide good indicators of biodiversity, but they do not fully capture the desired target property, if only because they fail to adequately account for the degree of difference between species. A measure that does take such degrees of difference into account must be based on a smaller currency unit; a unit such as character.

We can either interpret character richness as the true currency of biodiversity—the target property we want to measure—or as an indicator of some more fundamental target property yet to be specified. Faith, for example, adopts the former position, although he is by no means alone (e.g. Vane-Wright *et al.*, 1991; Williams & Humphreys, 1996). According to Faith,

[the] fundamental level of organismal variation corresponds to features or attributes of species. When biodiversity is defined at this level, the oft-stated conservation goal of ‘protecting biodiversity’ translates into protecting as much of this feature-diversity as possible. (Faith, 1994, p. 46)

Faith's rationale for this stance is that feature diversity corresponds to option value, as defined by

an attempt to keep options open, for realizing values of species in future, by saving as much biodiversity as possible now. (p. 46)

If we want to maximise option value, we should, according to Faith, maximise feature diversity, since

the greater the number of different features represented by a protected subset of taxa, the greater the option value. (p. 46)

So the important attribute of biodiversity that we call 'distinctness'²⁵ is rooted in character richness, not the other way around. How distinct a species is from other species will depend on the sort of characters it has and the degree to which these characters are unique or otherwise differ from those of other species. Measures of biodiversity based on character richness are a sensible way of quantifying distinctness, but it is the character richness itself that we are interested in.

Despite its attractions, there are several reasons why we might reject character richness as a target property. One major reason is that 'characters', 'traits' or 'features' are simply not well-defined. What exactly counts as a trait? Moreover, does it make sense to describe an organism by simply reducing it to a collection of traits? The debate over adaptationism (Gould & Lewontin, 1978) clearly demonstrates the difficulties with such a notion. Because organisms are highly integrated systems, delineating one trait from another is problematic. Even when certain organs can be delineated, should we think of them as a single trait or a collection of traits? In other words, should a character richness measure apply more weight to a highly complex

trait such as an eye than to a relatively simple trait such as skin colour? And how should a character richness measure account for variations within traits across groups of organisms? Given that lungs in terrestrial vertebrates are homologous to swim bladders in teleost fishes, do we count lungs and swim bladders as a single trait or as two different traits? There is a strong intuition that the evolution of some characters such as lungs is highly significant, whereas the evolution of others, such as a new skin colour, is of little significance (Müller, 1991).

Given these problems, it seems that character richness is not a robust property. Despite its distinctly cladistic origins, it sets out to quantify biodiversity in largely phenetic terms, where something like ‘overall similarity’ or ‘overall distinctness’ is assessed in terms of character state changes. The same criticisms that were levelled at phenetics similarly apply in a more limited form to character richness. In particular, we need to determine what counts as a character, and how characters can be compared with one another.

May (1994) offers an alternative target property of phylogenetic diversity measures; a property termed ‘independent evolutionary history’ (IEH). Under this view, phylogenetic diversity is not a measure of the richness of character state changes, but of evolutionary *history*. Many historical human artefacts are valuable. This is partly due to their rarity, but also in part because of what they tell us about the past. A 16th century map, for example, is valuable in part because it says something about the way people thought about the world in the 16th century. Organisms also tell us something about the past. They have a history that has been shaped by evolution. A character state change tells us something not only about the way an organism is today but about the evolutionary history of its lineage. One species represents an historical

²⁵ Although ‘distinctness’ is usually thought of morphologically, behavioural or biochemical forms of distinctness are also recognised as important.

record of long dead organisms, and the more phylogenetically distinct it is, the more it tells us.

Knowing something about evolutionary history does have some utility. For example, knowing that *teosinte* and domestic maize shared a recent common ancestor tells us that wild strains of *teosinte* can probably be crossed with strains of domestic maize to improve crop yields. Nevertheless, the value associated with independent evolutionary history seems predominantly like inherent value. Our desire to preserve evolutionary history is largely akin to our desire to preserve a 16th century map. Both stem, in large part, from some irreducible passion.

Despite this limitation, many of the phylogenetic diversity measures discussed above appear to bring us closer to capturing what we want from a *true* measure of biodiversity. Measures such as character richness appear to match many of our intuitions about biodiversity: that its unit of currency is smaller than individual species, that it encompasses diversity not only between but within species, and that it corresponds to something like the overall richness of biological structure. Nevertheless, none of these measures, character richness included, seems to fully capture what we want biodiversity to mean. Biodiversity remains an elusive concept.

All of the measures of biodiversity discussed above have been developed by biologists. Unsurprisingly then, there is an emphasis on the practicality and usefulness of these measures, and a lack of serious philosophical investigation into the metaphysical foundation of biodiversity itself. The main objective of my project is to fill this gap. Maclaurin (1998*b*) also addresses this problem. In doing so, Maclaurin proposes and develops a novel concept of biodiversity; a concept which he labels ‘raw biodiversity’.

3.3 Raw biodiversity

Maclaurin's concept of raw biodiversity is intended to achieve two things. On one hand it is intended to measure option value; the value associated with 'hedging our bets' (p. 175). On the other hand, it is intended to avoid the arbitrary nature of existing biodiversity definitions; to make biodiversity a non-anthropocentrically defined feature of the world. Maclaurin therefore defines it as the

... notion of diversity ... not based upon antecedent beliefs about what it is that makes a property or an entity valuable [nor] upon antecedent beliefs about what it is that makes a property or an entity scientifically important. So you might think of this as diversity without reference to the value or importance of properties or entities. (p. 175)

Hence raw biodiversity should both maximise option value and incorporate a range of properties without regard to how important we think they might be.

Many biologists' and conservationists' notions of biodiversity correspond to something like raw biodiversity—an all-encompassing concept that encapsulates every possible aspect of biological variety. However, there is a problem concerning its dual definition that I need first to dispel before looking at the concept in more detail. Maclaurin assumes a congruence between a biodiversity measure that maximises option value and one not based on antecedent beliefs about which properties are important or valuable. Such a congruence has not been demonstrated, and there are good reasons for rejecting it.

Recall that when culling from my jar of assorted hardware in Chapter 2, I was forced to make decisions about what to keep and what to discard in a way that would maximise option value. I concluded that a rational strategy would be first to discard those items that were unlikely to be useful in future, and to keep those that were more likely to be useful. The notion of likely future usefulness was based on what we know about the sort of properties that we think might be important, so my strategy

was not completely bereft of such considerations. It may well be that these properties do not fully correspond to the sort of properties that are eventually important. So the black nail I discarded because it was otherwise identical to the silver ones may eventually prove to have a useful property, namely blackness, despite my considering such a property unimportant on the basis of present knowledge. Nevertheless, option value can only ever be a 'best guess' of likely future value. We know that some properties are important, so it seems sensible to give them a greater weight than properties whose importance we can only guess at. In other words, whatever knowledge we currently have about what sorts of properties are important is a vital component of option value. The lack of knowledge under which an option value assessment must be made need only be partial, not complete. Wheat is valuable now and will almost certainly be valuable in future, so it has lots of option value. Yet this option value has everything to do with what we think of and are able to do with wheat. Although option value may correlate to some extent with some non-anthropocentric property of the biosphere on which some robust concept of biodiversity is based, option value cannot be congruent with such a property.

Henceforth I will consider raw biodiversity only in its second guise; as an all-encompassing concept that includes biological variety at many levels, without consideration of why we think such variety is important. As will be shown in Chapter 5, such a definition more closely corresponds to the sort of biodiversity concept that I, as well as many biologists and conservationists, seek—a robust intrinsic property of the world that encompasses many levels of organisation. Such notions of biodiversity typically ignore option value, and where option value is included (e.g. Kunin & Lawton, 1996), typically treat it as a secondary argument for protecting biodiversity, not as a definition of what biodiversity is.

Maclaurin clearly rules out the possibility of a theory-free concept of biodiversity. Any two objects are both similar and different with respect to infinite sets of

properties. Regardless of their differences, they may share any number of ‘abundant properties’ such as ‘being 4.2 ± 1.7 light years from *Proxima Centauri*’. Regardless of the number of similarities between the two objects, such abundant properties can be used to construct an equal or greater number of differences. Hence measuring the degree of diversity among a set of objects requires that we introduce some means of limiting what counts as a property. Maclaurin explores three ways of achieving this, which give rise to three types of raw biodiversity: process biodiversity, structural biodiversity and anthropocentric biodiversity. Process biodiversity is based on properties that can be defined in terms of biological processes, structural biodiversity is based on Lewis’ notion of *natural properties*, whereas anthropocentric biodiversity is based on Quine’s notion of innately-acquired pretheoretical categories in which human beings perceive the world.

In all three cases, Maclaurin’s approach to limiting properties is based on a consideration of natural kinds. As he notes, classifying the world on the basis of natural kinds should generate “... some large set of well-motivated, non-gerrymandered properties that allow us to talk about similarity in general” (p. 201). Unfortunately, as Maclaurin notes, natural kinds alone do not generate a single well-defined measure of biodiversity. One immediate problem is simply that we have no way of justifying the commensurability of different properties. Does ‘possessing an eye’ count as one property and ‘possessing red feathers’ count as another? Or is ‘having an eye’ actually a collection of properties such as ‘having a lens’, ‘having light sensitive cells’, ‘having an optic nerve’ and so on. In other words, at what level do we count properties? This is exactly the same problem that beset character richness measures, discussed above. Making an inventory of properties does not in itself give us a well-defined measure, although it may usefully illustrate where we should begin to look for one.

Maclaurin raises some interesting critiques of each of the three ways of formulating raw biodiversity that he proposes. Consider, for example, raw process biodiversity.

Recall that this is defined in terms of natural kinds arising from some set of biological processes. One such process is natural selection, so we can use process diversity as a means of comparing organisms on the basis of adaptations. As Maclaurin notes, however, limiting process diversity only to adaptations rules out properties that we might think are useful but which are mere 'biological epiphenomena', such as the ability of dung beetles to moderate blow-fly populations. Moreover, casting our net wider to include such epiphenomena by broadening the definition of process biodiversity appears only to return us to where we started, with nearly every conceivable property counting as some sort of epiphenomenal property.

A different set of objections can be raised against raw anthropocentric biodiversity, which is defined in terms of Quine's concept of natural kinds as stemming from the way in which humans pretheoretically perceive and classify the world into classes of similar objects. This notion of biodiversity is explicitly anthropocentric, since it is based on our 'innate' propensity for classification. Maclaurin illustrates some of the problems stemming from this notion of raw anthropocentric biodiversity, such as its inability to reflect non-obvious features that do not form part of pan-cultural attributes of classificatory schemes. For example, many past (and some present) cultures classified whales as fish, largely because of their fish-like appearance. Of course, we now know that whales are mammals. Whales have not changed, but our fish-category has. Clearly, our classificatory schemes are not pretheoretically fixed. We can change them in the light of new knowledge. More fundamental objections may also be raised against Quine's conception of natural kinds, both on philosophical grounds (Dupré, 1993, p. 277) and on biological grounds concerning the concept of innateness on which it is based (Lehrman, 1970; Lickliter & Berry, 1990; Bateson, 1991; Maclaurin, 1998*a*).

The third of Maclaurin's three approaches to developing raw biodiversity is based on what he labels raw structural diversity. Of his three approaches, this is the one of

which Maclaurin is most dismissive. It is also the approach that comes closest to the notion of biodiversity developed in Chapters 4 and 5. Raw structural biodiversity is based on Lewis' notion of *natural properties*. Of these Lewis says

Sharing of them makes for similarity, they carve at the joints, they are intrinsic, they are highly specific, the sets of their instances are *ipso facto* not entirely miscellaneous, there are only just enough of them to characterise things completely and without redundancy. (Lewis, 1986, p. 60)

As Maclaurin notes, the requirement of characterising without redundancy automatically excludes disjunctive properties such as 'being a mammal or a milk bottle'. A sensible scheme of classifying the world already includes mammals and milk bottles, and so a further classificatory rule based on 'being a mammal or a milk bottle' is redundant. Furthermore, raw structural biodiversity offers the possibility of including important epiphenomenal properties without also including 'unnatural' properties such as 'being 4.2 ± 1.7 light years from *Proxima Centauri*'. Hence it can overcome at least one of the major objections raised against raw process biodiversity.

The main objections against raw structural biodiversity that Maclaurin raises concern the way it treats ecological categories such as 'being a predator'. Lewis' natural properties are intrinsic. According to Maclaurin, based on Lewis' notion of what counts as an intrinsic property, ecological categories such as 'being a predator' do not qualify. Loosely speaking, Lewis defines a property of an object as intrinsic if it would still belong to that object unaccompanied—that is, if the object were in a *lonely* world (Lewis, 1986; Langton & Lewis, 1998). Since 'being a predator' is a property of an organism that depends on its relationship with other organisms, it is not an intrinsic property in Lewis' strict sense and therefore can not be included among Lewis' natural properties.

One possible solution to this problem is to relax Lewis' definition of intrinsic to include functional properties such as 'being a predator'. There are two main theories

of biological function (Sterelny & Griffiths, 1999). According to the *etiological theory*, biological functions are explained in terms of their evolutionary origins. Under the rival *propensity theory*, they are explained in terms of their present adaptive effects. So the etiological theory implies that a lion possesses the predator property because its ancestors were selected to kill and eat other animals, whereas the propensity theory implies it is because the lion is well-suited to killing and eating other animals. According to the propensity theory, a lion in a lonely world would no longer possess the predator property because there would be no other animals in that world for the lion to kill and eat. It would have no propensity for killing and eating other animals. According to the etiological theory, however, a lion can still possess the predator property in certain types of lonely worlds. These would be lonely worlds in which the lion has a history; in other words, worlds in which the lion persists through time. Admittedly, such worlds are not strictly lonely worlds, since the lion is accompanied by its past and future selves (Langton & Lewis, 1998). Nevertheless, although properties such as ‘being a predator’ are not strictly intrinsic according to Lewis’ definition, they are certainly not arbitrary. A suitably relaxed definition of intrinsic can accommodate functional properties without admitting such ontological detritus as ‘being 4.2 ± 1.7 light years from *Proxima Centauri*’.

Maclaurin raises a second objection against ecological categories such as ‘being a predator’. Because such categories are a matter of function rather than structure, they are multiply realisable. This means that the single category ‘being a predator’ might include a host of realisations that together represent a significant level of diversity not represented by the category alone. As Griffiths (1997) describes the problem,

... the causal homeostatic mechanism of each ecological category is a particular set of adaptive forces [which] are sensitive only to properties at the level of task description, so properties at [lower] levels are not projectable in ... categories derived from the ecological level ... The same

task can be performed in many ways at lower levels of description. (p. 234)

Here Maclaurin raises a serious and interesting issue about biological classification. Ecological categories are generally more weakly projectable than, say, phylogenetic categories.²⁶ However, they are projectable nonetheless. For example, if A and B are closely related species, and C and D are not closely related but share an ecological category such as ‘being a predator’, then we can generally be more confident about predicting the attributes of B from knowledge of the attributes of A than we can be about predicting the attributes of D from knowledge of the attributes of C. Nevertheless, despite the confidence level of predictions based on ecological categories being lower, they are not zero.

Maclaurin’s objection can be further countered by considering ecological-historical categories. If A and B are closely related *and* share an ecological category such as being a predator, then the properties of A are more highly projectable onto B than if they shared only the phylogenetic or ecological category alone. Ecological categories may be multiply realisable, but history imposes constraints on the possible space of forms for realising any particular ecological function (Griffiths, 1997). Predators come in a wide range of forms, but the features of avian predators are quite consistent despite their phylogenetic diversity—good eyesight, well-developed flying ability, sharp claws and hooked beaks.²⁷ On their own, ecological properties may be only weakly projectable, but when combined with phylogenetic properties, they give rise to highly projectable categories.

²⁶ For arguments in support of this assertion, see Griffiths & Sterelny (1999) Ch. 11 and de Queiroz & Good (1997).

²⁷ As a more detailed empirical example, Faith (1989) demonstrates the repeated evolution of certain morphological traits, in response to feeding mode, for many species of wading bird.

A final objection to raw structural biodiversity is simply that Lewis' natural properties are not sufficiently abundant for characterising biological diversity. There are many more intrinsic properties than there are natural properties, given that natural properties occur only in fundamental physics (assuming everything else supervenes). Hence biologically similar things are unlikely to share many more natural properties than are biologically dissimilar things. Natural properties simply apply at the wrong level of detail to be useful in characterising biological diversity.

