The Chances of Evolution:
An Analysis of the Roles of Chance in Microevolution and Macroevolution

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Roberta Lynn Millstein

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Advisers:
C. Kenneth Waters
John Beatty

August 1997

Introduction

Charles Darwin's *On the Origin of Species* profoundly changed the way we view the natural world and the way we view ourselves within that world. Ever since Darwin, we have seen species not as static entities specially created by a higher being, but as entities which change in response to environmental pressures. According to Darwin, such changes are primarily due to natural selection: the differential reproduction of organisms as a consequence of differences in fitness in a given environment. Thus, on Darwin's view, humans and other organisms are the result of a causal process that has taken place over millions of years, a causal process which has led to the adaptation of organisms within their environments.

In spite of this causal focus, Darwin's theory of natural selection was sometimes referred to as the "law of higgledy-piggledy" in the period immediately following the publication of the *Origin*. Indeed, Darwin incorporated various concepts of chance into different aspects of his theory. For example, Darwin refers to "mere chance" as causing "one variety to differ in some character from its parents" (Darwin [1859] 1964: 111); to "individuals having any advantage, however slight, over others" as having "the best chance of surviving and of procreating their kind" (Darwin [1859] 1964: 81); and to the proportions of plants and bushes along a bank as not being the result of "chance" but of "action and reaction of the innumerable plants and animals" (Darwin [1859] 1964: 74-5). Further compounding matters, at the outset of Chapter V, Darwin notes that he has "sometimes spoken as if the variations" in domesticated and non-domesticated organisms "had been due to chance," conceding that "[t]his, of course, is a wholly incorrect expression, but it serves to acknowledge plainly our ignorance of the cause of each particular variation" (Darwin [1859] 1964: 131). In other places, Darwin emphasizes that variations are "chance" in the sense that they may be beneficial to the organism, or harmful, or neither (e.g., Darwin [1859] 1964: 81). This is anything but an exhaustive list of Darwin's usages of chance, and yet it illustrates that Darwin used the term 'chance' in a number of different areas and with a number of different meanings. We shall see three of

these meanings arise in our discussions of present-day evolutionary theory: chance in the context of organisms which are fitter, but not more reproductively successful than their less fit counterparts (Chapter 1); chance as ignorance of the real causes (Chapter 2); and chance in the context of variations which are beneficial, harmful, or neither (Chapter 5).

In spite of Darwin's copious use of the term 'chance', the charge that Darwin's theory of natural selection is the law of higgledy-piggledy is not really justified. For Darwin, evolution by natural selection is a two-step process. The first step is the *production* of chance variations. The second step is the *perpetuation* of the variations, and this, according to Darwin, takes place by the process of natural selection: variations that are beneficial have the best chance of being preserved and passed along to offspring. Both of these steps involve "chance" to some extent. However, the second step, perpetuation of variation, is not really a chance process. *Sometimes* the fittest organisms may not be the most reproductively successful, but for the most part (and particularly, over the long run), they will. Thus, at best, Darwin's theory of evolution by natural selection is only in part a chancy process. Furthermore, Darwin's invocation of the Laplacean notion of chance as "ignorance of the real causes" means that for Darwin, evolution is a *deterministic* process.

However, recent developments in evolutionary biology raise futher questions about the role of chance in evolution. The theory of random drift raises the possibility that a large part of evolution may *not* be due largely to the greater reproductive success of fitter organisms (perhaps even over the long run). Stochastic theories of macroevolution suggest that speciation and extinction may not be the result of protracted acts of natural selection, but that speciation and extinction may occur randomly across taxa and over time. Neither of these theories were a part of Darwin's evolutionary theory, and together, they raise the question of whether the perpetuation of variation is more of a chancy process than Darwin ever thought. On the other hand, recent experiments with bacteria call into question the deeply held belief that all mutations are random, so that with respect to the production of chance variation, evolution may be *less* chancy than Darwin thought. Just as Darwin's theory altered the way we see ourselves and the natural world, the question of the role of chance in evolution has the potential to further alter those views. Perhaps not

only is our presence here *not* due to a special creation event – perhaps it is not even largely due to the causal process of natural selection. Or, if it turns out that non-random mutations are prevalent in the natural world, perhaps evolution is more of a directed process than we had previously thought. Darwin's theory of natural selection overturned our previous views of the natural world; the study of the role of chance in evolution has the potential to overturn our views yet again, and perhaps to cause us to reevaluate some of the many views and theories which have been based on the more causal theory of natural selection.

The study of the role of chance in evolution also has broader philosophical interest; the question of the nature of chance, and related questions of determinism and indeterminism, are longstanding philosophical problems that have occupied philosophers for centuries. More recently, philosophers of science in general and philosophers of physics in particular have looked to quantum mechanics to settle issues concerning the nature of chance and to settle questions concerning the fundamentally probabilistic nature of the universe. However, evolutionary theory itself has a decidedly probabilistic character, one which in some sense does not seem to rest on any new discoveries in quantum mechanics; evolutionary biology was given a probabilistic formulation prior to and independently of the development of quantum mechanics. Perhaps the probabilities inherent in evolutionary theory represent only our ignorance concerning the true causal structure of events – or perhaps not. Perhaps biological processes are inherently probabilistic in some deeper sense. And if this is the case, we may find that the tables are turned, and that the study of chance in biology sheds light on the study of chance in quantum mechanics and in other areas of science and philosophy. Thus one general goal of this project is to provide insight into the deeper philosophical questions concerning the nature of chance.

Apart from this general goal, the more specific goal of this dissertation is to explore and illuminate the nature and role of chance in evolution. This is accomplished by examining three distinct areas where chance enters into evolutionary theory (the areas mentioned above: random drift, stochastic macroevolutionary models, and random mutation). In each area, I describe the particular conception of chance in use, in part by

contrasting it with alternative, less chancy accounts. I also address issues of causality, explanation, and determinism, in order to explore the role of chance in each area of study. These discussions serve to draw further connections between issues in the philosophy of biology and issues in the philosophy of science more generally. Throughout, I discuss similarities between the use of chance in the different areas, and in the Conclusion to the dissertation I address the question of whether evolutionary theory is probabilistic in any *general* sense, or only in a variety of particular, distinct senses. This approach provides a comprehensive, synthetic account of the nature and role of chance in evolution that improves on previous accounts which have tended to focus only on one area at a time.

The first three chapters are devoted to issues connected to random drift; the fourth and the fifth chapters discuss stochastic processes in macroevolution and random mutation, respectively. Random drift receives greater attention than the other areas for three reasons: 1) Random drift is often considered to be *the* source of stochasticity within evolutionary theory, at least by American philosophers and biologists; 2) random drift has been the subject of much recent discussion by philosophers of biology, leaving a number of unresolved controversies; and 3) my discussion of random drift provides a template for exploring random mutation and stochastic models of macroevolution.

In Chapter 1, I describe seven different processes collectively known as random drift. The discussion of random drift processes is used to articulate the general conception of random drift implicit in each of the processes. (Chapters 2 and 3 also draw on the discussion of random drift processes). The conception of random drift is illustrated using a model in which colored balls are sampled from an urn (a widespread method of describing random drift). I then turn to an issue raised by John Beatty (1984), who asks whether the concept of random drift can be distinguished from the concept of natural selection, suggesting that in many cases, the two concepts are not distinguishable. If Beatty is right, then controversies in the latter half of the 20th century over the prevalence of natural selection versus random drift are largely misguided. I examine various candidate solutions to the problem of conceptually distinguishing natural selection from random drift, and show why they are inadequate. However, contra Beatty, I argue that the concepts of

natural selection and random drift *can* largely be distinguished, and I propose a solution that uses both causality and a widely used distinction between rare and common environmental events.

In Chapter 2, I discuss issues of determinism and indeterminism with respect to random drift, natural selection, and evolution as a whole. In the first part of Chapter 2, I examine the arguments of Alexander Rosenberg (1994), who maintains that an omniscient account of evolution does not need the concept of random drift, because, according to Rosenberg, explanations which appeal to random drift can be replaced by explanations which appeal to natural selection. I argue, contra Rosenberg, that any evolutionary theory, omniscient or otherwise, must take random drift into account. I examine Rosenberg's arguments for the claim that random drift is eliminable from an omniscient viewpoint, and show that the concept of random drift upon which they rest is an impoverished one. I then provide three reasons why random drift is not eliminable from an omniscient account of evolution.

Rosenberg's arguments for the eliminability of random drift occur in the context of an argument for the determinism of evolutionary biology. If Rosenberg is right that random drift is eliminable from an omniscient account of evolution, and if we assume that natural selection is deterministic and random drift is *the* source of stochasticity within evolutionary theory, it would follow that the evolutionary process is a deterministic one. However, my arguments against Rosenberg's claims suggest that in order to argue for the determinism of the evolutionary process, Rosenberg must show that random drift is *itself* deterministic. Barbara Horan (1994) also argues for the determinism of the evolutionary process; although she asserts that random drift is deterministic, she offers little proof for this claim. In a recent paper, Robert Brandon and Scott Carson (1996) examine the processes of random drift and natural selection, and mount a decisive challenge to Rosenberg's and Horan's claims concerning the determinism of the evolutionary process. If Brandon and Carson are right, a scientific realist should conclude that the evolutionary process is indeterministic, not deterministic. In the second part of Chapter 2, I examine Brandon and Carson's arguments against Rosenberg and Horan and argue that there is insufficient

reason to favor one side over the other. Instead, I will argue that a more philosophically defensible position argues neither for the fundamental determinacy nor indeterminacy of the evolutionary process. However, I suggest that even without making these kinds of empirical claims concerning the determinism or indeterminism of the evolutionary process, we can still make arguments concerning the *probabilistic character* of evolutionary theory. That is, it remains an open question as to whether evolution is inherently and unavoidably probabilistic. I show that even if one assumes that the evolutionary process is fundamentally deterministic, the status of natural selection and random drift as population-level processes implies that evolutionary theory is inherently and unavoidably probabilistic.

In Chapter 3, I address the subject of the explanatory nature of random drift. Chance is commonly held to be non-explanatory, for a number of different reasons that I outline. Yet biologists frequently use random drift in evolutionary explanations. This raises questions that have gone largely unexplored by philosophers of biology: how is it that chance, in the form of random drift, can be explanatory? What is it about the concept of random drift that makes it explanatory, when presumably, other conceptions of chance are not? I describe a recent account in the philosophy of science, due to Paul Humphreys (1989), which holds that chance is not explanatory. Humphreys proposes an account of causal explanation, in which a list of contributing causes can explain a property or a change in a property, and a contributing cause is a factor which increases the probability of the occurrence of an event (the effect). On Humphreys's account, chance lacks causal efficacy and is therefore non-explanatory. However, I argue that chance, in the form of random drift, *is* causal in Humphreys's sense, and therefore is explanatory under a causal account of explanation. This is done by showing how it is that random drift causes a number of different evolutionary phenomena, phenomena which evolutionary biologists invoke random drift to explain. Thus, in evolutionary theory, random drift explanations play the role of an alternative to explanations by natural selection. Furthermore, if my arguments are correct, we need to broaden our conception of chance to include chance that is both causal and explanatory. The study of random drift and explanation thus provides us with a reason to rethink some of our traditional philosophical ideas about chance.

In Chapter 4, I examine a stochastic model of the macroevolutionary processes of speciation and extinction. This is an area of evolutionary theory largely unexplored by philosophers.[1] This stochastic model was proposed as an alternative to traditional "deterministic" macroevolutionary explanations, which invoke specific causes as the explanation for the macroevolutionary phenomenon in question. (The term 'deterministic' is in scare quotes because such accounts are not Laplacean deterministic; both stochastic and deterministic explanations are consistent with either Laplacean determinism or indeterminism). I describe the stochastic model, contrasting it with deterministic explanations, and discuss the ways in which the stochastic model provides an alternative to deterministic explanations. Next, I explore the sense in which the stochastic account is stochastic, and suggest that the argument that I gave concerning the population-level nature of random drift applies to stochastic models as well. I then show how it is that the stochastic account can explain phenomena that have been traditionally explained deterministically, and demonstrate how stochastic explanations account for specific kinds of macroevolutionary phenomena. Throughout the discussion, parallels are drawn between stochastic macroevolutionary models and random drift. Last, I summarize the similarities between the random drift and stochastic macroevolutionary processes, and suggest some broader conclusions about stochasticity within evolutionary theory.

In Chapter 5, I analyze the topic of random mutation. It is a major tenet of present-day neo-Darwinian evolutionary theory that mutations, which lead to variation within a population, occur at random. From the time of a paper by S. E. Luria and M. Delbrück in 1943 (a paper that claimed to experimentally demonstrate that directed mutations do not occur) through 1988, few biologists doubted the claim that all mutations are random. However, in 1988, John Cairns, Julie Overbaugh, and Stephan Miller claimed to have demonstrated the occurrence of directed mutation in bacteria, setting off a series of experiments and papers by other biologists, some in an attempt to support their claims and others to refute them. While it is probably still true that most biologists believe that all mutations are random, it is perhaps less true than it once was; many respected biologists

[1] At a recent conference, I met another philosopher, Todd Grantham, who is also working on these issues, but to my knowledge, no one has published yet on this subject.

now take seriously the idea that some mutations might be directed, and some would say the issue is as yet unresolved. This recent debate has raised issues for how the concepts of random mutation and directed mutation are defined, and for how the opposing hypotheses ("all mutations are random" versus "some mutations are directed") can be distinguished empirically. I provide a short history of the directed mutation controversy. I then attempt to clarify and distinguish between the concepts of 'directed mutation' and 'random mutation', by examining three candidate sets of definitions, one from Sahotra Sarkar, one from Richard Lenski and John Mittler, and one from Eva Jablonka and Marion Lamb. Last, I describe the nature of the directed mutation controversy and suggest a way that the opposing hypotheses can be distinguished empirically. Note that, unlike the case for random drift and stochastic macroevolutionary models, it would be inappropriate to speak in terms of contrasting explanations; the topic of directed mutation is so contentious that it would probably be premature to refer to "directed mutation explanations". Thus, I explore the more fundamental question of how we would determine whether directed mutations occur *at all.* As with the stochastic macroevolutionary model, I draw parallels between the random drift case and the random mutation case throughout the chapter.

In the Conclusion, I explore the parallels between the three areas of chance under study with respect to the concepts in use in each area, their relationships to the question of Laplacean determinism or indeterminism, and their explanatory roles within evolutionary theory. I then offer some suggestions concerning the importance of chance in evolutionary theory and the evolutionary process as a whole.

Chapter 1: Random Drift and Natural Selection

1.1. Introduction

Sober characterizes random drift[2] as "the source of the stochastic element in evolution," in contrast to mutation and selection which "are treated as deterministic forces of evolution" (Sober 1984: 110). Rosenberg claims that, "[e]volutionary biologists seem to identify the source of unpredictability in evolution with the phenomenon of drift" (Rosenberg 1994: 67). Brandon and Carson (1996) suggest that the "crucial question" regarding indeterminism in evolutionary theory concerns random drift. Thus, while a philosophical examination of the role of chance in evolution might possibly examine mutation, recombination, or contingency, it is unavoidable that it would examine random drift.

If you pick up any population genetics textbook, you will inevitably find a discussion of random drift alongside the other mechanisms of evolution, such as natural selection, mutation, and migration. Yet the role of random drift in 20th century evolutionary biology has been a turbulent one. Biologists such as Sewall Wright (the person most often associated with random drift) and R. A. Fisher disagreed over the relative importance of random drift in the evolutionary process, with Fisher arguing that Wright's theory assigned too great a role to random drift. The debate continues today, with biologists such as Kimura and Crow (so-called "neutralists") claiming an even larger role for random drift than Wright did, while biologists ("selectionists" such as Ernst Mayr and Douglas Futuyma; see especially Mayr 1983 and Futuyma 1988) remain steadfast to the idea that it is natural selection that plays the preeminent role in phenotypic evolution. According to the selectionist camp, the role of random drift is at best a minor one. As Beatty (1984) has argued, this debate has taken place in specific contexts (over whether, for example, natural selection or random drift is more prevalent in a particular population) as well as more

[2] Authors refer variously to "drift","genetic drift", "random drift", "random walk", and "random genetic drift". For the sake of consistency, I will limit my usage to the term 'random drift' although authors I am discussing may be using an alternate terminology.

general contexts (over whether natural selection or random drift is more prevalent in evolution as a whole).

With all this contention concerning the relative importance of random drift and natural selection within evolution, one would hope that the concepts of random drift and natural selection could be clearly defined. Furthermore, one would hope that the concepts could be clearly and easily distinguished. After all, if one is to argue that random drift is very prevalent in evolution, and natural selection less so (or vice versa), it had better be the case that instances of natural selection can be distinguished from instances of random drift – and in order to do *that*, the concepts must be distinct. Even (or perhaps, especially) if one takes a non-empirical, theoretical approach to settling this debate, it is still necessary that one be able to distinguish the concept of random drift from the concept of natural selection. However, as Beatty (1984) has shown, it is not all that easy – perhaps even impossible – to distinguish the two concepts. But if that is true, then it seems as though the neutralist/selectionist debate has just been a tempest in a teapot, and that participants in the debate have been arguing about nothing.

While this is a possibility, it is clearly an undesirable conclusion that one would not want to accept until the subject had been thoroughly exhausted. It is with this motivation that I seek to reexamine the questions of how to characterize random drift, and whether it is possible to distinguish random drift from natural selection (without taking a stance on the neutralist/selectionist debate, an issue that is best left to biologists). My answers to these questions form the subject matter of this chapter, and provide the basis for Chapters 2 and 3.

I will take a three-step approach to characterizing random drift. In the first stage, I discuss the different processes (such as the indiscriminate sampling of gametes during reproduction and the founder effect) which are collectively known as random drift. This characterization provides the biological basis for the concept of random drift. In the second stage, I describe a model of random drift through an analogy to random sampling from an urn. Finally, in the third stage, I turn to the question of whether random drift can be distinguished from natural selection. After rejecting two solutions to this problem, I

propose a solution of my own, using the biological and theoretical characterizations of random drift that I develop in the first two stages. I argue that the definitions I provide are sufficient to make the distinction between natural selection and random drift, with some qualifications.

1.2. Kinds of Random Drift

The term 'random drift' actually refers to a number of distinct processes, not just one kind of process. Beatty (1992) isolates four kinds of random drift: indiscriminate parent sampling, indiscriminate gamete sampling[3] , the founder effect, and fluctuations in the rates of evolutionary processes (selection, migration and mutation). Related to the founder effect is the so-called "bottleneck effect." Another kind of random drift is the random assortment of genes into gametes (Dodson and Dodson 1985: 229). Brandon and Carson (1996) claim that Brandon (1990) identifies yet another "drift-like phenomenon". This leaves potentially at least seven different kinds of random drift. I say "potentially" because there has been some controversy as to which processes should be considered to be random drift (Beatty 1992); indeed, different authors will describe different subsets of this list as constituting the kinds of random drift.[4] I will describe each of the seven processes and try to draw out the general characteristics that are common to all of them.

"Parent sampling," according to Beatty, is "*the process of determining which organisms of one generation will be parents of the next, and how many offspring each*

[3] As far as I know, the terms "gamete sampling" and "parent sampling" were first used in Beatty 1984, though he does not make use of this terminology in his 1992 paper.

[4] Ridley (1993: 126-7) discusses indiscriminate gamete sampling, indiscriminate parent sampling, the founder effect, and the bottleneck effect. Hartl and Clark (1989: 61, 70, 83-4) also discuss these same four phenomena. Dodson and Dodson assert that random drift refers to indiscriminate gamete sampling and the assortment of genes into gametes (1985: 229), although bottlenecks and the similar founder effect are "corollaries" (1985: 231). Kimura (1983: 37) and Roughgarden (1996: 57) only discuss indiscriminate gamete sampling in their discussions of random drift; Roughgarden asserts that random drift "originally referred to all chance effects in evolution but has come to refer to a special source that is internal to the population," i.e. gamete sampling (1996: 57). Thus clearly there are some differences of opinion within the field of evolutionary biology as to which processes should be considered to be random drift. Indiscriminate gamete sampling seems to be the most uncontroversial form of random drift. Other than that, however, there appears to be no consensus. Although the question of which of these processes should properly be considered to be random drift is an interesting one, I will not be discussing it at any length.

parent will have" (Beatty 1984: 188; italics in original). When physical differences between organisms are causally irrelevant to this process, we can say that the parent sampling is *indiscriminate*. Conversely, when physical differences between organisms are causally relevant, we can say that the parent sampling is *discriminate*. Discriminate parent sampling is generally considered to be natural selection; indiscriminate parent sampling is random drift. For example, a population of grey and brown squirrels might be exposed to a forest fire which indiscriminately kills squirrels of both colors. Suppose that the grey squirrels died in greater numbers than brown squirrels – it just so happened, say, that more of the grey squirrels were in the area of the forest that caught fire. Consequently, the frequency of grey squirrels in the next generation would be less than the frequency in the parent generation. More grey squirrels than brown squirrels failed to survive to reproduce, but the color of the squirrels was causally irrelevant to this "sampling" process. Thus, the change in frequency in this squirrel population would be due to the kind of random drift known as indiscriminate parent sampling.

Beatty defines "gamete sampling" as "*the process of determining which of the two genetically different types of gametes produced by a heterozygotic parent is actually contributed to each of its offspring*" (Beatty 1984: 189; italics in original). Gamete sampling takes place during the process whereby gametes are united to form zygotes. In any given population, the number of gametes (the "gene pool"[5]) is very large relative to the size of the population. On average,[6] the genetic composition of the gene pool will be representative of the parents' generation. However, only a finite number of those gametes will become founders of the next generation; the number of zygotes is always finite. Thus the successful gametes (those that become zygotes) are a sampling of the total number of viable gametes in the gene pool. If each individual has an equal chance of contributing its

[5] The term 'gene pool' is metaphorical; it arises out of considering the phenomenon at the population level. At an individual level the phenomenon we are considering has to do with the alleles a heterozygote passes along to its offspring; the distribution of alleles in the offspring of a given heterozygotic individual may not be equal. There is of course no question as to which alleles a homozygote will pass along.

[6] Not all gametes may be viable, and in females only one quarter of the products of meiosis become eggs (Ridley 1993: 125). Also, there may be non-Mendelian production of gametes, i.e. meotic drive. All of these processes would cause the frequencies of alleles among the gametes to differ from the frequencies among the parents.

successful gametes to the next generation, i.e., if the success of gametes is not due to any physical differences between them, then there is *indiscriminate* gamete sampling. As a consequence, the successful gametes may form a non-representative sample of the gene pool (the gene frequencies among the former may differ from those among the latter). Indiscriminate gamete sampling, like indiscriminate parent sampling, is a form of random drift. Both are processes that sample indiscriminately with respect to physical characteristics; that is, processes in which physical differences are causally irrelevant to the sampling process. The difference is that with indiscriminate parent sampling, the sampling is of phenotypes, whereas indiscriminate gamete sampling involves the sampling of gametes.

The *founder effect*, usually attributed to Ernst Mayr, occurs when a small sub-population (the "founders") becomes spatially isolated from the rest of the population (the "parental" population). In the most extreme case, the founder population consists only of a single fertilized female (Mayr 1970). The isolation of the founder population can occur passively, as when a river might split a population, or actively, by the migration of a group of individuals away from the parental population. For example, suppose a few woodpeckers fly out of their usual range, away from the rest of the population, and colonize a new area. If the separation of the founder woodpeckers from the parental woodpecker population is not based on physical differences between the two populations (supposing, for example, that it is not the case that the woodpeckers left because they could fly farther than the other woodpeckers, but that any physical differences between the founder woodpeckers were causally irrelevant to their leaving), we can say that the founder woodpeckers were sampled indiscriminately from the parental woodpecker population. If a founder population is small, then it is likely that only a small fraction of the variation of the original population is represented amongst the founders. For example, if most of the woodpeckers in the parental population had large red tufts on their heads, it may be the case that a much smaller proportion of woodpeckers in the founder population has large red tufts, or perhaps none of the founders have them. We would then expect that subsequent generations of the new founder population would have different genotype

frequencies than that of the original parental population. The indiscriminate sampling by isolation of a small number of founders from a larger population is called the founder effect. Like parent sampling, the founder effect is sampling which is indiscriminate with respect to the physical differences of *phenotypes* (physical differences are causally irrelevant to the phenotype sampling). However, it differs from parent sampling in that the sampling occurs through a splitting of the population.

