

11 Metaphysical and epistemological issues in modern Darwinian theory

I A TWO-PART THEORY

Like Darwin's own theory of evolution, the modern Darwinian theory of evolution has two main elements:

The Tree of Life: All organisms now alive on earth trace back to a common ancestor.

Natural Selection: Natural selection has been an important cause of the similarities and differences that exist in the earth's biota.

The first of these propositions says that any two contemporary organisms have a common ancestor. Human beings are genealogically related to each other, but each human being also has a common ancestor with chimps, dogs, clams, daffodils, bacteria and yeast.¹ The second proposition, as I have formulated it, does not say that natural selection is the *only* cause of evolution. Indeed, it should be understood to leave open the possibility that there are traits for which natural selection is entirely irrelevant. This is the big picture, and evolutionary biology is devoted to filling in the details.

Although Darwinism is easy to describe, this simple theory gives rise to a rich range of metaphysical and epistemological questions. It is the purpose of this chapter to discuss some of them. In conformity with the structure of Darwinian theory, I have chosen one metaphysical and one epistemological problem from each of the two big ideas. I begin with a problem in the metaphysics of natural selection – the role of chance – followed by a problem in the metaphysics of the tree of life – the nature of a biological species. Turning from metaphysics to epistemology, the later sections of the chapter examine the testing

of hypotheses about genealogical relatedness (the tree of life) and the testing of adaptive hypotheses (natural selection).

II THE LOGICAL CHARACTER OF DARWINIAN THEORY

Before moving on to these four topics, it is useful to contemplate the logical character of the two propositions that comprise the Darwinian theory. Each is a historical claim, not a law of nature. Laws of nature are conventionally understood to be empirical generalisations that do not refer to any place, time or individual. In addition, they cannot be true accidentally; they are supposed to possess a kind of necessity (*nomological*, not *logical*). In contrast, the two propositions we are considering are expressed in singular statements about the organisms that happen to exist on earth.

In the days when philosophy of science was dominated by philosophy of physics, this feature of Darwinian theory was a matter of concern, if not embarrassment. With Newtonian mechanics, relativity theory and quantum mechanics as their paradigms of what a scientific theory should be like, the logical empiricists often equated science with the search for law. Since the Darwinian propositions are not laws, in what sense do they constitute a scientific theory at all? Now, in these post-positivist times, the impulse to make biology fit this physical ideal is less compelling. It now seems natural to recognise that sciences are of two types – *nomothetic* and *historical*. Nomothetic sciences aim at the discovery of laws; they use historical information about specific objects as a means to that end. Historical sciences aim to discover facts about the histories of specific objects; they use information about laws as a means to that end.²

This broader picture of what counts as science allows us to recognise that physics contains disciplines of both types, and so does evolutionary biology. The physical theories cited above belong to nomothetic disciplines. But physicists are also interested in the histories of stars and galaxies; as such, astronomy is an historical science. Indeed, the division of nomothetic from historical sciences need not be strict. Astronomers are interested in the histories of specific stars and also try to describe the laws that govern the development of stars. In the same way, biologists seek to understand the evolution of specific groups of organisms and also to describe the

laws that govern evolutionary change. A student of the social insects might also develop general models of sex ratio evolution.

Although theoreticians in evolutionary biology seek to formulate generalisations that are not true simply by accident (as is the case for many statements about 'evolutionary trends' – for example, that size increase has been more common than size reduction in the earth's evolving lineages), there is a feature of these generalisations that fails to conform to the logical empiricist concept of law. Whereas the logical empiricists held laws to be empirical rather than mathematical, models in evolutionary biology are 'if... then' statements that are mathematical truths. Consider, for example, elementary models in evolutionary genetics. They assign fitnesses to the various genotypes in a population, and assert that *if* those fitnesses have these values, *then* the population will evolve to certain future states. The 'if... then' statement that summarises such models is true *a priori*. No observations are needed to see that it is true; checking the algebra suffices. Of course, it is an empirical matter whether this or that natural population satisfies the conditional's antecedent. However, this empirical question concerns a *singular* statement – that *this population* exhibits certain properties.³

III CHANCE

The concept of chance features in evolutionary theory in two contexts. First, the variation on which natural selection operates is said to arise 'by chance'. Second, probabilities appear twice over in the characterisation of a selection process – the concept of fitness is defined probabilistically and finiteness of population size introduces a stochastic element into evolutionary trajectories.