Of the three forms of raw biodiversity discussed by Maclaurin, raw structural biodiversity has some potential for further development. Turning it into a workable concept would, however, require extending the range of properties used beyond Lewis' natural properties to include other intrinsic properties, as well as relaxing the definition of what counts as an intrinsic property to include functionally-defined properties such as 'being a predator'. The result, I suspect, would look much like the concept of biodiversity that I intend to develop in the following chapters. However, this is a point I do not wish to pursue in detail. My approach to developing an intrinsic concept of biodiversity will be quite different to that of Maclaurin. I plan to look first at a more general measurable property, and then show how biodiversity can be interpreted in terms of this property. Moreover, it is a universal, multilevel and, in Lewis' strict sense, intrinsic property: complexity.

4 Complexity

More is different.

— Phillip Anderson

Complexity is a concept widely applied in discussions of biology, although not without controversy (Castrodeza, 1978; Hinegardner & Engleberg, 1983; McShea, 1991, 1992, 1996*a*, 1996*b*; Yagil, 1985). Organisms and other biological systems are, it seems, complex entities. If only we could somehow quantify or otherwise understand biological complexity, we would perhaps better understand biology in general. The controversial aspects of complexity are manifold. First there is the issue of the origins of biological complexity. Why are living things complex compared to non-living things? How do some organisms become more complex than others? A second controversy relates to the notion of ‘progress’ in evolution; whether or not it occurs, and the sense in which progress can be equated with increasing complexity. Although there is little doubt that many present-day species are much more complex than early prokaryotes, there is much doubt over assertions that there is some evolutionary drive towards increasing complexity or that this increasing complexity can be characterised in some meaningful sense as progressive (Gould, 1989, 1991, 1993; McShea, 1991; Ridley, 1993).

A third controversy regarding biological complexity concerns mathematical and computational approaches to understanding complex systems. The so-called ‘sciences of complexity’ address problems in the biological sciences and other areas in which ‘complex adaptive systems’ play a fundamental role, and try to render difficult problems in these areas tractable through the use of newly developed mathematical techniques and modern computer technology. Much of this work has attracted significant criticism, partly because of the imbalance between its excessive hyperbole and modest successes (Horgan, 1995).

None of these criticisms concerns my project. I make no claims for the origins of biological complexity, nor for how complexity might change over time. Furthermore, I do not advocate any particular approach to understanding complex systems. My concerns are simply with the nature of complexity, its measurement, and its relationship to biodiversity. In this chapter, I hope to develop some tools for understanding complexity and to show that complexity is a well-defined and intrinsic property of any system. In the following chapter, I will reinterpret biodiversity in terms of complexity. In other words, in answer to the question *What is biodiversity?*, my answer is *biological complexity*. If complexity is an intrinsic property of any system, then biodiversity, construed as biological complexity, must also be an intrinsic property of any biological system. First, however, I want to discuss an intuitive concept of complexity.

4.1 An intuitive complexity concept

Virtually by definition, complex systems are difficult to describe. Indeed, the more complex something is, the more difficulty we have in describing it. That is, the description of a complex thing tends to be longer than the corresponding description of a simple thing. This observation provides an intuitive and surprisingly useful concept for thinking about complexity measures; namely, that the relative complexities of objects can be compared by comparing the lengths of some descriptions of those objects. So, for example, if we wish to compare the complexity of a Volkswagen Beetle with that of a Boeing 747 passenger jet, we could compare a service-manual description of the Beetle with a service manual description of the 747. By doing this we could then determine that the 747 is considerably more complex than the Beetle.

In carrying out such a comparison, it is obviously important that the descriptions be *at the same level of detail* to make valid conclusions about relative complexity. If we

compared a service-manual description of the Beetle with an encyclopaedia-entry description of the 747, the former would probably be longer. However, whereas the service-manual description of the Beetle would contain considerable detail about minor parts such as the bonnet release mechanism and the windscreen wipers, the encyclopaedia-entry description of the 747 might only contain details about the aircraft's dimensions, top speed, seating capacity and so on. It will ignore details about its altimeter, yaw-angle control system and a host of other features relevant to a service technician. In this case the description of the VW Beetle is longer than that of the 747 only because it contains more detail, not because the Beetle is more complex.

With two highly dissimilar objects, defining what we mean by *at the same level of detail* may prove problematic. Suppose we wish to compare the complexity of a VW Beetle with that of the Matisse painting *Les Poissons rouges* (1911). What level of description of this painting corresponds to the service-manual level description of a VW Beetle? We could try describing *Les Poissons rouges* using the steps a colour-by-numbers copy artist would use to reproduce the painting. Or we could describe it using the actual brushstrokes and colour choices Matisse made in creating the painting. We could even try scanning a photograph of the painting into a computer and applying the latest image compression software to give a description in terms of a computer file that we could send to a friend by e-mail. Clearly none of these is exactly analogous to the sort of engineering description used in a VW Beetle service manual. Indeed it appears that no such description exists, tempting us to conclude that the complexities of the two objects cannot be meaningfully compared in this way. Whereas comparing the complexities of similar types of objects, such as VW Beetles and Boeing 747s, is relatively straightforward, comparing the complexities of dissimilar objects appears problematic. Later on we shall see that this problem is not as intractable as it first appears. At this stage, it will suffice to observe that, provided the objects in question are not too dissimilar, then we can make meaningful comparisons

of their complexities by comparing similar level descriptions of the objects. Moreover, the more similar the two objects are, the more meaningful such a comparison becomes and the more meaningful is our measure of their relative complexities based on description length.

A further observation about this sort of complexity measure is that it tends to become more accurate as we include more detail in our descriptions. Superficial descriptions of a VW Beetle and Boeing 747 listing their colours, shapes and overall appearances might lead us to conclude that their complexities are similar. Even a description listing their major components (engine, wheels, windscreen, exhaust system etc.) might be misleading as a measure of complexity. Only when we come to describing each at a sufficient level of detail, such as the sort of descriptions required if we were to try to build each from scratch, do we begin to appreciate the significant differences in complexity between the two objects.

One of the features of complex objects is that they can usually be described in a number of different ways. Take, for example, copies of Shakespeare's *Julius Caesar* and the Barbara Cartland romance novel *Bride to a Brigand*. Since both of these are to some extent the same sort of thing—stories written in ink on paper—we should be able to make comparisons about their relative complexities. Which is more complex? If we use a length-of-description technique to determine the answer, the outcome will depend on the type of description we adopt. We could describe each in terms of the physical location of every atom of which they consist, or in terms of the two-dimensional patterns on their pages, or simply in terms of the texts themselves. In all three cases we might find that *Bride to a Brigand* is more complex, largely because it is longer and physically bulkier than our copy of *Julius Caesar*. What we might really be interested in, though, is the *meaningful content* of the stories. In this case we could try to describe each in terms of features such as the number of characters, details about each character, aspects of plot development, overall themes, interesting uses of

language and the other sorts of elements studied by students of English. Once we carry this out we will probably find that *Julius Caesar* is significantly more complex than *Bride to a Brigand*. Clearly this sort of description-length definition of complexity is not, on its own, well-defined unless we specify the sort of description to be applied. In other words, we cannot talk about the complexity of something unless we have some understanding of the sort of description-level attributes to which we wish our notion of complexity to refer.

At this stage it seems like complexity is not a very intrinsic property. If, for everything whose complexity we wish to measure, we first have to specify the type of description to be applied, then it looks as if we are not measuring any strongly intrinsic complexity-like property, only complexity with respect to some arbitrarily-selected set of qualities. This, however, is not what I want. I want a strongly intrinsic concept of complexity that can be used as a strongly intrinsic concept of biodiversity. Later I hope to demonstrate how we can solve this problem.

Leaving these metaphysical questions aside for the moment, let us think of complexity in terms of a metric—a mathematical function that we can use to measure the amount of some property (complexity) possessed by any *system*. So if x is some system, let $C(x)$ be the complexity of x and let $C(x, y)$ be the complexity of the *joint* system consisting of x and y . Well-defined mathematical measures tend to have properties such as the following:

1. $C(x) \geq 0$ (non negativity)
2. If $x = y$, $C(x) = C(y)$ (reflexivity)
3. $C(x, y) = C(y, x)$ (symmetry)
4. $C(x, y) \geq C(x)$ (accumulation)
5. $C(x) + C(y) \geq C(x, y)$ (convexity).

How then does description length fare in terms of the above properties? It can be clearly demonstrated that, properly understood, a description length definition of complexity satisfies all of them. It satisfies the test for non-negativity, since it is impossible to write a description of negative length. It satisfies the test for reflexivity, since identical objects can be given identical descriptions, which therefore have identical lengths. It also satisfies symmetry, given that a description of the system consisting of x and y is the same as a description of the system consisting of y and x . It satisfies accumulation, since the description of the system consisting of x and y is at least as long as the description of x alone. Finally, it satisfies convexity, since the description of the system consisting of x and y is no longer than the description of x together with the description of y .

This last point is interesting. The reason the description of the system consisting of x and y is generally shorter than the combined descriptions of each element considered separately is that x and y may share some similarities. In other words, a joint description of the two elements considered separately may contain some *redundancy*, which can be removed without loss of completeness to describe the system consisting of both elements. Using the expression $C(x:y)$ to specify the amount of redundancy in the system (x, y) —that is, the amount of complexity shared by x and y —we can measure redundancy by $C(x:y) = C(x) + C(y) - C(x, y)$. If x and y are identical, then $C(x:y) = C(x) = C(y)$, whereas if x and y share no similarities, $C(x:y) = 0$.

For example, consider once again our VW Beetle. Our Beetle is not more complex by virtue of its having two rear brake drums instead of only one. Description length can quite easily account for redundancy. A service manual description of a VW Beetle, for example, would not include a separate entry for each rear brake drum. Instead there would be merely one entry that covers both. This entry might take the form of a detailed description of the right rear brake drum, together with a short

statement such as “the left rear brake drum is a mirror-reversed copy of the right rear brake drum.” This additional sentence can be made arbitrarily short, thereby representing virtually no additional complexity. This suggests that we should not use just any description to assess complexity, but the shortest or at least nearly shortest possible description, given the existing language constraints. In other words, we should remove all redundancy. So instead of including in our description of a VW Beetle, “has one wheel with attributes $\{a_1, a_2, \dots, a_n\}$ and another wheel with attributes $\{a_1, a_2, \dots, a_n\}$ and another wheel ...”, we should simply say, “has four wheels with attributes $\{a_1, a_2, \dots, a_n\}$ ”. Whether there are four identical wheels or only one will therefore have little bearing on overall complexity.

This issue of convexity, however, is not so straightforward. In some cases, description length does not appear to satisfy convexity. The description of a system consisting of several elements can be longer than the combined descriptions of those elements alone, because the elements can have relational properties. This is a feature of complex systems in general. They are typically more complex than the complexity of their parts would suggest. According to Anderson (1972), more is not just more; more is different.

One of the ways around this problem is to consider relational properties when measuring complexity. So the system (x, y) consists only of the elements x and y without any relational properties, whereas the system consisting of x and y and including all their relational properties p_1, p_2, \dots, p_n should be written as something like $(x, y, p_1, p_2, p_3, \dots, p_n)$. For example, a working VW Beetle is in some sense more complex than a pile of junk consisting of the parts of a VW Beetle. The elements of a working VW Beetle comprise not just its parts but the relational properties between those parts. In contrast, the pile of Beetle parts contains no significant relational properties, only the single relational property ‘randomly assembled into a pile’. In terms of description length, the description of a working VW Beetle would have to

include more than simply a description of the parts that comprise that VW Beetle; it would also have to include a description of how those parts go together. A description of a disassembled Beetle is therefore shorter than a description of a working Beetle. By including all such relational properties, convexity is maintained.

This intuitive notion of measuring the complexity of an object by the length of a description of that object seems to be very useful. It is perhaps unsurprising, therefore, that description length, loosely defined, forms the basis for all complexity measures. The differences between measures arise in large part from adopting different types of descriptions, or more specifically, different *representations*. We will now look in more detail at some of these specific complexity measures. From a mathematical perspective, the simplest types of complexity measures to consider are those that apply to strings of symbols. And strings of symbols are exactly the sort of objects dealt with by the mathematical theory of information.

4.2 Information content as a measure of complexity

Information theory officially began in 1948 with the work of Bell engineer Claude Shannon (Shannon & Weaver, 1949). Shannon was concerned with quantifying information for the purposes of better understanding how messages could be efficiently transmitted over noisy communication channels. The technical concept of information applied in information theory is quite different to the everyday concept of information as ‘meaningful message content’. In information theory, the meaning of any message is explicitly ignored. Instead, information is identified more with the length of the message or the ‘cost’ of transmitting it. Hence the amount of information in a message depends not only on the message itself but on the communication system in which the message is transmitted.

A simple communication system considered by Shannon consists of an information *source* that transmits a message chosen from a finite set of characters in the form of a *signal* over a *channel* to a *receiver*. The channel can also introduce *noise*, which changes the signal at the receiver from that sent by the source. In the context of such an arrangement, Shannon defined information as a measure of the reduction in uncertainty at the receiver as a result of transmitting the message.

Suppose the source can only transmit two types of characters, say ‘0’ and ‘1’, and from the perspective of the receiver, each of these has equal probability of being transmitted. Assuming no noise, transmission of a single character therefore reduces the uncertainty at the receiver by half. Before transmission there were two possibilities—‘0’ or ‘1’—whereas after transmission there is only one possibility defined by whichever character was sent. If, instead of one character, we consider a message consisting of three characters, then the number of possible messages that can be sent is eight (000, 001, 010, 011, 100, 101, 110 and 111). Transmission of a particular message, say ‘110’, therefore reduces the uncertainty by a factor of eight. The information measure defined by Shannon takes the logarithm of the reduction in uncertainty as an estimate of the information content of the message. So in the above example, the quantity of information sent is $\log_2 8$. Although the logarithmic base used is arbitrary, it is conventional to use base two, in which case the unit of information is the ‘bit’ (a contraction of binary digit). Hence the above message ‘110’ consists of three bits of information. Indeed for any string of equiprobable binary digits, Shannon’s information measure corresponds exactly with the length of the string transmitted.

One of the nice aspects of considering information in terms of reduction in uncertainty is that it can account for redundancy. If I send you the e-mail message “I WILL SEE YOU INSIDE THE ROYAL RANDWICK HOTEL THIS EVENING AT 7 O’CLOCK”, I have transmitted just as much information as if I had sent the

much more compact “C U IN ROYAL 7 2NITE”. This is because both messages reduce your uncertainty about where and when I intend next to see you by the same amount. There are, however, some penalties associated with the second message form. One is that you have to do more work to understand the second message. For example, you have to know that ‘C’ and ‘U’ can be translated into the common English words ‘see’ and ‘you’, respectively, and you have to extract from the context of the message that this is how I intend you to understand them. A second penalty is that the second message is much more prone to being misunderstood, especially if noise happens to be introduced into our communication system. Suppose the e-mail system has the annoying habit of occasionally replacing individual characters by ‘N’ in all its messages. The first message might then be received by you as “I WILL SEE YOU INNIDE THE ROYAL RANDWINK HOTEL THIN EVENING AT 7 O’CLONK”. As a result you might exchange a few witty remarks on my poor spelling and/or typing skills upon meeting me at the intended time and place. If, however, I chose the more economic message format, you might receive a message such as “C U IN ROYAL 7 2NINE”, in which case I might end up waiting around for almost two hours while you drink with other friends at the Doncaster and discuss my odd habit of setting overly precise meeting times.

As the above example shows, redundancy in communication systems is not always a bad thing. It can reduce the amount of work required by the receiver to decode the message and provide a way of making the message robust to noise. Indeed ‘error correcting codes’ are used extensively in communication systems expressly for this purpose. In many cases, however, redundancy in a message can cost time and money. If I want to e-mail you a large text file, using a file compression program such as *WinZip* can reduce the size of the file by removing some of its redundancy, thereby reducing the amount of time and money it costs me to send it to you. Provided you also have a copy of *WinZip* on your computer, you should be able to decompress the file and read it. The key aspect of this strategy is it requires that we both have a copy

of *WinZip*. That is, not only must the source be able to encode the message, but the receiver must know how to decode it. Redundancy is therefore not a property of the message itself but an aspect of the entire communication system. There is a joke about a country pub in which the patrons have heard every joke so often that instead of bothering to tell an entire joke, they just refer to it by a number (upon which the other patrons laugh).²⁸ In other words, how redundant a message is depends not just on the message but on how much the receiver already knows.