A related phenomenon to the founder effect is what is sometimes called the *bottleneck effect*. A bottleneck occurs when there is a large reduction in numbers in a population (as occurs frequently in some populations which tend to fluctuate in numbers) and the reduced population carries only a small fraction of the variation of the original population. Some authors characterize the bottleneck effect as a kind of founder effect (see, e.g., Futuyma 1986, Ridley 1993), whereas others classify the founder effect as a kind of bottleneck effect (see, e.g., Hartl and Clark 1989). In both the founder effect and the bottleneck effect, the physical differences between the smaller population and the original population are not causally relevant to the formation of the smaller population (i.e., there is indiscriminate sampling). Furthermore, as a result of both the founder effect and the bottleneck effect, the smaller population may not have the same distribution of characteristics as the original population. The main difference between the founder effect and the bottleneck effect is that in the former, a small subpopulation splits from a larger population, while in the latter, the whole population undergoes a drastic reduction in size. To make matters even more confusing, the bottleneck effect is also quite similar to indiscriminate parent sampling, discussed above, since both processes involve sampling of parents which is indiscriminate with respect to physical differences. The primary difference seems to be that with the bottleneck effect, there is a drastic reduction in numbers, such that the "sample" is likely to be particularly unrepresentative of the original population, whereas indiscriminate parent sampling can be on a much smaller scale. The difference is thus more one of degree rather than of kind; the bottleneck effect can be seen as an extreme case of parent sampling.

Brandon claims that "there is another way, until now unrecognized, in which chance

can intervene to disassociate adaptedness from actualized fitness. By chance different types can be unequally distributed over a heterogeneous environment" (Brandon 1990: 46). (I will refer to this process as the "Brandon effect"). According to Brandon and Carson, this is a "drift-like phenomenon" which "certainly occurs and certainly has evolutionary consequences" (Brandon and Carson 1996: 321, n. 8). The example Brandon (1990) gives is of two seed types, *a* and *b*, which the wind distributes over a field containing patches of fertile soil and patches of toxic soil. He supposes that by chance, more *a*'s than *b*'s land in fertile soil, with the result that the proportion of *a*'s in the population increases. The process that Brandon describes, like the other random drift processes we have been discussing, is a process where there is sampling that is indiscriminate with respect to physical differences. The increase in the frequency of *a*'s relative to the frequency of *b*'s is not due to the physical differences between the *a* seed type and the *b* seed type (as would be the case under natural selection); rather, any physical differences between the two seed types are causally irrelevant to the change in frequency. Moreover, the Brandon effect is one where there is no splitting of the population and the sampling is of phenotypes. It is thus questionable as to whether Brandon has "recognized" a new form of drift; the characteristics of the process he describes fit the characteristics of indiscriminate parent sampling. There is little difference between the example we discussed above, where the frequency of brown squirrels increases in the population because they happen to be in the area of the forest that does not catch fire, or seeds of type *a* which happen to land on fertile soil (with no causal connection between the physical differences and the location of the organisms in either case). Brandon has succeed in drawing our attention to one kind of parent sampling, where there is unequal distribution of types in a heterogeneous environment, but he has not revealed a new kind of random drift.

However, a kind of random drift, which, if not new, is seldom discussed in the literature, is mentioned by Dodson and Dodson (1985). Dodson and Dodson assert that random drift "refers to the accidental fluctuations in the proportion of a particular allele, which depend upon the fact that the assortment of genes into gametes and the combination

of gametes to form zygotes are random processes" (Dodson and Dodson 1985: 229). Now, the random process whereby gametes combine to form zygotes is what we have been calling indiscriminate gamete sampling. But what is the random assortment of genes into gametes? Here Dodson and Dodson seem to be referring to the processes of independent assortment and recombination, both of which occur during meiosis (the process which creates gametes). *Independent assortment* refers to the independent alignment of the maternal and paternal pairs of homologous chromosomes; when the pairs segregate, each member of the pair has an equal chance of facing one pole or the other on the equatorial plate (Gardner, Simmons, and Snustad 1991: 66). So, in other words, for each pair of chromosomes, it is a chance process as to which member of a pair of chromosomes will be part of which gamete, with the result that different assortments of chromosomes are possible. *Recombination*, on the other hand, is a possible result of crossing-over (the exchange of chromosome segments) whereby a new combination of genes is formed (Gardner, Simmons, and Snustad 1991: 165). So, when gametes are formed from a pair of parental chromosomes, not only are different combinations of chromosomes possible (as a result of independent assortment), but, within each chromosome, different assortments of *genes* are possible (as a result of recombination). Together, the processes of independent assortment and recombination determine the assortment of genes into gametes, and thus determine the genotypes that are to be represented among the gametes. If the physical differences between genes are causally irrelevant to which gene assortments get represented among the gametes, then recombination and independent assortment, taken together, are an indiscriminate sampling process, analogous to indiscriminate gamete sampling, indiscriminate parent sampling, and the founder effect. We might call this process "gene sampling".

The last kind of random drift we will examine was described by Sewall Wright. Although initially Wright considered only indiscriminate parent and gamete sampling as random drift (even though he did not use those terms), he later broadened his conception of random drift to include random fluctuations in selection pressure, mutation rate, and migration rate (Beatty 1992: 277). For example, changes in weather may cause selection

pressure to fluctuate randomly, because some organisms in the population may be better adapted for warmer temperatures, whereas others are better adapted for cooler temperatures. Beatty argues that Wright's inclusion of these sorts of random fluctuations in his account of random drift is justified when we look at it in the context of Wright's shifting balance theory, since the effects of the fluctuations in selection, mutation, and migration rates have the same kind of effect as random drift and could therefore play the same role as random drift within Wright's theory (Beatty 1992: 279). However, as Beatty acknowledges, Wright's extension of the term to include fluctuations in the rates of other evolutionary processes is "questionable in certain respects." (Beatty 1992: 278). In particular, it is difficult to see how fluctuations in the rates of these processes could be construed as indiscriminate sampling.[7] Thus, while fluctuations in selection, mutation, and migration rates may share the same kind of *outcome* as random drift, they are not the same kind of *process* that we have been discussing, processes whereby sampling is indiscriminate with respect to physical differences. We will discuss the distinction between random drift as a *process* and random drift as an *outcome* further below.

Each of the processes described above, with the exclusion of fluctuations in selection pressure, mutation rate, and migration rate, can be characterized as an indiscriminate sampling process in the following way. For each kind of random drift, differences in the relevant "entity" are causally irrelevant to differences in the "success" of the relevant entity. With indiscriminate parent sampling, the founder effect, the bottleneck effect, and the "Brandon effect," the relevant entity is the individual organism. With indiscriminate gamete sampling, the pertinent entity is the gamete; with gene sampling, it is the gene. Moreover, each of these indiscriminate sampling processes leads us to expect certain kinds of outcomes. In the next section, we shall look at a general model for these indiscriminate sampling processes, and see why certain outcomes are likely given that kind of process (i.e., a random drift process).

[7] Moreover, as Cain and Currey argue, including these processes under the definition of drift will only lead to confusion: "Whatever may be the conveniences for mathematical geneticists of considering together all processes regarded as random, irrespective of their biological significance, the worker on actual examples must classify processes according to their biological significance" (Cain and Currey 1963: 59).

1.3. A Model for Random Drift

Population genetics textbooks standardly outline a model for random drift that maps on to the processes described above. Here I follow Hartl and Clark (1989), but the account is similar across textbooks; an ideal population model is introduced, with subsequent population models increasing in complexity. As we shall see, changing just one of the assumptions of the ideal model provides us with a simple model of random drift.

In an ideal population (diploid, sexually reproducing organisms; non-overlapping generations; random mating; infinite population size; negligible migration and mutation; no fitness differences between genotypes) genotype frequencies will not change from generation to generation. Instead, they will remain in a what is called a Hardy-Weinberg equilibrium. To see why, consider a population where the genotype frequencies of three genotypes *AA*, *Aa*, and *aa* are p^2, $2pq$, and q^2 respectively, where p represents the allele frequencies of *A*, q the allele frequency of *a*, and $p + q = 1$, since *A* and *a* are the only two alleles in the population. Assuming that gametes for each genotype will be produced with the same genotype frequencies already existing in the population, we would expect the probability that a sperm or an egg carries *A* is q, whereas the probability that a sperm or an egg carries *a* is p. Our ideal, initial assumptions (in particular, the assumption of random mating) imply that the combination of gametes will be random, so that the chance that an *A* sperm will fertilize an *A* egg is $p \times p = p^2$. Thus, the frequency of *AA* genotypes in the next generation will be p^2. Similarly, the chance that an *a* sperm will fertilize an *a* egg is q^2, so the frequency of *aa* genotypes in the next generation will be q^2. Finally, the probability that an *A* egg will fertilize an *a* egg is pq and the probability that an *a* sperm will fertilize an *A* egg is qp, so the frequency of *Aa* genotypes in the next generation will be $2pq$. Thus, the genotype frequencies in the next generation are the same as in the initial generation: the frequency of *AA* is p^2, the frequency of *Aa* is $2pq$, and the frequency of *aa*

is q^2. These frequencies are the Hardy-Weinberg equilibrium frequencies for two alleles. Therefore, in an ideal population, genotype frequencies will remain constant from generation to generation in a Hardy-Weinberg equilibrium.

If we change our population slightly to be a finite one, however, we introduce the potential for the genotype frequencies to deviate from their Hardy-Weinberg equilibrium by chance, that is to say, we introduce the possibility of random drift. Why and in what sense this should be the case is often explained by analogy to random sampling from an urn of colored balls.

First let us consider what "random sampling" amounts to. Giere argues that the randomness of a sample is not a feature of the sample *itself*, but rather a feature of the sampling *process* (Giere 1991: 149). According to Giere, a random sampling process must satisfy two criteria: 1) each item has an equal chance of being sampled, and 2) there is no correlation between the outcome of one sampling event and any other (Giere 1991: 184). Suppose we have a large urn filled with equal numbers of black and white balls which we sample according to Giere's two criteria, so that each ball has an equal chance of being picked and there is no correlation between the colors of the balls that get picked. If a large sample of balls is taken, say, 100 balls, we would expect the frequencies of colored balls in the sample to be very close to the frequencies in the urn. So, for example, one likely result might be that we pick 49 white balls and 51 black balls. In this case, the frequencies of our sample (0.49 white and 0.51 black) would be almost equal to the frequencies in the urn (0.5 white and 0.5 black). On the other hand, if we only take a small sample of colored balls, say 10 balls, our sample may very well have strikingly different proportions of colored balls than the urn does. For example, one likely result is that we draw only 3 black balls and 7 white balls. In this case, the frequencies of our sample (0.3 black and 0.7 white) do *not* reflect the frequencies in the urn very well. Furthermore, on a subsequent draw, we might get frequencies of 0.8 black and 0.2 white, or 0.6 black and 0.4 white. Thus, when we take a large sample, we expect the frequencies of the sample to be fairly representative of the frequencies of the urn, but when we take a smaller sample, we expect the frequencies to be less representative of the frequencies of the

urn. Therefore, the frequencies of small samples may fluctuate from sample to sample (See Chapter 3 for further discussion).

The phenomenon of drawing nonrepresentative samples is known as sampling error. If we are drawing 10 balls, we can calculate the probability for each possible outcome: 0 black balls, 1 black ball, 2 black balls,... through 10 black balls These probabilities provide us with a probability distribution of the all the possible outcomes (in this example, a bell-shaped distribution). Of course, the representative sample is always the most probable outcome, but the smaller the sample, the more probable non-representative outcomes become.

Just as larger urn samples have frequencies which tend to better reflect those of the urn from which they are picked, in an ideal population (of infinite size), infinite samples will always be representative (at least in theory). Thus, in our ideal population, genotype frequencies remain constant from generation to generation. However, if we allow our population to be finite, we introduce the possibility that genotype frequencies will not remain constant from generation to generation, just as the frequencies in our small sample of colored balls do not reflect the frequencies of the urn as well as the large samples do and thus are liable to change from sample to sample. But why is this, i.e., what is the random sampling occurring in the population? It should be the same as random sampling in the urn: each individual has an equal chance of being picked, and there is no correlation between one individual's being picked and another's.

This talk of "picking" individuals is rather abstract, however, so let us return to our example of indiscriminate parent sampling discussed above. In that example, a population of grey and brown squirrels was exposed to a forest fire. More of the grey squirrels happened to be in the area of the forest that caught fire, so a greater proportion of brown squirrels survived to reproduce. In terms of our random sampling model, we would say that each squirrel had an equal chance of being in the fire, and there was no correlation between the fate of any one squirrel and the fate of any other squirrel. The result of this sampling was that the frequency of squirrels in the subsequent generation changed from that of the parent generation; there was in increase in the frequency of brown squirrels.

Our random sampling model fits the squirrel scenario fairly well, yet there are two apparent problems with the way I have set things out. One is that it doesn't seem quite right to say that each squirrel had an equal chance of being caught in the fire. Weren't some of the squirrels simply not in the area of the forest that caught fire, and so did not have a very great chance of perishing before being able to reproduce? Of course, we could say that our urn model suffers from a similar problem; aren't the balls on the top more likely to be picked than the balls on the bottom? However, this response doesn't seem completely satisfactory, perhaps because (as I shall suggest momentarily) it serves to overlook an important issue regarding the sampling of both the balls and the squirrels.

The second apparent problem is that the model, as I have described it, doesn't seem to capture what I drew out as the general characteristic of a random drift process: differences in the phenotypes (gametes, genes) are causally irrelevant to differences in the "success" of phenotypes (gametes, genes). So, perhaps our model needs to be modified to take this into account. Instead of claiming that each ball has an equal chance of being selected, we can make the lesser claim that the differences in the colors of the balls are causally unrelated to which ball gets picked, as would be the case if we were blindfolded and picking balls out of the urn. We could deliberately pick balls from the top, but so long as ball color is causally irrelevant to a ball's being *on the top* , it would be the case that ball color is causally irrelevant to which ball gets picked. Likewise, an organism's location may make it more likely to be sampled, but so long as the given physical characteristics not causally relevant to that location, we can say that the physical characteristics are not causally relevant to the organism's reproductive success. We might say that under random drift, organisms' reproductive success occurs randomly with respect to the given physical characteristics, meaning that physical differences are causally irrelevant to organisms' reproductive success.

At this point a terminological note is in order. In describing random drift this way, I am focusing on random drift as an *indiscriminate sampling process*. However, it should be noted that while random drift is sometimes spoken of as a *process* (as I have done here), at other times (and by other people) it is spoken of as an *outcome.* That is, one

could focus on the outcome of the processes that I described in section 1.2 above, rather than on the processes themselves. Previously, we discussed why it is that the outcome of an indiscriminate sampling process may be one where the frequencies of the sample do not reflect the frequencies of the parent generation. We can associate with each frequency a certain probability of occurrence, producing a generally bell-shaped probability distribution of gene frequencies. Representative samples are the most probable outcomes, but the smaller the sample, the more probable the non-representative samples become. Thus, in a small population undergoing an indiscriminate sampling process over a number of generations, genotype frequencies may fluctuate, or drift, from generation to generation. This fluctuation in gene or genotype frequencies from generation to generation (generally as the result of an indiscriminate sampling process) is what some authors refer to as random drift. However, in this chapter and in subsequent chapters, I will use the term 'random drift' only as an indiscriminate sampling process, never as an outcome. This will make it easier to contrast and compare random drift with other evolutionary processes, such as natural selection.

1.4. Conceptual Distinction Between Natural Selection and Random Drift

We have been discussing random drift in the absence of natural selection, i.e., in populations where there are no fitness differences between genotypes. However, as Beatty (1984) notes, in any real, finite population where natural selection is occurring, random drift will be occurring as well. So, to thoroughly treat the subject of random drift, one must analyze natural selection at the same time as random drift. Furthermore, as we discussed in section 1.1, there has been (and continues to be) a great deal of debate over the question of whether, in particular instances and in general, natural selection is more prevalent than random drift. For these reasons, it is important that the concept of random drift be distinguishable from the concept of natural selection.[8]

[8] The distinct question of how one distinguishes between natural selection and random drift *empirically* (as opposed to *conceptually*) is a complex one, which I will not be addressing here.

1.4.1. The Problem

However, according to Beatty, it is "conceptually difficult to distinguish natural selection from random drift, especially where *improbable results of natural selection* are concerned" (Beatty 1984: 196). On Beatty's account, natural selection is *discriminate* sampling, or sampling with regard to fitness differences, whereas random drift is *indiscriminate* sampling, or sampling *without* regard to fitness differences (Beatty 1984: 190-1). Yet Beatty maintains that these concepts are not conceptually distinct. His conclusion is based on the following example. Consider a population of dark moths and light moths, both having the same color-sensitive predator. The moths inhabit a forest of light colored trees (40%) and dark colored trees (60%) – trees which provide camouflage from the predator for the light and dark colored moths, respectively. Since there are more dark colored trees than light colored ones, the dark moths have a greater likelihood of landing on a tree which will prevent their detection by a predator. Thus we would expect the dark colored moths to have greater reproductive success than the light colored moths, i.e., the dark colored moths are fitter than the light colored moths in the given environment.

However, as Beatty emphasizes, the expected reproductive success of each type of moth (dark or light) can be represented by a probability distribution where the probability of the given type for leaving each of the possible quantities of offspring is displayed. Now make three further suppositions: 1) there is an overlap in the probability distributions of the dark and light moths, such that it is possible for the light moths to have greater *actual* reproductive success than the dark moths, even though the dark moths overall have a greater *expected* reproductive success; 2) the possible becomes actual, and the lighter moths *do* have greater reproductive success; and 3) the greater actual reproductive success of the lighter moths is due to the fact that, by chance, the dark moths more frequently landed on light colored trees than dark colored trees, despite the fact that there was a greater proportion of dark colored trees. Given these three suppositions, Beatty raises the following question: should we consider this change in gene and genotype frequencies to be

the result of natural selection, or of random drift? He states:

> That is, is the change in question the result of sampling discriminately or indiscriminately with regard to fitness differences? It is not easy to maintain that the sampling was entirely indiscriminate with regard to differences in survival and reproductive ability. At least it is difficult to maintain that the death by predation of *conspicuously* dark moths in this environment is indiscriminate sampling, whereas the death of *conspicuously* light moths in the same environment is selection. On the other hand, it is also difficult to maintain that selection alone is the basis of the change. At least, it is difficult to maintain that the fittest were selected (Beatty 1984: 195-6; italics in original).

Beatty seems to be arguing as follows. We don't want to attribute the changes in the population entirely to random drift, since the light moths who perished did so on the basis of their lightness; the lightness of these moths made them less fit than the darker moths in an environment where the majority of trees was dark. In other words, it seems as though we would want to say that the light moths were sampled *discriminately* with regard to their light color and thus, we would say that they died as a result of natural selection. However, when we consider the dark moths who died by predation, it seems as though we would want to say that their deaths occurred *indiscriminately* with regard to fitness differences, since it is just by chance that so many of the dark moths landed on light colored trees. Since the dark moths were "sampled" without regard to their darkness, we would want to attribute their deaths to random drift. But how could it be that the deaths of the dark moths are due to random drift, whereas the deaths of the light moths are due to natural selection? After all, is there anything that is different about the situations of the two moths? In both cases, the color of the moth was distinct against its background, leaving it vulnerable to the predator. Why should one case be considered natural selection and the other random drift? So perhaps it is the case that the change in the population was entirely due to selection. Yet this seems counterintuitive as well – how can it be that there was discrimination on the basis of fitness differences, yet it is the *less* fit organisms which had the better reproductive success?

So, it doesn't seem to make sense to attribute the change entirely to random drift, nor

entirely to natural selection, and it certainly does not make sense to attribute the change in frequency of one type of moth to natural selection and the change in frequency in the other type to random drift. Given these considerations, Beatty concludes that it is conceptually difficult to distinguish natural selection from random drift.

But why do the difficulties involved in determining whether the change in the moth population was due to natural selection or random drift constitute a *conceptual* problem? Brandon and Carson suggest that Beatty is really putting forth a question about how to *empirically* distinguish natural selection from random drift rather than how to make the *conceptual* distinction (Brandon and Carson 1996: 325; n. 13). On the surface, the example as Beatty has described it does appear raise an empirical (if hypothetically empirical) question: we are looking at the "results" for this population, i.e. a particular change in gene frequencies, and attempting to determine whether that result is due to natural selection or random drift. This would be analogous to what would be done in attempting to answer the very same question in a real population; in a real population, as in a hypothetical population, determining whether an outcome was (in fact) due to random drift or to natural selection is an empirical task. However, it does not seem plausible that the problem that Beatty raises is an empirical one, for this reason: we know all the facts of the case (since it is a hypothetical situation) and we *still* have difficulty in answering the question of whether the cause is random drift or natural selection. Our concepts should be rich enough to cover any empirical situation. That we have difficulty answering the question indicates that our conceptual apparatus is not sufficient to account for the empirical situation. But then what is the conceptual question?

I think that Beatty uses the moth example to demonstrate the conceptual problem, but that should not confuse us into thinking that determining whether given instances of evolutionary change are natural selection or random drift *is* the problem. The conceptual problem lurks beneath the surface of the example. As Beatty suggests, natural selection and random drift are both probabilistic. On Beatty's account, natural selection is a process whereby organisms are discriminated on the basis of their fitnesses, but the discrimination is probabilistic. It is not that the fittest always have the greatest reproductive success; but

that they are more likely to do so than their less fit conspecifics. We might expect longer legged gazelles to have more success in outrunning their predators, but they may be "unlucky" and come into contact with more predators than the shorter legged gazelles. Thus, natural selection doesn't give us a specific prediction of an evolutionary outcome, but gives different predictions (assigns probabilities) for a range of outcomes. However, as we have seen, random drift is also a probabilistic process that assigns probabilities for a range of outcomes. According to Beatty, random drift is *indiscriminate* sampling, in that organisms are sampled without regard to differences in their fitnesses. So, the question is, when you have a probabilistic process occurring in the population, is it the probabilism of random drift or the probabilism of natural selection? To go back to the moth example, when the dark moths happen to land on a greater number of light colored trees, is that the probabilism of natural selection (discriminate sampling) or is it the probabilism of random drift (indiscriminate sampling)? It is hard to know how to answer these questions. Thus it is the probabilism of natural selection and random drift which leads to the conceptual problem that Beatty outlined; the fact that natural selection and random drift are both probabilistic seems to leave us with a conceptual overlap.

This is a problem for evolutionary biology given arguments over the relative importance of natural selection and random drift. How can natural selection and random drift be used as alternative explanations if the concepts of natural selection and random drift cannot be distinguished? Beatty suggests that some evolutionary changes are "to some extent, or in some sense, a matter of natural selection *and* to some extent, or in some sense, a matter of random drift" (Beatty 1984: 196), acknowledging that, at least in some cases there *is* no hard and fast distinction between random drift and natural selection. This is not a solution to the problem he has raised, but rather a concession to the problem. Perhaps this is the best that we can do. However, I think it is important to try to make sense of biologists' empirical claims about whether random drift or natural selection is more prevalent in given population. These claims are difficult enough to substantiate, but they are impossible without the conceptual apparatus. If Beatty is right, there will be some cases in which we cannot answer the empirical question, because our concepts don't

provide the distinction (this is different from, say, not being able to answer the question because we can't obtain the relevant information). However, if the concepts are not distinct, then it seems as though *all* evolutionary explanations which invoke natural selection and/or random drift will suffer from vagueness, not just certain fuzzy cases. As this is clearly an undesirable result, I will explore alternative solutions to the problem Beatty has raised.

1.4.2. Dissolution of the Problem

One way of "solving" the problem that Beatty has raised is to attempt dissolve it by criticizing the conception of natural selection on which Beatty's argument rests. On Beatty's account, natural selection is a probabilistic process, and as we saw above, it is that very probabilism which produces the conceptual overlap between natural selection and random drift. However, if the concept of natural selection were not a probabilistic one, there would be no overlap, and therefore no problem in making the conceptual distinction.

But what would a non-probabilistic conception of natural selection look like? One way of characterizing it would be to claim that natural selection is "deterministic" in the sense that in an infinite population, organisms will (with a probability asymptotically approaching unity) achieve their expected reproductive success (Sober 1984 puts forth such a view).[9] On this kind of account, any variation from that expected reproductive success in a real, finite population is attributable to random drift rather than natural selection (Sober 1984: 110-111). Although Sober does not make this argument, one could use his characterization of natural selection to claim that there was no conceptual overlap; any probabilism is attributed to random drift. In the case of Beatty's moth example, a Soberian would claim that the population experienced both natural selection and random drift. To the extent that the change in the population was in accordance with the expectations of natural selection, the change was due to natural selection. To the extent that the change deviated from those expectations, the change was due to random drift.

[9] Note that this sense of determinism is distinct from Laplacean determinism.

However, on this kind of view, the concepts of random drift and natural selection are distinct; natural selection is a process whereby organisms achieve their expected reproductive success (based on their fitness in a given environment), random drift is a process which probabilistically deviates the population from those expectations.

This argument is appealing in that it provides a clean distinction between natural selection and random drift. It also fits with the intuition that in most real populations, both natural selection and random drift are occurring; using Sober's conceptions of random drift and natural selection, this will be true of all real populations (since no real population is infinite).