Beginning with the first of these usages, we can discern one of its meanings in a remark of Darwin's: 'I have hitherto sometimes spoken as if the variations... had been due to chance. This, of course, is a wholly incorrect expression, but it serves to acknowledge plainly our ignorance of the cause of each particular variation.'⁴ Here Darwin echoes the French astronomer Pierre-Simon Laplace, who claimed that a demon with complete knowledge of the relevant laws and initial conditions, and who had limitless computational powers, would never need to talk about what would *probably* occur. Rather, for such

a being, 'nothing would be uncertain, and the future, and the past, would be present to its eyes'.⁵

A second meaning that modern biologists attach to the idea that variation arises by chance came to the fore only after Darwin's time. This is the doctrine, due to the German biologist August Weismann, that beneficial variations do not arise *because* they would be beneficial.⁶ This doctrine amounts to a rejection of the Lamarckian idea that there is inheritance of acquired characteristics. Applied to the distinction between genotype and phenotype, Lamarckism requires that a phenotype acquired by parents should change the genes that parents transmit to their offspring. Whereas the blacksmith gets big muscles because and only because he works at the forge, his son develops big muscles whether he exercises them or not – an acquired character is transformed into one that is 'innate'. When modern biologists say that mutations occur 'by chance', one thing they mean is that this Lamarckian causal pathway does not exist.

I now turn to the question of whether there is a 'chance element' in the process of natural selection itself. Modern biologists define natural selection in terms of the concept of fitness – a selection process occurs precisely when there is variation in fitness. An organism's fitness is its 'ability' to survive and reproduce. This ability is represented probabilistically, in terms of a fertilised egg's probability of reaching adulthood and the adult organism's expected number of offspring.

We may begin with the point that fitness is a theoretically interesting property because it is a property of *traits*. It is traits that evolve through multi-generational selection processes, whereas individual organisms are here today and gone tomorrow. Biologists care about the fitness of dorsal fins, not about the fitness of individual tunas. That said, evolutionary theory does not *reify* traits; the fitness of a trait does not float free of the fitnesses of the individuals that have the trait. The two are linked by a simple formula – the fitness of a trait is just the average fitness of the individual organisms that possess the trait.⁷

Let us apply this framework to a concrete example. Suppose that running speed is evolving in a population of zebras. Some zebras run fast while others run slowly. If the frequencies of these traits are changing because there is natural selection, the two traits *running*

fast and *running slowly* must differ in fitness. This means that fast zebras, on average, have a different fitness value from slow ones. Let us suppose that this is because fast zebras, on average, are better able to avoid being killed by predators.

Fast zebras differ among themselves in countless ways, so it is a mistake to think that there is a single fitness value that they have in common. Perhaps fast zebras have a thousand different probabilities of surviving to adulthood. Or maybe the lifetime of each zebra is a deterministic process wherein the organism is fated to die before reaching adulthood, or fated not to do so. This choice does not matter, because whether we average a thousand different probabilities, or average a thousand different 1's and 0's, the upshot is the same – we represent the fitness of the trait *running fast* as being between 0 and 1.

Once fitness values are assigned to the two traits, the fundamental question concerning what natural selection can be expected to produce depends on a simple comparative question – which trait is fitter? The absolute values of the fitnesses do not matter. If *running fast* is fitter than *running slowly*, then it is more probable than not that *running fast* will increase in frequency (assuming that the traits are heritable). But *how* probable is this outcome?

This is where the size of the population becomes relevant. The larger the population, the more certain it is that the fitter trait will increase in frequency. Consider an analogy – two coins that differ in their biases. The first has a probability of landing heads when tossed of 0.8, while the second has a probability of landing heads of 0.6. If I toss each coin a number of times, I expect the first to land heads more often than the second. However, the strength of this expectation depends on how many times the coins are tossed. If each is tossed twice, there is a considerable probability that the first coin will *not* yield the larger number of heads. But if I toss the coins a thousand times, this probability shrinks. The Law of Large Numbers says that as sample size increases, the probability increases that the frequency of heads produced by a coin will be close to its probability of landing heads. In the limit, the probability approaches unity (that is, certainty) that the first coin will land heads $80\% \pm \varepsilon$ of the time and the second will land heads with a frequency of $60\% \pm \varepsilon$, no matter how small ε is.

In coin tossing, small sample size gives 'chance' an enhanced opportunity to show itself. In evolution, it is small population size that

has this effect. This is the idea that Motoo Kimura exploited in his 'neutral theory of molecular evolution'.⁸ If traits differ only a little in fitness, and if population size is small enough, then traits will evolve by random walk. Modern Darwinians either reject the neutral theory or restrict their Darwinism to changes at higher levels of organisation; random walk is not evolution by natural selection.