A system's 'information content' can be used as a measure of its complexity. If the system is itself a message string, then assuming equiprobable digits and ignoring redundancy for the moment, we need only consider the length of the string. The string '1011' contains four bits of information, whereas the string '01101010' contains eight bits. The second string is in some sense twice as complex as the first because it is twice as hard for us to specify. Considerations of redundancy may lead us to reduce the complexity of some highly ordered strings. So, for example, the string '11111111' consists of eight bits, but it is in some sense less complex than the eight-bit string '01101010', because it can be replaced by the simple instruction "write '1' eight times". To account for this sort of redundancy, we need a different complexity measure, one based not simply on information content but on some notion of 'true information content'; that is, on information content with all redundancy removed. However, we have just seen that redundancy is context dependent, so how should we remove all redundancy in a consistent and well-principled manner?

²⁸ In one version of this joke, a visitor to the pub calls out '63' and nobody laughs because he didn't tell it right. In another version, everybody laughs hysterically because it was one they hadn't heard before.

Algorithmic information content (AIC) was introduced in the 1960s as a means of measuring the ‘true information content’ of a message string. It is a measure that can account for all redundancy in a message string without making any assumptions about the knowledge of a receiver. Consider a so-called *universal computer*—an idealised all-purpose computer with infinite storage capacity. Consider also a message consisting of a string of symbols taken from a finite alphabet. The AIC of the string is then defined as the shortest program that can be run on the computer which will cause the computer to print out that string and then stop.

This definition of AIC is quite useful for measuring the complexity of certain computational problems. This type of complexity measure is known as *Kolmogorov complexity*, after the Russian mathematician instrumental in its invention. The advantage of Kolmogorov complexity is that it effectively addresses the issue of redundancy. Instead of a program that says “print ‘1’ then print ‘1’ then ... then print ‘1’ then stop”, we could write a program that says “print ‘1’ eight times then stop”. The Kolmogorov complexity of a string of eight 1s is something like the length of the second program rather than the former. More typically, this complexity measure is not used as an absolute measure applied to individual strings, but as a way of understanding how the complexity of a class of computational problems increases as the size of the problem increases. For example, if we make the string of 1s longer and longer, its Kolmogorov complexity remains almost constant. Thus, regardless of its length, a highly ordered string can be described by a short program and has little Kolmogorov complexity, whereas the Kolmogorov complexity of a random string increases in proportion to the length of the string.

The definition of Kolmogorov complexity in terms of the length of a computer program seems to raise a problem of language dependence. A program written in one programming language will typically differ in length from a program written in another language. Indeed, the same criticism seems to apply to any sort of descrip-

tion-length definition of complexity. How can we be sure that, in trying to compare the complexities of different systems, the outcome does not depend on the choice of language?

The solution to this problem involves modifying the programming language in each case so as to minimise the expected value of the code word lengths. It has been proved that for any universal computer, there exists a slightly modified universal computer such that the average algorithmic information content over a class of systems is essentially equal to the Shannon information content over that class of systems (Gell-Mann & Lloyd, 1996). Once we have specified the class of systems we are dealing with, the residual algorithmic information content of a particular member of that class is then given by the length of the code assigned to it by the modified universal computer. In short, there is a technique for removing language-dependence from the overall measure, thereby reducing the Kolmogorov complexity of any system to a universal metric. Just as we can measure the information content of a message with a universal metric, so too can we measure its Kolmogorov complexity. In both cases the same unit applies, namely the bit.

To illustrate how Kolmogorov complexity can be applied to physical systems, consider a box containing helium atoms. A simplified example of this is illustrated in Figure 7a below. One way of measuring the complexity of this system is to measure the amount of information I would be required to send to you in order for you to reproduce this system almost exactly. I would therefore have to specify the location of every atom in the box to some specified degree of accuracy and using some agreed coordinate system, and then transmit this information to you. If the atoms were in a gaseous state and randomly distributed within the box, then my message would look like a string of mostly random numbers. Now consider a second box in which every helium atom is crammed into one corner, as shown in Figure 7b. I could again specify the location of each atom using the same technique as before, but my

message would contain a high degree of redundancy. This is because I can use the fact that all the atoms are crammed into one corner of the box to reduce the length of the message I need to send to you. For example, if the corner into which all the atoms were crammed corresponded to the origin of my coordinate system, then I could simply reduce the number of digits required to specify the location of each atom. In a sense I have specified a much shorter algorithm for you to use in reconstructing the box of atoms.




Figure 7: Examples of highly disordered (*a*) and ordered (*b*) systems.

The measure of complexity used in the above example is akin to a measure of the entropy of the physical system. This is far from coincidental, since there are deep links between information theory and the thermodynamic concept of entropy. Indeed the logarithmic equation for measuring information proposed by Shannon is analogous to an equation used to measure the entropy of a physical system. What this implies is that highly entropic systems—those with a lot of entropy—also have, in the sense discussed above, a lot of information content and hence a lot of complexity. So if we took a cockroach and heated it up in the absence of oxygen, thereby turning it into a gas, the gaseous cockroach would have much more entropy and therefore much more Kolmogorov complexity than the living cockroach.

There is an obvious problem here. This is that complexity as normally understood does not usually apply to systems with a highly random arrangement of elements.

The gaseous cockroach in the above example is more entropic than our living cockroach, but it is in some sense much less complex. There is something important about the way the components of a living cockroach are arranged that is lost when we turn the cockroach into a gas. A cockroach has lots of internal relational properties between its elements. When we turn the cockroach into a gas, we lose these relational properties. Indeed neither highly ordered systems such as crystals nor highly disordered systems such as gases are complex in the sense that a cockroach is complex. The sort of measure we require to differentiate between simple and complex systems is not the same as the measures we can use to differentiate between ordered and disordered systems. Unfortunately, the more random a string, the higher is its Kolmogorov complexity; which is almost the opposite of what we want.

There is a second, less obvious problem with Kolmogorov complexity as a practical complexity measure. Consider the following example. There is a relatively short program that will print out digits 1,000,001 to 2,000,000 of the decimal expression of the number π . The Kolmogorov complexity of these one million digits is therefore quite low. However, if I simply gave you these digits and asked you to find the shortest algorithm that would print them out, you would have little idea of where to start looking. The digits themselves appear truly random. They will pass any statistical test for randomness you might subject them to in order to discern a pattern. You might therefore conclude that the only suitable algorithm is “print ...” followed by the string itself; in other words that the Kolmogorov complexity for this string is maximal. This problem is an illustration of a curious property of Kolmogorov complexity: its uncomputability (Gell-Mann, 1994). What this means is that, for a given string, there is no algorithm that will produce the shortest possible program that will output that string and hence return a value for its Kolmogorov complexity. In other words, even if we devise a program that efficiently prints out a string by accounting for some of the string’s redundancy, we can never be sure that this is the most efficient program that exists.

There are other types of complexity measure that overcome the uncomputability problem by placing limits on the set of algorithms from which the minimum length algorithm is to be selected. Despite this advance, both Kolmogorov and these other types of complexity measure are ‘entropic measures’; that is, they are maximal for truly random strings. Although they are useful for assessing the difficulty of certain computational problems, they fail to capture the sense of complexity we want to use as a measure of biodiversity. What we need is a different sort of complexity measure; one based on a different sort of system representation.

4.3 Alternative complexity measures

To overcome the above limitation of informational measures of complexity, several alternative measures of complexity have been proposed for use in a range of disciplines (Cornacchio, 1977; Grassberger, 1988; Löfgren, 1977; Papentin, 1980). The key feature of these alternative complexity measures is that the random parts of the system in question—that is, those highly entropic parts that do not contribute to what we consider to be the system’s overall complexity—are dealt with in such a way that they contribute little to the final complexity measure. As mentioned previously, all complexity measures, even those based on measuring information content, require us to measure the length or some other size aspect of a representation of the system in question. For a purely informational measure of a string of symbols, the representation is the string itself. For Kolmogorov complexity, the representation is the shortest computer program that will output the string. Neither of these representations proved appropriate. However, if we adopt a different representation, we can devise a complexity measure that captures the sense of complexity in which we are interested.

Consider the previous example of a box filled with helium. In this example I was trying to devise a message that I could send to you in order for you to recreate the

system. Since I wanted you to recreate the system almost exactly, I went about measuring the position of each helium atom and putting this information into my message. Suppose instead that I was not interested in the exact position of each helium atom. Suppose that all I wanted you to reproduce was a box filled with helium. In this case the message I would have to send to you in order for you to reproduce such a box would be very small. All I would have to say is “build a box of dimensions $x \times y \times z$ and fill it with h amount of helium”. I could then be pretty sure that you would produce a box of helium atoms very much like mine. It would not be *exactly* the same, since the positions of the atoms in my box would be quite different to those of the atoms in your box. Nevertheless, the *representations* of each box as “a box of dimensions $x \times y \times z$ filled with h amount of helium” would be the same. Moreover, this representation is now much smaller than the previous representation that specified the position of every atom. Hence the complexity of such a system measured as the length of this representation is also quite small.

Consider an example of a complexity measure based on the length of a simplified representation. The type of complexity is known as ‘regular language complexity’ (RLC), and is applied to character strings. The measure is applied not to particular strings but to the class of strings to which a particular string might belong. Any string can be classified in terms of the rules governing its generation. That is, the string may be constrained by certain ‘grammatical’ features. For example, in a portion of English text, the character following the character ‘Q’ is almost always ‘U’, whereas the character ‘V’ is commonly followed by ‘E’ or ‘O’ but almost never by ‘N’. Similarly, the phrase ‘Once upon a’ is commonly followed by the word ‘time’. These ‘grammatical’ rules can be represented using a ‘deterministic graph’. To illustrate the idea of a deterministic graph, consider the example of a string of 0s and 1s whose only grammatical rule is “the sequence ‘11’ is always followed by ‘0’”. In other words, the string will never contain three or more 1s in a row. An example string of this class is

‘1011001101011000’. The deterministic graph for all strings of this class is illustrated in Figure 8a below.




Figure 8: Deterministic graphs for sequences of binary digits (a) with the grammatical rule “‘11’ is always followed by ‘0’” and (b) with no grammatical rules.

To understand how a deterministic graph encodes a grammatical rule, try writing out a sequence of 1s and 0s using the graph as a guide. Using any node as a starting point, select one of the paths leading from the node at random to generate the first digit, using the path label indicated. Following this path takes you to another node (or in some cases the same node). The second digit is generated by taking one of the available paths from this second node to a third node and so on. Applying this process using the graph of Figure 8b will generate a sequence of digits in which the grammatical rule is obeyed; that is, in which ‘11’ is always followed by ‘0’. The class of sequences in which no grammatical rules apply; that is, in which the selection of digits is random and completely unconstrained, can be represented by the graph of Figure 8a.

The regular language complexity of a class of sequences is measured by the number of nodes in its deterministic graph. As the above graphs show, under this measure, strings of binary digits constrained only by the rule “‘11’ is always followed by ‘0’”

are three times as complex as purely random strings. Strings of binary digits representing ASCII²⁹ coded English sentences would be constrained by a large set of rules pertaining to permissible letter combinations, words, word combinations, English grammar and so on. Such strings would thus require a very large graph for their full representation and would have a correspondingly large measure of RLC.

Although RLC only captures the complexity of the string's 'grammar', not of its meaningful content, it provides a practical example of how the right sort of representation can adequately account for the entropy problem that afflicted complexity measures such as Kolmogorov complexity. This sort of simplified representation is exactly what was being used in the previous example comparing the complexity of a VW Beetle with that of a Boeing 747 aircraft. If I wanted you to recreate a VW Beetle I could, as in the box of helium example given previously, measure the position of every atom in the Beetle and transmit this information to you. However, the vast majority of this information is not required to produce a working VW Beetle. The exact position of every iron atom in the steering column is not particularly useful information. From the perspective of how difficult it is to reproduce a working VW Beetle, this information is simply not relevant, because many combinations of atoms are adequate. What would be more useful and efficient is to specify the dimensions of each part, its materials, how all the parts fit together and so on. This brings me back to the sort of service-manual description considered previously.

Suppose now that instead of a VW Beetle I wanted you to recreate a cockroach. Again I would not need to transmit the position of every atom in the cockroach, since the exact position of every atom is in many cases not relevant to producing a working cockroach. Nevertheless, the amount of information that would have to be

²⁹ American Standard Code for Information Interchange.

transmitted in order for you to reproduce a working cockroach would be very large. Even if I disassembled the cockroach and accurately measured every part I could find and transmitted this vast amount of information to you, it is doubtful in the extreme that you could successfully reproduce a working cockroach. You might successfully produce an inanimate replica of a cockroach, but to reproduce a working cockroach from scratch would require at least a vast extension to existing scientific knowledge about basic insect physiology. Even supposing that such knowledge were available and reproduction from scratch of a working cockroach were feasible, the amount of information that I would need to send to you to effect this task, compared to that for a VW Beetle, would be truly mind boggling. This is the sense in which biological organisms are complex. The complexity of a VW Beetle is impressive compared to that of a bicycle, and that of a Boeing 747 is more impressive still, but all of them are dwarfed by the complexity of a single cockroach.

4.4 The metaphysics of complexity

Is complexity an intrinsic property? Is it, like mass, a property that something possesses even in lonely worlds? Is it, like beauty, in the eye of the beholder? Or is it something in between? As we have seen, to measure the complexity of a system, we first have to generate a representation of the system, and then measure the size of that representation. Complexity is therefore representation-dependent; the sort of measure we end up with depends on the sort of representation we choose. Unless I can specify some way of ‘privileging’ one type of representation over all others, we are free to choose whatever representation we want. What we are left with is not complexity as a strongly intrinsic property, but some arbitrarily-defined version of complexity. How much complexity something has would not only depend on the thing itself but on the way we choose to look at it. To an entomologist, a cockroach might be complex, but to an echidna it is just one more type of food.

Many authors accept the thesis that complexity is not a strongly intrinsic property, or at least that there are many different types of complexity, none of which can be singled out in some principled manner. Löfgren sums up this position in terms of the representation-dependency of complexity.

... complexity [should] be associated with descriptions, rather than being thought of as an intrinsic property of objects. (Löfgren, 1977, p. 197)

Kurths and Witt also emphasise the contextual nature of complexity ...

It is important to recognise that complexity measures are contextual, i.e. they cannot be defined universally but depend on some context. (Kurths & Witt, 1994, p. 191)

... while Cornacchio emphasises its observer-dependence.

the complexity of a system is relative to an observer's knowledge of the system ... to the neurophysiologist the brain ... is certainly complex ... while to a butcher, a sheep brain is simple, since he has to distinguish it from only, say, 30 other meats. (Cornacchio, 1977, p. 221)

McIntyre bluntly denies complexity any intrinsic status.

... complexity is derivative rather than inherent. Complexity exists not merely as a feature of the world, but as a feature of our attempts to understand the world. Complexity, in short, is inextricably bound up with your point of view. (McIntyre, 1998, p. 28)

Should we accept this position, or is there some way of formulating complexity that makes it a strongly intrinsic property of some systems? My aim is to show the latter. Although complexity is representation-dependent, such that different representations result in different complexity measures, some of these representations may be 'better' than others when judged according to some sensible criteria. If we can pick out a single measure from the alternatives using these criteria, the choice of representation is no longer arbitrary, and we can perhaps establish that the resulting measure is based on some strongly intrinsic property. To demonstrate that complexity is a

strongly intrinsic property, therefore, I need to specify some sensible criteria for picking out a single privileged representation and ruling out the alternatives.

4.5 Effective complexity

The sort of intrinsic complexity measure I have in mind is one based on a description of a system that captures all its perceivable regularities. The property defined by such a measure is intrinsic because there is no room for making subjective decisions about what counts as a regularity and what does not—all regularities are counted. To illustrate what I mean by ‘all perceivable regularities’, consider the previous example of generating strings of symbols based on pre-defined grammatical rules. For these systems, we can exhaustively list every rule governing their regularities, if only because we specified such rules prior to creating the system.³⁰ There is no room for arbitrary decisions about what does and does not count as a regularity in such systems, because the systems have been created with all their regularities already specified.

To define a measure of complexity based on ‘all perceivable regularities’ for less contrived examples, we need to specify a well-defined procedure for determining what does and does not count as a regularity of any given system. Gell-Mann and Lloyd (1996) propose such a procedure. What this procedure gives us is a way of privileging some representations over others in terms of how much of the regularities of the system each describes. Complexity is then measured by the length of the representation that most concisely describes the maximal set of regularities. This they call ‘effective complexity’.