However, excluding chance from our conception of natural selection does not seem to fit with longstanding usage of the term in biology. Darwin, in describing natural selection, asserts: "But if variations useful to any organic being do occur, assuredly individuals thus characterized will have the *best chance* of being preserved in the struggle for life" (Darwin [1859] 1964: 127). On Darwin's conception of natural selection, chance, in this case the chance that organisms' possessing useful variations may not be successful, is *part* of the concept of natural selection, not separate from it. Ernst Mayr also utilizes a probabilistic conception of natural selection, distinct from the probabilism imparted by random drift:

> Natural selection is a statistical phenomenon; it means merely that the better genotype has a 'better chance of surviving' (Darwin). A light-colored individual in a species of moth with industrial melanism may survive in a sooty area and reproduce, but its chances of doing so are far less than those of a blackish, cryptically colored individual. It happens not infrequently in nature that, for one reason or another, a superior individual fails to reproduce while an inferior one does so abundantly … Natural selection, being a statistical phenomenon, is not deterministic; its effects are not rigorously predictable (Mayr 1970: 107-8).[10]

In more recent writings, he asserts that "selection and chance are not two mutually exclusive alternatives…there are stochastic perturbations ("chance events") during every stage of the selection process" (Mayr 1983: 332).

[10] Mayr is clearly not considering random drift here; he brings random drift and other forms of drift into his account at a later point in the chapter (Mayr 1970: 120-8).

The conception of natural selection inherent in these quotes from Darwin and Mayr is one which many biologists and philosophers of biology would grant, I think: we acknowledge that natural selection is a process where organisms may not be as reproductively successful as one would expect, based upon their traits – they may "by chance" fail to achieve that success. But is random drift this chance element or is the chance element an inherent part of natural selection? Or, to shift the burden of proof to Sober, why speak of natural selection as "deterministic"? I think the reason natural selection is sometimes described in this way is to provide a relatively simple model that can be analyzed quantitatively. In the context of population genetics, it is easiest to deal with a model of natural selection where population size is infinite and expected values are realized. Then, when one "complicates" the model by introducing finite population size, random drift suddenly comes into play. (This is the model we discussed in section 1.3 above, but with fitness differences). Further "complications", such as migration and mutation, can also be added to the model. So the population models serve as step-wise descriptions of evolutionary processes, with increasing complexity. Considering natural selection under the conditions of an infinite population allows biologists to better isolate the role that *other* factors play, such dominant and recessive genes, without complicating factors "getting in the way".

However, I think it is a mistake to take the population genetics model of natural selection and apply it to the *concept* of natural selection, for the very simple reason that in the real world, there are no populations of infinite size. If concept of natural selection is a deterministic one, then it is a concept which can't be applied to the real world. In the real world, organisms don't always achieve their expected reproductive success; they come across predators, forest fires, diseases. That in itself shouldn't mean that they are any less subject to natural selection. Our concept of natural selection should be a "real world" concept, one that incorporates finite populations and the Darwinian intuition that fitness only imparts a likelihood of reproductive success, not a guarantee. In other words, in the real world, the process of natural selection is probabilistic; our concept of natural selection should reflect this reality and be probabilistic as well. Unfortunately, we are still left with

the problem of how to distinguish this probabilistic conception from random drift. So let us turn to another candidate solution to the problem, that of M. J. S. Hodge.

1.4.3. Hodge's Solution

Hodge claims that a proper distinction between natural selection and random drift relies on *causation*: "… differential reproduction in selection is distinguished from any in drift by its causation; by contrast with drift, it is occurring because the physical property differences constituting the hereditary variation that is being differentially produced are not merely correlated with differences in reproduction – they are causally relevant to them" (Hodge 1987: 251). This distinction uses essentially the same definition of random drift I developed in sections 1.4.1 and 1.4.2 above; random drift is a process where physical differences are *not* causally relevant (i.e., are causally irrelevant) to differences in reproductive success. However, Hodge doesn't seem to think that causation can completely solve the problem that Beatty raises. Citing Beatty, Hodge states:

> … explicating selection by contrast with drift allows for – even encourages – the admission that in real life the ecology and genetics of butterflies and their predators may include some changes that are equally plausibly categorized as selection or drift. Once one considers such familiar complications as linked genes, correlated responses to selection, patchy environments, frequency dependent effects, habitat preferences, and so on, it is possible to think of scenarios for which it is impossible to draw a sharp line between [differential reproduction where physical property differences are causally irrelevant and differential reproduction where physical differences are causally relevant]. But here, as always, it is to be emphasized that such impossibilities do not nullify the rationales for making the conceptual distinction (Hodge 1987: 253).

Here Hodge seems to be confusing distinguishing between natural selection and random drift *empirically* (deciding whether real life cases are natural selection and random drift) and distinguishing between natural selection and random drift *conceptually* (developing definitions of the two concepts which are non-overlapping). Without further explanation as to how the various "complications" he cites make the distinction impossible to draw, it

is difficult to determine whether he is making this mistake. In any case, I think that Hodge's distinction (with two modifications I will discuss below) is more clear-cut than he realizes – that it can be used to solve the problem that Beatty raises. However, before I can support that claim we must first examine Hodge's arguments in further detail.

Hodge notes that textbook definitions of natural selection are standardly in agreement on two counts. The first area of agreement is concerning the necessary conditions for natural selection:

1. Variation
2. Heritability of variation
3. Differential reproduction of heritable variation

The second area of agreement concerns the need for an additional, fourth condition to make the conditions jointly sufficient (Hodge 1987: 250). Where textbook definitions tend to disagree is over what the additional condition should be. A fourth condition is needed because differential reproduction of heritable variation could describe either natural selection or random drift. Thus, finding this additional condition addresses the very problem we have been discussing: how to distinguish natural selection from random drift. If a population satisfied the first three conditions, as well as the fourth condition, we would say that the population was undergoing natural selection. However, if the population satisfied the first three conditions, but failed to satisfy the fourth condition, we would say that the population was undergoing random drift. According to Hodge, we can divide most of these "extra" conditions into two categories: the statistical (which relies on terms such as 'consistent', 'systematic', or 'nonrandom') or the biological (which relies on 'fitness' or 'adaptation') (Hodge 1987: 251). Hodge rejects both of these proposed conditions in favor of causation.

As the reader might guess from the account of random drift I gave above, I am very sympathetic to this solution. Along with Hodge, I think that causality captures the essence of the distinction between natural selection and random drift, and I agree that random drift

is a physical, rather than mathematical process, making a purely statistical condition an inadequate discriminator. The processes of random drift we have been discussing, such as indiscriminate gamete and parent sampling and the founder effect, are all physical processes, and as we have seen, causality (or rather, *lack* of causal relevance) is at the root of each one. I also agree that a biological condition by itself cannot serve to distinguish between natural selection and random drift. Hodge considers a number of different definitions of adaptation and concludes that none of them compromise "the appropriateness of a physicalist, causalist explication of selection" (Hodge 1987: 257). As for fitness, defined as expected reproductive success, Hodge asserts that it is a mistake to ascribe causal properties to fitness so defined, because an expectancy is neither causal nor explanatory (Hodge 1987: 257). Thus, according to Hodge, adaptation is not necessary in defining natural selection, whereas fitness is not sufficient. However, I don't think that either of these arguments gets at the heart of the reason why a biological condition *alone* is inadequate in distinguishing between natural selection and random drift. Since many authors use this kind of criterion as a discriminator, I think it is important to understand why a biological condition alone fails to do the job. In addition, by examining the failure of the biological condition, we will see why a causal criterion is necessary.

A biological condition would be something along the lines of:

4. Variants must differ in fitness (adaptiveness).

A hypothetical example will serve to show why this condition is inadequate. Consider a population of pink and brown snails living in a region where most of the surroundings are brown. In this environment, the brown snails are fitter (and better adapted) than the pink snails, because the predator which feeds on the snails has a greater ability to discern the pink snails, giving the pink snails a lower viability. Now suppose that one summer there is a drought which causes both brown and pink snails alike to die in great numbers. Let us further suppose that more brown snails are able to find moist shelter and thus survive the drought than pink snails (and that the brownness of the brown snails

was causally unrelated to their greater ability to find shelter). As a result, the next generation of snails includes more brown snails than pink snails. However, with so few snails in the population, the difference between the number of brown snails consumed and the number of pink snails consumed is minimal.

Assuming the pink and brown colors are heritable in the snail population, this scenario fits the three necessary conditions that Hodge outlined: variation, heritability of variation, and differential reproduction of heritable variation (the brown snails were more reproductively successful than the pink snails). The scenario also satisfies the "biological" condition: there is a difference in fitness (adaptiveness) between the variants, and the variants with the greater fitness are more reproductively successful. So, on an account which takes the fitness condition to be that which distinguishes natural selection from random drift, the snail scenario would be described as an instance of natural selection.

Yet it doesn't seem to make sense to describe these events as natural selection. Yes, the brown snails are fitter than the pink snails, but that is not what led to their greater reproductive success in this instance. Rather, the reproductive success of the brown snails is due primarily to an incidental (and non-heritable) property: the location of brown snails near shelter. Intuitively, we want to say that to the extent that the *color* of the snails is not relevant to their reproductive success, it shouldn't be a factor in determining whether or not this a case of natural selection. Yet nothing about the definition (differential reproduction of heritable variants differing in fitness) instructs us to ignore the color of the snails in making this determination. Furthermore, it fails to direct us to pay attention to the location of the snails. If we do consider location, we note that it is not a heritable difference between the snails, and so to the extent that the differential reproduction of the snails is due to their location, the case should not properly be classified as natural selection.

It should be fairly evident that what the definition in question needs to direct our attention appropriately is *causation*. The definition fails precisely because it lacks this concept. However, if we expand the fourth condition from "variants must be differ in fitness", to "variants must differ in fitness, and that difference must be causally relevant to the difference in reproductive success" we will get the desired result. Using this expanded

fourth condition, we could argue that the population contains variants which differ in fitness (brown snails are fitter than pink snails) but that the color difference is *not* causally relevant to the differential reproduction. Based on this expanded fourth condition, we would argue that the change in the population was not due to natural selection. Thus, in the absence of other factors which could explain the change (such as migration) we would say that the change is due to random drift. Causation allows us to make this distinction, whereas fitness or adaptation alone failed to make the proper distinction.

However, this expanded condition, essentially the one that Hodge gives, doesn't give a completely satisfactory result either. As I have described the scenario, the color difference between the snails is largely causally irrelevant, but need it be completely so? Suppose, for example, that even though most of the brown snails survived simply because they were able to find shelter, that some small portion of them survived because their color made them difficult for the predator to spot. In other words, it does seems as though to some extent natural selection *is* occurring, although not to a very great extent. Here we have the (biologically common) situation where both random drift *and* natural selection are occurring in the population, yet Hodge's causal relevance criterion will specify only one or the other of these processes (in this case, only random drift). The concepts do not cover the phenomena correctly; there is natural selection occurring in the population, but Hodge's criterion will specify that the changes in the population are entirely due to random drift. What we need is an additional criterion, one which will let us describe to what extent each process (natural selection and random drift) is occurring in the population. In the next section, I will propose just such a criterion.

The careful reader will note that my expanded fourth condition differs somewhat from Hodge's in that it refers to differences in fitness, rather than to differences in physical properties. I won't go into any detail concerning this issue, as it would take this discussion into a long (although interesting!) tangent about the correct way to define fitness. I will say only that if we distinguish the way fitness is *measured* from what fitness *is*, we can understand fitness as expected reproductive success (this is how fitness is measured), while identifying fitness with its underlying physical properties (this is what

fitness is). If fitness is defined as the capacity of an organism to produce offspring in a particular environment, where the capacity of an organism is identified with the underlying physical properties, then fitness can play a causal role in our definition of natural selection. This definition of fitness allows us to avoid jettisoning 'fitness' from the definition of natural selection, and to retain the generalization across populations that a term like 'fitness' provides. With regard to random drift, when we wish to distinguish it from natural selection, we will similarly want to spell out random drift in terms of fitness differences. In other cases, particularly when there are no fitness differences in the population, we will wish to speak of physical differences alone (as we did in sections 1.4.1 and 1.4.2 above).

Since under this definition of "fitness" fitness is identified with its underlying physical properties, its use makes my dissimilarity with Hodge's view primarily terminological; my "differences in fitness" are essentially equivalent to Hodge's "differences in physical properties."

1.5. Millstein's Solution

I am proposing that we define natural selection and random drift as follows. To the extent that heritable fitness differences are causally relevant to differences in reproductive success, the process in question is natural selection. To the extent that the heritable fitness (or physical[11]) differences are not causally relevant to differences in reproductive success, the process in question is random drift. As a test of this proposal, we will apply these definitions to Beatty's example and see if they succeed.

Recall that Beatty's example deals with a population of dark moths and light moths, both having the same color-sensitive predator. The moths inhabit a forest containing 40% light colored trees and 60% dark colored trees, conferring a greater fitness on the dark moths in this environment. In spite of the greater fitness of the dark moths, the lighter moths have greater actual reproductive success due to the fact that the dark moths happen

[11] I make this qualification here to allow us to describe random drift in populations where there are no fitness differences, i.e., populations in which the variations are "neutral".

to land more frequently on light colored trees than dark colored trees. Beatty asks, is this a case of natural selection, or of random drift?

On the account I am proposing, it is primarily a case of natural selection. The color difference between the light and the dark moths is causally relevant to the death of the light moths; the light moths died because the predator could distinguish them more easily than the dark moths in an environment where there are greater numbers of dark trees (similarly, the dark moths died in an environment where the predator could distinguish them less easily). But how causally relevant is the color difference? There are, broadly speaking, two ways of estimating the causal relevance: mathematically and biologically. I will first explore the mathematical, and show why that fails; I will then argue for the biological measure.

One mathematical measure of causal relevance is that of causal effectiveness. According to Giere (1980), the causal effectiveness of a cause C for an effect E in a stochastic system is:

$$Ef\ (C,E) = Df.\ \mu c - \mu \neg c$$

where μc is the mean value of the probability distribution of possible outcomes of C, and μ¬c is the mean value of the probability distribution of possible outcomes of ¬C. If we just focus on viability, we can take the effect E to be survival and C to be the property of being dark. The forest is composed of 60% dark trees, so the expected value for the frequency (the mean value of the probability distribution of possible outcomes) of dark moths is 0.6. Likewise, the forest is composed of 40% not-dark (i.e., light colored) trees, so the expected value for the frequency of not-dark (i.e., light-colored) moths is 0.4. Thus, the causal effectiveness of the dark color for survival in the given population is 0.6 - 0.4 = 0.2. Since causal effectiveness for a positive causal factor (being dark is a positive causal factor since it increases the probability of survival in this population) ranges from 0 to 1, we can see that the causal effectiveness of the property of being dark is not very high. Thus, we can say that the population is subject to natural selection, but not to any large

extent. This would suggest that (in the absence of other processes operating in the population) that random drift is predominant.

Since the causal effectiveness is calculated solely from fitness values (0.6 for the dark moths and 0.4 for the light moths), we can see that we obtained the result that there is strong drift solely from the knowledge that there was weak selection. Likewise, if there had been strong selection (say 0.9 dark moths and 0.1 light moths), we would have gotten the result that there was weak drift (with a causal effectiveness of 0.7, selection would be strong and random drift relatively weak). This, of course, echoes the mathematical criterion for assessing the relative roles of natural selection and random drift standardly found in population genetics textbooks:

> If $4N_e s >> 1$, then natural selection is predominant.
> If $4N_e s << 1$, then random drift is predominant.

where N_e is the effective population size and s is the selection coefficient. Beatty does not give the effective population size for his example; however, the value of s would be the relative fitness of the dark moths minus the relative fitness of the light moths, or 1 - 2/3 = 1/3. Without knowing the population size, we cannot make the determination, but clearly with a small enough effective population size, we would get the same result that we got with causal effectiveness: random drift has a stronger influence than natural selection. And, as with causal effectiveness, with a higher selection coefficient, you would expect that selection has a stronger influence than random drift.

However, both the causal effectiveness measure and the population genetics measure will sometimes produce incorrect results, for reasons similar to those already discussed in section 1.4.3 above. Since both these measures rely on fitness without taking the causal story into account, they run the risk of overestimating (or underestimating) the actual causal contribution of fitness. Suppose you had a population with large fitness differences (high selection coefficient) where the fitness was causally relevant, but not very much so; say most of the organisms died in a flood rather than as a result of inferior

fitness. In this case, you would want to say that even though the selection coefficient was high, random drift predominated. So, the mathematical calculations overestimate the causal relevance of fitness and give you the wrong determination: strong selection. On the other hand, suppose you had a population where the selection coefficient was low, but that the reproductive success of the organisms in the population very strongly reflected what you would expect from the fitnesses. In this case, you would want to say that even though the selection coefficient was low, natural selection had a stronger influence than random drift. Here the mathematical calculations underestimate the causal relevance of fitness and say that random drift predominates over natural selection, when it should be the other way around.

Since mathematical solutions to the question of how to determine the strength of causal relevance don't seem to work, let us turn to a biological solution. Suppose we knew exactly what was going on in a population: which organisms lived and reproduced, which organisms died, and why. We would then know which organisms were reproductively successful (or not) as the result of their fitness traits, and for which organisms fitness was not causally responsible for their success (or lack thereof). Unfortunately, even this (completely unrealistic) point of view will not help us very much in trying to determine how to determine the relative roles of natural selection and random drift; as we have already seen, if we take the probabilism of natural selection seriously, it will sometimes be the case that, as part of the selection process, the fittest organisms will not always be the most reproductively successful. We can't say that all of the organisms whose fitness did not cause their reproductive success were subject to random drift, because then it turns out that natural selection isn't probabilistic. So, what we need is another criterion which would sort the cases of natural selection and the cases of random drift from instances where fitness differences are not the cause of differences in reproductive success.[12]

One such criterion is that of *rarity*. As Beatty has emphasized, fitness must be

[12] Just to clarify: whereas my definition uses the concept of causal relevance (*C* is causally relevant if it increases the probability of the effect), here we are considering a different issue – what the actual causes in the population are.

specified relative to a particular environment (Mills and Beatty 1979, Beatty 1984), since while phenotype *A* might be fitter than phenotype *B* in environment *E1*, phenotype *B* might be fitter than phenotype *A* in environment *E2*. However, specifying an organism's environment is a difficult, and less than straightforward task.[13] How does one draw the physical boundaries of the environment? What are the time boundaries of the environment (given that most environments change over time)? Assuming you can give reasonable answers to these two questions, which characteristics of the environment should be included? Do you include rare events (the impact of a meteor, a one-time flood) or do you just include common events (prevalent predators, common diseases, yearly floods)?

It is this last question that concerns us here. Biologists do often distinguish between events which commonly occur in an environment and events which rarely occur. In some cases, this allows biologists to further refine their models of evolutionary change to account for the differing prevalence of events. In other cases, it lets them focus on what sort of change will most likely be occurring, given the events that are common in the environment. And sometimes, the distinction is simply a practical one, since some events, such as the impact of a meteor, may be completely unforeseen, so that it would be impossible to take such events into account even if biologists wished to do so. Biologist Jonathan Roughgarden makes just such a distinction, between "regular" and "unusual" conditions:

> It is easy to observe that the environment of virtually every population changes from time to time. Presumably, if the environment changes, then the selective values of the various genotypes also change. The presence of variation in the environment becomes particularly interesting if the selective values actually change *direction* from generation to generation. For example, most places have an unusual type of weather every now and then. A late frost may strike, or severe rains and flooding, and so forth. Presumably, there are genes that are favored during these unusual conditions that are not favored during regular conditions (Roughgarden 1996: 259; italics in original).

In the context of random drift, this distinction was used by Wright to characterize one of the kinds of random drift: "I have held that random processes (sampling drift, *unique*

[13] See Brandon 1990 for one solution to this problem.

events, fluctuations in the coefficients of selection and immigration and even systematic local differences in selection) bring about differences among local populations which are random with respect to the general course of selection" (Wright to Kimura, June 2, 1971; quoted in Provine 1986; emphasis added). So, on Wright's account, random drift included not only sampling drift (what we have been calling indiscriminate parent and gamete sampling) and fluctuations in selection and migration, but also *unique events,* i.e., "events that are unique or nearly so in the span of history under consideration" (Wright 1955: 18) such as "swamping by mass immigration" or "unique reduction in numbers" (Wright 1955: 19).

What I am suggesting is that we use this distinction between rare and common events to help sort reproductive success that is not caused by fitness into "reproductive success due to natural selection" and "reproductive success due to random drift", but only once we have made the initial determination that fitness differences are causally relevant. In Beatty's example, as we have already noted, the color difference between the light and the dark moths is causally relevant to the death of the light moths; the light moths died because the predator could distinguish them more easily than the dark moths in an environment where there are greater numbers of dark trees. So, there is natural selection occurring in the population. Then, to determine the *extent* of the natural selection, we want to consider the events which were not caused by fitness; should we attribute them to natural selection or random drift? Under Beatty's scenario, the reason more dark moths died than light moths was that "dark moths chanced to land on light trees more frequently than on dark trees, even though the frequency of dark trees was greater" (Beatty 1984: 195). This, I would maintain, is no uncommon or unique environmental event, such as the impact of a meteor or a once-in-one-hundred-years flood; this is a regularly occurring environmental event concerning the patterns of moth flight. Therefore, I would argue, using the biological criterion I have put forth, that the death of the moths is, to a large extent, due to natural selection.

Now this might seem like an odd result for a population where dark moths are fitter than light moths, yet light moths are reproductively successful. However, it should not be

too surprising a result if we *take seriously the proposition that natural selection is probabilistic*. If the fitness of an organism is represented by a probability distribution of possible outcomes, then certain outcomes are more likely than other outcomes, but why should the *less likely* outcomes be any less natural selection than the *more likely* outcomes? All are predicted by natural selection; it's just that the predictions are of differing strengths. More tellingly, consider the following: let's suppose that the population initially consisted of 50% light moths and 50% dark moths, but in the next generation (after dark moths unluckily land on too many light trees, and the surviving moths reproduce in equal numbers) the dark moths are 45% of the population and light moths are 55%. However, since the most likely result (the mean or expected value) is 60% dark moths and 40% light moths, and since the probability distribution is bell-shaped, the outcome of 55% dark moths and 45% light moths is *has the same probability* as the outcome of 45% dark moths and 55% light moths. Yet, in this case (55% dark moths and 45% light moths), where the fitter organisms *do* have the greater reproductive success, there is less difficulty in seeing that natural selection is the operative process. However, since both results are equally likely, there is no reason to consider the one outcome (55% dark moths and 45% light moths) to be natural selection, and the other outcome (45% dark moths and 55% light moths) not to be.

There is a further aspect to consider in evaluating the result that a population in which the fittest are not the most reproductively successful can be a population undergoing natural selection. We are only looking at a change in the population from one generation of moths to another. Natural selection, as Darwin conceived it, is a process which occurs "in the course of many generations" (Darwin [1859] 1964: 114). In the short run, the unlikely can occur (the fittest organisms may be exposed to a disease, or encounter more predators), so that in the occasional generation, the fittest may not prevail. However, in the long run, we would expect the fitter organisms to be more successful. When we consider what would most likely happen to Beatty's moth population in succeeding generations, the claim that what has occurred in the generation in question is (to a large extent) natural selection is more intuitive.

Let us now consider how my proposed definition would apply to the scenario discussed in section 1.4.3. In that population, pink and brown snails live in a region where most of the surroundings are brown. As a result of their ability to camouflage and avoid the predator, the brown snails are fitter than the pink snails. However, one summer there is a drought which causes both brown and pink snails alike to die in great numbers. More brown snails are able to find moist shelter and thus survive the drought than pink snails (and that the brownness of the brown snails was causally unrelated to their greater ability to find shelter). As a result, the next generation of snails includes more brown snails than pink snails. With so few snails in the population, there is little difference between the number of brown snails consumed and the number of pink snails consumed.

As we discussed above, if it were the case that some small portion of the brown snails survived because their color made them difficult for the predator to spot, the fitness difference between the snails would be causally relevant to the difference and reproductive success, and so, we would want to say that to some extent, the population was undergoing natural selection. But again, the question is, *to what extent*? Let us suppose that the drought that the snails are subjected to is a rare occurrence in the population. Applying the biological criterion, we would determine the reproductive success which was not caused by fitness differences to be largely random drift. Since most of the difference in reproductive success was not due to differences in fitness, we would say that the random drift predominated in the population, although natural selection occurred to some (small) extent. (It follows, of course, that if the drought were a common, rather than a rare, occurrence in the environment, we would consider that it was natural selection rather than random drift that predominated.)

These two scenarios should serve to illustrate how my proposed definition would function in allowing us to distinguish natural selection from random drift in populations of

interest.[14] However, I wish to clarify that in populations where there are no fitness differences between the entities in question (i.e., if the traits are neutral), then (in the absence of mutation and migration) changes in the population are due entirely to random drift. Natural selection requires that variants in a population differ in fitness; without fitness differences, natural selection has nothing upon which to act. Conversely, in populations where fitness differences are very strong, in the absence of evidence to the contrary, we would expect that changes in the population would be due primarily to natural selection.[15] So, it is not always difficult to know what to label a particular instance of evolutionary change; there are clear-cut cases of both natural selection and random drift.