I hope this brief discussion gives the reader a feeling for the fact that modern evolutionary theory is saturated with probability concepts. Probabilities are used to describe mutations, they are used to characterise the fitness values of traits, and they are used in models that allow one to calculate the outcomes of specified initial conditions. Some of these models are said to be 'deterministic'; they apply only to populations that are infinitely large. Such models may be suitable idealisations when the finite populations under study are big enough, but these deterministic models are a special case. The body of theory, taken as a whole, is probabilistic to its core.

What do these probability concepts mean? To begin with, they do not entail that determinism is false. This is not a problem on which biology has any purchase. When a biological model assigns a probability to a given event, there may be factors influencing the process leading up to that event that are not acknowledged in the biological model. These *hidden variables* may be biological in character (and so a more complex biological model can be constructed to capture them), or they may involve events that cannot be described in biological language. Either way, the theory is said to be causally incomplete. It is at this point that physics may have to take over – the buck has been passed. It is interesting that the buck never gets passed in the opposite direction – when physicists think that a physical model is incomplete, they do not turn to biologists for help. This asymmetry arises because there is no reason to think that biology is causally complete, but the idea that physics is causally complete is taken very seriously indeed.⁹

In order to investigate how the probability concepts used in evolutionary biology should be interpreted, let us assume that determinism is true. The Laplacean interpretation (with which the quotation from Darwin agrees) is that probability concepts must therefore be placeholders for ignorance; either determinism is false or probabilities must describe subjective degrees of belief.

There is a third possibility. Consider, first, the fact that the mathematical formalism of the probability axioms can be interpreted in terms of *actual frequencies*. Under this interpretation, 'the probability is $1/2$ that the next toss of this coin will land heads' means that the coin's history of tosses (past, present and future) yields 50 per cent heads. I do not claim that this interpretation does justice to much of what we want to say in probability language – after all, a fair coin can be tossed an odd number of times – but it does bring out the possibility that probability statements can describe objective features of the world even if determinism is true.

The question of whether probability statements can be objectively true in a deterministic world needs to be separated from the pragmatic question of which statements we should use in making a prediction. If we toss a coin and determinism is true, full information will allow us to predict with certainty whether the coin will land heads. If we had this full information, we would not use the fact that the coin landed heads half the time in past tosses to infer that the probability of heads on the next toss is $1/2$. However, this is a pragmatic point, not a semantic one. The fact that we would not *use* the probability statement to make our prediction does not mean that it is not objectively *true*.

Scientists introduce probability models to describe repeatable processes that exhibit different outcomes with different frequencies. The probability of an outcome is not the same as the observed frequency, but rather is a theoretical quantity introduced to explain and predict that observed frequency. Like all theories, probabilistic theories are inferentially connected to observations. Values for probabilities are estimated from observed frequencies, and postulated probabilities make predictions about which observations will (probably) occur. When we ask whether nonprobabilistic theoretical postulates are objectively true, all we can do is point to the confirmation that those theories have received. This is why we are entitled to think that electrons objectively exist – they are not figments of our imagination. Precisely the same standard should be applied to the question of whether various probability concepts are objective. We know that uranium has a given half-life; this is an objective feature of that substance. The same holds true of the mutation probabilities and fitness values discussed in evolutionary biology.¹⁰

This point about the interpretation of probability concepts in a deterministic world has implications for how probabilities should be understood if determinism is false. Suppose that a complete physical theory were to assign a probability of x to a given event (where $0 < x < 1$). What does this imply about the probability that some other (perhaps biological) theory must assign to that event? Must the other theory also assign a value of x , on pain of being dismissed as 'merely subjective' (or just outright false)? The answer is *no*. The physical theory assigns a value of x by conditionalising on a set of (true) propositions P . A different theory can conditionalise on a different set of (true) propositions Q , and thereby assign a value of y . The probability statements do not conflict, since they conditionalise on different propositions. Laplace assumed that determinism is true, and concluded that all probability statements are mere confessions of ignorance. But the deeper position he defended goes beyond the assumption of determinism – this is the idea that the only objective probabilities are the ones provided by a theory that is causally complete. Here is a reductionist thesis that we should reject.¹¹

IV ESSENTIALISM AND THE SPECIES CONCEPT

Species have long been a favourite example that philosophers cite when they discuss *natural kinds*.¹² For example, John Stuart Mill claims that *human being* is a natural kind, but the class of *snub-nosed individuals* is not, on the grounds that 'Socrates is a human being' allows one to predict many of the other characteristics that Socrates has, but 'Socrates is snub-nosed' does not.¹³ Aristotelian essentialism endows the concept of natural kind with a more burdensome characterisation. Natural kinds not only have predictive richness; in addition, they have *essences*. The essence of a natural kind is the necessary and sufficient condition that all and only the members of the kind satisfy. Indeed, it is a necessary truth that the members of the kind, and they alone, have this essential property. Further, the essence is explanatory; the fact that an individual has this species-typical essence explains many other features that the individual possesses.