³⁰ Even though a particular string might happen, out of pure chance, to possess a small number of regularities that cannot be accounted for by the rules governing its generation, such regularities play a

The idea that some representations are better than others can be made precise through the definition of ‘total information content’. This is the sum of the Kolmogorov complexity of the representation we are using to approximately describe the system, and the information left over that is required to fully specify the system (much of which is residual entropic or Shannon information). The best representation is the one that minimises the information content of the representation with the constraint that it also minimises total information content. In other words, the best representation uses all law-like regularities of the system to be both maximally descriptive and maximally concise. It contains no redundancy and is otherwise no longer than it needs to be (otherwise the complexity measure would not be minimal), but neither is it overly simple in the sense that it fails to describe certain regularities of the system (otherwise the total information content would not be minimal). Expressed another way, the complexity of a system can be thought of as the amount of compressible information possessed by the system, once that information has been maximally compressed using a Kolmogorov-type procedure.

For complexity measures of physical systems, the constraint on minimising total information content automatically rules out spurious representations, such as specifying the system by the distance of every component from *Proxima Centauri*. This automatically overcomes the problem encountered in the previous chapter when discussing raw biodiversity, namely how to decide what properties to use in determining differences among groups of entities. The constraint on minimising total information content means that a representation based on a set of scientific theories that concisely represent certain regularities in the system is better than an atheoretical representation or one based on only a crude understanding of the system. Indeed the more we know about the system in terms of being able to concisely express some of

small part in any complexity measure and become vanishingly small as we extend the measure to

its attributes, the closer is our complexity measure to an idealised effective complexity measure derived from a representation based on everything there is to know about the system.

The effective complexity of any object is an intrinsic property of that object in Lewis' strict sense. It does not matter what we think of that object, nor what its relationships are with any other objects; its effective complexity is an intrinsic property that can be objectively measured. Our complete and maximally concise representation of the object, and hence our measure of its effective complexity, would be the same length regardless of whether or not the object existed in a lonely world. Any representation that is not the same length is either not complete or not maximally concise. Of course, we might not exist in such lonely worlds to construct a representation and so measure the object's effective complexity, just as we might not exist to measure its mass, but the property itself would still exist.

Such an abstract idea as a complete and maximally concise representation can be better understood with an example. Let us call C our complexity measure, which is simply the information content of our representation, and S the residual information required to fully specify those aspects of the system not already expressed in C . The total information content, say T , is defined simply as $T = C + S$. Now take our VW Beetle. A measure for C based on a representation in which the position of every atom in the Beetle is specified would leave nothing left to specify, so then we would have $S = 0$ and $T = C$. This representation includes no theories about any regularities in the positions of the atoms of the VW Beetle. Hence C is very large, and so is T . Now we ask ourselves, can we specify a representation that reduces C and/or T ? Clearly we can. We know that the steering column, for example, is made from steel

strings of infinite length.

and has a specified set of dimensions. We can therefore predict the approximate location of a large number of iron and carbon atoms, and therefore reduce the amount of information required to exactly specify their locations. Using this knowledge allows us to produce a new representation that reduces C . We still have $S = 0$ and $T = C$, so we have simultaneously reduced C and T .

If we keep modifying our representation by including everything we know about VW Beetles, we will continue to decrease C while simultaneously decreasing T . All the time, $S = 0$. Eventually, we will reach the limit of our knowledge and be unable to further reduce C while keeping $S = 0$. Now we ask ourselves, can we further reduce C without making T any bigger? Clearly we can. We have no theory that allows us to concisely express the exact location of every air molecule in the left front tyre, because these molecules are distributed randomly, at least within a bounding toroidal envelope. Their locations are still specified in C , but only because we had to measure them. In other words, most of the information required to specify the location of every molecule in the left front tyre is incompressible. If we now take this information out of C and put it in S , we can further reduce C without increasing the size of T . Hence we again have a better representation. We could not do this for aspects of the VW Beetle that have a high degree of regularity because doing so would decrease C while keeping T constant, whereas if we know about the regularities, we can actually do better than this and decrease C while also decreasing T . If we keep moving all the incompressible information from C to S , we eventually come up with a minimised value for C , with the constraint that T is also minimised. Moreover, the more knowledge we include about VW Beetles in our representation, the closer is our complexity measure to this ideal measure based on a representation that includes all the predictive theoretical knowledge about the structure of VW Beetles that could ever exist.

How would such a complexity measure apply to organisms? Obviously, we could carry out the same sort of analysis for a cockroach, but in that case we are not even close to possessing a full theoretical understanding of cockroaches. Nevertheless, it seems more than plausible that such a full theoretical understanding is possible, at least in principle. If this is indeed the case, the notion of an intrinsic complexity measure of cockroaches is well-defined, even though our existing complexity measures are at best only poor approximations of this intrinsic complexity. And despite the fact that we can measure such complexity only very approximately, we can nevertheless choose between competing complexity measures on the basis of how much theoretical knowledge they embody. Even on the basis of such crude measures, we can easily recognise that a cockroach really is much more complex than an equivalent mass of gas or crystalline solid. The amount of possible theoretical knowledge about either gases or crystals is relatively small. Gases have a lot of entropy and therefore a lot of total information content, but they have little complexity because there is virtually no theoretical knowledge that can be used to specify the state of their constituents more concisely than a complete description. Crystals, on the other hand, are easily described, so both their total information content and complexity are low. In contrast, there is an enormous amount that can be said about a cockroach that enables us to compress much of its information content. A cockroach may possess less entropy than a gas, but it possesses vastly more complexity.

5 Biocomplexity

It is interesting to contemplate an entangled bank, clothed with many plants of many kinds, with birds singing on the bushes, with various insects flitting about, and with worms crawling through the damp earth, and to reflect that these elaborately constructed forms, so different from each other, and dependent on each other in so complex a manner, have all been produced by laws acting around us.

— Charles Darwin

5.1 Complexity as a measure of biodiversity

At the core of my project is a simple proposal: that biodiversity be equated with complexity—in particular effective complexity. I have called this type of biodiversity *biocomplexity*. Biocomplexity is simply the effective complexity of any biological system. In this and the following chapter, I hope to flesh out this concept of biocomplexity and illustrate how it can be used to resolve some of the problems that beset biodiversity. Principally, biocomplexity is real. The biocomplexity of any living system is a strongly intrinsic property of that system. It does not depend on what we think of it.

To justify the proposal that biodiversity be reformulated in terms of biocomplexity, I need to do more than simply list its advantages. I need at least to show that biocomplexity somehow corresponds to what we want biodiversity to mean. First, consider the similarities between the two concepts. Both diversity and complexity capture notions of heterogeneity; the unlikeness or richness of constituent parts. The inability to adequately account for heterogeneity is central to the criticisms already raised against species richness as a concept of biodiversity. Dissimilar species should constitute more diversity than similar species. Likewise, the notion of redundancy embodied by the concept of effective complexity accounts for the degree of heterogeneity within a system by excluding multiple descriptions of any similarities.

Of course, when biologists (among others) use the terms ‘diversity’ and ‘complexity’, they rarely describe exactly the same thing. Lawton (1997), for example, uses ‘ecological complexity’ to denote something like the extent and integrity of ecological functions, and ‘biodiversity’ to denote species richness. These are two quite distinct concepts. Nevertheless, something like effective complexity is the kind of property that many biologists and conservationists seem to be referring to when they attempt to define biodiversity. A few examples from both scientific and popular publications can be found where the link between the two is explicit, such as:

Biodiversity has been seen as the total (and irreducible) complexity of all life, including not only the great variety of organisms but also their underlying behaviours and interactions. (NHM, 1994)

Similarly, Levin (1996) relates the appropriate target property for a biodiversity measure to the ‘functional complexity’ of ecosystems.

In the measurement of biodiversity, one must recognise the diversity within species as well as the diversity in terms of number of species, or do away with the notion of species entirely in favor of ‘continuum’ measures of the genetic and functional diversity of communities. Such continuum measures are more robust and probably more nearly represent the functional complexity of the system and its ability to respond to perturbations. (p. 279)

Hence, at least in some cases, the concepts of biodiversity and complexity are not particularly distinct.

A survey by Takacs (1996, p. 46–52) of eminent biologists’ opinions on what ‘biodiversity’ means to them is revealing. Many of the given definitions of biodiversity are vague, covering “the living resources of the planet” (Ehrlich), “the diversity of living things on the face of the earth” (Ilitis), “the diversity of life in all its manifestations” (Erwin) and “the richness of life” (Noss). More explicit definitions, however, sound a lot like effective complexity. Lovejoy describes biodiversity as “diversity at

all levels of organization,” while for Falk it is “the dimension of difference at multiple levels of organization.”

The most explicit definitions of biodiversity focus on these multiple levels of biological organisation, with particular emphasis on the three levels already mentioned: genes, species and ecosystems. Wilson (1996), for example, offers

Biodiversity is defined as all hereditarily based variation at all levels of organization, from the genes within a single local population or species, to the species composing all or part of a local community, and finally to the communities themselves that compose the living parts of the multifarious ecosystems of the world. (p. 1)

whereas May (1994) asserts that

Biological diversity exists at many different levels, from the genetic diversity within local populations of a species, or between geographically distinct populations of the same species, all the way up to communities or ecosystems. (p. 13)

Several consistent themes emerge from these attempted definitions. Biodiversity, it seems,

1. includes notions of organisation and interaction as well as heterogeneity,
2. covers many (perhaps all possible) levels of biological organisation, and
3. is an intrinsic property.

Whereas the first and second points are made explicit in many of the above quotations, the third is implied by phrases such as May’s “exists at many different levels.” Not only is biodiversity something that *exists* independent of what we think of it, it manifests itself at multiple levels, some of which we have yet to discover. Biodiversity therefore has, in the words of Michael Polanyi, “the capacity ... to reveal itself in

unexpected ways in the future.” Biodiversity is, in short, an intrinsic property of the world.

Biocomplexity, as I have shown, has all of the above three attributes we want from a concept of biodiversity. Recall that the biocomplexity of a biological system is measured by the length of a complete and maximally concise description of the regularities of that system. As such, a measure of biocomplexity must include all possible heterogeneity, organisation and interaction within the system. Otherwise the description on which it is based would be either incomplete or not maximally concise. Moreover, biocomplexity is an intrinsic property of any system, since it is based on a single, well-defined³¹ and privileged system representation. Hence biocomplexity seems like the right sort of concept for understanding biodiversity. In the following section, I hope to show why any reasonable biodiversity measure based on a more weakly defined target property could not be strongly intrinsic.

5.2 The uniqueness of biocomplexity

I have adopted the term *biocomplexity* for several reasons. One is to provide a concise means of referring to the concept of biodiversity based on effective complexity. Another is simply to delineate the concept from the glut of alternative biodiversity concepts already discussed. There is, however, a third and stronger reason. This is that *biodiversity*, as I see it, is a misnomer. *Biodiversity* is not about *diversity* but about *complexity*. The concepts of *diversity* and *complexity*, although similar in many respects, are not exactly the same. In particular, biodiversity is an intrinsic property of any biological system (or at least that is the sort of property we want it to be), and yet, as

³¹ Well-defined in the sense that there exists a well-defined procedure for generating the appropriate system representation, even though full execution of this procedure may be hopelessly impractical.

I hope to show, no concept of diversity weaker than effective complexity itself could possibly be intrinsic.

Despite asserting that *biodiversity* is a misnomer, I am not advocating that it be expunged from the popular lexicon and replaced by *biocomplexity*. *Biodiversity* is too deeply embedded for that, whereas *biocomplexity* has already been coopted for use in the application of complex systems theory to biological systems.³² *Biodiversity* is a useful and acceptable term. As a concept for use in conservation efforts, however, we should recognise that what we mean by biodiversity, if we are talking about some intrinsic feature of the world, is necessarily biocomplexity.

Recall the objection raised in Section 3.3, when discussing raw biodiversity, against a theory-free measure of diversity. Without any theory to discount ‘unnatural’ properties such as ‘being 4.2 ± 1.7 light years from *Proxima Centauri*’ or ‘being neither an aardvark nor an aeroplane’, any two objects are both similar and dissimilar in infinitely many ways. Hence a theory-free measure of diversity is impossible. To construct a sensible diversity measure, we need to limit what counts as a property by considering only those properties that stem from theoretical knowledge about the objects in question. We were still left with a problem, however; that of deciding which theories to use. How do we decide which theories are unnatural and which are not? How can we decide between ‘appropriate’ theories like evolutionary biology, under which we might derive a property such as ‘sharing a recent common ancestor’, and inappropriate theories like astrology under which we might derive a property such as ‘born under the sign of Scorpio’? An obvious solution is to limit ourselves to theories that have some predictive value; that is, those that give rise to projectable

³² A recent Internet search generated several hundred hits, with the term now featuring in several journal and conference titles.

properties. After all, such projectable properties are what allow us to derive compressed descriptions of complex biological systems. So ‘sharing a recent common ancestor’ counts because it allows us to predict some of the features of one object by studying those of others that share the same property. If one bird builds its nest from mud, we should not be surprised if other birds of that species do the same. Spatial relationships with distant celestial objects, in contrast, do not give rise to such projectable properties. Scorpion horseshoe crabs share no more properties with me than do those born under the sign of Sagittarius.

Although we have now limited ourselves to theories with predictive value, we have yet to fully specify which theories to include and which to exclude from our diversity measure. Unless we can find some way of privileging one set of theories over all others, we will fail to demonstrate that our measure of diversity is strongly intrinsic. That is, we would have only diversity with respect to some arbitrarily chosen set of theories.

I suggest that there are several privileged sets of theories that can be used to generate a description of a system. One is the empty set. This is the set that would be used to generate a theory-free description of the system. This description would contain all statements of the form ‘is 4.2 ± 1.7 light years from *Proxima Centauri*’ or ‘has never been found in a ferret’s stomach on a Tuesday’. The empty set is privileged, because it is unlike any other set in a very special way. It is the only set containing no members. There is nothing arbitrary about the theories contained in the empty set, because it contains no theories. However, as we have already seen, it is impossible to measure diversity with respect to this set; that is, to measure theory-free diversity. Clearly, the list of properties possessed by any system with respect to the empty set of theories would be infinitely long, and a measure based on the length of such a description is therefore ill-defined.

Another privileged set of theories is that represented by a complete and maximally concise description of the system. In other words, the set consists of *all* possible descriptive theories pertaining to the system in question. This description is akin to the complete description of our VW Beetle encountered previously, where all possible theoretical knowledge of VW Beetles has been used to compress the information content of the description, but all incompressible portions of the complete description remain. As we have seen, measuring the length of such a description is akin to measuring Kolmogorov complexity, which is akin to measuring entropy.

Using this same maximal set of theories pertaining to VW Beetles, we could also generate another description in which incompressible information has been removed. This would then be a complete and maximally concise description of all the *regularities* of the VW Beetle. The regularities of the system can be distinguished from the irregularities by the ability to describe them more concisely than with a complete (uncompressed) description, so no arbitrary choice need be made about what counts as a regularity. As we have also seen, the length of this description can be used as a measure of effective complexity.

Any set of theories that falls between the two extremes described above—the empty set and the maximal set—must necessarily be arbitrary. We must either include no theoretical knowledge of the system, or all possible theoretical knowledge of the system. Anything in between must involve an arbitrary selection. Even the set containing all *current* theoretical knowledge involves an arbitrary choice, namely the moment of time defining *current*. Since we can discount the empty set as not generating a useful measure, it follows that any intrinsic diversity-like measure must be based on a description using all possible theoretical knowledge of the system.

As we have just seen, there are at least two intrinsic measures that can be generated using the maximal set of theoretical knowledge pertaining to a system. One is an entropy measure and the other a measure of effective complexity. Are there others? I suggest not. To demonstrate why, take a particular system such as our VW Beetle. Now consider the set containing all the descriptions based on the maximal set of theoretical knowledge about our VW Beetle. Such descriptions need not be complete; they need only be based on the maximal set of theoretical knowledge about our VW Beetle. Suppose we now choose one such description from this set at random. What can we conclude about this description?

First, we know that the description must be maximally compressed—if it were not maximally compressed, then there would be some piece of theoretical knowledge about our VW Beetle that could potentially be used to compress the description further, meaning that the given description would not be based on the maximal set of theoretical knowledge about our VW Beetle. Secondly, since the description uses the maximal set of theoretical knowledge about our VW Beetle, it must at least contain a complete description of the compressible aspects of that VW Beetle. If it did not, the set of theories on which it is based would once again not be maximal.

Finally, we cannot say anything about the overall completeness of the description. The overall description would contain a complete description of the regularities of the system, but need not contain a complete description of the non-regularities. In other words, any description based on the maximal set of theoretical knowledge about a given system must contain a complete and maximally compressed description of the regularities of that system and some possibly partial description of the non-regularities.