1.6. Conclusion

The concept of random drift is a ubiquitous element in present day accounts of evolutionary theory. We have seen that a complete characterization of random drift is a three-fold process, involving a description of the physical processes of random drift (such as indiscriminate gamete sampling and the founder effect), an analogy to random sampling from an urn of colored balls, and a definition that represents the key elements of the physical processes and the sampling model and provides a distinction between the concepts

[14] Again, I wish to emphasize here that I am not proposing a way of empirically distinguishing random drift from natural selection in a given population; rather, I am proposing that if we define random drift and natural selection in the way that I have specified, they will be distinct enough so that we have the ability say to what extent a population is undergoing one process or the other. Determining whether fitness is causally relevant to differential reproduction is in many populations a difficult task – but how to do so is part of the task of distinguishing between natural selection and random drift empirically. Furthermore, what we *know* about a population is independent of what *occurs* in the population. As Hodge notes: "For a process to be [one where physical differences are causally relevant] rather than [one where physical differences are causally irrelevant], in the requisite sense, is for it to be so quite independently of how far we have succeeded in observing the causal bases for the causal relevance" (Hodge 1987: 254). So, the fact that it is difficult to determine whether fitness differences are causally relevant to differences in reproductive success should not preclude us from using such a criterion for our definition. Rather, the criterion should serve to focus, and perhaps provide an ideal for, our empirical investigations.

[15] Note that there is an asymmetry between random drift and natural selection. If traits in a population are entirely neutral (no fitness differences) then the population will undergo random drift in the complete absence of selection. However, since all real populations are finite, we cannot entirely rule out the possibility that random drift may be occurring, even in populations where fitness differences are large (although it is likely to be the case that in large populations, the effects of random drift are negligible).

of random drift and natural selection. While the first two of these tasks are relatively straightforward, the third, unfortunately, is not. In spite of the importance of a clear cut distinction for understanding and resolving debates over the prevalence of random drift in evolution, it should be evident from the above discussion that providing such a definition is far from trivial.

I have offered definitions of random drift and natural selection which rely on causality to make the distinction between natural selection and random drift: natural selection is a process where heritable fitness differences between organisms are causally relevant to differences in reproductive success, whereas random drift is a process where heritable fitness differences (or heritable physical differences, in the absence of fitness differences) are causally irrelevant to differences in reproductive success. The extent to which a process should be considered to be natural selection is the extent to which fitness differences are causally relevant to reproductive success, where "extent of causal relevance" is determined by the rarity of the cause within the specified environment. To the extent that a process is not natural selection, it is (in the absence of other processes such as mutation and migration) random drift.

These definitions are not entirely satisfactory, for three reasons. First, although they do distinguish between natural selection and random drift, because they rely on a fuzzy distinction between rare and common environmental events, the distinction that they provide is somewhat fuzzy as well. Second, although they give a way to qualitatively say which is more prevalent, random drift or natural selection, they do not give a way to quantify the degree to which one or the other process is prevalent. Third, there is a possibility that the definitions favor natural selection over random drift. Some biologists may want to argue that common as well as rare environmental causes can lead to random drift.

However, the definitions I have proffered have a number of advantages. They permit us to keep a realistic, probabilistic conception of natural selection, while providing a way of differentiating the probabilism of natural selection from the probabilism of random drift. They are rich enough to allow for the commonly held view that both natural selection and

random drift often occur simultaneously in the same population. They respect various views concerning random drift, furnishing us with a way to understand the view that random drift is an indiscriminate sampling process (in that differences between variants being sampled are causally irrelevant to which variants are sampled) as well as the view that random drift events are unique events. And by making causality prior to rarity, the definition of random drift avoids labeling rare, yet selective processes as "random drift". (In Chapter 2, we will examine a case where Rosenberg uses the concept of drift in this confusing way). My hope is that despite their deficiencies, the definitions, in providing a way to distinguish and weigh between natural selection and random drift, can serve as useful tools for both biologists and philosophers alike.

Chapter 2: Random Drift, Determinism, and Indeterminism

2.1. Introduction

Present-day biologists and philosophers describe Darwin's theory of natural selection as a *deterministic* theory, although it was referred to as the "law of higgledy-piggledy" during its own time. There are at least two reasons for this deterministic interpretation of Darwin's views. One is Darwin's assertion that by "chance" he meant only to acknowledge "our ignorance of the cause of each particular variation" (Darwin [1859] 1964: 131).[16] A second reason Darwin's theory is viewed as deterministic is that it lacks the concept of random drift, commonly held to be the most important source of stochasticity within twentieth century evolutionary theory.[17] Yet Alexander Rosenberg claims that the presence of random drift in evolutionary explanations is itself merely an "admission of ignorance" (1994: 82), thus applying this Darwinian concept of chance to random drift. In so doing, Rosenberg extends the deterministic interpretation of Darwin's evolutionary theory to encompass current evolutionary theory.

Random drift is a process in which physical differences between organisms are causally irrelevant to their survival and reproductive success, in contrast to natural selection, a process where physical differences between organisms *are* causally relevant to their survival and reproductive success (see Chapter 1). Rosenberg claims that an omniscient account of evolution would have no need for the concept of random drift – that all instances of random drift can be explained in terms of natural selection. Rosenberg uses this claim to argue that although evolutionary *theory* is statistical, the evolutionary *process* is a deterministic one. According to Rosenberg, evolutionary theory is statistical purely for instrumental reasons; random drift serves merely as a useful fiction. A similar claim for

[16] Darwin also used the term 'chance' in other ways, e.g., to mean "not designed," to contrast with the theological arguments of his day.

[17] This statement is not meant to imply anything about the relative roles of random drift and natural selection in evolutionary theory. The English tradition, in particular, attributes a very small role to random drift (see Turner 1987 for discussion). Thanks to Michael Ruse for pointing out my American ethnocentrism.

the determinism of the evolutionary process is made by Barbara Horan (1994). Brandon and Carson (1996) challenge Rosenberg's and Horan's claims; instead, they maintain that a scientific realist should conclude that the evolutionary process is fundamentally indeterministic.

This chapter will first address the issue, raised by Rosenberg, of the necessity of drift from the omniscient viewpoint. I argue, contra Rosenberg, that any evolutionary theory, omniscient or otherwise, must take random drift into account. I will begin by examining Rosenberg's arguments for the claim that random drift is eliminable from an omniscient viewpoint, and show that the concept of random drift upon which they rest is an impoverished one. I will then provide three reasons why random drift is not eliminable from an omniscient account of evolution.

The second issue this chapter addresses concerns the (in)deterministic character of the evolutionary process; is it deterministic, as Rosenberg and Horan argue, or indeterministic, as Brandon and Carson maintain? I will argue that a more philosophically defensible position argues neither for the fundamental determinacy nor indeterminacy of the evolutionary process. However, even without making these kinds of empirical claims concerning the determinism or indeterminism of the evolutionary process, we can still make arguments concerning the probabilistic character of evolution. That is, it remains an open question as to whether evolution is inherently and unavoidably probabilistic. Brandon and Carson (1994), as well as Sober (1984), maintain that even if evolution is deterministic at the individual level, population-level probabilistic explanations are necessary for providing complete evolutionary explanations. Although I am essentially in sympathy with these arguments, I don't think they make their case as strongly as they might. I seek to show that even if one assumes that the evolutionary process is fundamentally deterministic, the status of natural selection and random drift as population-level processes implies that evolutionary theory is inherently and unavoidably probabilistic.

2.2. Rosenberg's Argument

Rosenberg (1994) constructs a fictional situation in which poachers have removed the longest-necked members from a small giraffe population and transported them to zoos, where they fail to reproduce. For several generations afterward, the short-necked giraffes predominate in the population. Biologists are given the task of determining why the population shifted towards short-necked giraffes even though long-necked giraffes are fitter than short-necked giraffes in this environment. Since the biologists don't know about the poachers, the explanation they would provide is a non-omniscient one. According to Rosenberg, their non-omniscient explanation would invoke random drift, since to do otherwise would be a "disconfirmation of the theory of natural selection" (1994: 71).

However, since *we* know the reason for the increase in short-necked giraffes, we *can* provide an explanation from the omniscient viewpoint. Would *we* attribute the change in the giraffe population to random drift? Rosenberg's answer is no; rather, "we, who know the facts … will say that for a short time the environment changed, making long-necks maladaptive and therefore shifting gene frequencies through selection" (1994: 73). Rosenberg thus reinterprets a case of random drift as a case of selection by switching from a non-omniscient viewpoint to an omniscient one. Random drift is eliminated from the explanation, eliminating the element of stochasticity along with it (assuming, of course, that natural selection is itself non-stochastic). Generalizing from this particular example, Rosenberg concludes that, "from a position of omniscience, there is no need for the notion of drift" (1994: 73). Elaborating, he suggests that "a deterministic theory in which drift plays no role" would either be "so generic in its claims as to have little predictive content" or extremely "detailed in its enumeration of selective forces" (1994: 83).

According to Rosenberg, because we are non-omniscient, finite creatures, this "detailed" theory is "beyond our powers to discover and express" and is "hopelessly unwieldy" (1994: 83). Rosenberg argues that since the omniscient account is beyond our reach, we *do* need probabilities and the concept of random drift in our evolutionary theory.

However, Rosenberg maintains that it is solely for reasons of utility that random drift appears in our evolutionary accounts.

Rosenberg's claim that random drift is just a way of referring to selective forces of which we (as non-omniscient beings) are unaware may strike some as extreme selectionism. However, it is important to note that this claim places Rosenberg squarely in the tradition of such prominent biologists as Ernst Mayr and Arthur J. Cain. For example, Mayr asserts that only after all attempts to explain biological processes through natural selection have failed, is a biologist "justified in designating the unexplained residue tentatively as a product of chance" (1983: 326). Even more strongly, Cain asserts that the real basis for every random drift hypothesis is that "[t]he investigator finds that he, personally, cannot see any [evidence of selection], and concludes that, therefore, there is none" (1951: 424). The reason some biologists have come to hold these selectionist views is that the accepted explanations for a number of phenomena, such as the frequencies of chromosome shapes on Drosophila, changed from random drift to natural selection (Gigerenzer et al. 1989: 156-157). It would seem, however, that while Rosenberg shares the selectionist views of these biologists, he offers additional reasons for adopting them. It is to these reasons, and the conception of random drift that lies beneath them, that we now turn.

2.3. Rosenberg's Random Drift

The poacher scenario is one of two examples of random drift upon which Rosenberg's arguments rest. Both of these examples prove to be problematic. They are arguably cases of natural selection, not random drift, and they fail to reflect all the different kinds of random drift. I will discuss each example in turn, explore its implications for Rosenberg's arguments, and elaborate a fuller conception of random drift.

In the first example, discussed above, poachers remove the longest-necked members of a giraffe population and take them to zoos, where the giraffes fail to reproduce. Over the next several generations, the short-necked giraffes come to predominate in the small

population. In what sense is this a case of random drift? Rosenberg seems to be suggesting that non-omniscient biologists would *consider* the change in the giraffe population to be random drift because 1) long-necked giraffes are normally fitter than short-necked giraffes in this environment; 2) the fittest organisms, the long-necked giraffes, were not the most reproductively successful; and 3) the biologists don't know the reason the long-necked giraffes were less successful than the short-necked giraffes. Since the generational change in the population does not *appear* to be due to the relative fitnesses of the giraffes in their environment, but rather due to chance, Rosenberg can claim that biologists would consider random drift as the explanation.

But in actuality (despite appearances to non-omniscient biologists), the change in the population *is* due to the differential fitnesses of the giraffes. The selective agent in this case is the poachers themselves; they have specifically picked out the longest-necked giraffes in the population for removal. This is clearly a case of artificial selection. In essence it is no different from natural selection; the organisms were sampled on the basis of their physical traits. Since neck length determined the giraffes' future reproductive success, the poacher scenario is an instance of selection, not random drift. Rosenberg admits as much when he asserts that "we, who know all the facts" will attribute the change in gene frequencies to selection (1994: 73).

Thus, Rosenberg is examining hypotheses of *random drift* that biologists would offer as explanations of a clear-cut case of *selection*. This is problematic in itself. The real problem, however, lies in the conclusions that he draws from the poacher example. Rosenberg asserts that the omniscient viewpoint would not need the concept of drift in this example; generalizing, he asserts that random drift is not necessary for the omniscient viewpoint at all. That the omniscient viewpoint would have no need for drift in the poacher case should not be a surprising result, however, since the example was a case of selection to begin with. If you start with a case of non-drift and conclude that, knowing all the facts, you would not need the concept of drift to describe the case, you haven't proved much. You certainly have not laid the basis for a broader generalization. Rosenberg needs to consider a clear-cut case of *random drift*, rather than a clear-cut case of *selection*. Then,

and only then, can conclusions be made about the necessity of drift from the omniscient viewpoint.

The only other specific example which Rosenberg provides of random drift is also arguably an instance of natural selection. Rosenberg suggests that biologists, in ignorance of the poachers' actions, might hypothesize that some singular conditions intervened to change the proportion of long-necked giraffes within the population. As an example of such a hypothesis, Rosenberg asks us to imagine that "through freakish, never-to-be-repeated wind conditions, the two tallest trees … became so twisted that they accidentally trapped the heads of most of the tallest giraffes and broke their necks" (1994: 72-3). This example is arguably a case of natural selection, not random drift, since the difference in neck lengths is causally relevant to the giraffes' survival. For this to be a clear-cut case of random drift, long-necked and short-necked giraffes would have to be killed indiscriminately. Having a long neck would not be a *cause* of the giraffes' failure to survive to reproduce, but would instead be irrelevant.

Since the action of the wind in this example is a singular, unique event, there is some basis for Rosenberg's claim that this is a case of random drift. As I noted in Chapter 1, Sewall Wright maintained that unique events were one kind of random drift, and indeed, I used the criterion of uniqueness to help differentiate between natural selection and random drift (although it was not the primary criterion). However, since the singular action is one that so clearly produces a difference in fitness between long and short-necked giraffes, to call it random drift can only serve to muddle the distinction between natural selection and random drift. At best, Rosenberg's example is ambiguous; it is hard to know whether to classify it as natural selection or random drift. At worst, it is simply mistaken; it seems likely that most biologists would label this a case of natural selection, not of random drift. As with the first example we discussed, Rosenberg has used an example that is Phanerozoic more clearly natural selection than random drift, casting doubt on any conclusions concerning random drift he draws from it.

Furthermore, even if the wind example is an instance of random drift, it represents only one kind of random drift. The kind of random drift represented by the action of wind

conditions in Rosenberg's example is probably best characterized as parent sampling. As we discussed in Chapter 1, parent sampling, as defined by Beatty, is the process of determining which organisms of one generation will be parents of the next, and how many offspring each parent will have. In Rosenberg's example, the action of the wind determines which giraffes will survive to become parents, and so it is a case of parent sampling. However, on the account that we presented there, Rosenberg's example would not be considered to be random drift, because it is a case of discriminate parent sampling, rather than indiscriminate parent sampling; the difference in neck length is causally relevant to the giraffes' survival. For it to be an instance of indiscriminate parent sampling, it would have to be the case that the difference in neck length was causally irrelevant to the difference in survival rates. However, let us grant Rosenberg's claim that his example does describe a case of random drift. If so, then he has only described one kind of random drift, the kind having to do with parent sampling. As we saw in Chapter 1, there are potentially at least 6 other kinds of random drift: indiscriminate gamete sampling, the founder effect, the bottleneck effect, the Brandon effect, indiscriminate gene sampling, and fluctuations in selection pressure, mutation rate, and migration rate. The existence of these other kinds of random drift points to a further deficiency within Rosenberg's account; namely, that he has failed to discuss and take into account the different kinds of random drift in his arguments.

Consider, for example, indiscriminate gamete sampling and the founder effect. Gamete sampling occurs during the process where gametes are joined to form zygotes. *Indiscriminate* gamete sampling occurs when physical differences between gametes are causally irrelevant to differences in fertilization success (see Chapter 1 for further discussion). So, indiscriminate gamete sampling operates on an entirely different level than indiscriminate parent sampling; it concerns gametes rather than organisms. The founder effect, on the other hand, occurs when a small population becomes spatially isolated from the rest of the population (see Chapter 1 for further discussion). Like parent sampling (and unlike indiscriminate gamete sampling), the founder effect operates on the level of the organism. However, they concern different sorts of processes: parent

sampling concerns issues of viability and reproductive success, whereas the founder effect concerns events such as the migration of small groups of individuals, or the splitting of a group due to a physical barrier. Thus, even though parent sampling[18] , indiscriminate gamete sampling and the founder effect are all instances random drift, each is a distinct phenomenon; they concern different organismic levels and describe different kinds of processes.

The minimal point I am trying to make here is that Rosenberg's conclusions about random drift are (at best) based on just one kind of random drift; it is unclear how they would work when one takes into consideration other kinds of random drift. Since the different kinds of random drift are different kinds of processes, it is not obvious that Rosenberg's arguments would extrapolate to other kinds of random drift. Can he develop an example similar to the giraffe example, showing that an instance of indiscriminate gamete sampling can be described in terms of natural selection? Can he show that gamete sampling, discriminate or indiscriminate, is deterministic? Without answering these questions (and other similar questions), Rosenberg's arguments for the determinism of the evolutionary process are incomplete. However, there are further implications to the existence of different kinds of random drift processes. In section 2.4.1, I will show how the existence of different kinds serves to argue against Rosenberg's claim that from an omniscient viewpoint the concept of random drift is unnecessary.

2.4. The Non-Eliminability of Random Drift

We have seen that Rosenberg's account uses an impoverished notion of drift in that it is limited to random drift as parent sampling only. Moreover, his conception is a disputable notion of drift, including cases more commonly classified as selection. Thus, to justify his claim that random drift is eliminable from an omniscient viewpoint, Rosenberg's arguments must use clear-cut cases of random drift, and they should hold for different kinds of random drift. In this section I will argue that such a demonstration is not

[18] Again, note that on my account, only indiscriminate parent sampling is random drift; here we are granting Rosenberg's claim that discriminate parent sampling can be random drift as well.

forthcoming – that random drift is not eliminable from evolutionary theory, even if one "knows all the facts." I offer three arguments in support of this position: 1) That biologists distinguish between different kinds of random drift suggests that random drift is not merely a cover for processes of which we are ignorant; 2) Instances of random drift cannot be spelled out in terms of selection; and 3) Random drift provides the means to explain certain evolutionary changes that cannot in principle be explained by natural selection. I will discuss each of these in turn.

2.4.1. The Importance of Different Kinds of Random Drift

I mentioned previously (above, and in Chapter 1) that biologists differentiate between a number of different kinds of random drift. For example, Cain and Currey emphasize the importance of distinguishing random fluctuations in selection, mutation, and migration from indiscriminate parent and gamete sampling (Cain and Currey 1963: 2). They also agree with Dobzhansky's (1959) claim that parent and gamete sampling should be distinguished from the founder effect. So not only do biologists consider a number of different processes to be random drift, they think that it is important to retain the distinctions between them. Dobzhansky emphasizes that while random drift phenomena are "closely related" they are "different enough to make some distinction useful" (1959: 85). In other words, different kinds of random drift account for different kinds of phenomena. Furthermore, as Dobzhansky suggests, one kind may be more prevalent in one population, another kind more prevalent in another population. Distinguishing between the different kinds of random drift gives biologists the means to explain the differences between the evolutionary processes occurring in different populations.

For example, suppose there exists a population of snails, half of which are pink and half of which are yellow. Let us further assume that yellow snails are twice as fit as pink snails, due to their greater tolerance to heat from the sun. The snails are observed for a generation, and it is discovered that the proportion of yellow snails in the subsequent generation has actually *decreased*, in spite of their superior fitness. In the absence of

random drift we would expect that the proportion of yellow snails would increase, not decrease; that the proportion of yellow snails would decrease is an unlikely result of natural selection alone. A reasonable explanation is that random drift caused the change in gene frequencies. However, this explanation in itself doesn't tell us very much; a deeper explanation would tell us *why* this change came about. For example, further study might determine that this particular summer there were fewer shade-providing bushes than normal, causing snails of both colors to die in large numbers – and it just so happened that a greater number of yellow snails were unable to find shade. Another possibility is that the original population has been split, leaving an unrepresentative sample of snails. Either of these scenarios might lead to a decrease in the proportion of yellow snails. However, in the first case, the change in gene frequencies would be due to indiscriminate parent sampling; in the second case, it would be due to the founder effect. Thus, distinguishing between the different kinds of random drift provides the tools to isolate and further specify the different kinds of processes that may be occurring in a population.

Yet the practice of distinguishing between different kinds of drift is puzzling on Rosenberg's account. According to Rosenberg, all drift explanations can, in principle, be put in terms of selection explanations; the problem is that non-omniscient biologists are unaware of the real selective events, and so must postulate random drift. Random drift, on Rosenberg's account, thus acts as a cover for our ignorance. However, biological practice belies this claim. Rather than being a cover for our ignorance, random drift is a general term for a number of distinct biological processes such as indiscriminate parent sampling and the founder effect. In practice, it may be difficult to determine which process has predominated. But that is not to imply that such a task is impossible, nor does it imply that biologists cite random drift when they are unaware of the real evolutionary processes. On the contrary, biologists seek to determine which of the kinds of random drift has occurred to in order to support and further expand upon their claims that random drift has been prevalent in a population. Since random drift is used to refer to a number of distinct biological processes, rather than being a cover for our ignorance, drift will not be eliminable from an omniscient account of the evolutionary process as Rosenberg claims.

Even an omniscient being would need to take into account the different processes leading to evolutionary change which the term 'random drift' represents.

2.4.2. Random Drift Is Not Selection

My second argument is a denial of Rosenberg's claim that instances of random drift can be spelled out in terms of selection. According to Rosenberg, if one were omniscient, one could give an account of the evolutionary process that was so detailed in its enumeration of selective forces that one would have no need for the concept of drift. Rosenberg illustrates this claim with the poacher scenario; he suggests that we who "know all the facts" would say that the changes in the giraffe population were due to selection, rather than random drift. Yet we saw that the poacher example was *not* a case of random drift, but of natural selection. If we examine a case of random drift, however, we will find that it cannot be replaced by an enumeration of selective processes.

Consider, for example, the process of indiscriminate parent sampling. If all the variants in a population have an equal chance of success – if, in other words, they are equal in fitness – there may still be a change in gene frequencies. The frequencies of offspring may by chance fail to reflect the frequencies in the parent generation. This is random drift, since the physical differences between organisms are causally unrelated to their survival and reproductive success. However, without fitness differences, selection cannot occur, since there is nothing upon which to select. Therefore, in populations where there are no fitness differences, it is not possible to interpret indiscriminate parent sampling in terms of selection. Indiscriminate parent sampling in such populations is a phenomenon wholly distinct from natural selection. It is not eliminable from evolutionary explanations; even if we were omniscient, we would still see it as random drift. Although I will not argue the point here, I maintain that the other kinds of random drift likewise cannot be accounted for in terms of selective processes. Thus the omniscient viewpoint needs the concept of random drift; it cannot be supplanted by natural selection.

2.4.3. The Indispensable Explanatory Role of Random Drift

If natural selection and random drift are distinct processes, as I have argued, it follows that there are phenomena that can be explained as random drift, but not as natural selection.[19] Yet if random drift can explain what selection cannot, then an omniscient being would not be able to explain all instances of evolutionary change without the concept of random drift. In other words, the being's arsenal of explanatory tools would be incomplete. Thus the concept of drift is a necessary part of evolutionary theory, even from an omniscient viewpoint; it provides explanations where selection cannot.

In response to these claims, Rosenberg might argue that an omniscient being would know all the details of any evolutionary event. Thus, even if the being could not explain a particular phenomenon in terms of *selection*, it could explain it by an elaboration of the various processes (the reproductive success of individual organisms, or even the fates of individual gametes) that constitute the indiscriminate sampling event. This would suggest that random drift *is* eliminable from the omniscient viewpoint.

Yet if random drift is eliminable in this sense, then so is natural selection. An omniscient being could just as easily discuss the processes that make up natural selection, through an elaboration of the fates of individual organisms (as Horan 1994 suggests). It would follow that natural selection is eliminable from an omniscient account of evolutionary theory, making its use purely instrumental. However, such a conclusion would be devastating for Rosenberg's overall argument for the instrumentalism of biology, as Rosenberg notes:

> … the argument that biology is an instrumental science rests mainly on the fact that biological processes are the result of the operation of random variation and natural selection. If our commitment to the theory of natural selection is itself in part heuristic, if the theory of evolution is itself a useful fiction for creatures no smarter than we are, then the linchpin of my argument is pulled away. After all, my strategy is to provide realist arguments for biological instrumentalism. I cannot afford to include as a premise a theory that itself must be interpreted instrumentally, lest the argument beg the question of why biology is an instrumental science

[19] This claim assumes that random drift is an explanatory concept, a point that will be argued for in Chapter 3.

(Rosenberg 1994: 16-7).

Thus, the argument that random drift can be replaced with an accounting of the fates of individual organisms is not one that is open to Rosenberg, because if random drift is eliminable in this way, then so is natural selection, "pulling away the linchpin" to Rosenberg's program. In order to keep that linchpin, Rosenberg must acknowledge that random drift plays as essential a role in evolutionary theory as natural selection does.