Besides citing biological species as examples, philosophers often point to the chemical elements as paradigm natural kinds. Gold is a kind of substance; its essence is said to be the atomic number 79.

This atomic number is what makes a lump of matter an instance of gold. And atomic number explains many other properties that gold things have. According to Saul Kripke and Hilary Putnam, science is in the business of empirically discovering the essences of natural kinds.¹⁴ Formulated in this way, essentialism is not established by the existence of trivial necessary truths. It is a necessary truth that all human beings are human beings, but this does not entail that there is an essence that human beings have. It also is important to separate the claim that kinds have essences from the claim that individuals in the kind have essential properties.¹⁵ The fact that different elements have different atomic numbers leaves open the possibility that an individual may persist through time as it changes from being made of one element to being made of another. Essentialism does not rule out the possibility of radioactive decay (nor of more mundane replacements, as in the constantly rebuilt ship of Theseus).

The example of the chemical elements illustrates a further feature that kind essences are supposed to have. Notice that 'atomic number 79' does not refer to any place, time or individual. What makes two things members of the same natural kind is that they are similar in the requisite respect. There is no requirement that they be causally connected to each other. The essence is intrinsic, not relational.

Although philosophers who accept this essentialist picture of the chemical elements usually think that chemistry has already discovered the essences that various chemical kinds possess, they must concede that biology so far has not done this for biological species. Is this simply because biology's work is not yet done? No – there are strong reasons to think that Darwinian theory undermines this essentialist picture of biological species.¹⁶ Species are not natural kinds, at least not on the usual essentialist construal of what a natural kind is.

The reasons for this conclusion need to be stated carefully. The fact that species evolve is, *per se*, not a conclusive argument against essentialism. Just as essentialists can agree that chemical elements undergo transmutation, so essentialists can agree that lineages undergo evolution, with ancestor and descendant belonging to different species. And the fact that there are vague boundaries between species is not, in itself, a refutation of essentialism, either. When an atom of uranium-235 gives rise to atoms of bromine and

lanthanum, there may be intermediate stages of the process in which it is indeterminate what the natural kinds are to which the matter belongs.¹⁷

Unfortunately, there still is disagreement in evolutionary biology about how the species category should be understood. The most popular definition is Ernst Mayr's¹⁸ *biological species concept*.¹⁹ Its anti-essentialist consequences are to a large extent also the consequences that other species concepts have, so we may examine it as an illustrative example. Mayr's idea is that a biological species is an ensemble of local populations that are knit together by gene flow. The individuals within local populations reproduce with each other. And migration among local populations means that there is reproduction between individuals in different populations as well. This system of populations is reproductively isolated from other such systems. Reproductive isolation can be a simple consequence of geographical barriers, or it can mean that the organisms have behavioural or physiological features that prevent them from producing viable fertile offspring even when they are brought together. Reproductive isolation allows two species to evolve different characteristics in response to the selection pressures imposed by their different environments. However, the different phenotypes that evolve are not what make the two species two; it is reproductive isolation, not physical dissimilarity, that is definitive.

Mayr initially allowed two populations to belong to the same species if there is *actual or potential* interbreeding between them, but he later changed the definition so that *actual* interbreeding is required. This raises the question of what the time scale is on which interbreeding must take place. How often must individuals in different local populations reproduce with each other for the two populations to belong to the same species? Indeed, the same question can be posed about individuals living in the same local population. Another detail that needs to be addressed concerns individuals that exist at different times. Human beings who are alive now are not having babies with human beings who lived thousands of years ago. What makes past and present human beings members of the same species? One necessary condition is that human beings now and human beings then are genealogically related. But this is clearly not sufficient; otherwise, we could not make the distinctions we do between a present-day species and a distinct ancestral species. Finally, I should note that

Mayr's definition excludes the possibility of asexual species; this is another feature that has made it controversial.

The important point about Mayr's definition is that similarity is neither necessary nor sufficient for conspecificity. Members of the same species may have very different characteristics. And if creatures just like tigers evolved independently in another galaxy, they would not belong to the species to which earthly tigers belong. Conspecificity is defined by the causal-historical connections that arise from reproductive interactions. Biological species and chemical elements are very different in this regard.