If the description of the non-regularities is maximal in size (that is, it describes all the non-regularities), then our total description is complete and the measure based on the

length of that description is an entropy measure. If the description of the non-regularities is minimal in size (that is, of zero length), then the total description only applies to the regularities of the system and the measure based on the length of that description is a measure of effective complexity. For anything in between, an arbitrary choice must be made as to which parts of the incompressible description to include. That leaves the two extremes as the only two non-arbitrary measures based on the maximal set of theoretical knowledge: the entropy measure and the effective complexity measure. Any more weakly-defined notion of diversity must involve an arbitrary choice about which theories to use. Hence, only by defining diversity in terms of entropy or effective complexity can it be considered an intrinsic feature of the world. And entropy is simply not the sort of diversity measure we want to use for assessing biodiversity. If that were the case, maximising bioentropy could be achieved by turning organisms into gases. That leaves one reasonable and intrinsic biodiversity measure: biocomplexity.

This discussion explains why biodiversity is such a problematic concept, and why I think the term *biodiversity* is a misnomer. From the outset we have been trying to develop a concept of biodiversity that is intrinsic—a non-arbitrary, non-anthropocentric and objectively measurable feature of the world. And yet no notion of diversity weaker than effective complexity itself could possibly be intrinsic as well as measure some diversity-like property. Only by looking at biodiversity in terms of complexity can we make it an intrinsic property while retain some sensible meaning.

5.3 Defending biocomplexity

Before I try to describe biocomplexity beyond the claim that it can be measured using the length of a complete and maximally concise representation of the regularities of a given biological system, I need to address some immediate potential objections to biocomplexity.

One apparent problem with effective complexity (and hence biocomplexity) is that the choice of language used in any description is arbitrary, which seems to counter the claim that effective complexity is strongly intrinsic. We have already encountered this problem when discussing Kolmogorov complexity. Recall that the Kolmogorov complexity of a string of symbols is defined as the length of the shortest computer program that will output that string and then stop. The definition says nothing about the programming language that should be used. As we saw, this problem can be overcome through some mathematical tricks to minimise the size of the programming language used for each class of systems being measured. Although this procedure might be near impossible to apply in practice, in principle at least, Kolmogorov complexity is not language-dependent. Since the effective complexity of any system can be equated to its Kolmogorov complexity less the information content of its ‘noise’, effective complexity, when measured at its most fundamental level, is also not language-dependent. In practice, of course, we might use natural languages and other less perfect forms of representation when trying to construct an approximate and relative measure of effective complexity, but this is a problem only with our measurement techniques, not with the underlying property itself.

A second possible objection with my definition of biocomplexity concerns what I mean by ‘regularities’. It might seem that what counts as a regularity to one person might not be a regularity to another. I have, however, already countered this objection by giving a well-defined procedure for deciding what counts as a regularity. Regularities are those elements whose complete descriptions are compressible. A gas, for example, has no regularities beyond its bounding envelope, density and so on. A complete description of the gas would require that the position of every molecule within the bounding envelope be specified. No more concise description exists, and hence such a description is incompressible. A perfect crystal of silicon, on the other hand, can be concisely described by simply specifying the position, shape and orientation of the crystal, and specifying the structure of its lattice. The position of

every atom in the crystal can then be predicted with precision, because each atom exhibits a well-defined and fixed spatial relationship with its neighbours.

Another possible objection concerns the temporal aspects of biocomplexity. To date my discussions of effective complexity have been very much focussed on the structure of static physical systems, whereas for biological systems we are much more interested in functions and processes; things with a clear temporal dimension. Functions and processes are an important component of biocomplexity. There is no reason why biocomplexity cannot be applied to dynamic systems; I have simply limited my discussions to static systems to make the examples easier to understand.

Finally, there is the objection that a complete and maximally concise description of any physical system is simply not available. This is particularly the case for biological systems, which are so complex as to defy complete description. Although true, this is a practical rather than conceptual limitation. Effective complexity is almost impossible to measure precisely in practice, but that does not prevent us from using it to define a non-operational measure of biodiversity. Moreover, we can still devise reasonable indicators of biocomplexity based on course-grained descriptions. The indicator that best approximates the biocomplexity of a given system will be the one that uses the most knowledge at our disposal to characterise all law-like regularities.

5.4 Biocomplexity—a summary

Biocomplexity—the measure of biodiversity based on effective complexity—captures the main attributes that we want from a concept of biodiversity. First, it incorporates the notions of heterogeneity, organisation and interaction that underlie many biologists' discussions of biodiversity. Complex objects such as organisms and ecosystems have many different parts, intricately organised and interacting among

themselves. The greater the internal heterogeneity and organisation of such objects, the greater is their measure of effective complexity.

Secondly, biocomplexity is a multilevel measure. It is not limited to diversity at a single level such as species, characters or genes. It incorporates all these levels and more, accounting for diversity both between and within species and ecosystems. It incorporates what we might previously have considered different types of biodiversity, such as phylogenetic diversity and ecological diversity. We can therefore use biocomplexity to quantify more specific conservation targets such as species diversity, intra-species diversity or ecosystem integrity. A change in any of these conservation targets has a quantifiable effect in terms of biocomplexity.

Thirdly, and perhaps most importantly, biocomplexity is a well-defined, measurable and intrinsic feature of the world. In contrast, a measure of biodiversity such as higher taxon richness is based on some arbitrary choice of higher taxon levels, making it inherently anthropocentric. Unless we can formulate biodiversity in some strongly intrinsic sense, reaching agreement on a target property for conservation may prove problematic. Just as the lack of a stable target property led to the failure of phenetic taxonomy, agreement can never be reached on what constitutes the underlying target property of any non-intrinsic biodiversity measure. There simply is no single privileged target property underlying such a measure on which we can all agree.

Biocomplexity solves this problem outright. It is an intrinsic property of any biological system. Moreover, it is the only intrinsic property that at least approximately corresponds with what we want from a biodiversity concept. In other words, not only is biocomplexity an intrinsic property of biological systems, it is the only intrinsic property of such systems on which we can possibly base a concept of biodiversity. This gives us some hope that we can at least all agree on what biodiver-

sity is. Just as there is one true tree of life, ensuring stability in attempts to derive a complete cladistic taxonomy, there is also one true measure of biodiversity, namely biocomplexity, that can form the foundation on which we can agree about how to make rational conservation decisions.

6 Applied biocomplexity

Nothing in biology makes sense except in the light of evolution.

— Theodosius Dobzhansky

If the concept of biocomplexity is to have any impact on the process of making real conservation decisions, then we must be able to apply it in real situations. In this chapter, I hope to develop the concept of biocomplexity further and show what a practical approach to measuring biocomplexity might look like. I then return to the various measures of biodiversity discussed in Chapter 3, and discuss their relevance as indicators of biocomplexity.

6.1 Estimating biocomplexity

The effective complexity of any system, you will recall, is the information content of a complete and maximally concise description of that system. A true measure of the effective complexity of any biological system—in other words of biocomplexity—would therefore require complete knowledge about that system. Despite its conceptual advantages, developing a true measure of biocomplexity is almost certainly impossible in practice. Even a single organism such as a cockroach is too complex to describe in complete detail. Hence we will never, it seems, be able to measure with complete precision the biocomplexity of even a single cockroach, let alone an entire ecosystem.

Clearly, any practical approach to assessing biocomplexity will be based not on a true and complete measure of biocomplexity but on some representational estimator. The most straightforward way of estimating biocomplexity is to forgo a complete description of the given biological system, and instead use a partial ‘coarse-grained’ description. The estimated biocomplexity of the system would then be based on the length of this coarse-grained description. Although this would not produce an

absolute estimate of biocomplexity, it could still provide a useful estimate of relative biocomplexity. Fortunately, for nearly all practical problems involving biocomplexity, such as assessing conservation value or determining whether biocomplexity correlates with ecosystem stability, a relative measure is sufficient. Although a crude estimator of biocomplexity may not be able to tell us, for example, how much biocomplexity is possessed by a particular area of land, it may be able to tell us whether that area of land has more or less biocomplexity than some other area.

As we saw in Chapter 4, the more detailed we make our description of something, the more accurate it tends to become as a measure of effective complexity. The most accurate estimators of biocomplexity will therefore tend to be those that include the most detail. Conversely, less detailed descriptions will tend to give rise to less accurate estimators of biocomplexity, although such limited accuracy may be sufficient for many applications.

What sort of description is appropriate for estimating biocomplexity? In Chapter 4, we compared the complexity of a VW Beetle and a Boeing 747 by comparing the lengths of their descriptions written at the level of a technical service manual. In comparing the complexities of biological systems, we could adopt a similar approach. Consider a single organism such as a cockroach. Although our existing biological knowledge is limited, it could still be used to give a highly detailed description of a cockroach by focussing on its various levels of internal organisation. For example, we know that multicellular animals such as cockroaches can be described in terms of their constituent parts or organs, each of which has one or more specialised functions. The organs of cockroaches, for example, include a mouth, two eyes, six legs and two antennae. Most of these organs, in turn, also consist of numerous subcomponents. Each leg, for example, comprises both a chitinous shell and numerous hairs. We can generally classify each of these subcomponents in terms of its tissue type. Each tissue type is in turn composed of multitudes of microscopic cells, which can

themselves be divided into numerous subcomponents or organelles—chromosomes, ribosomes, mitochondria, vacuoles, lysosomes and a host of other structures. Even the organelles are not the most fundamental unit of organisation, as each can be described in terms of its constituent enzymes, which in turn consist of various amino acids.

It seems as though the more we learn about the internal organisation of cockroaches, the longer our description of them becomes. However, there is a sense in which our knowledge of cockroaches reduces their potential description length. That is, if we were to describe our cockroach by specifying the exact position and type of every atom in the cockroach's body, as well as how these positions change over time as the cockroach moves, eats, digests, reproduces and carries out its other everyday functions, we would be confronted with an extremely long description. Such a description might be complete, but it is not very concise, and hence unsuitable for assessing biocomplexity. Using knowledge of the cockroach's internal organisation reduces the length of this description. We can give a pretty good description of where a large part of the cockroach's atoms reside by describing how many legs the cockroach has, what their morphologies are, what cell types these legs are made of, what those cell types look like and so on. Only with this detailed knowledge of cockroach physiology can we have some hope of concisely describing the internal organisation of a cockroach in detail, rather than relying on the hopelessly impractical approach of describing the exact position of every cockroach atom.

Whereas the sort of detailed description of a machine such as a VW Beetle or a Boeing 747 results in something that resembles a service manual, a detailed descriptions of an organism such as a cockroach results in something that resembles a monograph. Of course, a typical monograph takes for granted much basic biology, and thereby omits much of the sort of detail discussed above. Nevertheless, we can think of the sort of detailed description of an organism being considered here as a

sort of idealised monograph. Such a description, although not complete, is still detailed enough for making reasonably accurate relative estimates of biocomplexity. We could, for example, compare the length of the idealised monograph of one organism such as a cockroach with that of another organism such as a coral to estimate the relative biocomplexity of one with respect to the other.³³

Writing idealised monographs of individual organisms such as cockroaches and corals is still a long way from contributing to the practical process of making conservation decisions. Conservation decisions generally apply at larger scales of biological organisation such as whole ecosystems. To assess the biocomplexity of an ecosystem using the above methodology would require us to write an idealised monograph of each species in the ecosystem, combine them, and then take away all the redundant sections, such as might occur with closely related species. Such an exercise is still far too impractical.

The way around this problem is to omit more detail; to carry out the coarse-graining process at a higher level. For most practical applications of biodiversity in conservation, we are typically only interested in determining whether ecosystem A has more or less biodiversity than ecosystem B. In such cases, we can generally ignore many of the lower levels of biological organisation discussed above. Since DNA is a component of all earthly organisms, it is not necessary to describe the structure of DNA when comparing the biocomplexities of each system. The complete descriptions of each system will both contain a common description of the structure of DNA. Similarly, nearly all areas of conservation value will contain a mixture of plants, animals, bacteria and fungi, so there is no need to describe the cell structures of each

³³ Such an exercise would presumably demonstrate that cockroaches are substantially more complex than corals.

of these types of organism if they are common to each system being evaluated. By omitting these lower level details, we can focus instead on the larger scale patterns.

If we only wish to decide which of ecosystem A or B has more biocomplexity, then even descriptions with a modest level of detail may no longer be necessary. In many cases, the answer is obvious. A tropical rainforest really is more complex (more biodiverse) than a desert. There are not only more species in a tropical rainforest than a desert but more morphological disparity between species and more ecological relationships. Without even attempting to describe either in detail, we can be reasonably certain from the outset that a detailed description of one will be much longer than that of the other.

In cases where the outcome is less obvious, existing estimators of biodiversity may often be adequate. All other things being equal, a forest containing nine endangered mammal species is probably more biocomplex than another forest containing seven. Of course, a more detailed assessment may show this not to be the case. The set of seven mammal species may include several unique members, while the set of nine mammal species may all be from the same genus. However, such cases can still be adequately addressed by using a more sophisticated measure than simple species richness, such as higher taxon richness.

As suggested by the above discussion, many existing measures of biodiversity such as species richness are reasonable, albeit imperfect, indicators of biocomplexity. This is a topic to which we shall return in the following section. Before doing so, however, let us consider in more detail the type of information we might expect to use in devising practical indicators of biocomplexity.

Many of the various measures of biodiversity discussed in Chapter 3 were classified as phylogenetic measures. Such measures generally involved taking a set of species

and then measuring the relative amount of biodiversity of each subset by applying some set of criteria based on phylogeny. Hence phylogenetic knowledge is a good candidate for inclusion in an indicator measure of biocomplexity. Phylogenetic knowledge is useful because it explains a large amount of the pattern we see in nature. Since closely related species generally share a host of features, phylogenetic knowledge is predictive, in the sense that we can guess a lot of facts about a particular species from knowledge about its close relatives. In other words, phylogenetic knowledge allows us to concisely express a large amount of information. Once we have described the phylogeny of a group of species, we have generally gone a long way to describing the patterns of similarity between them.

As we saw in Chapter 3, the patterns of similarity between species arise from two different sources: phylogeny and adaptation. Although phylogeny accounts for much of the larger-scale patterns of biological similarity, it does not account for all of it. Some is attributable to adaptation. So while the similar egg-laying habits of birds and lizards owe their origins to a common ancestor, the similarities in wing shape between birds and pterosaurs are due to convergent adaptation.

Given that phylogenetic knowledge can be used to devise indicator measures of biocomplexity, we might expect that ecological factors could be used to devise alternative indicator measures. Doing so would require a scheme for classifying and describing organisms on the basis of their ecological categories. One way of doing this is to use Hutchison's notion of a *niche space* (Sterelny & Griffiths, 1999, Ch. 11). A niche space is a multidimensional space in which each axis corresponds to some environmental variable such as temperature or food size. Organisms 'inhabit' this space, in the sense that they can only survive within a certain range of values for each variable. For example, Figure 9 illustrates a two-dimensional portion of niche space inhabited by three species, with dimensions representing temperature and food size. In this example, species 1 inhabits a distinct portion of niche space with respect to

the other two species, whereas the regions inhabited by species 2 and 3 significantly overlap.

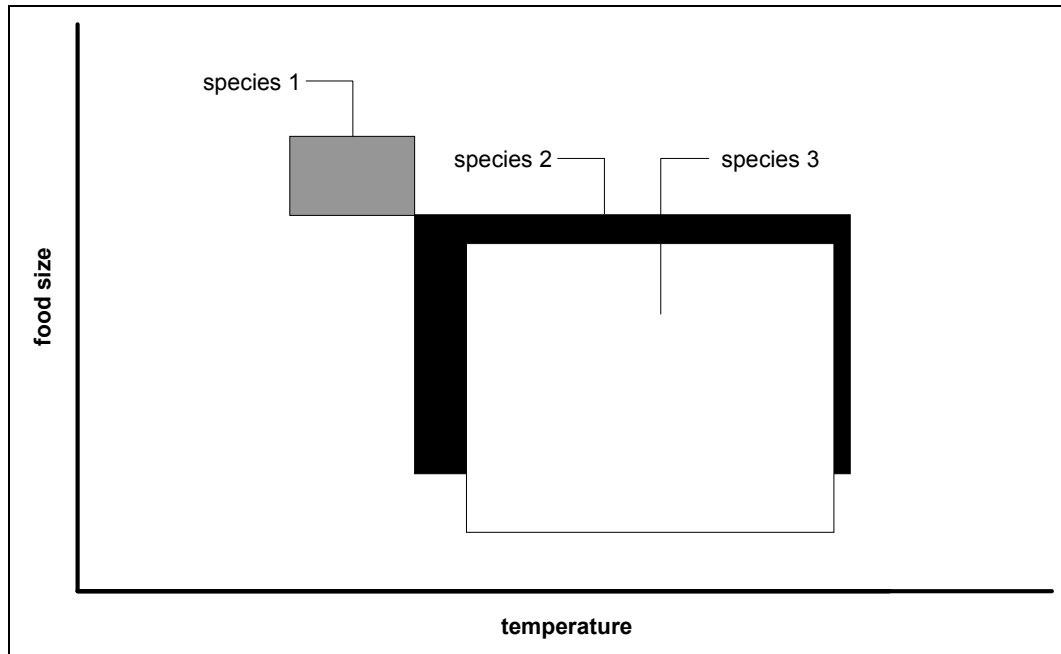


Figure 9: A two-dimensional view of a small portion of niche space inhabited by three species.