However, there is a deeper reason why random drift cannot be supplanted with an accounting of the fates of individual organisms (or gametes, or genes). Evolution is a population-level process, not an individual-level process. By focusing on the individual level, the omniscient being would miss the processes that occur at the population level. This will be discussed further below.

2.5. Determinism Versus Indeterminism

Rosenberg uses the claim that instances of random drift can be replaced with instances of natural selection to argue that the evolutionary process is a deterministic one. (Of course, this argument will only work if one assumes that natural selection is itself deterministic). However, as we have seen, random drift cannot be eliminated in this fashion. So, if Rosenberg is to argue for the determinism of the evolutionary process, he will have to do so in another way. But if Brandon and Carson (1996) are right, the evidence points to an indeterministic evolutionary process, not a deterministic one. In this section, I shall argue that the evidence is not strong enough to support either conclusion.

Brandon and Carson's claims for the indeterminism of evolutionary biology rest on two arguments: the "percolation" argument, and an argument from experiments on cloned organisms. I will examine each of these arguments, and argue that they are not sufficient to settle the disagreement between the determinist and the indeterminist.

The percolation argument amounts to the claim that the quantum mechanical indeterminism of the micro-level can "percolate up" to the macro-level described by

evolutionary biology (Sober 1984, Brandon and Carson 1996). As formulated by Sober, the percolation argument asserts:

> If *enough* elementary particles had behaved differently, the behavior of the macro-object (the organism, the population) that they compose would have also been different. And there is no deterministic guarantee that the ensemble of particles *must* have behaved the way it did. The most that the ensembles of particles we call organisms can do is exhibit an impressive degree of predictability. But, so long as they are made of particles that have an irreducible chance component in their behavior, they too must be indeterministic systems. If chance is real at the micro-level, it must be real at the macro-level as well (Sober 1984: 121; italics in original).

Sober's argument here is a direct denial of Rosenberg's claims; Rosenberg maintains that, "[i]n general, the quantum probabilities are so small, and the asymptotic approach to determinism of everything physical above the level of the chemical bond is so close, that the quantum mechanical probability could never explain the probabilistic character, if any, of either evolutionary phenomena or evolutionary theory" (Rosenberg 1994: 61). (This is in spite of Rosenberg's admission that the micro-level can "infect" – i.e., percolate up to – the macro-level).

Both Sober's and Rosenberg's arguments are "in principle" arguments. Sober claims that in principle, micro-level indeterminism can percolate up to the macro-level, while Rosenberg maintains that in principle, by the time we reach the macro-level we will have asymptotically approached determinism. Since these "in principle" arguments are in direct conflict, they leave us at a philosophical impasse.

Brandon and Carson attempt to break down the impasse by presenting an example where a quantum mutation would have a population-level effect. However, the example they present, rather than being a general one, is of a very specific situation – a haploid population with an unstable equilibrium point, in which a mutation occurs. Thus Brandon and Carson's example does not settle the question of whether the evolutionary process is *generally* indeterministic, although it does suggest that it may be indeterministic *to some extent.* Brandon and Carson's example shows that the micro-level *can* "percolate up" to the macro-level, but without further evidence that this "percolation" occurs across different

types of evolutionary scenarios, the most we can say is that evolution may be indeterministic to some small extent. The percolation argument thus fails to answer the larger question of whether evolutionary biology is generally indeterministic.[20] Brandon and Carson's attempt to break down the philosophical impasse does not succeed, leaving us with differing philosophical intuitions about the implications of the percolation argument, and no resolution of the determinism versus indeterminism question.

Brandon and Carson's second argument analyzes experiments performed on (in theory at least) identical organisms in identical environments. As Brandon and Carson note, experimental setups that use cloned organisms in controlled environmental settings are quite common in biology. The results are equally commonplace: organisms that are genetically identical differ in reproductive success. For example, cloned plants grown in identical environments will be of differing height, weight, flower number, etc.[21] I take it that these sorts of results are not in dispute. What *is* in dispute is how to interpret the results; do they provide evidence for the indeterminism of the evolutionary process, or not?

Brandon and Carson claim that such results *are* evidence for the indeterminism of evolution. However, there is an obvious response that the determinist can make. The determinist can maintain that either the *organisms* were not truly identical (a mistake occurred during the cloning process), or the *environments* were not truly identical (a mistake was made in constructing the environmental settings). Note that the determinist can claim either that there is one variable that appears in some preparations, but not others (e.g., an undetected difference in soil among the preparations), or that there are *numerous*

[20] It should be noted that Brandon and Carson don't think that the percolation argument is decisive for the indeterminism of evolutionary biology: "For ET to be autonomously indeterministic it must be indeterministic in a way that does not depend on QM" (Brandon and Carson 1996: 320). However, I don't think this is the right way of looking at the issue; if the evolutionary process is in fact indeterministic, it is genuinely and autonomously indeterministic; it is only our description of the events that ascribes the indeterminacy to quantum mechanics. My point here is that Brandon and Carson fail to show the general indeterminacy of the evolutionary process not because they look to quantum mechanics to make their arguments, but because they only demonstrate it for one limited case. In other words, in one sense, I take Brandon and Carson's argument to be potentially more important than they do; I think that the *kind* of argument that they give might ultimately prove fruitful in settling the debate, even if I don't think that this particular argument succeeds.

[21] The implicit, and fairly uncontroversial claim that Brandon and Carson are making here is that these differences will translate into differences in reproductive success.

hidden variables which differ across the different preparations.

Brandon and Carson suggest that the experimental procedure used in these kinds of experiments *assumes* an indeterministic response; "it assumes that different copies of the same genotype in the same treatment will give different results; otherwise the experiment could be made much smaller with single copies of each genotype for each treatment" (Brandon and Carson 1996: 330). However, there is a very natural *deterministic* interpretation of this practice: that of the randomized experiment. As R. A. Fisher notes: "…the uncontrolled causes which may influence the result are always strictly innumerable" (Fisher 1953: 55); to overcome this problem, we use the method of a randomized experiment. In a randomized experiment, every characteristic possessed by members of the population is, on average, matched in the different treatments (Giere 1991: 241; Cartwright 1989: 64). The determinist can argue that we use multiple preparations for each treatment so that if there *are* any hidden differences between preparations, those differences on average occur equally in the all the treatments. This prevents us from mistakenly attributing a particular effect to the difference in question when it is really due to a hidden difference between the treatments. If we used only *one* replicate for each treatment, a hidden variable (rather than the difference under study) might be the cause of the given result. By using *multiple* replicates for each treatment, the hidden variables are randomized across treatments. Thus, the determinist can argue that we use this kind of experimental design because we cannot control for every possible factor.

So, both the determinist and the indeterminist have a way to account for the experimental results and a way to interpret the experimental procedure. Yet Brandon and Carson claim that the results favor an indeterministic interpretation. The question is, can they make this argument without simply begging the question against the determinist? In other words, do they offer any additional reasons to favor an indeterministic interpretation over a deterministic one?

Brandon and Carson propose two criteria for determining when it is appropriate to posit a theoretical entity (in this context, a hidden deterministic variable). The first criterion is that "the positing of the entity aids the development of theory"; the second criterion is

that "the available empirical evidence supports the posit" (Brandon and Carson 1996: 331). According to Brandon and Carson, the determinist's hypothesis fails to satisfy either of these criteria: "… the positing of deterministic hidden variables serves no theoretical purpose at all, and, insofar as it is allowed to be addressed by data is contradicted by empirical data" (Brandon and Carson 1996: 331).

However, the claim that the positing of deterministic hidden variables serves no theoretical purpose is not a view that is shared by at least some evolutionary biologists. If we can view natural selection hypotheses as deterministic, in the sense that they posit specific causes for a given evolutionary change, and random drift hypotheses as indeterministic, in the sense that they do not posit such causes, then we can understand some biologists to be arguing that a deterministic, natural selection explanation is to be favored over an indeterministic, random drift explanation. These biologists argue that not only does the determinist's hypothesis (that a particular change in the population was due to natural selection) direct us to keep looking for a selective factor, the indeterminist's hypothesis (that the change was due to random drift) may cause us to stop looking for a selective factor, and possibly miss one. For example, Mayr and Cain hold views of this sort. Mayr asserts that only after all attempts to explain biological processes through natural selection have failed, is a biologist "justified in designating the unexplained residue tentatively as a product of chance" (1983: 326). Cain states that the real basis for every random drift hypothesis is that "[t]he investigator finds that he, personally, cannot see any [evidence of selection], and concludes that, therefore, there is none" (1951: 424). Lest we think that it is only "selectionists" who favor deterministic hypotheses over indeterministic ones, it should be noted that Sewall Wright was also concerned about the dangers of the indeterminist's outlook; he states: "A certain danger for science must be squarely face [sic]. The acceptance of statistical description as ultimate may lead sometimes to premature abandonment of analysis in cases in which analysis would be pushed farther by one who believes firmly that there is a deterministic mechanism to be found" (Wright 1964: 288).

Here the indeterminist might concede that in field studies of evolution, where we cannot rigorously control conditions, there is an advantage to the determinist's hypothesis,

while arguing that in laboratory studies (particularly controlled studies of the kind that Brandon and Carson describe) there is no such advantage. But even to grant this point is to grant that the positing of hidden deterministic variables has been useful in evolutionary biology, for certainly the theory should account for natural as well as experimental contexts.

Let me clarify that my point here is not to argue that the determinist's hypothesis is theoretically superior to the indeterminist's. The indeterminist could counter the arguments of Mayr, Cain, and Wright by pointing out that the determinist may end up wasting time fruitlessly looking for possible hidden variables. What I am arguing is that neither view is superior to the other; both the determinist's hypothesis and the indeterminist's hypothesis have their advantages and disadvantages. Furthermore, I think Brandon and Carson clearly go too far in claiming that the positing of deterministic hidden variables in evolutionary theory serves no theoretical purpose whatever.

Let us turn then to Brandon and Carson's second claim, the claim that the deterministic hypothesis is contradicted by the empirical data insofar as it is allowed to be addressed by data. I must admit to being somewhat puzzled by this claim. If they mean that the determinist makes claims that go beyond the data, then they are just begging the question against the determinist, since the determinist's primary claim is exactly that there are data which we are missing. If, on the other hand, they mean only that the indeterminist's hypothesis is inconsistent with the data, it is hard to see how this could be the case. In fact, it would seem as though one could more easily make the opposite charge, that the deterministic hypothesis is cooked up simply to fit the data. However, since determinism and indeterminism are both long-standing philosophical traditions, it is difficult to dismiss one or the other *purely* on the basis of *ad hoc*ness. Furthermore, it seems to me that the kind of argument one would make, say, for the reality of electrons, is inappropriate for general theses about the nature of reality. I think that the criteria for believing in such general theses should be stronger than the criteria for believing in electrons, or even phenomena such as directed mutation. In any case, the question here is whether the data distinguish between the two hypotheses, and it seems clear that they do

not.

My arguments here have focused on rejecting Brandon and Carson's arguments for the indeterminacy of the evolutionary process. The reader should not interpret this to mean that I favor the determinist's hypothesis over the indeterminst's. I think that Brandon and Carson succeed in showing us why we need to take the indeterminist's hypothesis seriously, but they do not succeed in persuading us to prefer the indeterminist's hypothesis over the determinist's. Brandon and Carson readily admit that their arguments cannot be on par with Bell's results in quantum mechanics – that they cannot force their opponent into a logical contradiction. However, the criteria they do offer, of theoretical fruitfulness and experimental confirmation, are not decisive either. The criteria do not give us a decisive reason to prefer one hypothesis over the other.

At this point, when no philosophical argument is able to decide between the determinist's hypothesis and the indeterminist's, the most philosophically defensible position would seem to be to remain agnostic between the two positions. The question of determinism versus indeterminism is a longstanding issue in philosophy; it is no less so here because we consider not global indeterminism and determinism, but only the determinism or indeterminism of the evolutionary process. It does not seem likely, therefore, that philosophical argument will settle the question. However, to grant this point is not to give up all hope of providing an account of the probabilistic nature of evolution. That is to say, without settling the debate between the determinist and the indeterminist, we can still ask whether evolutionary theory is inherently and unavoidably probabilistic, or whether it is probabilistic solely for instrumental reasons – because we find probabilities more useful and tractable in evolutionary contexts. On the one hand, if the evolutionary process is *indeterministic*, then the answer to this question is clear; evolutionary theory is inherently and unavoidably probabilistic. On the other hand, if the evolutionary process is *deterministic*, the answer is not as obvious. Sober (1984), as well as Brandon and Carson (1996), give arguments which attempt to demonstrate that even if evolution is deterministic at the individual level, one must give population-level probabilistic explanations in order to provide complete evolutionary explanations. I am

essentially in sympathy with these arguments, and the argument I will make here is similar in certain respects to theirs. I will begin by assuming, for the sake of argument, that the evolutionary process is deterministic. The determinist's hypothesis (and indeterminist's as well) is consistent with the view that evolution is a population-level process. So, let us explore this population-level nature a little further, beginning with a discussion of what is meant by the claim that evolution is a population-level process.

2.6. The Inherent Probabilism of Evolution

A population can be defined as "a group of conspecific organisms that occupy a more or less well defined geographic region and exhibit reproductive continuity from generation to generation" (Futuyma 1986: 554-5). Even though a population is composed of individual organisms, it is important to distinguish between properties that apply to *individual organisms* and properties that characterize *the relationships among organisms* – that is, properties that apply to populations. For example, individual organisms have properties such as color, shape, and length.[22] Populations, on the other hand, have properties such as size (defined as the number of individuals), frequency (defined as the proportion of individuals of one type or another), and growth rate (defined as the rate of change in the number of individuals in the population). Thus, in a sense, population-level properties are properties that arise only given the collection and interaction of individuals. Nonetheless, a property can apply to a population without applying to the individuals in the population. For example, the growth rate of a population could be increasing while the growth rate of the individuals in the population remained constant.

Although an individual organism's properties may change over the course of its lifetime – the organism may grow, change color, become stronger, even lose a limb – it cannot be said to evolve. It is only *populations* which can evolve, through a change in the genetic composition of a population from one generation to the next. For example, suppose there existed a population composed of 50% dark moths and 50% light moths.

[22] Thanks to Carl Chung for suggesting that this distinction be made explicit in the following discussion.

The population is observed for a generation, and it is found that the next generation of moths consists of 40% dark moths and 60% light moths. We would say that an evolutionary change had taken place – perhaps by natural selection, perhaps by random drift, or perhaps by mutation or migration. In any case, the change which has occurred is a change in the *frequencies of the population*, which is a property of the population rather than of an individual. Thus, it is the *population* that has evolved. Some of the individual moths may change; some may die, some may reproduce, some may suffer injury, but none can be said to evolve.

This evolutionary focus on the properties of populations, with its consequent de-emphasis on the properties of individuals, is what is meant by the assertion that evolution is a population-level process. As Sober notes: "Darwin and Galton focused on the population as a unit of organization … The details concerning the individuals who are parts of this whole are pretty much irrelevant" (Sober 1994: 175). Darwin focused on populations, not because focusing on individuals was too unwieldy, but because he saw, and wanted to describe, processes *which were occurring at the population level*. If Darwin had written a book announcing merely that some organisms lived and reproduced, while other organisms died, he would have been forever buried in obscurity. Darwin's insight was to see not *individuals*, but *groups of individuals*, and to seek to describe the changes in populations by revealing the processes that occur at the population level. Darwin's theory of evolution is not a theory about individuals, but a theory about populations.

Natural selection and random drift are also population-level processes. This can be seen by examining how these terms are used and defined. As we discussed in Chapter 1, the three necessary conditions for both natural selection and random drift are:

1. Variation
2. Heritability of variation
3. Differential reproduction of heritable variation

(See Chapter 1 for a discussion of the "fourth condition" which distinguishes between

natural selection and random drift). These three conditions jointly describe a population-level process. The first condition, "variation", is a condition which applies to a population, not an individual. It does not refer to the variability (or the heterozygocity) of an individual; instead, it means that within a population, some individuals will differ from other individuals. The third condition, differential reproduction of heritable variation, is likewise a condition which applies to a population, not to an individual. It specifies that some individuals are more reproductively successful than other individuals, something that makes sense only with reference to a population, not with reference to an individual. Thus, since random drift and natural selection share these necessary conditions, they are both population-level processes.

We sometimes speak of one individual as having been selected, but this linguistic practice is somewhat misleading. It does not mean that the individual has undergone natural selection. Rather, it implies that one individual has been selected *over other individual(s)* in the population; the former lived to reproduce, whereas the latter did not. In other words, it is only within the context of a *population* that natural selection occurs. Indeed, biologists assign relative, rather than absolute fitnesses to organisms in a population. What is relevant with respect to natural selection is only how well an individual fares relative to other individuals in the population; how it fares in an absolute sense is irrelevant because it cannot constitute natural selection.

I take it that my characterizations of evolution, natural selection, and random drift as population level processes are fairly uncontroversial. That is, I am simply describing the way these terms are generally defined and used by evolutionary biologists; they are used to describe population level processes. What *is* controversial is the kind of explanation that should be given for these processes: an individual-level explanation or a population-level explanation? According to the determinist, it is possible to explain these population-level processes by reference to the individuals in the population. Horan (1994), for example, proposes this kind of focus on individuals, while admitting the difficulty of providing such an account. The indeterminist agrees that individual-level explanations can be given to some extent, although, of course, on the indeterminist's account such explanations are

indeterministic rather than deterministic. However, certain indeterminists (namely Sober, Brandon, and Carson) make the further claim that individual-level explanations fail to *fully* capture, or explain, the processes that occur at the population level; in their terms, there are population-level generalizations which Laplace's demon would miss.

As I argued above, the question of the determinism or indeterminism of evolution remains unresolved. However, there is hope that we can resolve this further issue that divides the indeterminist and the determinist, the issue of the appropriate level of explanation to be given for evolutionary processes. To that end, we will examine each kind of explanation. Since the argument for the superiority of the population-level explanation is supposed to hold whether evolution is deterministic or indeterministic, we will grant the determinist's claim that evolution is composed of deterministic processes.

First let us consider an individual-level explanation. Suppose again that there exists a population of light and dark-colored moths, and let us further suppose the light colored moths are fitter than the dark colored moths because their color provides them with better camouflage in a primarily light-colored environment. The frequency of light-colored moths is observed to change from 50% to 60% from one generation to the next. The question at issue is, what is the explanation of this change in gene frequencies? The individual-level explanation would describe the fate of each moth – which moths were able to avoid detection by the predator and thus survive to reproduce, and which ones were picked off. So, assuming we had the ability to provide such an account, our explanation might go as follows: moth 1, a light-colored moth, had two offspring; moth 2, a dark-colored moth, had one offspring; moth 3, a dark-colored moth had none, etc. We can imagine that such an account might further discuss the causal chain of events which led to each moth being in its particular location, leading to its success or failure. Sum up the fates of all of the moths, and you arrive at the individual-level explanation for the evolutionary change in the population.

Now let us consider a population-level explanation. The population-level explanation notes that there is a heritable fitness difference between the two types of moths, and that there is differential reproduction of the two types. Furthermore, it is claimed that the color

difference between the two moths is causally relevant to the the difference in reproductive success.[23] More specifically, the population-level explanation maintains that the greater success of the light-colored moths is due to their greater ability to avoid detection by the predator, given the primarily light-colored environment that the moths live in. Thus, according to the population-level explanation, the population was subject to natural selection; as a result of this natural selection process, the frequency of the light-colored moths increased relative to the frequency of the dark-colored moths. Therefore, the change in the population was the result of evolution by natural selection.

It is important to note there that the individual-level explanation makes no reference to natural selection, while the population-level explanation does. This is, of course, what you would expect given that natural selection is a population-level process. The individual-level explanation cannot make reference to natural selection, because once it does so, it has made a population-level argument.

What I want to argue is that the individual-level explanation cannot fully account for the population-level process; we must give a population-level explanation. While the individual-level account – impossible to provide in practice – would be informative in many ways, it fails to fully answer the question, a question which concerns a change in the population *as a whole*. We may know which of the moths lived and reproduced, and which of them died, but why is it that the number of light-colored moths increased *overall*? What is the explanation for that population-level phenomenon? To sum up the fates of these individuals merely reiterates that there was a change in frequency of light-colored moths from 50% to 60%, but it does not give us a general account of *why* there was such a change; that is, it does not tell us what it is about the light-colored moths that led to their increased numbers in the population (and conversely, what it is about the dark-colored moths that led to their decreased numbers). To know these things, we need to know not what has happened to the individual moths, but what were the causal influences in the population as a whole. And on the whole, the color difference of the moths was causally relevant to the difference in the reproductive successes of the two types. In other words,

[23] Thus satisfying the criteria for natural selection outlined in Chapter 1.

what has happened to the population is that it has been subject to the population-level process of natural selection.

Although it is individuals which are selected (over other individuals), it is not until we look at the population that we see what is changing. The individuals are not changing in any way that has to do with a change in gene frequencies; they are simply living, dying, and reproducing. What is changing is the *composition of the population* (a property of populations, not of individuals), and the reason it is changing is that the color of the light-colored moths is causally relevant to their survival and reproduction. To explain why the frequency of light-colored moths changed from 50% in the parent generation to 60% in the subsequent generation, one does not cite the fate of individual organisms; they are irrelevant. Rather, it is the processes occurring at the population level which explain the change in gene frequencies. *The individual-level explanation misses the overall process*, a process that occurs at the population level. It thus fails to provide a complete explanation of the evolutionary process occurring in the population.

I have made this argument using a case of natural selection as an example because it makes the example a little easier to follow. However, the same argument can be made using a case of random drift. In fact, you can take exactly the same moths, undergoing the same fates, but let us now assume that the predator is color blind. Again, we can talk about the fates of individuals, but to do so fails to give us a general explanation for why there was an increase in light-colored moths. In this population, the reason for the increase in the light-colored moths is that there was an indiscriminate sampling process occurring in the population; that is, a process whereby the color differences between the moths was causally irrelevant to their survival and reproductive success. In other words, what has happened to this population is that it has been subject to the population-level process of random drift. The individual-level explanation completely misses this point, and so fails to provide a complete explanation of the evolutionary change.

There is a further point to be made, along the same lines as arguments made by Sober, Brandon, and Carson. Sober (1984) contends that fitness and probability are two concepts which generalize *across* populations. But the concepts of random drift and

natural selection themselves play this role as well. We use the concepts of natural selection and random drift to describe what is similar from population to population. If you only focus at the individual level, you lose the concepts of random drift and natural selection. And without those concepts, you don't fully see what is going on in a particular case (as I argued above), nor can you describe what is similar from case to case. The concepts of random drift and natural selection generalize across species and populations; they capture the similar processes that occur in different populations and in different species. To the extent that the individual-level explanation ignores the population-level concepts of random drift and natural selection, it fails to capture what is similar across populations and species.

If we must give population-level explanations in order to give complete evolutionary explanations, it follows that natural selection and random drift – and evolution – are inherently and unavoidably probabilistic. This means simply that in order to describe evolutionary processes completely, we are *forced* to use the language of probability theory; the use of probabilities is not merely convenient or instrumental. A non-probabilistic explanation would be an individual-level explanation that described the fates of individual organisms and determined one specific outcome. Population-level explanations, on the other hand, abstract away from the specific causal histories of individuals, and thus do not predict one specific outcome.

The population-level explanations provided by natural selection and random drift involve the determination of whether, for the population as a whole, fitness or physical differences in a population are causally relevant to differences in reproductive success. Thus, they consider only physical or fitness differences; other factors having to do with the specifics of individuals (such as the location of a given individual at a given point in time) are ignored. Given a particular distribution of genotypes differing in fitness, a natural selectionist explanation predicts not one outcome, but a probability distribution of possible outcomes. This is because given two genotypes differing in fitness (for example), we know that the most likely outcome is that the fitter type outreproduces the less fit type. However, the reproductive success of the two types may not reflect fitness values; in fact, the less fit genotype might even outreproduce the more fit genotype (although such an

outcome would be less likely). Similarly, a particular distribution of genotypes in a population undergoing random drift yields a probability distribution of possible outcomes. We can say what the most likely outcome is (for random drift, the most likely outcome is that there is no change in the genotype distribution), but other outcomes are possible, even relatively probable.

Thus, the population level explanations provided by natural selection and random drift express which outcomes are the most likely, but they do not designate any particular outcome. In this way, population-level explanations of evolution are probabilistic. So, if we are forced to give population-level explanations to capture evolutionary processes, as I have argued, we are forced into probabilism.

2.7. Conclusion

Rosenberg offers us an intriguing and provocative account of evolutionary theory, one where the concept of random drift is necessary only because of our limited capacities as human beings. However, his arguments use an impoverished concept of random drift, one that is imprecise and not sufficiently elaborated. Random drift is a term which biologists use to refer to a number of distinct, biological processes, not a term they use when they are unaware of the "real" processes. Furthermore, random drift is a real phenomenon in its own right, distinct from natural selection; random drift cannot be replaced with an elaboration of selective processes. Consequently, random drift can explain certain kinds of evolutionary change that selection cannot. Each of these considerations points to the same conclusion: random drift is an essential, ineliminable part of *any* complete evolutionary theory. Darwin's evolutionary theory is arguably deterministic, but Darwin didn't consider the possibility of random drift. No proof of the determinism of current evolutionary theory can go through without a demonstration of the determinism of random drift. You cannot simply eliminate random drift from consideration, as Rosenberg does.