Evolutionary biologists talk about species in the same way they discuss individual organisms. Just as individual organisms bear genealogical relationships to each other, so species are genealogically related. Just as organisms are born, develop and die, so species come into existence, evolve and go extinct. These considerations led Michael Ghiselin and David Hull to maintain that *species are individuals, not natural kinds*.²⁰ There is room to doubt, however, that species are as functionally integrated as individual organisms often are. The parts of a tiger depend on each other for survival; excise an arbitrary 30 per cent of a tiger, and the tiger dies. However, the extinction of 30 per cent of a species rarely causes the rest of the species to go extinct. This suggests that individuality (in the sense of functional interdependence of parts) comes in degrees, and that species are often less individualistic than organisms often are. Still, Hull and Ghiselin's main thesis remains; perhaps it should be stated by saying that species are *historical entities*.²¹

Similar points apply to broader classificatory groupings, that is, higher taxa. Although ordinary language may suggest that carnivores all eat meat, this is not how biologists understand *Carnivora*. Taxa are understood genealogically; they are *monophyletic groups*, meaning that they are composed of an ancestral species and all of its descendants. Pandas belong to *Carnivora* because they are descended from other species in *Carnivora*; the fact that pandas are vegetarians does not matter. Superspecific taxa, like species themselves, are conceptualised as big physical objects; they are chunks of the genealogical nexus. And just as species are often not very individualistic, superspecific taxa are even less so.²²

The chemical kinds do not comprise an *ad hoc* list. Rather, there is a theory, codified in the periodic table of elements, that tells us how

to enumerate these chemical kinds and how they are systematically related to each other. To say what the chemical kinds are, we can simply consult this theory; we do not, in addition, have to do fieldwork. No such theory exists in biology for species and higher taxa; fieldwork is the only method that biology has for assembling its list of taxa. The terms 'botanising' and 'beetle collecting' both allude to this feature of systematic biology. Species and higher taxa are things that happen to come into existence owing to the vagaries of what transpires in the branching tree of life.

It does not follow that there are no natural kinds in evolutionary biology. Perhaps sexual reproduction is a kind; perhaps being a predator is another.²³ What makes it true that two organisms each reproduce sexually, or that both are predators, is that they are similar in some respect; it is not required that they be historically connected to each other. The sexual species do not form a monophyletic group, and neither do the predators. These kind terms appear in models of different evolutionary processes; there are models that explain why sex might evolve and models that describe the dynamics of predator/prey interactions. Although Darwin's theory of evolution undermines essentialist interpretations of species and higher taxa, it is another matter whether essentialism is the right way to understand these other, nontaxonomic, theoretical categories.

V TESTING HYPOTHESES ABOUT COMMON ANCESTRY

Although a great deal of work in evolutionary biology is devoted to inferring phylogenetic relationships, almost none of it provides a test of the tree-of-life hypothesis. When biologists attempt to reconstruct the phylogenetic relationships that link a set of species, they usually *assume* that all the taxa under study are genealogically related. Whatever method one uses – cladistic parsimony, distance measures, or maximum likelihood methods – the typical question is *which* tree is the best one, not *whether* there is a tree in the first place.²⁴

This is not to say that biologists have totally ignored the issue of how the tree-of-life hypothesis might be tested. For example, Francis Crick, co-discoverer of the double helical structure of the genetic molecule DNA, argued that the genetic code is a 'frozen accident', meaning that the pattern by which nucleotide triplets code for amino

acids is functionally arbitrary.²⁵ Given Crick's thesis, the (near) universality of the genetic code among the earth's organisms provides strong evidence that all trace back to a common ancestor. If the tree-of-life hypothesis were true, we would expect the code to be universal; however, if lineages arose separately, we would not expect them to exhibit the same code. This argument is often repeated as if it constitutes a conclusive case for the tree-of-life hypothesis, but, in fact, the claim that all codes are equally fit raises subtle and ongoing questions. For example, it has been argued that the code we now observe is optimal.²⁶ If this turns out to be correct, the argument for the tree-of-life hypothesis that appeals to the universality of the genetic code must be rethought. If there is a selective advantage to the code we observe, the question of whether the tree-of-life hypothesis or the hypothesis of multiple start-ups is better supported will depend on quantitative considerations – how much of a selective advantage did the code we observe possess, how much time would there have been for selection to make over a lineage that initially exhibited an alternative, and how deep and wide is the 'valley' that separates a code on one adaptive peak from a code on another?