We can think of an ecological description of a species as a description of the region of niche space it inhabits. An indicator measure of biocomplexity based on niche space would simply equate the relative biocomplexity of a set of species with the total amount of niche space inhabited by those species. So a set of similar species all with similar niches would not occupy too much of niche space and would have a relatively low measure of biocomplexity, whereas a set of disparate species, each occupying a largely disjoint region of niche space, would have a relatively high measure of biocomplexity.

An ideal measure of biocomplexity would account for all biological structure, regardless of its origin. Hence we might expect a sophisticated indicator measure of biocomplexity to integrate both phylogenetic and ecological knowledge. However, the balance between the two need not be equal. Indeed, there are good reasons for

such a measure placing more emphasis on phylogenetic rather than ecological patterns. As discussed previously, both phylogenetic and ecological categories are projectable, but they are not equally so. Knowing that two species share a phylogenetic category such as *Aves* allows us to project much of the knowledge of one species onto the other. Ecological categories such as *pelagic macroorganism*, however, are not as strongly projectable, since ecological functions are multiply realisable. Aquatic propulsion is achieved by horizontal tail movement in sharks, vertical tail movement in cetaceans and jet propulsion in cephalopods. Even where convergent adaptations do exist, the similarities between them are largely superficial. Vertebrate and arthropod eyes share a function, that of vision, but morphological they are very different. They both embody principles of optics, but otherwise there is little to be said about one that can be said about the other. The evolution of shared derived traits generates a degree of additional biocomplexity that the evolution of shared inherited traits does not.

There are also good reasons to suggest that the balance between phylogenetic and ecological knowledge in indicator measures of biocomplexity might vary depending on the evolutionary distance between the taxa being assessed. When assessing closely related taxa such as the Anseriformes,³⁴ we might expect homoplasy to be common. In the case of disparate taxa, homoplasy is much rarer. Two species of duck share a relatively recent common ancestor, but they may also share a range of ecological variables such as feeding mode. The two species may therefore share similar evolutionary pressures, as well as inherited evolutionary constraints. Hence many of the adaptations to a diving mode of feeding in one species will be shared by another species with the same feeding mode (Faith, 1989). Koalas and caterpillars also share a

³⁴ Ducks, geese and swans.

feeding mode, that of leaf eating, but the radically different evolutionary histories of these two taxa have led to radically different solutions to a leaf-eating lifestyle. Hence we should see an increased predominance of phylogenetic over ecological patterns as we move from measuring biocomplexity over smaller to larger scales. If the scale of analysis is sufficiently large, a totally phylogenetic focus may be valid.

6.2 Measures of biodiversity—a retrospective

In Chapter 3, I discussed some of the principal biodiversity measures that have been developed for assisting the process of making conservation decisions. We are now in a position to revisit these measures and assess their usefulness as indicators of biocomplexity.

6.2.1 Biocomplexity, species richness and higher taxa

The first biodiversity measure encountered in Chapter 3 was *species richness*—simply the number of species in a given sample. The main advantages of species richness are its conceptual and practical simplicity. Species are, for the most part, easily delineated and countable entities. However, as we saw in Chapter 3, species richness is unable to account for the degree of difference between species. Some clades such as beetles are millions of times richer in species than others such as tuataras and coelacanths, yet seem to possess far less than a million times as much biodiversity. Similarly, a group of organisms comprising two distinct species would appear to have more biodiversity than a group comprising two similar species.

As an indicator of biocomplexity, species richness adopts a fairly coarse-grained level of representing biological systems. The system is described simply in terms of a list of species, with each species counting as a single unit. As we have already seen, such coarse-grained analysis may be perfectly adequate for many applications. In most

cases, a larger set of species will tend to be more biocomplex than a smaller set, which is often all we need to know.

Although species richness does tend to correlate positively with biocomplexity, it does so rather imperfectly given its inability to account for the degree of difference between species. Species are not additive units in terms of their biocomplexity. There is much redundancy between the descriptions of two closely related species, and hence their joint biocomplexity is substantially less than the simple addition of their individual biocomplexities. Conversely, species richness will tend to be more accurate as an indicator of biocomplexity when the species in question are relatively disparate.

In Chapter 3, we also encountered the biodiversity measure of *higher taxon richness*. As we saw, this measure offers both practical and theoretical advantages over species richness. Higher taxa are more easily distinguished and less numerous than individual species, making them a more practical unit of study in field surveys. Higher taxa also tend to categorise similar species together, thus avoiding the problem of highly speciose clades having inordinately high measures of biodiversity. The main theoretical problem with higher taxon richness is the artificiality of higher taxa themselves. There simply is no well-principled means of determining where we should draw the boundary that delineates one higher taxon from another.

Despite these limitations, higher taxon richness is potentially a very good indicator of biocomplexity. Provided we delineate the higher taxa appropriately, we can think of each higher taxon as comprising a similar amount of biocomplexity. Some taxa will contain many similar species, while others will consist of one or two distinct species. The combined description of all the species in a speciose taxon will contain lots of redundancy, whereas the description of the single species in a unitary taxon will not. Provided we have properly determined the right levels for delineating each taxon, the concise descriptions of each taxon should be about the same length. Furthermore,

given that the degree of overlap between the descriptions of higher taxa is less than in the case of individual species, combining the biocomplexities of higher taxa will tend to approximate an additive measure more accurately, albeit still somewhat imperfectly.

Another nice feature of higher taxon richness is that it meaningfully combines both phylogenetic and ecological patterns. Species richness, in contrast, measures only phylogenetic patterns. Its limitations as an indicator of biocomplexity stem in large part from its inability to account for the degree of evolutionary change within each lineage. Australia's crows and ravens are all remarkably similar in appearance and habit, while Darwin's Galapagos finches are not. Species richness simply treats each species as a single unit of biocomplexity, regardless of diversity. In contrast, the degree of evolutionary change in response to ecological pressures is a central criterion for determining the levels at which higher taxa are defined. The species within a taxon may be unified by their phylogeny, but they are separated from other taxa in part by their ecological-determined disparity.

6.2.2 Genetic biocomplexity

In Chapter 3, I discussed genetic measures of biodiversity and rejected the idea that we should think of genetic diversity as any more fundamental than, say, phenotypic diversity. I argued that there is no meaningful sense in which we can reduce phenotypic diversity to genotypic diversity. How then does genetic diversity fair as an indicator of biocomplexity?

Thinking of genes in terms of complexity is not a new idea. "Genetic information" is a common phrase even outside scientific circles. A straightforward approach to measuring genetic information would be to count the number of bits of information encoded by a sequence of DNA. The human genome, for example, contains approximately 3 billion base pairs, each of which encodes two bits of information (2^2

= 4 possible base pairs). This gives a total information content of 6 gigabits, roughly the information storage capacity of a single CD-ROM. Other species have even more genetic information. The newt *Triturus cristatus*, for example, has a genome with 40 gigabits of information (Dawkins, 1999). It seems, however, that newts are not an order of magnitude more complex than humans, so it is commonly presumed that the genomes of species such as *Triturus cristatus* contain a lot of redundancy. To measure what Richard Dawkins calls the “true information content” of the genome (Dawkins, 1999), we first have to strip out this redundancy.

As we have seen, complexity is not simply measured by information content. To measure the complexity of a genetic sequence, we need to measure the information content of the compressible portions of the sequence, once that information has been maximally compressed. Such a procedure should be applicable to genetic sequences, at least in theory. Some genomes really are more complex than others in terms of the extent to which they encode lots of compressible information. Perhaps there is even a good correlation between the complexity of genotypes and the complexity of phenotypes; an empirical question of some interest.

Practically, however, there are major obstacles to measuring genetic complexity. First, we cannot easily distinguish signal from noise. In other words, we do not know which parts of a genome are compressible, because there is insufficient agreement as to what constitutes so-called “junk DNA” (Kimura, 1983). Secondly, we have little idea of what compression algorithm to apply to the remaining genetic information to determine its maximally compressed length. In the computer sciences, various compression algorithms have been developed for dealing with all sorts of data files. Some of these algorithms are well suited to compressing files containing images, while others are good at compressing files containing music. The better we understand the system, the more optimally we can compress its information content. So if we know that a file contains an image, we can write a really good algorithm to

achieve a good compression ratio based on what we know about the way information in image files tends to be organised. If the file contains a music track, on the other hand, we would be able to achieve greater compression by writing a different algorithm using what we know about the way information is organised in music files.

The problem with trying to quantify genetic complexity is that we simply do not know enough about the way genetic information is organised to write a reasonable compression algorithm. In time, though, we may develop such understanding. Doing so requires two-way feedback between molecular and developmental biology. To infer a phenotype from a genotype, we first have to understand how the phenotype is built and then read that understanding back into the genotype. Our understanding of molecular biology grows by exploring the genetic and phenotypic levels simultaneously and letting them incrementally illuminate one another. With such an understanding, we might eventually be able to measure genetic complexity, which in turn may prove to be a very useful indicator of biocomplexity. At present, however, genetic complexity seems like a very poor practical indicator of biocomplexity.

6.2.3 Biocomplexity and phylogenetic biodiversity

In Chapter 3, I discussed various phylogenetic biodiversity measures. The main feature of these measures is that they use taxonomic distinctness as a means of quantifying the degree of difference between taxa. So if species A and B are more closely related to one another than either is to C, a phylogenetic measure of biodiversity would typically consider the set containing only A and B to be less biodiverse than the sets containing either A and C or B and C.

The most sophisticated of the phylogenetic diversity measures examined was *character richness*. We can use this measure to estimate which subsets of a group of taxa have the greatest amount of phylogenetic diversity by constructing the cladogram for the

group of taxa, weighting each branch according to the inferred character state changes, and selecting those subsets which maximise total intervening branch length.

Character richness is likely to be a more accurate estimator of biocomplexity than simple measures such as species richness because it uses a smaller unit of currency, namely character. Rather than use a representation in which each species counts as a single unit, character richness represents each species in terms of its constituent characters. In theory, such an approach could not only give an indication of the complexity of each species in terms of the number of characters it possesses, it could also, as an indicator of biocomplexity among groups of species, resolve the additivity problem by ensuring characters are only counted once, thereby removing much of the redundancy between the representations of each species.

Several criticisms of character richness were raised in Chapter 3. In particular, not all characters seem to warrant equal weighting. A complex trait such as a vertebrate eye, for example, can be broken down into a suite of constituent traits. There are no clear rules about whether we should consider this a single trait or a collection of several traits. A conceptual solution to this problem is to weight each trait according to its complexity. Whereas an eye, for example, would be given a high weight, simple traits such as skin pigmentation would be given a correspondingly low weight. Again, for many applications, such detail may be unnecessary; an equally weighted character richness measure may be more than adequate. Nevertheless, the potential to weight traits according to their complexity counters the conceptual objection that traits are incommensurate.

There are several other ways in which existing measures of character richness can be improved to make them more suitable as indicators of biocomplexity. For example, most existing measures do not count root characters. The reasoning is that character richness should measure diversity *between* set members. Hence any character state

changes before the first speciation event separating members of a solution subset do not count. From a biocomplexity perspective, this assumption is not valid. If we consider character state changes simply as an indicator of biocomplexity, then any character state changes are relevant, no matter when they occur. This problem is easily resolved by simply including root characters in the procedure of selecting solution subsets.

A further way of making existing character richness measures more suitable as indicators of biocomplexity is to change the way they treat homoplastic characters; that is, similar characters that have evolved independently. Biocomplexity is measured by describing all regularity. There is no discrimination as to the origins of regularity. So when we described our VW Beetle to generate a measure of its complexity, it did not matter whether the left front wheel was similar to the right front wheel because they were made by the same manufacturer or because different manufacturers came up with the same solution to a common problem. All that mattered was that they were the same. Most of the description of one can be used as a description of the other, allowing us to compress our overall description and minimise its length without sacrificing completeness.

An analogous argument can be applied to species. From the perspective of biocomplexity, it does not matter whether a particular trait shared by species A and B evolved independently or not. All that matters is that they share the same trait. Existing character richness measures ignore homoplasy by counting shared derived traits along each lineage. The solution is simply to refrain from double-counting homoplastic characters. In other words, we should measure character richness in a literal sense, rather than measuring phylogenetic diversity weighted by character state changes.

Depending on how we interpret the relationship between our indicator measure of character richness and our target property of biocomplexity, the solution suggested above may not be completely appropriate. So far I have been interpreting character richness as a way of representing a species or group of species. Under this interpretation, character richness is a direct indicator of biocomplexity. An alternative way of interpreting character richness is as an indicator of some intermediate property such as evolutionary distance. So the more character state changes there are between species A and B, the more independent evolution has occurred between them. According to this interpretation, evolutionary distance is a direct indicator of biocomplexity, whereas character richness is an indirect indicator. Under such an interpretation, double-counting of homoplastic characters is probably valid, since it is likely to result in a better correlation with the intermediate property of evolutionary distance. This second interpretation is reasonable, given that character richness measures are usually applied to representative samples of character state change data, rather than to some vast set of data intended to describe a set of species completely. In other words, character richness often looks like the ‘sample model’ phylogenetic diversity measures illustrated by Figure 3 in Chapter 3.

Despite the enormous practical difficulties in measuring biocomplexity with precision, estimating biocomplexity using various indicator measures is practically feasible. Moreover, several existing measures of biodiversity adequately provide such indicator measures. By providing a definitive answer to the question ‘What is biodiversity?’, the concept of biocomplexity can allow the practical development and application of such measures to continue without metaphysical encumbrance.

7 Conclusion

God is in the details.

— Ludwig Mies van der Rohe

7.1 A rational argument for biocomplexity

How should we understand biodiversity? My answer is simple—complexity. This is, I believe, the answer that makes most sense of the concept. It is the only way of understanding biodiversity that both makes it an intrinsic property of biological systems and incorporates the main attributes we want from such a concept. If we want a well-defined, measurable and metaphysically strong concept of biodiversity that captures more than just species richness—that incorporates all forms of heterogeneity, organisation and interaction at every conceivable level—then biocomplexity is the only answer. It is the only reasonable concept of biodiversity that delineates some objective property of the world.

The advantages of this strongly intrinsic concept of biodiversity should not be underestimated. First, it matches our intuitions—that biodiversity is an independently specifiable property we can measure. Secondly, it ensures stability. Without some such firm foundation on which to base our measurement procedures, there is little chance of reaching agreement on what it is we are trying to measure. Thirdly, a non-anthropocentric concept of biodiversity is meaningful. It enables us to discover things we do not already know, such as whether ecosystem A really does have more biodiversity than ecosystem B. Fourthly, an intrinsic concept of biodiversity carries more political weight. If biodiversity were only defined anthropocentrically, those arguing against its protection would have two basic strategies available to them. One strategy would be to dismiss the importance of biodiversity. This strategy might admit the existence of biodiversity, but deny its relevance as a target of conservation efforts. The second strategy would be to deny biodiversity's very existence as an

independently-specifiable target property of conservation. By arguing that biodiversity is suspect, proponents of such an argument would not even have to address its relevance to conservation. With a metaphysically robust concept of biodiversity, this second strategy is cut-off, narrowing the ground on which such battles can permissibly be fought.

Finally, an intrinsic concept of biodiversity is essential to those wishing to use biodiversity in pursuing environmental ethics from a position of moral realism. If species have intrinsic value—if they have rights that transcend what we think of them—then those rights must be rooted in something more fundamental than our particular preferences and prejudices. As I have already stated, I do not wish to pursue such an argument. For those who do, however, an intrinsic concept of biocomplexity is a good place to start.

In Chapter 2, I pointed out that there are two distinct types of biodiversity concept. One type, on which I have focussed, is targeted at issues of conservation. This sort of biodiversity is something of value worth protecting. The other type of biodiversity concept is scientific. It is the sort of biodiversity that might play some useful role in scientific theories. Biocomplexity, I have argued, is useful for understand the conservation-based type of biodiversity concept. How useful, then, is it for understanding the scientific concept of biodiversity? Could biocomplexity play some useful role in scientific theories, just as the concept of mass plays a useful role in Newtonian mechanics?