Through an examination of random drift, Brandon and Carson challenge

Rosenberg's claims and argue that that evolution is indeterministic. However, neither of the arguments which they offer (the percolation argument and the argument from experiments on replicates) is decisive against the determinist. This leaves us at a philosophical impasse. Of course, it is possible that a philosophical argument will be proposed that can settle the question of the determinism or indeterminism of evolution. However, given the longstanding nature of this problem, such a solution is not likely to be forthcoming. Thus, our best hope lies in exploring the nature of evolution which follows from *either* thesis. On the thesis of indeterminism, evolution is clearly probabilistic. However, under thesis of determinism, evolution is probabilistic as well. In order to provide complete explanations of evolutionary processes, one must give population-level explanations, rather than individual-level explanations. Yet the population-level explanations provided by natural selection and random drift are probabilistic. So, even under the thesis of determinism, we are forced into probabilism, because we are forced to give population-level explanations. Evolution is thus inherently, and unavoidably, probabilistic.

Chapter 3: Random Drift and Explanation

> It is universally allowed that nothing exists without a cause of its existence, and that chance, when strictly examined, is a mere negative word, and means not any real power which has anywhere a being in nature (Hume, [1777] 1988 , §74).
>
> 'Chance' often serves as a powerful explanatory concept in evolutionary biology (Shanahan 1991: 267).

3.1. Introduction

On many formulations, chance is held to be non-explanatory. That is, statements which claim that an event occurred by chance are often seen as having failed to provide an explanation for the event's occurrence. Contrary to this claim, I will argue that in evolutionary biology chance *does* play an explanatory role. Biologists use the concepts of natural selection, mutation, and migration to explain changes in gene frequencies and other evolutionary phenomena. However, they also use chance, in the form of random drift, to explain these same phenomena. Random drift is a process in which physical differences between organisms are causally *irrelevant* to their survival and reproductive success, in contrast to natural selection, a process where physical differences between organisms *are* causally relevant to their survival and reproductive success (see Chapter 1). Random drift is different from more traditional philosophical notions of chance because it is an explanatory concept. It thus gives us reason to rethink our traditional conceptions of chance within the context of scientific explanation.

In general, philosophers of biology have paid little explicit attention to the explanatory nature of random drift. Sober (1984) discusses it indirectly by suggesting that the use of probabilities captures important biological distinctions that Laplace's demon would miss, a claim denied by Rosenberg (1994). Sober (1984) also considers the issue in the context of his account of explanation in evolutionary biology more generally, but only in passing. For the most part, remarks concerning the explanatory nature of random drift are brief, and simply point out the kinds of phenomena which random drift explains, or has been taken to explain. Shanahan (1991) is an exception to this general trend; he

discusses at some length the essential explanatory role that chance (construed more broadly, but including random drift) has played in the history of evolutionary biology. Accounts such as these suggest the importance of the explanatory role of random drift, but they do not elaborate on that role, leaving unanswered questions. How is it that random drift can be explanatory, if it is a form of chance? What is it about the concept of random drift that makes it explanatory, when presumably, other conceptions of chance are not? What kinds of phenomena does random drift explain, and how do these explanations proceed? And finally, should random drift lead us to rethink our traditional conception of chance as non-explanatory? My goal in this chapter is to shed some light on the answers to these questions.

The structure of the chapter is as follows. First, I discuss what is meant by various claims that chance is non-explanatory. Second, in response to one of these claims, I show how it is that chance in the form of random drift *can* be explanatory. In so doing, I explore the causal role played by random drift and show how random drift can explain fluctuations in gene frequency in a population, as well as how it can explain the divergence of small populations from one another. Third, I examine yet another phenomenon (fixation) that random drift has been invoked to explain, in order to further explore the explanatory nature of random drift and to give an indication of the scope of its explanatory role within evolutionary theory.

3.2. The Counter-Claim: Chance is Not Explanatory

However, before we begin, we need to understand what is meant by the claim I will be arguing against, the claim that chance is not explanatory. Here are three possibilities:

> **1. When we "explain" an event by saying that it is due to chance, it is because we are ignorant of the real cause(s) of the event.**
>
> On this view, if scientific explanations are causal explanations, citing chance will fail to provide an explanation of a phenomenon because the very thing that would provide an explanation, the cause, is admittedly absent from the explanation. Thus,

on this account, chance is not explanatory because to say that an event occurred "by chance" is to fail to provide an explanation of the event. In the context of evolutionary biology, any purported explanation using random drift is suspected of having missed the real cause (the real cause being an instance of natural selection) and thus fails to provide a genuine explanation. Biologists who hold this view include Cain (1951) and Mayr (1983); philosophers such as Rosenberg (1994) and Horan (1994) make similar claims.[24]

2. Chance is not sufficiently prevalent or efficacious to bring about the event in question.

Therefore, it cannot be used to explain the phenomenon in question. Fisher, in his disagreement with Sewall Wright over the prevalence of random drift, held that random drift is at best just random noise that serves to spoil the effects of natural selection. Proponents of this view do not claim that chance is completely irrelevant to the event, only that it is insufficient to account for, and therefore to explain, the phenomenon.

3. Chance lacks causal efficacy altogether; chance cannot be the cause of an event.

This view differs from **2** in that chance has no efficacy whatsoever, as opposed to merely insufficient efficacy. On this view, if to give a scientific explanation is to cite a cause, chance cannot be cited because it cannot *be* the cause of another event. Therefore, chance is non-explanatory. This view, held by Humphreys (1989), will be examined in further detail momentarily.

The first two of these arguments have been thoroughly examined, if not resolved, in the biological and philosophical literature; they raise issues about the relative importance of the

[24] It is interesting to note one difference between the views of Mayr and Cain on the one hand, and Rosenberg and Horan on the other. Mayr and Cain do not go so far as to assert that there is always a cause underlying any claim of chance, as a Laplacean would; rather, the claim is that natural selection is much more prevalent than chance (random drift). Any purported explanation using random drift is suspected of having missed the real cause (an instance of natural selection) and thus fails to provide a genuine explanation. Rosenberg and Horan, however, *do* make the Laplacean claim (see Chapter 2).

roles played by selection and drift. In the context of evolutionary biology, argument **1** implies that most, if not all evolutionary events are the result of natural selection, rather than random drift. Argument **2**, on the other hand, implies that chance has played a minimal and unimportant role in these events (although perhaps a factor in all or most of them). Since these issues have been covered extensively elsewhere, I will not discuss them here. My focus therefore will be on argument **3**.

Paul Humphreys, in his recent work *The Chances of Explanation,* claims that chance is not explanatory because chance is not causally efficacious. The context for this claim is Humphreys's account of causal explanation. Humphreys maintains that in order to explain the occurrence of a property or change in property Y in a system S at trial t, one must cite:

> 1. A non-empty list of terms referring to the contributing causes of *Y*
>
> 2. A possibly empty list of terms referring to the counteracting causes of *Y* (Humphreys 1989: 101)

We would then say that *Y* occurred because of the contributing causes, despite the counteracting causes.

Thus, on Humphreys's account, an essential part of an explanation is to provide the contributing causes. However, according to Humphreys, chance cannot be a contributing cause, because: "Chance by itself has no causal properties or powers, just because there is no such thing as a universal property of chance in the world, any more than there is a universal deterministic tendency in the world…chances themselves are acausal" (Humphreys 1989: 20-1). In other words, Humphreys believes that chance is not something which produces effects; it cannot itself be a cause: "…once we have described the contribution of *B* to the occurrence of *A*, there is no further explanation of *A* to be had. After citing the causal factor…the event *A* either just happens or just does not happen. That is the nature of chance" (Humphreys 1989: 35). Moreover, "…after all the causal factors have been cited, all that is left is a value of sheer chance, and chance alone explains nothing" (Humphreys 1989: 113). So, on Humphreys's view, chance cannot be a cause, and it cannot be a contributing cause as part of a causal explanation. Thus, according to Humphreys, chance is non-explanatory.

Humphreys's account is interesting because it provides an opportunity to examine a recent claim made about scientific explanation *in general*, from the point of view of explanation *in evolutionary biology*. The claim that chance is non-explanatory appears in biological contexts as well as broader scientific contexts. Thus an argument concerning the explanatory nature of chance has implications for both philosophy of biology *and* philosophy of science. What remains to be seen is whether upon closer examination, the claim that Humphreys has made for scientific explanation in general holds for explanations in evolutionary biology. I shall argue that it does not – that in evolutionary biology, chance *is* explanatory.

3.3. How Can Chance Be Explanatory?

Humphreys's argument creates a challenge for philosophers of biology: How is it that chance (or more specifically, random drift) can be explanatory? I shall argue that Humphreys's own account of explanation and causality can be used to show that chance is explanatory. The argument has the dual purpose of arguing *that* random drift is explanatory, while providing an account of the *way* in which it is explanatory. The key to this argument lies in showing that random drift has causal properties; this will allow us to bring chance in under the umbrella of Humphreys's theory of explanation.

Humphreys's claim that chance is not explanatory follows from his claim that chance is not causally efficacious, given his causal account of explanation. However, there is reason to think that chance in the form of random drift *is* a causal agent. Certainly many prominent evolutionary biologists speak of random drift in causal terms. For example, Sewall Wright, the biologist most often associated with random drift, refers to the "effects of fluctuations in mutation, selection, and immigration pressures and of sampling in small populations" (in other words, the effects of random drift) as one of three factors contributing to evolution (1949: 474). Kimura, in some ways Wright's "successor" as a proponent of random drift, stresses "the importance of random drift as a major cause of evolution" (1991: 5972). Another biologist, Cavalli-Sforza, claims that "[random] drift

can cause the frequency of a gene to vary markedly from one population to another" (1969: 32).

That evolutionary biologists speak of random drift as a causal agent suggests that the conception of random drift within evolutionary theory is a causal one. But in what sense is random drift a causal agent? Here we can turn to Humphreys's own view of causality for an answer. Humphreys asserts that a factor X is a contributing cause of Y iff the occurrence of X results in an increased chance of Y. (Correspondingly, a factor Z is a counteracting cause of Y iff the occurrence of X results in an decreased chance of Y). (Humphreys 1989: 16).

A couple of points of clarification are needed here. First, by "chance" Humphreys means probability. For Humphreys, probability is "not reducible to relative frequencies"; instead, it is a propensity, "grounded in structural bases and affected by multiple exogenous factors" (Humphreys 1989: 36). The second point to remember here is that when Humphreys says that a contributing cause is that which increases the chance of an outcome, he is not claiming that chance in any way causes the outcome, but rather that the cause increases the probability of the outcome, after which the outcome "either just happens or just does not happen" (Humphreys 189: 15).

It can be shown that random drift is a contributing cause in Humphreys's sense. Suppose there exists a population of giraffes not undergoing mutation, migration, or selection – only random drift. As we saw in Chapter 1, random drift is an indiscriminate sampling process. In other words, differences in organisms' physical traits are causally irrelevant to their survival and reproduction. Thus, if the long-necked giraffes in the population become more prevalent than short-necked giraffes, it is by chance (and not because of their long necks). In an infinite (or very large) population, we would expect the proportion of long-necked to short-necked giraffes to remain relatively constant from generation to generation. However, we would not expect the proportions to remain constant in a small population. To understand this, imagine an urn filled with colored balls where balls are sampled without respect to color. If a large sample of balls is taken, we would expect the frequencies of colored balls in the sample to be very close to the

frequencies in the urn. (See Figure 3.1). On the other hand, if we only take a small sample of colored balls, our sample may very well have different proportions of colored balls than the urn does. (See Figure 3.2). In the same way that the color difference between the balls is causally irrelevant to which ball gets picked, in a population undergoing random drift, the physical differences between organisms are causally irrelevant to their survival and reproductive success. So in large populations, as with large urn samples, we would expect gene frequencies to be representative of the parent generation, but in small populations, as with small urn samples, gene frequencies may or may not be representative. Thus, when random drift occurs over a number of generations in a small population, gene frequencies may fluctuate randomly from generation to generation. (See Figure 3.3).

In this (theoretical) manner, one could argue that random drift increases the probability that a population will experience gene frequencies which fluctuate from generation to generation, and use that argument to claim that chance, in the form of random drift, is causally efficacious and therefore explanatory on Humphreys's account. However, two practical problems intervene in providing examples where biologists invoke random drift to explain fluctuations in gene frequencies. The first problem is that studies of populations over many generations are difficult and costly. The second problem is that fluctuations in selection due to a changing environment may produce the same pattern of gene frequency fluctuations. (See Figure 3.3). Thus biologists cannot (with any justification) claim that random drift, and not natural selection is the relevant causal factor without additional evidence beyond patterns of gene frequency change. The phrase "without additional evidence" is key here; it is *possible* to explain gene frequency fluctuations by reference to random drift, but it is also possible to cite natural selection as the cause of those fluctuations, so that without further evidence, such claims would be unjustified. Thus, biologists often look to other kinds of evidence when trying to determine whether evolutionary change within that population is due primary to random drift or natural selection.

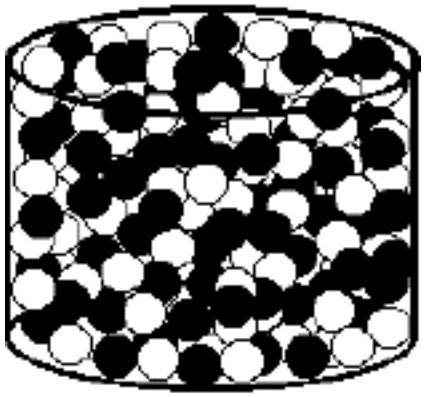

Frequency of white = 0.5
Frequency of black = 0.5

100 balls are sampled without respect to color (ball color is causally irrelevant to the sampling)

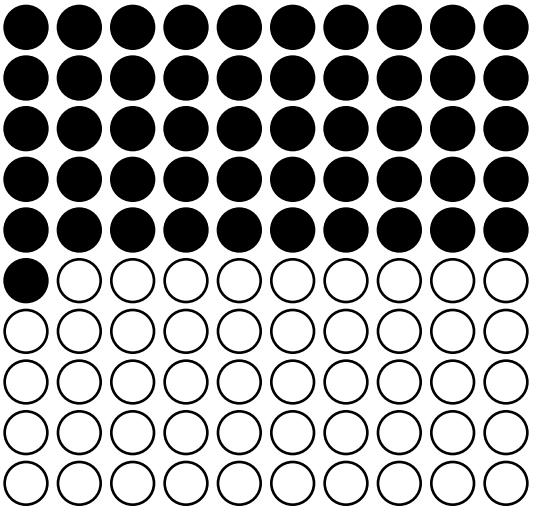

Frequency of white = 0.49
Frequency of black = 0.51

Figure 3.1: An example of sampling 100 colored balls from an urn. Since this is a relatively large sample, it is likely that the frequencies of the sample will be fairly close to the frequencies in the urn.

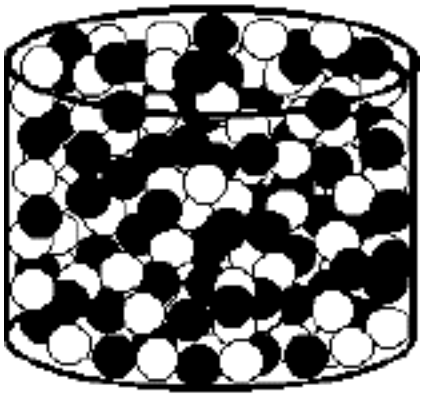

Frequency of white = 0.5
Frequency of black = 0.5

10 balls are sampled without respect to color (ball color is causally irrelevant to the sampling)

●●●○○○○○○○

Frequency of white = 0.7
Frequency of black = 0.3

Figure 3.2: An example of sampling 10 colored balls from an urn. Since this is a relatively small sample, it is likely that the frequencies of the sample will not be very reflective of the frequencies in the urn.

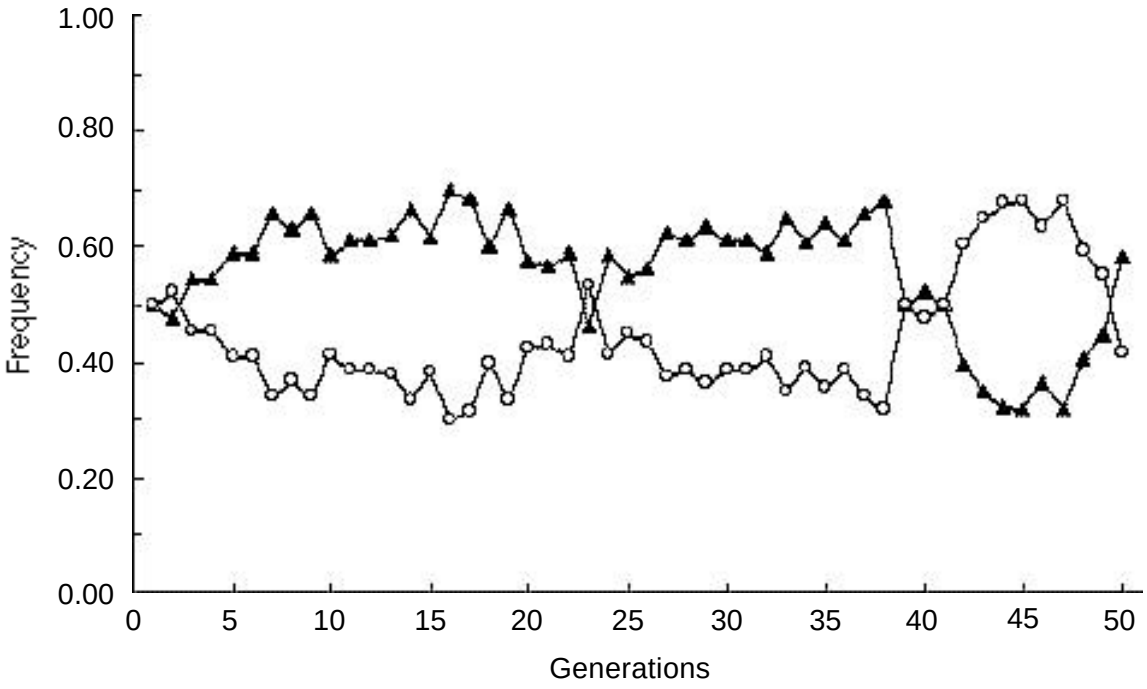

Figure 3.3: Sample plot of gene frequencies for a population with two alleles (represented by a circle and a triangle, respectively). Intially, the allele frequencies are equal, but over the course of 50 generations, the allele frequencies fluctuate. Fluctuations like these can be caused by random drift, but they can also be caused by natural selection in a fluctuating environment (as the environment fluctuates, fitness values fluctuate, causing the differerent genotypes represented by the different alleles to have varying reproductive success).

For example, biologists often study many populations at once rather than one population alone. Consider a number of small, isolated populations, each consisting half of yellow snails and half of pink snails, all undergoing random drift in the absence of selection. Each would experience random fluctuations in gene frequencies from generation to generation. But would they undergo the *same* random fluctuations? Most likely not. Some populations would end up containing both pink and yellow snails, but the proportion of pink to yellow snails would vary; pink would be more prevalent in some, yellow would be more prevalent in others. Other populations might consist solely of yellow snails, or solely of pink snails. In this way, random drift can lead to the genetic divergence of small,

isolated populations. (See Figure 3.4). Now consider a number of large populations under the same conditions. Because the "samples" (the new generations) are large, frequencies will tend to remain relatively constant from generation to generation, and thus populations will be similar to one another. (See Figure 3.5). In this manner, random drift tends to produce a pattern whereby small populations have considerable genetic divergence, and large populations have limited genetic divergence.

A classic study of this kind was performed by Cavalli-Sforza on blood type frequencies in Italian villages (Cavalli-Sforza 1969). In this study, the pattern of divergence we have been discussing was found: there was less divergence between the largest populations, and a greater divergence between the smaller populations. Cavalli-Sforza argued that natural selection is unlikely to have caused the pattern, because within each village, the blood type frequencies matched the frequencies of other genetic traits. For this to be the result of natural selection, all the different traits would have to have identical fitnesses, an unlikely coincidence. This would be akin to saying that blood types aided humans' survival and reproduction to the same degree as hair color, to the same degree as height, etc. Thus, while natural selection cannot be ruled out entirely as the cause of the pattern, there is reason to think that it is unlikely to be the cause. Moreover, if there were no selection and no random drift (as well as no mutation and migration) in these populations, we would not see this pattern. Populations would simply reproduce in the same proportions from generation to generation.[25]

[25] I point this out to clarify that if none of these processes are operating, the outcome in question would not occur. I do not mean to suggest that natural selection, migration, and mutation are completely absent. In both studies, the small populations were relatively isolated, meaning there would be little or no migration. Migration may have been a factor in the large populations, which would also tend to produce the uniformity seen in these populations. However, because migration tends to produce uniformity, it cannot have been responsible for the overall outcome. Mutations are rare and their effects (in the absence of the other factors) are considered to be negligible.

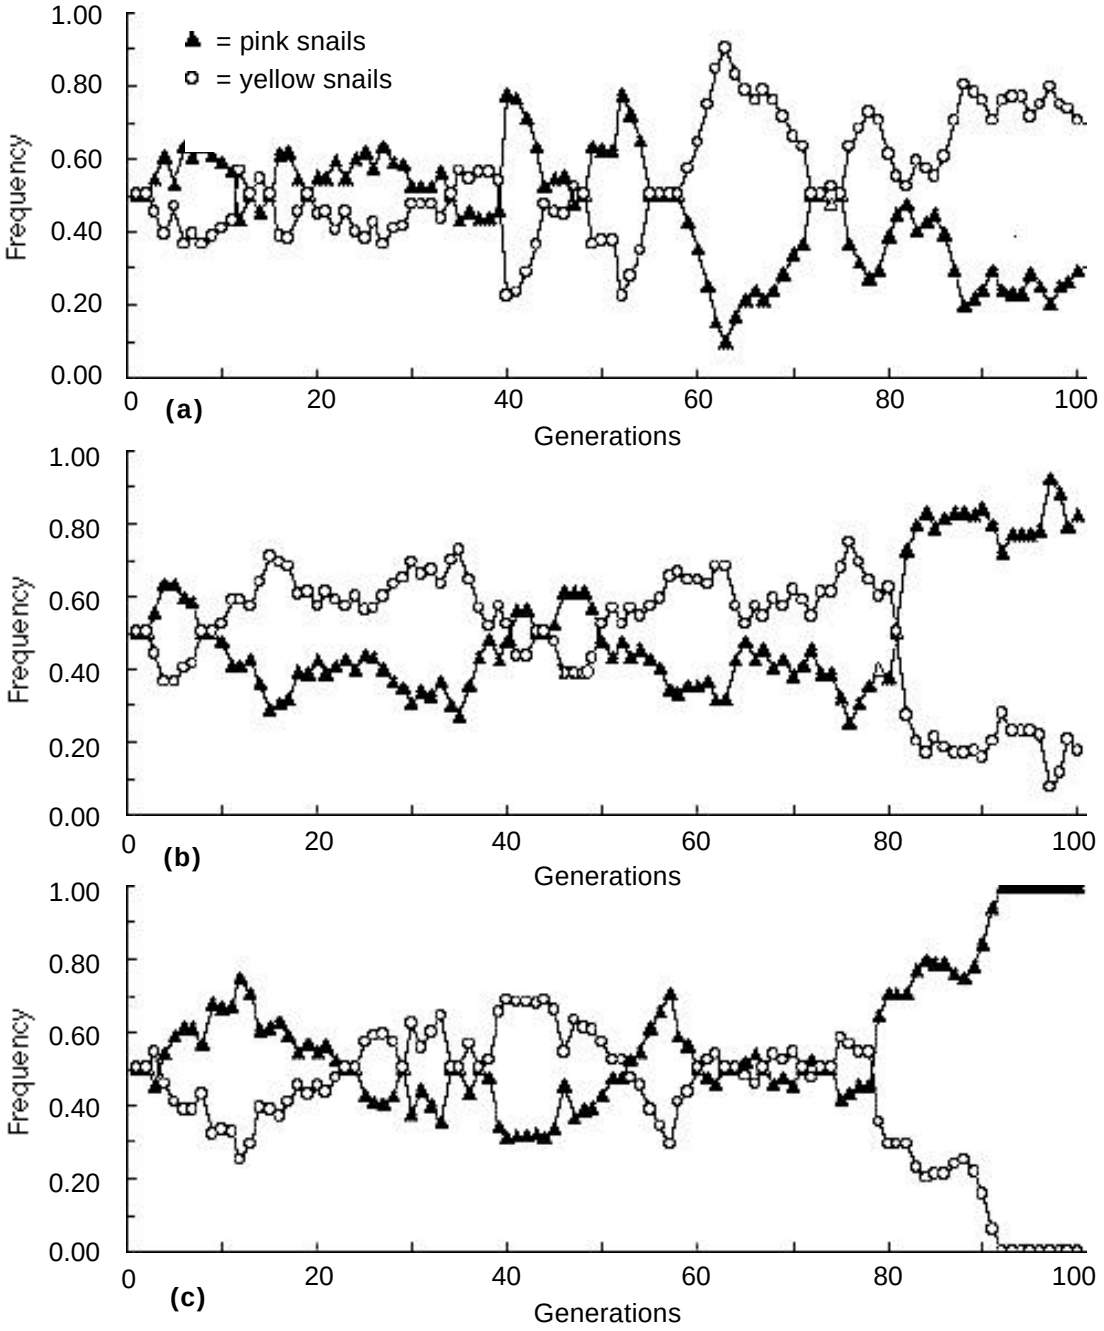

Figure 3.4: Three sample plots of small populations undergoing random drift in the absence of selection. Intially, all three populations contain equal numbers of pink and yellow snails. After 100 generations, yellow snails are more prevalent in population (a), pink snails are more prevalent in population (b), and pink snails have gone to fixation in population (c). (See below for a discussion of random drift and fixation).