I now turn to the methods that biologists use to evaluate competing phylogenetic hypotheses that all assert that the taxa under study have a common ancestor. Parsimony is the method most often used. The basic idea can be understood by considering the two trees depicted in Figure 11.1. Suppose we observe that Sparrows and Robins both have wings, but that Crocodiles do not. On the assumption that winglessness is the ancestral condition (that it is the character state




Figure 11.1. Parsimony favours the (SR) C tree on the left.




Figure 11.2. Parsimony does not discriminate between these trees.

of the common ancestor at the root of the tree),²⁷ the (SR)C tree can explain the observations by postulating a single change in character state (from no wings to wings) in the tree’s interior; the S(RC) tree, on the other hand, requires two such changes. Thus, the (SR)C tree provides the more parsimonious explanation of the data.

Now consider the problem represented in figure 11.2, which also involves evaluating two trees. We observe that Lizards and Crocodiles lack wings, but that Sparrows have them. If winglessness is the ancestral condition, the (LC)S tree and the L(CS) tree each can explain the observations by postulating a single change in character state. If parsimony is our guide, we will conclude that this character distribution does not discriminate between the two phylogenetic hypotheses.

As this pair of examples illustrates, parsimony treats some similarities, but not others, as evidence of common ancestry. Notice that the similarity considered in figure 11.1 is derived, while that in figure 11.2 is ancestral. The principle of cladistic parsimony regards only the former as evidentially significant. Parsimony therefore is a different methodology from that of phenetic clustering, which counts *all* similarities (ancestral as well as derived) as evidence of relatedness.

In addition to parsimony and phenetic clustering, there is a third approach to phylogenetic inference, which is explicitly statistical. The maximum likelihood approach is to find the phylogenetic tree that maximises the probability of the observations.²⁸ In terms of the problem depicted in figure 11.1, the question will be whether the (SR)C hypothesis makes the observations more probable than the S(RC) hypothesis does. This question cannot be answered until a probabilistic model of character evolution is provided.

Unfortunately, biologists who do not already know the genealogy of a set of taxa will often also be in the dark as to the rules of character evolution that those taxa obeyed. And if one simply *assumes* that a given process model is correct, maximum-likelihood inference of phylogenetic trees can lead one seriously astray, if that model is mistaken.

Although biologists usually use parsimony and phenetic clustering without stating an explicit process model, this does not mean that these methods perform well regardless of how the evolutionary process proceeds. For example, a tree of the form (XY)Z can follow rules of character evolution that lead X and Z to exhibit far more similarities (both ancestral and derived) than X and Y. When this happens, parsimony and phenetic clustering will both mislead; each will converge on the wrong tree as more and more data are gathered.²⁹

The problem of evaluating competing methods of phylogenetic inference is an active area of current investigation. The debate is by no means over. One central line of enquiry is the investigation of what parsimony and phenetic clustering presuppose about the evolutionary process. Another is the development of more realistic process models that can be used in maximum-likelihood inference.

VI TESTING ADAPTIVE HYPOTHESES

How can hypotheses about the effects of natural selection be tested? If you catch natural selection in the act, you can observe the process of replacement unfold, and empirically determine whether the trait that is increasing in frequency allows its bearers to survive better and reproduce.³⁰ If the zebras in the herd you are observing differ in running speed, you can check whether fast zebras are killed by lions less often than slow ones. But suppose you come on the scene too late; the variation has disappeared, and so you cannot directly compare the fitness values of different traits. If all the zebras in the population you observe run fast, how are you to test the hypothesis that fast zebras replaced slow ones, and that this happened because slow zebras were more vulnerable to lion attack?

In fact, comparison still is possible, but you must conceive it on a wider scale. Rather than compare one zebra with another, you should compare one population to another. If running speed is an adaptive response of prey organisms to predator attack, then you should find

that the running speeds of prey species differ in the same direction as the running speeds of their predators. If species A preys on species X, and species B preys on species Y, then if A runs faster than B, X should run faster than Y.³¹ This is a modest deployment of what biologists call 'the comparative method'.³² The comparison, of course, is *across* species, not within them.

Suppose the running speeds of A and X are 35 and 33 miles per hour respectively, and that the running speeds of B and Y are 22 and 19 miles per hour. This is evidence that running speed in predators and running speed in prey are not independent. It does not tell you whether predators evolved to catch their prey, or prey evolved to evade their predators, or both. Still less do these data tell you that the running speeds of the four species are *optimal*. After all, your verdict would have been the same if you had observed that the four running speeds are 50, 45, 10 and 7. The observations you made do not settle whether the observed running speeds are the best ones that the different species could deploy.