As I pointed out in Chapter 2, there are at least three features a concept must have for it to be scientifically valuable. It must be well defined, based on some natural properties of the world and useful in some scientific theory. Biocomplexity meets the first two criteria. Moreover, it is the only well-defined and strongly intrinsic property that roughly corresponds to the popular meaning of biodiversity. How well biocom-

plexity meets the third criterion—how useful it is in scientific theories such as the species-area law—is an empirical question, and therefore not mine to answer. However, as an intrinsic property, biocomplexity seems the most appropriate concept to use in attempting to answer such questions. Are complex ecosystems more or less stable than simple ecosystems? This is a question already being asked in ecology. Complexity is seen as the sort of property that might have some causal relationship with ecosystem stability. Like complex machines with cybernetic control systems, ecosystems have lots of feedback loops. A feedback loop with a small amount of gain can enhance system stability, while a feedback loop with a large amount of gain can create instability. Perhaps the complexity of an ecosystem determines how many and what type of feedback loops it has, which in turn determine its stability. Does biodiversity reflect the evolutionary time elapsed without major disturbance? Does it reflect the frequency of major disturbances? Again, these are not unreasonable questions. Answering them may be difficult, but at least with a solid concept of biodiversity we can be sure what questions we are trying to answer.

7.2 An impassioned plea for biocomplexity

In Chapter 1, I raised a question about the metaphysical status of biodiversity. I have, hopefully, provided a satisfying answer to this question. I also raised a quite different question—Why should we protect biodiversity? From the outset, I made a clear choice to avoid any attempt to answer this question definitively. And yet the question lingers. Why protect biodiversity? In the light of biocomplexity, this question has taken on a new significance. Why should we protect biocomplexity? What is so special about complexity that makes it worth having?

There is no simple answer to this question. Correlations between complexity and the various types of instrumental value possessed by the natural world are in many cases poor. The instrumental value of a wheat field is typically greater than the instrumen-

tal value of the forest that was cleared to produce it, and yet the wheat field is undoubtedly less complex. Some types of instrumental value, however, do correlate reasonably well with complexity. A forest is certainly more likely to harbour a cure for cancer or a natural predator for an introduced crop-eating insect than a wheat field. If we care about protecting the world's living resources for purely instrumental reasons, biocomplexity seems like a pretty good indicator of such instrumental value.

In the end, most arguments for protecting nature on the basis of instrumental value are, I believe, weak rationalisations of a deeper passion. In a thought experiment in which we strip nature of all its instrumental value, nature is not left valueless. Nature has non-instrumental (but non-intrinsic) value. It has inherent value. I have never seen Da Vinci's *Last Supper*. It was being restored when I visited Milan. Call me a philistine, but I have never read *The Tempest*. And I have never seen a kakapo. Yet all of these things matter. Why celebrate Shakespeare? Why preserve the *Last Supper*? Why protect the kakapo from extinction? None of these things are of much instrumental value to me, and yet I value them. I do not really understand why. Their value seems rooted in a desire that simply springs forth without rational foundation.

Complexity, I believe, is a pretty good indicator of inherent value. Why do we celebrate Shakespeare? Because his work is complex. Look at how difficult it can be to fully understand him, at how many books have been written on his work. We find Shakespeare fascinating because his work is complex; because there is so much to say about it. Barbara Cartland may have written more words than Shakespeare, but she essentially wrote the same story several thousand times with different names and locations. Unlike Shakespeare, her life's work is full of redundancy.

Why do we (ought we) try to protect the kakapo from extinction? Because its continued existence adds a significant amount of complexity to the world. Of course, a world without kakapos would still have lots of other parrots, but it would have no

nocturnal flightless parrots. There is a lot to be said about kakapos that cannot be said about any other creatures. They represent several million years of interesting and irreplaceable evolutionary design. Complexity matters. In the end, though, my argument for protecting biocomplexity is not easily defended. I can give no rational justification for protecting the inherent value of biocomplexity, because it is not reducible to more basic human desires. I can only hope that you accept my impassioned plea.

8 Bibliography

1. Anderson, P.W. (1972), 'More is different', *Science*, **177**, pp. 393–396.
2. Atkinson, I. (1989), 'Introduced animals and extinctions', in D. Western & M. Pearl (Eds.), *Conservation for the Twenty-first Century*, Oxford University Press, New York, pp. 54–69.
3. Atlan, H. (1988), 'Measures of biologically meaningful complexity', in L. Peliti & A. Vulpiani (Eds.), *Measures of Complexity*, Springer-Verlag, Berlin, pp. 128–138.
4. Bateson, P. (1991), 'Are there principles of behavioural development?', in P. Bateson (Ed.), *The Development and Integration of Behaviour*, Cambridge University Press, Cambridge UK, pp. 19–39.
5. Beardsley, T.M. (1988), 'Complexity counted? Physicists ponder a new way to measure an elusive concept', *Scientific American*, **259**, (2), pp. 16–17.
6. Bonner, J.T. (1988), *The Evolution of Complexity by Means of Natural Selection*, Princeton University Press, Princeton.
7. Bovet, D.P. & P.L. Crescenzi (1988), 'An introduction to the theory of computational complexity', in L. Peliti & A. Vulpiani (Eds.), *Measures of Complexity*, Springer-Verlag, Berlin, pp. 102–111.
8. Boyd, R. (1991), 'Realism, anti-foundationalism and the enthusiasm for Natural Kinds', *Philosophical Studies*, **61**, pp. 127–148.
9. Brooks, D.R. *et al.* (1995), 'Systematics, ecology, and behavior', *BioScience*, **45**, (10), pp. 687–695.
10. Brooks, D.R. & D.A. McLennan (1991), *Phylogeny, Ecology, and Behaviour: A Research Program in Comparative Biology*, University of Chicago Press, Chicago.
11. Brooks, D.R. & E.O. Wiley (1988), *Evolution as Entropy: Toward a Unified Theory of Biology (2nd ed.)*, University of Chicago Press, Chicago.
12. Callicott, J.B. (1992), 'The ethical view', *Defenders*, November/December.
13. Callicott, J.B. (1996), 'On Norton and the failure of monistic inherentism', *Environmental Ethics*, **18**, pp. 219–221.
14. Campos, D.H. (1995), 'Defining complexity', *Complexity*, **1**, (6), p. 5.
15. Casti, J.L. (1995), 'Complexity and simplicity, in the eye of the beholder', *Complexity*, **1**, (2), pp. 2–3.
16. Casti, J.L. (1998), 'Complexity and aesthetics', *Complexity*, **3**, (5), pp. 11–16.
17. Castroseza, C. (1978), 'Evolution, complexity, and fitness', *Journal of Theoretical Biology*, **71**, pp. 469–471.
18. Cohen, J. & I. Stewart (1994), *The Collapse of Chaos: Discovering Simplicity in a Complex World*, Penguin, New York.
19. Colwell, R.K. & J.A. Coddington (1994), 'Estimating terrestrial biodiversity through extrapolation', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 101–118.

20. Cornacchio, J.V. (1977*a*), 'Maximum-entropy complexity measures', *International Journal of General Systems*, **3**, pp. 215–225.
21. Cornacchio, J.V. (1977*b*), 'Systems complexity—a bibliography', *International Journal of General Systems*, **3**, pp. 267–271.
22. Cronquist, A. (1987), 'A botanical critique of cladism', *The Botanical Review*, **53**, pp. 1–52.
23. Crozier, R.H. (1992), 'Genetic diversity and the agony of choice', *Biological Conservation*, **61**, pp. 11–15.
24. Crutchfield, J.P & K. Young (1989), 'Inferring statistical complexity', *Physical Review Letters*, **63**, (2), pp. 105–108.
25. Csuti, B. *et al.* (1997), 'A comparison of reserve selection algorithms using data on terrestrial vertebrates in Oregon', *Biological Conservation*, **80**, pp. 83–97.
26. Daily, G.C. & P. Ehrlich (1992), 'Extinction and the biodiversity crisis', manuscript, quoted in Takacs (1996, p. 62).
27. Darwin, C. (1859), *On the Origin of Species by Means of Natural Selection*, John Murray, London.
28. Davies, P. (1987), *The Cosmic Blueprint*, Unwin, London.
29. Dawkins, R. (1976), *The Selfish Gene*, Oxford University Press, Oxford.
30. Dawkins, R. (1986), *The Blind Watchmaker*, Penguin, London.
31. Dawkins, R. (1999), 'The information challenge', *The Skeptic*, **18**, (4) pp. 21–25.
32. De Queiroz, K. (1988), 'Systematics and the Darwinian revolution', *Philosophy of Science*, **55**, pp. 238–259.
33. De Queiroz, K. & D.A. Good (1997), 'Phenetic clustering in biology: a critique', *The Quarterly Review of Biology*, **72**, (1), pp. 3–30.
34. Dennett, D.C. (1995), *Darwin's Dangerous Idea: Evolution and the Meaning of Life*, Penguin, London.
35. Donoghue, M.J. & P.D. Cantino, (1988) 'Paraphyly, ancestors, and the goals of taxonomy: a botanical defense of cladism', *The Botanical Review*, **54**, (2), pp. 107–128.
36. Drake, J.A. *et al.* (1996), 'Diversity and higher levels of organization', in K.J. Gaston (Ed.), *Biodiversity: A Biology of Numbers and Difference*, Blackwell Science, Oxford, pp. 149–166.
37. Duncan, T. & T.F. Stuessy (Eds.) (1984), *Cladistics: Perspectives on the Reconstruction of Evolutionary History*, Columbia University Press, NY.
38. Dupré, J. (1993), *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Harvard University Press, Cambridge MA.
39. Edgar, G.J. *et al.* (1997), 'The conservation-related benefits of a systematic marine biological sampling programme: the Tasmanian reef bioregionalisation as a case study', *Biological Conservation*, **79**, pp. 227–240.

40. Ehrlich, P.R. (1986), 'The loss of diversity: Causes and consequences', in E.O. Wilson (Ed.), *Biodiversity: National Forum on Biodiversity*, National Academy Press, Washington DC, pp. 21–27.
41. Ehrlich, P.R. & A.H. Ehrlich (1981), *Extinction: The Causes and Consequences of the Disappearance of Species*, Ballantine Books, New York NY.
42. Einarsson, N. (1993), 'All animals are equal but some are cetaceans: Conservation and culture conflict', in K. Milton (Ed.), *Environmentalism: The View from Anthropology*, Routledge, London.
43. Embley, T.M. *et al.*, 'Biodiversity at the molecular level: the domains, kingdoms and Phyla of life', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 21–33.
44. Erwin, T.L. (1982), 'Tropical forests: their richness in Coleoptera and other arthropod species', *Coleopterists Bulletin*, **36**, pp. 74–75.
45. Erwin, T.L. (1996), 'Biodiversity at its utmost: tropical forest beetles', in M.L. Reaka-Kudla *et al.* (Eds.), *Biodiversity II: Understanding and Protecting our Biological Resources*, Joseph Henry Press, Washington DC, p.1.
46. Fagan, W.F. & P.M. Kareiva (1997), 'Using compiled species lists to make biodiversity comparisons among regions: a test case using Oregon butterflies', *Biological Conservation*, **80**, pp. 249–259.
47. Faith, D.P. (1989), 'Homoplasy as pattern: multivariate analysis of morphological convergence in Anseriformes', *Cladistics*, **5**, pp. 235–258.
48. Faith, D.P. (1992), 'Conservation evaluation and phylogenetic diversity', *Biological Conservation*, **61**, pp. 1–10.
49. Faith, D.P. (1994a), 'Genetic diversity and taxonomic priorities for conservation', *Biological Conservation*, **68**, pp. 69–74.
50. Faith, D.P. (1994b), 'Phylogenetic pattern and the quantification of organismal biodiversity', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 45–58.
51. Fink, W.L. (1979), 'Optimal classifications', *Systematic Zoology*, **28**, pp. 371–374.
52. Freitag, S. *et al.* (1997), 'Ranking priority biodiversity areas: an iterative conservation value-based approach', *Biological Conservation*, **82**, pp. 263–272.
53. Gaines, B.R. (1977), 'System identification, approximation and complexity', *International Journal of General Systems*, **3**, pp. 145–174.
54. Gaston, K.J. (1996a), 'What is biodiversity', in K.J. Gaston (Ed.), *Biodiversity: A Biology of Numbers and Difference*, Blackwell Science, Oxford, pp. 1–9.
55. Gaston, K.J. (1996b), 'Species richness: measure and measurement', in K.J. Gaston (Ed.), *Biodiversity: A Biology of Numbers and Difference*, Blackwell Science, Oxford, pp. 77–113.
56. Gell-Mann, M. (1994), *The Quark and the Jaguar: Adventures in the Simple and the Complex*, Abacus, London.
57. Gell-Mann, M. (1995), 'What is complexity?', *Complexity*, **1**, (1), pp. 16–19.

58. Gell-Mann, M. & S. Lloyd (1996), 'Information measures, effective complexity, and total information', *Complexity*, **2**, (1), pp. 44–52.
59. George, L. (1977), 'Tests for system complexity', *International Journal of General Systems*, **3**, pp. 253–258.
60. Godfrey-Smith, P. (1996), *Complexity and the Function of Mind in Nature*, Cambridge University Press, Cambridge.
61. Goodwin, B.C. (1994), *How the Leopard Changed its Spots: The Evolution of Complexity*, Charles Scribner & Sons, NY.
62. Gould, S.J. (1989), *Wonderful Life: The Burgess Shale and the Nature of History*, Norton & Co., New York.
63. Gould, S.J. (1991), 'The disparity of the Burgess Shale arthropod fauna and the limits of cladistic analysis: why we must strive to quantify morphospace', *Paleobiology*, **17**, pp. 411–423.
64. Gould, S.J. (1992), 'Bligh's bounty' in *Bully for Brontosaurus*, Penguin, London, pp. 281–293.
65. Gould, S.J. (1993), 'How to analyze Burgess Shale disparity—A reply to Ridley', *Paleobiology*, **19**, pp. 522–523.
66. Gould, S.J. & R.C. Lewontin (1979), 'The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme', *Proc. Royal Society of London B*, **205**, London, pp. 581–598.
67. Grassberger, P. (1988), 'Complexity and forecasting in dynamical systems', in L. Peliti & A. Vulpiani (Eds.), *Measures of Complexity*, Springer-Verlag, Berlin, pp. 1–21.
68. Griffiths, P.E. (1997), *What Emotions Really Are: The Problem of Psychological Categories*, University of Chicago Press, Chicago.
69. Hacking, I. (1991a), 'A tradition of Natural Kinds', *Philosophical Studies*, **61**, pp. 109–126.
70. Hacking, I. (1991b), 'On Boyd' (A response to Boyd, 1991), *Philosophical Studies*, **61**, pp. 149–154.
71. Hammond, P.M. (1992), 'Species inventory', in B. Groombridge (Ed.), *Global Biodiversity: Status of the Earth's Living Resources*, Chapman and Hall, London.
72. Hammond, P.M. (1994), 'Practical approaches to the estimation of the extent of biodiversity in speciose groups', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 119–136.
73. Hampicke, U. (1994), 'Ethics and economics of conservation', *Biological Conservation*, **67**, pp. 219–231.
74. Harper, J.L. & D.L. Hawksworth (1994), 'Biodiversity: measurement and estimation', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 5–12.
75. Harvey, P.H. & M.D. Pagel (1991), *The Comparative Method in Evolutionary Biology*, Oxford University Press, Oxford.