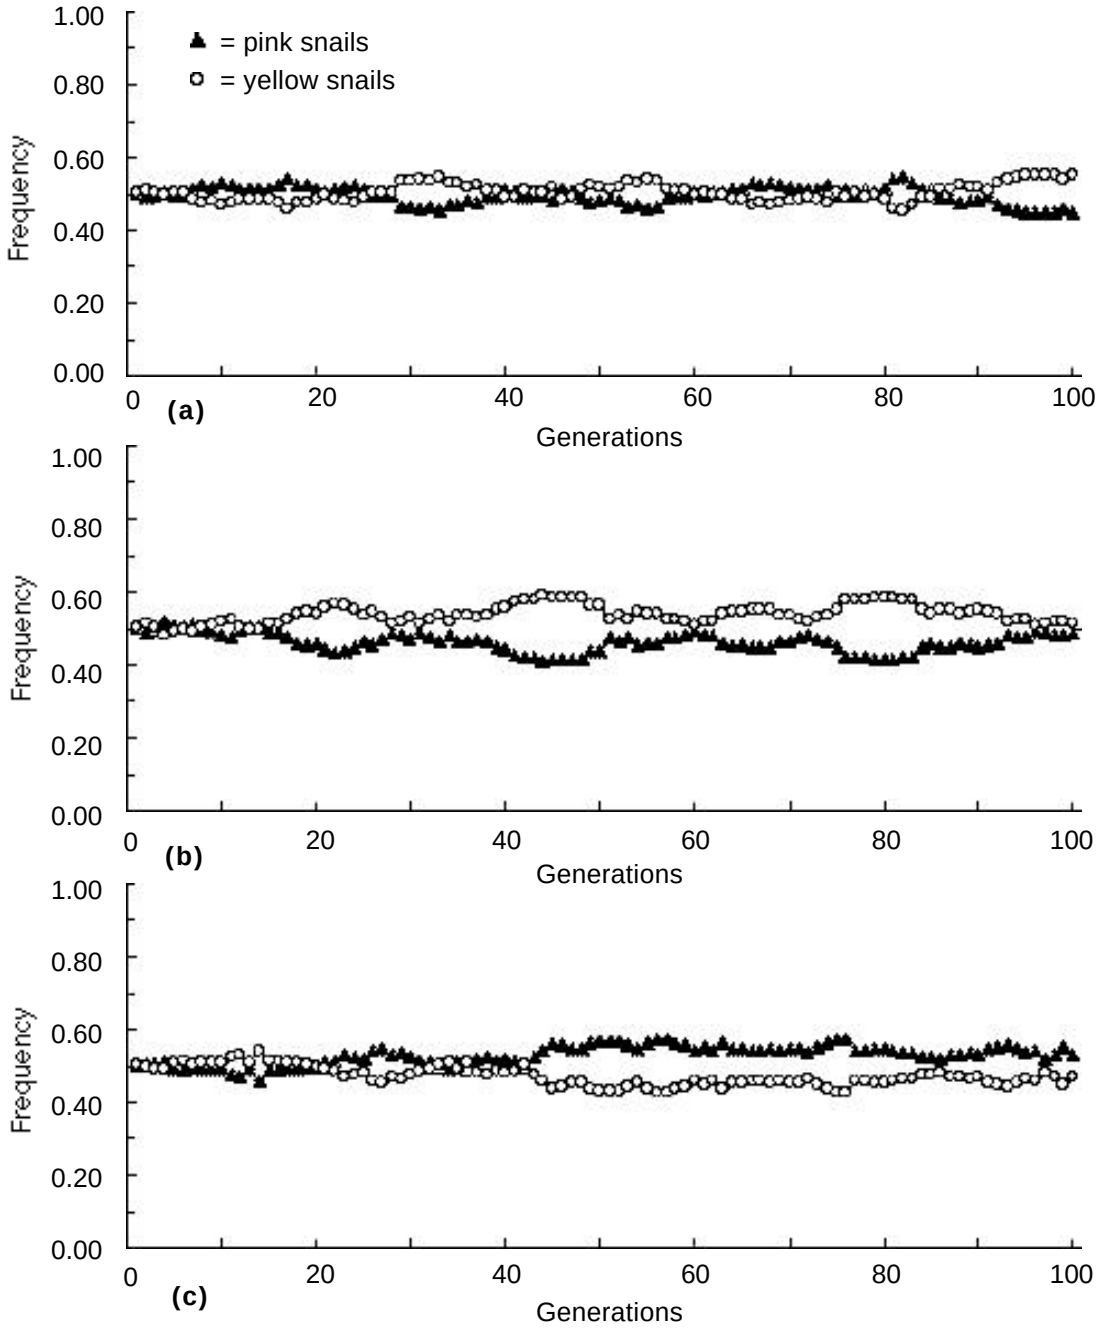

Figure 3.5: Three sample plots of large populations undergoing random drift in the absence of selection. Intially, all three populations contain equal numbers of pink and yellow snails. After 100 generations, yellow snails are more prevalent in populations (a) and (b) and pink snails are more prevalent in population (c). Overall, however, there is very little change in any of the populations, and they are relatively similar to one another.

These examples show that random drift leads us to expect a pattern whereby smaller populations are less uniform and larger populations are more uniform. Moreover, consideration of other possible causes shows that they are unlikely to produce this outcome. Thus we can conclude that it is random drift which has increased the probability of the outcome. If this is the case, then chance, in the form of random drift, *is* a contributing cause according to Humphreys's definition. Ironically, this claim contradicts Humphreys's belief that chance lacks causal efficacy. If a list of terms referring to contributing causes of an outcome is an explanation of that outcome, then random drift is a possible member of that list when the outcome is of an evolutionary nature. In other words, chance in evolutionary biology *is* explanatory, despite Humphreys's claim that chance (in general) is not.

At this point, it is appropriate to consider two possible counter-arguments to my contention that random drift is causally efficacious and explanatory. The first counter-argument is that random drift is not the appropriate "level" at which to provide an account of evolutionary change. Rather than citing random drift as the cause, according to this counter-argument, one should cite the causal history regarding the lives, deaths, and reproductive successes of individual organisms. However, as I argued in Chapter 2, this "individual-level" explanation fails to completely explain the inherently "population-level" process of evolution. Only population-level explanations will suffice – and random drift provides population-level explanations. Thus citing random drift as the population-level cause *is* invoking causality (and explanation) at the right level.

The second counter-argument maintains that there is a way of understanding random drift that is not causal, although it may be explanatory.[26] On this view, Humphreys's account of causality may be too permissive, so that whereas random drift may be considered to be causally efficacious under Humphreys's account, it would not be considered so under some other (more restrictive) account. Thus, one might argue random

[26] I owe thanks to Robert C. Richardson, a commentator on version of this chapter presented at the APA Pacific meetings in March, 1997, for pointing out this possible counter-argument. Richardson holds that while random drift is explanatory, it is not (under a proper conception of causality) causal (personal communication).

drift is a simply theoretical framework for dealing with changes in the frequency of genes (or genotypes or phenotypes) in finite populations; it describes general patterns of change, or expected patterns of change. Given certain parameters of a population (such as its effective size and the selection coefficient) we can determine whether random drift increases the chance of a given change, and by how much. However, we can do this without invoking specific causes in a given case, i.e., without giving details of the causal history. On this view, random drift, as a theoretical framework, may be explanatory, but it is not causal.

It is certainly true that my argument for the causal nature of random drift is dependent on the soundness of Humphreys's account of causality. As a defense of Humphreys would take us too far afield, I will point out only that 1) Humphreys's account is one which has some initial plausibility; that is, it is a reasonable account of causality, and 2) One cannot hold an account of causality such as Humphreys while maintaining that random drift is not causally efficacious or explanatory. However, the second counter-argument as it stands is not successful against my position. Random drift is more than a theoretical framework – more than just a model of population genetics. As we saw in Chapter 1, random drift is an actual process, or rather, set of processes, such as indiscriminate parent sampling, indiscriminate gamete sampling, and the founder effect. Each of the kinds of random drift we discussed constitutes a physical, causal process, involving individuals (indiscriminate parent sampling, founder effect, bottleneck effect, Brandon effect), gametes (indiscriminate gamete sampling), or genes (gene sampling). So when we say that random drift has been a factor in a given change in a population, we are implying that a particular process has been occurring in the population. In so doing, we are, I would argue, implicitly making a causal claim, even though it is a very vague causal claim. Furthermore, biologists often do make their causal claims less vague, by specifying the kind of random drift that is occurring in the population, and providing details about the environment and the organisms in question (see Chapter 2 for discussion). To the extent that their causal claims are less vague, I would maintain, more explanatory. It may not be necessary to provide such details, but to do so provides us with a deeper explanation, by

telling us *why* the changes occurred as they did.

Because random drift is more than just a theoretical framework – because it is a set of real processes operating in nature – random drift is a causal process, and thus should be considered to be a explanatory under a causal account of explanation.

3.4. The Explanatory Use of Random Drift

In section 3.3, we saw that random drift can explain two kinds of phenomena: the fluctuation of gene frequencies from generation to generation, and a pattern where small populations have considerable genetic divergence, and large populations have limited genetic divergence. However, evolutionary biologists invoke random drift to explain a number of related phenomena. Shanahan elaborates on three of these: 1) nonadaptive traits, 2) rate of evolutionary change, and 3) the rate of evolution of neutral alleles (Shanahan 1991). Since Shanahan's focus is primarily historical, he discusses the first two in the context of Sewall Wright's shifting balance theory. The third phenomenon is discussed in a section providing an overview of the use of chance in explanation in contemporary biology, and is explored in the light of Motoo Kimura's neutral theory. Shanahan's work is important because it demonstrates the extent of the role played by random drift in evolutionary theory (that is, random drift is invoked to explain a number of different kinds of phenomena) and illustrates the explanatory nature of random drift. In this section, I will take up where Shanahan left off by examining yet another phenomenon random drift can explain: the fixation of an allele within a population. Understanding the connection between random drift and fixation is important because it shows the potential for random drift to produce long-lasting evolutionary effects. The particular experiment that I analyze in this section is not a recent one; as the careful reader will have noted, neither is the Cavalli-Sforza study discussed above. I do this in order to focus on "classic" studies, the kind that present-day biologists refer to in their textbooks, experiments, and studies of natural populations. The goal is to show how random drift experiments proceed, rather than to provide a snapshot of the current "state of the art". Thus, in this

section, I seek to show the canonical way in which random drift is used to explain the fixation of an allele within a population.

Previously, we discussed the way in which random drift in small populations leads to genetic variability *between* populations. However, random drift can also cause the *loss* of genetic variability *within* a small population (that is, a decline in the heterozygocity of the population). Sometimes, random drift can cause the complete fixation of an allele within a population; that is, it comes to occupy 100% of all the available sites, so that all the members of the population are homozygous for the same allele. For example, in a laboratory study performed by Buri (1956) of just such a loss of genetic variability in populations of *Drosophila melanogaster* was observed and attributed to random drift. The study examined two culture series, the first consisting of 107 populations, and the second consisting of 105 populations. Each initial population was quite small, consisted of 16 heterozygote flies, bw^{75}/bw (8 male and 8 female). Subsequent generations were formed by choosing 8 males and 8 females at random from the progeny of the parental generation, so that the size of each population remained constant. Over the course of nineteen generations, an increasing number of populations became fixed for one allele or the other (see Table 3.1). By Generation 19 (considering the initial generation as Generation 0), the first culture series contained 30 populations which were fixed for the *bw* allele (consisted solely of bw/bw homozygotes) and 28 populations were fixed for the bw^{75} allele (consisted solely of bw^{75}/bw^{75} homozygotes), leaving only 49 out of 107 populations which were still heterozygotic. The second culture series produced similar results, although somewhat less striking. After 19 generations, 13 populations were fixed for the *bw* allele and 30 generations were fixed for the bw^{75} allele, leaving 62 out of 105 unfixed populations.

Generation	Total fixed b_W	Total fixed b_W^{75}	Total fixed b_W	Total fixed b_W^{75}
0	0	0	0	0
1	0	0	0	0
2	0	0	0	0
3	0	0	0	0
4	0	1	0	0
5	0	2	0	1
6	1	3	0	2
7	3	3	0	4
8	5	5	1	8
9	5	6	2	9
10	7	8	3	10
11	11	10	3	11
12	12	17	3	13
13	12	18	5	15
14	15	21	5	20
15	18	23	7	21
16	23	25	7	23
17	26	26	11	24
18	27	28	13	30
19	30	28	13	30
		Series I		**Series II**

Table 3.1: Number of fixed populations over the course of 19 generations. Series 1 consists of 107 populations total; Series II consists of 105. From Buri (1956).

Buri argued that these results were due to the effects of random drift (Buri 1956). But why did he come to this conclusion – how is it that random drift can account for the fixation of a type within a population? After all, it might seem as though random drift would always produce results where the frequencies fluctuated from generation to generation, as we saw above. How is it that random drift can produce results that are

seemingly *directional* (without being in any *particular* direction)?

The answer has to do with the kind of process that random drift exemplifies. Random drift is an example of a *Markov process*. A Markov process is a process in which "future development is completely determined by the present state and is independent of the way in which the present state has developed" (Feller 1968: 420). Thus, the particular history that leads to a given outcome is irrelevant; a given future outcome depends only on the present outcome and not on any of the proceding outcomes. Random drift is a particular kind of Markov process, called a random walk. The easiest way to understand a random walk is through an ideal coin-tossing game. Let heads represent an increase of one unit, tails a decrease of one unit, and the probability of heads or tails be equal. (This probability is know as the *transition probability*, or the probability of changing from the current state to a subsequent state). We can then describe a sequence of "winnings" (we start at zero) $s_1, s_2, \ldots, s_n$, where each s_i is the total amount "won" (positive or negative) after the *i*th trial (Feller 1968). This is a Markov process because it does not matter what sequence of events has led to the present outcome; for example, the sequence *heads*, *heads*, *heads*, *tails* is equivalent to the sequence *heads*, *tails*, *heads*, *heads*; in both cases, our winnings have the value 2, and in both cases, our possible values of our winnings after the next toss are either 1 (if we get tails) or 3 (if we get heads).

Note that in our ideal coin tossing game, for each trial *i*, s_i depends on the previous outcome s_{i-1} and the transition probability, which is 0.5 for an increase of 1 (heads) and 0.5 for a decrease of 1 (tails). So, for example, if our winnings are currently 7 (which would imply that we had turned up heads 7 more times than we had turned up tails), there are two possibilities for our next toss of the coin: 6 (if the coin turns up tails) and 8 (if the coin turns up heads). However, since heads and tails are equally likely, it is just as likely that we will have winnings of 8 as it is that we
will have winnings of 6; there is no tendency to "equalize out" the number of heads and tails. The outcome of the previous toss thus constrains the outcomes of future tosses. Each *toss* is independent, but the *outcome* of each toss, since it is added to the outcome of the previous toss, is not independent. Thus it can happen that outcomes can "walk" in one

direction or another; the values of our coin tosses will probably fluctuate somewhat, but the value may continue to increase overall (although presumably not indefinitely).

In fact, our winnings may not return to 0 as quickly as one might (intuitively) think. Our intuition suggests that, in a long series of tosses, our winnings would fluctuate from negative to positive values frequently, so that in any series of tosses, it is likely that our winnings recently had a value of 0. However, as Feller has demonstrated, in any sequence of $2n$ trials, where the $2k$th trial is the last trial where the number of heads and tails is equal, it is just as likely that $k = i$ as $k = n - i$.[27] So, for example, in sequence of 100 coin tosses, it is just as likely that the last point our winnings equalized was at the 98th toss as it is that the last point our winnings equalized was at the 2nd toss. This somewhat non-intuitive result implies that "it is quite likely that in a long coin-tossing game one of the players remains practically the whole time on the winning side, the other on the losing side" (Feller 1968: 81). For example, for a large number of tosses, the probability that our winnings are positive for 96.7 percent of the time is 0.2; the probability that our winnings are positive for 99.4 percent of the time is 0.1 (Feller 1968: 82).[28]

Random drift is analogous to the ideal coin tossing game in the following way. Recall that random drift is a process in which physical differences are causally irrelevant to differences in reproductive success. So, when random drift occurs in the absence of selection, in any given generation it is equally likely that one genotype will be as successful as another genotype (just at it is equally likely that we would get heads or tails). Furthermore, if we consider each generation as a trial and the number of alleles as the outcome of that trial, the number of alleles in each generation depends only on the number of alleles of the previous generation and the various transition probabilities T_{ij}, (where i is the present number of A alleles, j is the number of A alleles in the next generation, and T_{ij}, is the probability of going from state i to state j). The transition probability is defined by the binomial distribution:

[27] We use a value of $2k$ since, of course, our winnings can only be equal when there has been an even number of tosses.

[28] Likewise, the probability that our winnings are *negative* for 96.7 percent of the time is 0.2; the probability that our winnings are negative for 99.4 percent of the time is 0.1.

$$T_{ij} = [(2N)!/(2N - j)!j!] (i/2N)^j (1 - i/2N)^{2N - 1}$$ (Roughgarden 1996: 65-6)

To get a general idea of the properties of this distribution, let us look at a transition probability matrix for a population of size 3 ($2N$ = 6 genes) containing the allele *A*:

Number of *A* alleles in generation *t*

Number of *A* alleles in generation $t + 1$

	0	1	2	3	4	5	6
0	1.000	0.000	0.000	0.000	0.000	0.000	0.000
1	0.335	0.402	0.201	0.054	0.008	0.001	0.000
2	0.088	0.263	0.329	0.219	0.082	0.016	0.001
3	0.016	0.094	0.234	0.312	0.234	0.094	0.016
4	0.001	0.016	0.082	0.219	0.329	0.263	0.088
5	0.000	0.001	0.008	0.054	0.201	0.402	0.335
6	0.000	0.000	0.000	0.000	0.000	0.000	1.000

Table 3.2 : Transition probability matrix for a population of size $N = 3$ ($2N$ genes). Numbers generated by computer program and rounded to the third decimal place.

Note that for each state *i* at *t*, the most probable state at $t + 1$ is *i*. (That is, the most likely result is that the number of alleles will not change from one generation to the next). However, other transitions occur with a relatively high probability. In particular, note that when allele *A* is predominant in the population, it is still fairly likely that it will increase in frequency in the next generation. For example, in state 4 at time *t*, the probability of changing to state 5 at time $t + 1$ is 0.263, whereas the probability of changing to state 3 at time $t + 1$ is 0.219. Similarly, when allele *A* does *not* predominate in the population, it is still relatively likely that it will decrease in the generation. For example, in state 2 at time *t*, the probability of going to state 1 at time $t + 1$ is 0.263, whereas the probability of going to state 3 at time $t + 1$ is only 0.219. (States 0 and 6 are special cases, called *absorption states*, and will be discussed further below).

Thus, as with the ideal coin tossing game, there is no tendency to return to the equalization point, which for random drift is the point at which all alleles are represented equally (in a population where there are 2 alleles at a locus, this is the point at which there are N copies of each allele, or half of the total of $2N$ alleles). When an allele predominates

in the population, it may still continue to increase in the population (or it may decrease). Also, since any future generation will be formed from the previous generation, the frequencies of the present generation will tend to constrain the values of the future generation. Since frequency values, unlike coin toss values, can change by variable amounts at each time *t*, future values are not as constrained by past values in the random drift case as in the coin tossing case. However, the smaller the changes in frequency from generation to generation, the more future values will be constrained by the value of the previous generation. Larger frequency changes in a population mean that frequency values may return towards the equalization point in just one generation, but they also mean that the frequency may diverge even farther from the equalization point. Thus, as with the ideal coin tossing game, it is possible for values to "walk" in one direction or the other. Furthermore, since random drift is a random walk process like the coin tossing game, the chances of one allele predominating for a long period of time are reasonably high.

There is another important difference between the coin tossing game and the process of random drift. In the coin tossing game, our "winnings" did not have an upper or lower bound; if we were very lucky, our winnings could increase indefinitely (or, if we were very unlucky, decrease indefinitely). However, the frequencies of organisms in a population do have an upper (and lower bound). As one allele comes to predominate in a population, the population becomes increasingly homozygous. Once all the members of a population are homozygous for an allele, the frequency of that allele cannot increase any further (it has a frequency of 1, i.e. there are $2N$ copies of the allele), while the frequency of the other allele (in a population with two alleles at a locus) cannot decrease any further (it has a frequency of 0, i.e., there are no copies of the allele remaining). More importantly, once an allele is fixed, new alleles can only be introduced through mutation or migration. So, once gene frequencies walk to the point where one allele is fixed (and the other allele is eliminated entirely), they will remain in a fixed state unless a new allele is introduced through a mutation or the migration of an individual with a different allele into the population. In Table 3.2, fixation is represented by the absorption states 0 and 6. If a population is a state 0 at time *t*, it has a probability of 1 of staying in the same state, and a

probability of 0 of going to any of the other states (and similarly for state 6). This property of population frequencies, combined with the possibility that a gene frequency may "walk" to fixation, implies that eventually, a population undergoing random drift in the absence of selection will become fixed for one allele or the other. If one considers a suite of populations, all undergoing random drift in the absence of selection, over the course of generations an increasing number of the populations will become fixed for one allele or the other, until eventually all alleles are fixed (Roughgarden 1996: 58).

We are now in a position to understand how random drift can explain the results of the Buri study discussed above. In the Generation 0, all populations contained equal numbers of *bw* alleles and bw^{75} alleles; no populations were fixed. By the 10th generation, 15 populations were fixed in Series I (7 for *bw* and 8 for bw^{75}) and 13 populations were fixed in Series II (3 for *bw* and 10 for bw^{75}). By the 19th generation, as we already noted, 58 Series I populations were fixed and 43 Series II populations were fixed (See Table 3.1). If we think of random drift as a process where gene frequencies fluctuate from generation to generation, it is hard to see why the experimenters would attribute these results to random drift; instead, we might expect that random drift would produce populations which differed from each other, but where gene frequencies continued to fluctuate indefinitely. However, once we understand the properties of a random walk process such as random drift – that one allele can predominate over the course of many generations without the population returning to the equalization point – combined with the biological fact that populations can only leave a state of fixation through mutation or migration, it is easy to see how random drift can explain the results of the Buri study. In fact, the results are qualitatively exactly what you would expect from the properties of random drift that we discussed above; over the course of generations, an increasing number of populations become fixed for one allele or the other. Thus random drift *can* explain results that are seemingly directional (i.e., not just random, fluctuating noise), although clearly random drift is not directional in the sense that one direction is expected over the other.

It is important to note here that random drift is not the only explanation for fixation within a population. For example, selection can also explain the loss of heterozygocity within a population, leading to the fixation of a type.[29] Natural selection, a process where fitness differences between organisms are causally relevant to differences in reproductive success, tends to lead to an increase in the fitter type over the course of generations. Over time, the fitter type may become fixed in the population. Indeed, Buri (1956) is careful to consider whether natural selection can explain all or part of his results, although he rejects this hypothesis. In fact, it would be unlikely that natural selection were wholly responsible for Buri's results, given that large numbers of populations were fixed with the *bw* allele as well as the bw^{75} (if the results were due primarily to natural selection, we would expect that the same allele would have been fixed in most of the populations).[30] My goal, then, is not to show that random drift can replace natural selection as an explanation for fixation, but rather to point out that random drift *can* explain fixation, and that certain patterns of fixation are more likely to be caused by random drift than by natural selection. In addition, once we understand that random drift can cause fixation, it becomes clear that random drift can have concrete evolutionary consequences. These considerations point to the importance of the explanatory role played by random drift.

3.5. Conclusion

I've argued here that random drift has causal efficacy, and thus, under an account of causal explanation, can be considered to be explanatory. This implies that we need to broaden our conception of chance to include chance that is both causal and explanatory. Evolutionary biology thus provides us with a reason to rethink some of our traditional

[29] Natural selection is also a Markov process. Thus, the discussion of Markov processes does not serve to distinguish random drift from natural selection (that was done in Chapter 1). Rather, the goal is simply to explain how it is that random drift can explain fixation, since, at first glance, it seems counter-intuitive that random drift would lead to fixation. On the other hand, the idea that natural selection might lead to the fixation of the fitter type seems perfectly straightforward, and thus is in no need of a special explanation.