The attentive reader will have detected a change in subject in the preceding three paragraphs. I began by asking whether zebras run fast in order to avoid lions. I then shifted to the more general question of whether prey species run fast in order to avoid predators. These questions are not the same, and it is conceivable that the zebra–lion relationship differs from the relationship that obtains between most prey organisms and their predators. Though the questions are different, the shift is forced on us if all zebras run fast and we want to test adaptive hypotheses empirically. Adaptive hypotheses assert that natural selection played a specified *causal role*. And causal claims assert that one factor *makes a difference* in the expression of another. For example, the hypothesis that smoking causes lung cancer predicts that smokers should get cancer more often than nonsmokers, once one has controlled for other causal influences. If everyone smokes, the hypothesis cannot be tested.

Although these remarks may sound humdrum, they in fact have implications about a controversy that has stirred strong feelings in biology. This is the debate about *adaptationism*. Stephen Jay Gould and Richard Lewontin criticised biologists for uncritically espousing 'just-so' stories about natural selection.³³ They even went so far as to claim that adaptationism is unfalsifiable; since the defeat of one adaptive hypothesis allows you to invent another, there is no way to

refute adaptationism as a claim about nature. Gould and Lewontin also defended a 'pluralistic' view of the evolutionary process, according to which natural selection is one, but not the only, important influence on trait evolution. John Maynard Smith responded by defending adaptationism – although he conceded that observations never test the claim that a trait is an adaptation.³⁴ This appeared to confirm the worst fears that critics of adaptationism harboured: adaptationism seems to be an undefended and indefensible assumption.

One clarification that this debate sorely needs concerns what adaptationism asserts as a claim about nature. Here it is useful to distinguish the following two propositions:

- (I) Natural selection has been an important cause of the evolution of most phenotypic traits in most populations.
- (O) Natural selection has been the only important cause of the evolution of most phenotypic traits in most population.

Gould and Lewontin say that they agree with Darwin that (I) is true. What they deny is that nonselective influences on trait evolution can be ignored. In this light, it does no good to point out that natural selection is the only resource that evolutionary theory has for explaining complex features like the vertebrate eye.³⁵ This is not at issue. The question is whether the features of the eye are optimal – whether natural selection has sifted through a rich array of variation and provided organisms with the best of the available phenotypes. The debate concerns the *hegemony* of natural selection, not whether selection is *important*; (O) is the heart of the matter.

To understand the debate about adaptationism, it is important to distinguish methodological claims from claims about nature. Gould and Lewontin advanced both – they criticised the inferential practices of their colleagues, and they advanced a pluralistic conception of how traits evolve. These points are separate. Their critics sometimes responded by claiming that the concept of adaptation is an indispensable tool in investigating nature.³⁶ The point is correct and important; both adaptationists and anti-adaptationists need optimality models if they are to determine empirically the degree to which an organism's traits are optimal. However, this observation does not establish that adaptationism is correct as a claim about nature, nor does it show that adaptationists have tested their hypotheses with sufficient rigour.

At the same time, it needs to be recognised that Gould and Lewontin overstated their contention that adaptationism is untestable. They are right that if one adaptationist explanation of a trait is refuted by observations, another can be constructed. However, the same can be said of a pluralistic model. Adaptationism and pluralism are both *isms*. Each describes the *kind* of explanation that most traits have without saying anything very specific about why any given trait evolved. It is *specific* optimality models and *specific* pluralistic models that, in the first instance, can be brought into contact with data. This does not mean that the *isms* are untestable, but just that they can be evaluated only in the long run.³⁷ Each embodies a large-scale generalisation about trait evolution; case studies of individual traits are the vehicles by which these larger generalisations can be judged.³⁸

The controversy about adaptationism has been heated, but nowhere more so than in discussions of human evolution. Gould and Lewontin criticised adaptationism because they saw it as the deep problem afflicting E. O. Wilson's sociobiology.³⁹ For Gould and Lewontin, sociobiology was the symptom, naive adaptationism the disease. I began this section by discussing the methodological problems that need to be addressed if all zebras run fast. The very same problems arise in sociobiology when one considers a trait that all human beings share. Why are human beings able to speak a language? Why do human societies have religious practices and ethical norms? If a trait is a human universal, how can an adaptive explanation of the trait be tested? When we reach for the comparative method to answer this question, we run into a problem. The nice feature of running speed is that it is a quantitative characteristic; there is no difficulty in comparing the running speeds found in different species. But how can 'speaking a language' and 'having a religion' be redescribed, so that they become quantitative characters that render cross-species comparisons intelligible? This is the challenge that faces those who want human evolution to be part of the larger story.