76. Hinegardner, R. & J. Engleberg (1983), 'Biological complexity', *Journal of Theoretical Biology*, **104**, pp. 7–20.
77. Horgan, J. (1995), 'From complexity to perplexity', *Scientific American*, **272**, (6), pp. 104–109.
78. Hull, D.L. (1988), *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of a Science*, University of Chicago Press, Chicago.
79. Humphries, C.J. & J.A. Chappill (1988), 'Systematics as science: a response to Cronquist', *The Botanical Review*, **54**, (2), 1988, pp. 129–144.
80. Humphries, C.J. *et al.* (1995), 'Measuring biodiversity value for conservation', *Annual Reviews of Ecology and Systematics*, **26**, pp. 93–111.
81. Hutchinson, G.E. (1959), 'Homage to Santa Rosalia, or Why are there so many kinds of animals?', *American Naturalist*, **93**, (870), p. 146.
82. Iltis, H.H. (1986), 'Serendipity in the exploration of biodiversity: What good are weedy tomatoes?', in E.O. Wilson (Ed.), *Biodiversity: National Forum on Biodiversity*, National Academy Press, Washington DC, pp. 98–105.
83. James, E.A. & G.R. Ashburner (1997), 'Intraspecific variation in *Astelia australiana* (Lilaceae) and implications for the conservation of the Australian species', *Biological Conservation*, **82**, pp. 253–261.
84. Kamps, G. & V. Csányi (1987), 'Notes on order and complexity', *Journal of Theoretical Biology*, **124**, pp. 111–121.
85. Kauffman, S. (1995), *At Home in the Universe: The Search for Laws of Self-Organization and Complexity*, Oxford university Press, Oxford.
86. Keller, E.F. & E.A. Lloyd (Eds.) (1992), *Keywords in Evolutionary Biology*, Harvard University Press, Cambridge, MA.
87. Kennedy, L. (1994), 'Museum leaps to monitor biodiversity', *Geotimes*, **39**, (3), p. 5.
88. Kimura, M. (1983), *The Neutral Theory of Molecular Evolution*, Cambridge University Press, Cambridge.
89. Kingsland, S.E. (1985), *Modelling Nature: Episodes in the History of Population Ecology*, University of Chicago Press, Chicago.
90. Kokol, P. *et al.* (1996), 'Software complexity—An alternative view', *SIGPLAN*, **31**, (2), pp. 35–41.
91. Kolmogorov, A.N. (1965), 'Three approaches to the quantitative definition of information', *Problems of Information Transmission*, **1**, pp. 1–7.
92. Kunin, W.E. & J.H. Lawton (1996), 'Does biodiversity matter? Evaluating the case for conserving species', in K.J. Gaston (Ed.), *Biodiversity: A Biology of Numbers and Difference*, Blackwell, Oxford, pp. 283–308.
93. Kurths, J. & A. Witt (1994), 'On complexity measures', *World Futures*, **42**, pp. 177–192.
94. Langton, R. & D. Lewis (1998), 'Defining "intrinsic"', *Philosophy & Phenomenological Research*, **58**, pp. 333–345.

95. Lauder, G.V. *et al.* (1995), 'Systematics and the study of organismal form and function', *BioScience*, **45**, (10), pp. 696–704.
96. Leakey, R. & R. Lewin (1995), *The Sixth Extinction: Patterns of Life and the Future of Humankind*, Doubleday, New York.
97. Lehrman, D.S. (1970), 'Semantic and conceptual issues in the nature/nurture problem', in L.R. Aronson *et al.* (Eds.), *Development and Evolution in Biology*, Freeman, San Francisco, pp. 17–52.
98. Levin, S.A. (1996), 'Biodiversity: interfacing populations and ecosystems', in T. Abe *et al.* (Eds.), *Biodiversity: An Ecological Perspective*, Springer-Verlag, New York.
99. Lewin, R. (1992), *Complexity: Life at the Edge of Chaos*, Collier, New York.
100. Lewin, R. (1993), 'The right connections', *New Scientist*, **137**, (1859), pp. S4–5.
101. Lewis, D. (1986), *On the Plurality of Worlds*, Basil Blackwell, Oxford.
102. Li, W. (1997), 'The complexity of DNA', *Complexity*, **3**, (2), pp. 33–37.
103. Lickliter, R. & T. Berry (1990), 'The phylogeny fallacy', *Developmental Review*, **10**, pp. 348–364
104. Löfgren, L. (1977), 'Complexity of descriptions of systems: a foundational study', *International Journal of General Systems*, **3**, pp. 197–214.
105. Lovejoy, T.E. (1986), 'Species leave the ark', in B.G. Norton (Ed.), *The Preservation of Species: The Value of Biological Diversity*, Princeton University Press, Princeton NJ.
106. Lovejoy, T.E. (1994), 'The quantification of biodiversity: an esoteric quest or a vital component of sustainable development?', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 81–87.
107. Maclaurin, J. (1998*a*) 'Reinventing molecular Weismannism: information in evolution', *Biology and Philosophy*, **13**, (1), pp. 37–59.
108. Maclaurin, J. (1998*b*), *The Metaphysics of Biodiversity*, PhD thesis, Australian National University.
109. Maddox, J. (1990), 'Complicated measures of complexity', *Nature*, **344**, (6268), p. 705.
110. Mallet, J. (1995), 'A species definition for the Modern Synthesis', *TREE*, **10**, (7), pp. 294–299.
111. Mallet, J. (1996), 'The genetics of biological diversity: from varieties to species', in K.J. Gaston (Ed.), *Biodiversity: A Biology of Numbers and Difference*, Blackwell, Oxford, pp. 13–53.
112. Mann, C.C. & M.L. Plummer (1995), *Noah's Choice: The Future of Endangered Species*, Knopf, New York NY.
113. Martinez, N.D. (1996), 'Defining and measuring functional aspects of biodiversity', in K.J. Gaston (Ed.), *Biodiversity: A Biology of Numbers and Difference*, Blackwell Science, Oxford, pp. 114–148.
114. May, R.M. (1990), 'Taxonomy as destiny', *Nature*, **347**, pp. 129–130.

115. May, R.M. (1994), 'Conceptual aspects of the quantification of the extent of biological diversity', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 13–20.
116. Mayr, E. (1976), *Evolution and the Diversity of Life: Selected Essays*, Belknap Press, Cambridge, MA.
117. McQuillan, A.G. (1998), 'Passion and instrumentality: Further thoughts on the Callicott-Norton debate', *Environmental Ethics*, **20**, (3), pp. 317–324.
118. McIntyre, L. (1998), 'Complexity: a philosopher's reflections', *Complexity*, **3**, (6), pp. 26–32.
119. McShea, D.W. (1991), 'Complexity and evolution: what everybody knows', *Biology and Philosophy*, **6**, pp. 303–24.
120. McShea, D.W. (1992), 'A metric for the study of evolutionary trends in the complexity of serial structures', *Biological Journal of the Linnean Society*, **45**, pp. 39–55.
121. McShea, D.W. (1996a), 'Complexity and homoplasy', in M.J. Sanderson & L. Hufford (Eds.), *Homoplasy: The Recurrence of Similarity in Evolution*, Academic Press, San Diego.
122. McShea, D.W. (1996b), 'Metazoan complexity and evolution: is there a trend?', *Evolution*, **50**, (2), pp. 477–492.
123. Meyer, J.-A. (1988), 'Complexity in ecological systems', in L. Peliti & A. Vulpiani (Eds.), *Measures of Complexity*, Springer-Verlag, Berlin, pp. 139–145.
124. Mikkelsen, G.M. (1997), 'Methods and metaphors in community ecology: the problem of defining stability', *Perspectives on Science*, **5**, pp. 481–498.
125. Miller, D.R. & A.Y. Rossman (1995), 'Systematics, biodiversity, and agriculture', *BioScience*, **45**, (10), pp. 680–686.
126. Minelli, A. (1993), *Biological Systematics: The State of the Art*, Chapman and Hall, London.
127. Minter, B.A. & R.E. Manning (1999), 'Pragmatism in environmental ethics: Democracy, pluralism, and the management of nature', *Environmental Ethics*, **21**, (2), pp. 191–207.
128. Myers, N. (1991), 'Biological diversity and global stability', in F.H. Bormann & S.R. Kellert (Eds.), *Ecology, Economics, Ethics: The Broken Circle*, Yale University Press, New Haven.
129. Naeem, S. & S. Li (1997), 'Biodiversity enhances ecosystem reliability', *Nature*, **390**, pp. 507–509.
130. Nash, R. (1989), *The Rights of Nature: A History of Environmental Ethics*, University of Wisconsin Press.
131. Natural History Museum (1994), Biogeography and Conservation Laboratory World Wide Web site, 'Measuring biodiversity value', <http://www.nhm.ac.uk/science/projects/worldmap/diversity/index.html>.
132. Norton, B.G. (1991), *Toward Unity among Environmentalists*, Oxford University Press, New York.

133. O'Donnell, A.G. *et al.* (1994), 'Theoretical and practical aspects of the quantification of biodiversity among microorganisms', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 65–73.
134. O'Neil, R. (1997), 'Intrinsic value, moral standing, and species', *Environmental Ethics*, **19**, (1), pp. 45–52.
135. Oyama, S. (1985), *The Ontogeny of Information*, Cambridge University Press, Cambridge.
136. Papentin, F. (1980), 'On order and complexity I: general considerations', *Journal of Theoretical Biology*, **87**, pp. 421–456.
137. Parisi, G. (1988), 'On complexity', in L. Peliti & A. Vulpiani (Eds.), *Measures of Complexity*, Springer-Verlag, Berlin, pp. 22–27.
138. Pattee, H.H. (1977), 'Dynamic and linguistic modes of complex systems', *International Journal of General Systems*, **3**, pp. 259–266.
139. Patterson, C. (1977), 'The contribution of paleontology to Teleostean phylogeny', in M.K. Hecht *et al.* (Eds.), *Major Patterns in Vertebrate Evolution*, Plenum Press, New York.
140. Pearson, D.L. (1994), 'Selecting indicator taxa for the quantitative assessment of biodiversity', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 75–79.
141. Prance, G.T. (1994), 'A comparison of the efficacy of higher taxa and species numbers in the assessment of biodiversity in the neotropics', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 89–99.
142. Pressey, R.L. *et al.* (1997), 'Effectiveness of alternative heuristic algorithms for identifying indicative minimum requirements for conservation reserves', *Biological Conservation*, **80**, pp. 207–219.
143. Preston, C.J. (1998), 'Epistemology and intrinsic values: Norton and Callicott's critique of Rolston', *Environmental Ethics*, **20**, (4), pp. 409–428.
144. Raisbeck, G. (1963), *Information Theory: An Introduction for Scientists and Engineers*, MIT Press, Cambridge MA.
145. Randall, A. (1986), 'What mainstream economists have to say about the value of biodiversity', in E.O. Wilson (Ed.), *Biodiversity: National Forum on Biodiversity*, National Academy Press, Washington DC, pp. 217–223.
146. Rheinberger, H.-J. (1997), 'Experimental complexity in biology: some epistemological and historical remarks', *Philosophy of Science*, **64**, pp. S245–S254.
147. Richardson, R.C. (1997), 'Natural and artificial complexity', *Philosophy of Science*, **64**, pp. S255–S267.
148. Ridley, M. (1986), *Evolution and Classification: The Reformation of Cladism*, Longman Scientific and Technical, Harlow.
149. Ridley, M. (1993), 'Analysis of the Burgess Shale', *Paleobiology*, **19**, pp. 519–521.
150. Rosen, R. (1977), 'Complexity as a system property', *International Journal of General Systems*, **3**, pp. 227–232.

151. Routley, R. (1973), 'Is there a need for a new environmental ethic?', Proc. 15th World Congress on Philosophy, Sophia Press, Sophia, pp. 205–210.
152. Ruthen, R. (1993), 'Adapting to complexity', *Scientific American*, **268**, (1), pp. 130–137.
153. Sarkar, S. (1996), 'Biological information: a skeptical look at some central dogmas of molecular biology' in S. Sarkar (Ed.), *The Philosophy and History of Molecular Biology: New Perspectives*, Kluwer, Dordrecht, pp. 187–231.
154. Sarkar, S. (1999), *Cultures of Nature: Biological Conservation under Philosophical Scrutiny*, Draft manuscript, Department of Philosophy, University of Texas at Austin.
155. Saunders, P.T. & M.W. Ho (1976), 'On the increase in complexity in evolution', *Journal of Theoretical Biology*, **63**, pp. 375–384.
156. Savage, J.M. (1995), 'Systematics and the biodiversity crisis', *BioScience*, **45**, (10), pp. 673–679.
157. Schneider, E.D. & J.J. Kay (1995), 'Order from disorder: the thermodynamics of complexity in biology', in M.P. Murphy & L.A.J. O'Neill, *What is Life? The Next Fifty Years: Speculations on the Future of Biology*, Cambridge University Press, Cambridge.
158. Schrödinger, E. (1945), *What is Life? The Physical Aspect of the Living Cell*, MacMillan, New York.
159. Shannon, C.E. & W. Weaver (1949), *The Mathematical Theory of Communication*, University of Illinois, Urbana.
160. Simpson, B.B. & J. Cracraft (1995), 'Systematics: the science of biodiversity', *BioScience*, **45**, (10), pp. 670–672.
161. Sober, E. (1988), 'The conceptual relationship of cladistic phylogenetics and vicariance biogeography', *Systematic Zoology*, **37**, (3), pp. 245–253.
162. Sober, E. (Ed.) (1994), *Conceptual Issues in Evolutionary Biology*, MIT Press, Cambridge, MA.
163. Sogin, M.L. & G. Hinkle (1996), 'Common measures for studies of biodiversity: molecular phylogeny in the eukaryotic microbial world', in M.L. Reaka-Kudla *et al.* (Eds.), *Biodiversity II: Understanding and Protecting our Biological Resources*, Joseph Henry Press, Washington DC, pp. 109–122.
164. Steneck, R.S. & M.N. Dethier (1994), 'A functional group approach to the structure of algal-dominated communities', *Oikos*, **69**, pp. 476–498.
165. Sterelny, K. & P.E. Griffiths (1999), *Sex and Death: An Introduction to the Philosophy of Biology*, University of Chicago Press, Chicago.
166. Stewart, I. (1993), 'A new order', *New Scientist*, **137**, (1859), pp. S2–3.
167. Stoltz, K. (1999), *Interactive Construction: Organizational Constraints on Evolution and Cognition*, PhD thesis, RUG–Faculteit Letteren en Wijsbegeerte.
168. Stork, N.E. (1996), 'Measuring global biodiversity and its decline', in M.L. Reaka-Kudla *et al.* (Eds.), *Biodiversity II: Understanding and Protecting our Biological Resources*, Joseph Henry Press, Washington DC, pp. 41–68.

169. Takacs, D. (1996), *The Idea of Biodiversity: Philosophies of Paradise*, John Hopkins University Press, Baltimore & London.
170. Templeton, A.R. (1994), 'Biodiversity at the molecular genetic level: experiences from disparate macroorganisms', *Philosophical Transactions of the Royal Society of London B*, **345**, pp. 59–64.
171. Tsonis, P.A. & A.A. Tsonis (1997), 'Simplicity and complexity in gene evolution', *Complexity*, **2**, (5), pp. 23–30
172. Vane-Wright, R.I. *et al.* (1991), 'What to protect?—Systematics and the agony of choice', *Biological Conservation*, **55**, pp. 235–254.
173. Waldrop, M.M. (1992), *Complexity: The Emerging Science at the Edge of Order and Chaos*, Simon & Schuster, New York.
174. Wicken, J.S. (1987), *Evolution, Thermodynamics, and Information: Extending the Darwinian Paradigm*, Oxford university Press, Oxford.
175. Wilcox, B.A. (1984), 'In situ conservation of genetic resources: determinants of minimum area requirements', in J.A. McNeeley & K.R. Miller (Eds.), *National Parks, Conservation and Development: The Role of Protected Areas in Sustaining Society*, Proc. World Congress on National Parks, Bali, Indonesia, 11–22 October 1982, Smithsonian Institution Press, Washington DC.
176. Williams, P.H. *et al.* (1994), 'Do conservationists and molecular biologists value differences between organisms in the same way?', *Biodiversity Letters*, **2**, pp. 67–78.
177. Williams, P.H. & K.J. Gaston (1994), 'Measuring more of biodiversity: can higher-taxon richness predict wholesale species richness?', *Biological Conservation*, **67**, pp. 211–217.
178. Williams, P.H. & C.J. Humphries (1996), 'Comparing character diversity among biotas', in K.J. Gaston (Ed.), *Biodiversity: A Biology of Numbers and Difference*, Blackwell, Oxford, pp. 54–76.
179. Wilson, E.O. (1984), *Biophilia*, Harvard University Press, Cambridge MA.
180. Wilson, E.O. (1986), 'The current state of biological diversity', in E.O. Wilson (Ed.), *Biodiversity: National Forum on Biodiversity*, National Academy Press, Washington DC, pp. 3–18.
181. Wilson, E.O. (1992), *The Diversity of Life*, Penguin, London.
182. Wilson, E.O. (1993), 'Building an ethic', *Defenders*, Spring.
183. Wilson, E.O. (1997), introduction to M.L. Reaka-Kudla *et al.* (Eds.), *Biodiversity II: Understanding and Protecting our Biological Resources*, Joseph Henry Press, Washington DC, p. 1.
184. Yagil, G. (1985), 'On the structural complexity of simple biosystems', *Journal of Theoretical Biology*, **112**, pp. 1–23.