VII CONCLUDING COMMENTS

It is remarkable that philosophical questions about meaning and methodology engage the attention of evolutionary biologists. Like Molière's Monsieur Jordan, who spoke prose without realising it, biologists will not always describe their research as philosophical

in character, but the fact remains that this is part of what they are doing. Here is a case in which philosophy is continuous with the science it studies. However, it would be a mistake to conclude from the fact that philosophical questions are live issues in this science that something is amiss. Enquiry does not proceed with clear concepts and well-justified methods all laid out at the outset. Rather, the methods of science and the results of science both develop, with each informing the other.

In the previous sections on the testing of genealogical and adaptive hypotheses, I outlined some of the methodological questions that this two-part theory raises. Indeed, a good deal of current scientific work seeks to bring these two components – the tree of life and natural selection – into more intimate contact with each other. Hypotheses about phylogenetic relationships cannot be tested in isolation from models of the processes governing trait evolution. And adaptive hypotheses about trait evolution are increasingly being examined against the background of our knowledge of phylogenetic relatedness. Darwinism is a two-part theory, but the two parts are methodologically connected. The metaphysical picture is that life-on-earth is a large physical object, extended through space and time. Biological taxa are pieces of this branching tree, with characters evolving on branches according to rules that need to be described in the language of probabilities. So novel is this framework for describing nature that science is still developing methods for testing hypotheses concerning the details of the evolutionary world picture.

NOTES

My thanks to James Crow, Anthony Edwards, Ellery Eells, Chris Lang, Michael Steel, Christopher Stephens and the editors of this volume for useful comments on earlier drafts.

1. The tree-of-life hypothesis, thus stated, does not assert that life forms a phylogenetic tree in the strict sense of that term. As one goes from root to tips in a tree, lineages split but never join. Plant species formed by hybridisation do not form a tree, and the same is true when there is pervasive horizontal transfer, as is the case in some bacteria. See also Gayon, this volume.
2. See Sober 1993.
3. See also Beatty 1987; Lloyd 1988; and Thompson 1988.

4. In C. Darwin [1859] 1964, 131.
5. See Schweber 1983.
6. Mutation gives rise to novel *alleles*, but recombination is another source of variation, in that it generates novel *combinations* of already existing alleles.
7. See Mills and Beatty 1978 and Sober 1984.
8. Kimura 1983. See also Gayon, this volume.
9. See Sober 1999.
10. This brief discussion is not intended as a defence of scientific realism; the point is just that the standards we use for deciding whether electrons are objective should be the same as the standards we use for deciding whether probabilities are objective.
11. For a different interpretation, compare Brandon and Carson 1996.
12. This section is drawn from Sober 2002.
13. Mill 1872.
14. Kripke 1980 and Putnam 1975.
15. Enç 1986.
16. Hull 1965 and E. Mayr 1976.
17. See Sober 1994b.
18. See E. Mayr 1963, 1970.
19. For other species concepts, see Ereshefsky 1992.
20. Ghiselin 1974 and Hull 1978, 1987.
21. See Wiley 1981.
22. See Ereshefsky 1991.
23. On adaptations as forming a natural kind, see Radick, this volume.
24. It is widely held that if a given tree-selection method (e.g., parsimony) singles out the same tree as best when different data sets are considered, this is evidence that the taxa considered have a common ancestor. Penny *et al.* 1982 have made this argument rigorous. I suggest that the test is flawed – a tree can generate characters that are incongruent with each other, and a set of unrelated species can generate characters that all lead parsimony to the same (erroneous) tree.
25. Crick 1968.
26. In Freeland *et al.*, 2000.
27. Why think that winglessness is the ancestral condition? Characters are usually polarised by the method of outgroup comparison. See Sober 1988 for discussion.
28. See Lewis 1998.
29. Felsenstein 1978.
30. Endler 1986.
31. See Burt 1989 and Orzack and Sober 2001.
32. See Harvey and Pagel 1991.

33. In Gould and Lewontin 1979.
34. In Maynard Smith 1978.
35. See Dawkins 1983.
36. See Dennett 1995.
37. See Sober 1993 and also Orzack and Sober 1994.
38. In just the same way, the generalisation 'most speciation is allopatric' can be tested, but only indirectly, by looking at a range of case studies.
39. E. O. Wilson 1975.