[30] On the other hand, it is possible that natural selection could be partially responsible for the results; in particular, the fact that in Series II only 13 *bw* alleles went to fixation by the 19th generation, whereas 30 bw^{75} populations went to fixation. However, Buri (1956) rejects this hypothesis.

philosophical ideas about chance.

The phenomena I have discussed here – the random fluctuation of gene frequencies from generation to generation, the pattern of large genetic divergence between small populations and small genetic divergence between large populations, and the fixation of alleles within a population – as well as other kinds of phenomena I haven't discussed, such as the production of nonadaptive traits, the rate of evolution of neutral alleles, and molecular evolution, are all phenomena which should be viewed as the effects of random drift. (See Shanahan 1991 for discussion of these other phenomena). As a result, they are all phenomena which random drift can explain; furthermore, they are phenomena which biologists have invoked random drift to explain. These explanations fill out our picture of the explanatory nature of random drift; they show the way in which random drift is explanatory. Moreover, the multitude of the phenomena explained by random drift is an indication of the importance of the explanatory role of random drift within evolutionary theory.

What this chapter has shown is that random drift plays a positive, important role in evolutionary theory. According to Hartl and Clark, two of the empirical goals of population genetics are "…determining how much genetic variation exists in natural populations and…explaining this variation in terms of its origin, maintenance, and evolutionary importance" (Hartl and Clark 1989: 4). In this chapter, we have seen that random drift can explain how variation can be maintained in a population (at least in the short term, by producing fluctuations in gene frequencies from generation to generation) and that random drift can explain the genetic divergence between small, isolated populations. We have also seen how random drift can explain the loss of variation within a population, which raises questions for the maintenance of variation in any population undergoing a significant amount of random drift. Thus, the explanatory role of random drift acts to further the overall goals of population genetics.

Chapter 4: Stochasticity in Macroevolution

> Stochastic explanations have been offered for evolutionary patterns, but why be defeatist when there are credible deterministic alternatives? (Thayer 1979: 460).
>
> The purpose of this present article is to support the view that chance (and its statistical counterpart when large numbers of chances are involved, namely stochastic processes) forms the first-order pattern of organization for the history of life (Schopf 1979: 337).

4.1. Introduction

Whereas microevolution concerns changes*within* populations and species, macroevolution concerns the *origination* of new species, genera, families, orders, etc. (usually through the splitting of lineages) and the *extinction* of species and higher categories (i.e., the termination of lineages). Paleontologists (biologists who study fossils) are often interested in the processes which have led to the formation and extinction of species and higher categories as exhibited in our fossil record (part of the sub-discipline of paleontology known as paleobiology). In 1972, a group of paleontologists held two "informal meetings" at the Marine Biological Laboratory in Woods Hole, Massachusetts. The outgrowth of these and other subsequent meetings was a series of papers suggesting that stochastic processes might play a large role in the evolution of species and higher orders. The main tool of the researchers (who came to be known as the "Woods Hole Group"[31]) was a computer program known as the MBL program, which used a Monte Carlo approach to model the macroevolutionary processes of speciation and extinction. Comparing the results of the model to real world paleontological data, the researchers called into question the ubiquity of so-called "deterministic" explanations for macroevolutionary phenomena, arguing that the same data could be accounted for by

[31] The first paper produced by this group was written in 1973 by David Raup, Stephen Jay Gould, Thomas Schopf, and Daniel Simberloff. Subsequent papers include Raup and Gould (1974); Schopf, Raup, Gould, and Simberloff (1975); Gould, Raup, J. John Sepkoski, Jr., Schopf, and Simberloff (1977); Raup (1977); and Schopf (1979).

stochastic processes.[32]

The stochasticity of macroevolutionary processes has gone largely unexamined by philosophers of biology; analyses of the role of chance in evolution focus primarily on random drift and mutation, quintessentially microevolutionary phenomena. Yet the research of the Woods Hole group is striking in what it reveals about the concept of chance used in evolutionary theory, about the extent of the role of chance in evolution, and about the explanatory nature of chance. Furthermore, there are many interesting parallels between random drift and stochasticity at the macroevolutionary level. In this chapter, I will explore these issues, and use the parallels between the stochasticity of macroevolution and random drift to support previous conclusions concerning random drift.

The structure of the chapter is as follows. First, I describe the stochastic model, contrasting it with deterministic explanations, and discuss the ways in which the stochastic model provides an alternative to deterministic explanations. Second, I explore the sense in which the stochastic account is stochastic. Third, I show how it is that the stochastic account can explain phenomena that have been traditionally explained deterministically, and demonstrate how stochastic explanations account for specific kinds of macroevolutionary phenomena. Last, I summarize the similarities between the random drift and stochastic macroevolutionary processes, and suggest some broader conclusions about stochasticity within evolutionary theory.

4.2. The Stochastic Model: An Alternative to Deterministic Models

Before we examine stochastic macroevolutionary models, it will be helpful to have an understanding of traditional, deterministic explanations. Peterson (1983) provides an example of this kind of explanation. Peterson argues that four genera of iguanid lizards

[32] I say "so-called" because deterministic explanations are not deterministic in the Laplacean sense. Similarly, stochastic accounts do not deny Laplacean determinism. Rather, both deterministic and stochastic explanations are consistent with either Laplacean determinism or indeterminism. The terms "deterministic" and "stochastic" will be discussed further below. Throughout this chapter, I will use the term "deterministic" when referring to the kind of macroevolutionary explanation under discussion and "Laplacean determinism" when discussing determinism of a philosophical kind.

(*Anolis*, *Chamaeleolis*, *Chamaelinorops*, and *Phenacosaurus*, referred to as anoline lizards or anoles) have been more successful than other iguanid lizard genera due to their possession of a subdigital adhesive pad. Peterson states:

> The pad complex appears to be a key innovation (*sensu* Liem, 1973) in the successful radiation of the group. Compared to animals that must rely on their claws and toe position for grip, anoles appear to have greater range of usable locomotor substrates. They can grip or adhere to a perch of almost any diameter (or radius of curvature), spatial orientation, or surface texture (e.g., surfaces as smooth as glass or as rough as bark). It may become possible to state the relationship between the pad complex and physical aspects of the habitat more precisely in the future, but the adhesive pad seems to be among the morphological adaptations that permit anoles to adopt highly acrobatic locomotor behavior and to exploit the spatial heterogeneity of arboreal niches (Peterson 1983: 245).

In this explanation, the phenomenon being explained is the successful radiation of the anoline lizards, i.e., the greater number of anoline lizards species as compared to other species of iguanid lizards. (In macroevolutionary contexts, a "radiation" is a marked increase in the number of sub-groups within a group; in this case, the radiation in question is the increase in the number of species in the four anoline lizard genera). According to Peterson, there are more anoline lizard species than other iguanid species because the possession of a subdigital adhesive pad provided the anoline lizards with a superior grip, giving them greater agility and enabling them to more fully explore their habitat as compared to lizards which must rely on claws and toe position. This is referred to as a "deterministic" explanation because it invokes a specific cause (the presence of the sub-digital adhesive pad) as an explanation for the macroevolutionary phenomenon in question (the greater radiation of anoline lizards). Such an explanation is "taxonbounded" in that it does not treat all taxa alike; rather, it claims that the physical characteristics of some taxa (the anoline lizard genera) caused them to have a higher rate of branching than other taxa (non-anoline iguanid lizard genera).

Taxonbounded explanations are one kind of deterministic explanation. Another kind of deterministic explanation is a "timebounded" explanation (Gould et al. 1977). Timebounded explanations do not treat all times alike. For example, one could give a

timebounded explanation for why mammals were a very small group up until the Tertiary period; prior to that time, they were kept small by a successful competitor: the dinosaurs. Once the dinosaurs became extinct, times were "good" for mammals, and the number of genera expanded (Gould et al. 1977: 37). Because this explanation presupposes that some times are "better" for a group then others, it is considered to be timebounded. Other timebounded explanations claim that taxa have specifiable life histories; for example, that older taxa are diverse early in their history, whereas younger taxa in the same group are diverse late in their history (Raup 1985: 42). Since these explanations invoke specific causes (in this case, specific time periods) as the explanation for the phenomenon in question, they too are considered to be deterministic explanations.

In contrast to deterministic explanations, stochastic models of macroevolution are neither taxonbounded nor timebounded; that is, they "treat all times and taxa alike" (Gould et al. 1977: 25).[33] In the discussion that follows, I will show how the concepts of untimeboundedness and untaxonboundedness are embodied by the model used in the MBL program. Although elements of this model have been criticized since it was first proposed, and alternative models have been proffered, the MBL model (as described in Raup et al. 1973, Gould et al. 1977, and Raup 1977) continues to be cited in textbooks (see, e.g, Futuyma 1986, Ridley 1993) and recent papers (see, e.g., Uhen 1996, Guyer and Slowinski 1993). My focus will therefore be on the model used by the Woods Hole Group in order to: 1) describe the "classic" stochastic model from which other models are derived, and 2) to delineate the basic elements of stochastic models (the untimebounded, untaxonbounded, and Markovian qualities) without getting bogged down in the details of the differences between different models. With these goals in mind, in the description of the MBL program that follows, I will discuss the details of the program only insofar as they pertain to the discussion of stochasticity. For further details of the program, see Raup et al. (1973) and Gould et al. (1977).

The MBL programattempts to model the evolutionary process by producing

[33] It should be noted that I am contrasting a kind of explanation that has commonly appear in the literature (deterministic explanation) with a stochastic *model* that can be used for evolutionary explanations. In other words, there is no specific "deterministic model", although such models have been developed (see Raup 1985 for discussion).

simulated phylogenetic trees. (A phylogenetic tree is a tree that attempts to depict the genealogical history of a group of organisms). It begins with a single lineage, from which a "life history" is generated. At each time interval, and for each current lineage, one of three things occurs: 1) The lineage becomes extinct, 2) The lineage splits into two daughter lineages, or (3) The lineage persists to the next interval without branching. Which of these three fates occurs is determined by a computer-generated random number, based on predetermined probabilities. There are two main versions of the MBL program. In one version, the "freely floating" model, the probabilities of branching and extinction are equal (probability of branching = probability of extinction = 0.1) and unchanging throughout the simulation (Gould et al. 1977). In the other version, the damped-equilibrium model, there is an optimum number of co-existing lineages for any interval (the same optimum number applies to all intervals), and the probabilities of branching and extinction are adjusted to obtain this optimum value.[34] This adjustment occurs by raising the probability of branching when the number of lineages is below the optimum number, and by raising the probability of extinction when the number of lineages is above the optimum number. Thus, in the initial phase of the program, where the number of lineages is few, the probability of branching will exceed the probability of extinction, causing an "explosion" of lineages (Gould et al. 1977).

Both of these stochastic models can be said to be "untaxonbounded" and "untimebounded" to some extent. Both are untaxonbounded because the same probabilities of branching and extinction apply to all current, endpoint lineages in the simulation. Thus, all taxa are treated alike. This contrasts with the taxonbounded, deterministic explanation given by Peterson, where, the rate of extinction was less for some taxa (the anoline lizards) than it was for other taxa (non-anoline lizards), and the rate of branching was greater for some taxa (the anoline lizards) than it was for other taxa (the

[34] It might seem odd that one of the stochastic models is an equilibrium model. The reason for the use of an equilibrium model is to see if the successful use of equilibrium models in ecology and population biology can be replicated in paleontology (Raup et al. 1973). Furthermore, such a model may be more biologically realistic: "The maintenance of an equilibrium diversity in the present work implies that an adaptive zone or a geographic area becomes saturated with taxa and remains in a dynamic equilibrium determined by the opposing forces of branching (speciation) and extinction" (Raup et al. 1973: 529).

non-anoline lizards).

The freely floating model is completely untimebounded; the probabilities of branching and extinction remain the same throughout all the intervals in the simulation.[35] The damped-equilibrium model is untimebounded in the sense that the optimum number of lineages applies to all time intervals. However, it is not completely untimebounded because not all time intervals are alike; as previously mentioned, in the initial phase of a simulation the probability of branching exceeds the probability of extinction. However, for the rest of the simulation the model is untimebounded; even though the probabilities of branching and extinction change, they do not do so with respect to any particular phase of the simulation. Thus, all times are treated alike. This contrasts with the timebounded, deterministic explanation given above, where prior to the Tertiary period, mammals were not very successful (high extinction rate, low speciation rate), but once dinosaurs became extinct, times were "better" for mammals (low extinction rate, high speciation rate).

Thus, explanations which are *both* untimebounded and untaxonbounded are referred to as stochastic; explanations which are *either* timebounded *or* taxonbounded are referred to as deterministic. The use of the term "deterministic" is somewhat unfortunate here; it does not refer to the traditional philosophical concept of Laplacean determinism, but rather to the kind of macroevolutionary explanation which refers to a specific cause as the explanation of an event. Furthermore, the contrast between "deterministic" and "stochastic" seems to imply that the stochastic account rests on a notion of ontological randomness, but this is not necessarily the case (as will be explained below). So-called "deterministic" explanations might be better termed "determinate cause" explanations; they make reference to specific causes related to the properties of particular lineages or time periods. Although this usage of determinism may be somewhat confusing for philosophers, I retain the usage here for consistency with the paleontological literature.

Whatever one calls them, deterministic explanations have been (and perhaps continue to be) the predominant mode of explanation within the field of paleontology. Stochastic

[35] In the MBL model, the probabilities of branching and extinction are equal as well as constant. Other stochastic models allow the probabilities of branching and extinction to differ from one another (Raup 1985).

models thus present a challenge to this traditional method of explanation, and provide a new way to look at preexisting data. Since the synthesis of the 1940's and 1950's, it has generally been argued that the processes responsible for microevolution, such as natural selection and random drift, are responsible for macroevolution as well. Deterministic explanations are completely consistent with this thesis; the explanation given above for the success of the anoline lizards is simply natural selection, writ large. In one sense, stochastic models of macroevolution do not break with this general thesis either. Thus, the Woods Hole Group asserts that, "[t]he model makes no assumptions about the evolutionary mechanisms involved, such as natural selection and mode of speciation" (Raup et al. 1973: 528). In another sense, however, stochastic models provide a distinctively macroevolutionary explanation, one that cannot simply be extrapolated from microevolutionary processes, according to Gould (Gould 1995). The Woods Hole group sought to create a "nomothetic" paleontology, i.e., "an approach to historical science favoring the study of 'cases and events as universals, with a view to formulating general laws'" (Raup et al. 1973: 526; quoting the Random House Dictionary). In so doing, they hoped to achieve the same success that similar approaches have had in other sciences:

> We confess to a larger, ulterior motive in this and related studies (Raup and Gould 1974; Schopf et al., 1975). We believe that paleontology – the most inductive and historical of sciences – might profit by applying some deductive methods commonly used in the non-historical sciences (without sacrificing its important documentary role for the history of life). We may seek an abstract, timeless generality behind the manifest and undeniable uniqueness of life and its history. We take as our guide the recent success of simple, general models in the other branch of natural history most celebrated for the complexity and uniqueness of its subject – ecology (Gould et al. 1977).

Thus, the stochastic models which the Woods Hole Group proposed are untimebounded and untaxonbounded (true for all time and all taxa) in order to provide models having the character of scientific laws, similar to "laws" proposed in ecology, and, presumably, in sciences such as physics. On a small scale, the effect is to provide an alternative to traditional deterministic explanations. On a larger scale, if Gould's motives can speak for the rest of the group, the stochastic models serve to sever the dependence of

macroevolutionary processes on microevolutionary processes by proposing distinctively macroevolutionary "laws". So, while stochastic macroevolutionary models may be *consistent* with microevolutionary processes such as natural selection and random drift, their existence means that paleontology has a theory "…which is derived from internal standards of the field rather than borrowing from biology or geology" (Schopf 1979: 350).

However, it should be noted that while stochastic models provide an alternative to deterministic explanations, the two kinds of explanation are not mutually exclusive. That is, it may be the case that *some* macroevolutionary phenomena are best explained stochastically, whereas other macroevolutionary phenomena are best explained deterministically. This is similar to the natural selection and random drift case, as we have seen in earlier chapters; although natural selection and random drift can sometimes explain the same phenomena, it is not the case that in accepting a random drift explanation one must reject all natural selection explanations (and vice versa). Beatty, in discussing the debate between neutralists (proponents of random drift) and selectionists (proponents of natural selection) suggests that there are both specific and general issues at stake in the debate. Beatty states:

> The more specific issues concern whether or not a change in frequency of a *specific set of genetic alternatives* can be ascribed to their selective neutrality and the consequent random drift of their frequencies. Complimentarily, one might expect that the more general issues concern whether *all or no* evolutionary changes – or whether all of a certain kind or none of a certain kind of evolutionary change – are due to random drift alone. It is important to recognize, however, that even the most general issues surrounding this version of the importance of random drift do not boil down to questions of all or none, but to questions of *more or less* (Beatty 1984: 199).

That is, the debate between proponents of random drift and proponents of natural selection is sometimes over whether a specific instance – a specific change in a population, for example – is due to natural selection or random drift. Or, such debates might discuss whether the specific evolutionary change was *primarily* due to random drift or *primarily* due to natural selection (as we discussed in Chapter 1). At other times, the debate between neutralists and selectionists is over whether, in general, most evolutionary changes (or, evolutionary changes of a particular kind) have been due primarily to natural selection, or

primarily to random drift.

With regard to debates between proponents of deterministic explanations and proponents of stochastic explanations, the situation is strikingly similar. Sometimes, disputants disagree over the proper explanation for a specific phenomenon. For example, Flessa and Imbrie (1973) and Flessa and Levinton (1977) argue that common diversity patterns across taxa during the Phanerozoic call for a deterministic explanation, whereas Smith (1977) maintains that these patterns can be accounted for by a stochastic model. (This example will be discussed further below). At other times, the debate concerns the more general issues. As the Woods Hole Group asks:

> How different, then, is the real world from the stochastic system? How, in other words, is the real world "taxonbound" and "timebound" – i.e., in need of specific, causal explanations involving uniqueness of time and taxon at various stages of earth history. The answer would seem to be "not very" (Gould et al. 1977: 32).

However, immediately following this claim that the "real world" can be largely explained stochastically, the Woods Hole Group proceeds to outline three instances in which deterministic explanations *are* required. Similarly, one group of deterministic proponents admits: "There are, of course, stochastic elements in the genesis and development of all clades, including those that we designate as [the consequence of some condition or event]" (Stanley et al. 1981: 117). However, these same proponents are careful to circumscribe the extent of those stochastic elements: "Only within small taxa, such as genera comprising a small number of species, do chance factors have more than a small probability of prevailing" (Stanley et al. 1981: 125). Thus, proponents on both sides assert that *their* form of explanation should predominate, while carefully circumscribing a small area where the other form of explanation is appropriate. Proponents of stochastic models do not assume that all macroevolutionary phenomena should be explained stochastically; the same is true of proponents of deterministic explanations. Beatty argues that "disputants [in evolutionary biology] defend the importance of their favorite modes of evolution without ruling the others entirely out of the question" (Beatty 1984: 207); the disputants in this debate do just that.

4.3. The Stochasticity of Stochastic Models

Above, I suggested that deterministic explanations in macroevolutionary theory are not deterministic in the Laplacean sense. So, if stochastic models are not a denial of Laplacean determinism, in what sense can stochastic accounts be said to be stochastic? In discussing the nature of stochasticity presupposed by stochastic accounts, one might think that the proper way to proceed would be to examine the phenomena which underlie the stochastic models in order to determine whether the underlying processes were themselves stochastic. Yet the proponents of the stochastic models deliberately avoid discussing the processes which underlie them; instead (as we noted above), they assert that, "[t]he model makes no assumptions about the evolutionary mechanisms involved, such as natural selection and mode of speciation" (Raup et al. 1973: 528). Raup and Gould, who propose a stochastic model that incorporates morphological changes (changes in the features of organisms), explain that they are "quite consciously avoiding the subject of evolutionary mechanisms"; they insist that their model is consistent with selection in randomly fluctuating environments as well as the fixation of mutations by random drift (Raup and Gould 1974: 309).[36] What their model does exclude, Raup and Gould assert, is unidirectional selection, but it makes no assumptions other than that. (Unidirectional selection would constitute a deterministic explanation). These claims suggest the somewhat radical conclusion that the proper way to go about examining the stochastic nature of the models is not to examine their underlying phenomena; to do so is to miss the point of the models entirely. The models do not consider underlying phenomena at all; they are concerned only with the macroevolutionary phenomena. Thus, the stochasticity of the models does not depend on the stochasticity of the underlying the phenomena; the underlying phenomena could be ontologically random, or Laplacean deterministic.

How, then, to proceed? As we saw above, stochastic models are characterized by

[36] It is interesting to note that because selection *may* underlie the stochastic model, the debate between neutralists and selectionists does not fall along the same lines as the debate between proponents of deterministic explanations and proponents of stochastic models.

explanations which are untimebounded and untaxonbounded. These characteristics themselves delineate a conception of chance; they imply that differences between taxa are causally irrelevant to differences in rates of branching and extinction within the taxa (this is the untaxonbounded element) and that different time periods are causally irrelevant to differences in rates of branching and extinction (this is the untimebounded element). In other words, branching and extinction occur randomly with respect to time interval and taxon. For example, differences between genera do not confer different rates of speciation and extinction upon different species of the genus (as would be the case in a deterministic, taxonbounded explanation, such as the one given for anoline lizards above); rather, origination and extinction occur randomly among the different *species* of the taxa.

The analogy between the "untaxonbounded" criteria of stochastic models and random drift here is striking; however, they are not identical phenomena. In random drift, physical differences between organisms are causally irrelevant to differences in reproductive success; that is, reproductive success of organisms occurs randomly with respect to their physical traits. An untaxonbounded explanation, on the other hand, claims that differences between taxa are causally irrelevant to differences in rates of branching and extinction within the taxa (differences between the properties of the higher taxonomic group are not causally relevant to the branching and extinction of the taxa of the lower category).[37] That is, branching and extinction occur randomly with respect to the taxonomic groups to which taxa belong. The phenomena are similar, but they are occurring at different levels; random drift occurs at the microevolutionary level, whereas the stochastic model we have been discussing pertains to the macroevolutionary level. Also, there does not seem to be an analog of "untimeboundedness" in the random drift case. The issue of whether some times are "good" or "bad" for organisms does not seem to arise in the microevolutionary case as it does in the macroevolutionary.

[37] Note that whereas the "success" of organisms is measured by the number of offspring they produce, the "success" of a taxa is measured by the number of taxa at the lower level. For example, one genus is more successful than another genus if it contains a greater number of species (implying that the species of the genus have had either a greater speciation rate, a lower extinction rate, or both). This is why the definition of untaxonbounded is multi-level (refers to different taxonomic categories), while the definition of random drift is not (in a given instantiation, refers only to organisms, or only to gametes, etc.).

So, stochastic models of evolution are stochastic in the sense that branching and extinction occur randomly with respect to time interval and taxon. But are stochastic models *inherently* stochastic (or probabilistic), especially when one considers that the processes which underlie them may be Laplacean deterministic? In Chapter 2, I argued that the status of natural selection and random drift as population-level processes implies that evolution is inherently a probabilistic process. Here I argue that a similar case can be made for stochastic models of macroevolutionary theory.

In denying that macroevolutionary events are determined by specific time intervals or membership in a taxonomic group, proponents of stochastic models do not thereby mean to imply that such events are uncaused or ontologically random (Raup et al. 1973); however, they do not deny that possibility either (Gould et al. 1977). Rather, they suggest that evolutionary events such as multiple extinctions are likely to be the result of *many unrelated or unconnected causes* (Slowinski and Guyer 1989), not a single cause as would be proposed under a deterministic explanation. For example, Raup and Marshall ask: "Did the order Creodonta die out because of the 'creodontness' of its genera (meaning small cranial volume and presumably low intelligence), *or* did the order die out just because all its genera went extinct through independent and widely varying causes?" (Raup and Marshall 1980: 10). Thus, the extinction of *each genus* of an order might have a single cause, but the causes would be unrelated ("independent" events), so that there would be no single cause of the extinction of the order *as a whole*. However, as with the random drift case, it is at the "higher" level where the stochastic process occurs, *not* at the level of the individual originations and extinctions.[38] If the causes of origination or extinction are independent, then it is solely by "coincidence", or chance, that the higher level macroevolutionary event occurs. It is in this sense that the claim that macroevolutionary events are due to chance should be understood.

By saying that a macroevolutionary event is the result of chance, in spite of the fact

[38] This can get confusing, since the stochastic model is supposed to apply to different taxonomic levels. Thus, a stochastic explanation might claim that the demise of a specific order was due to stochastic processes, although extinctions of its genera were due to independent and widely varying causes. A different application of a stochastic model might claim that the demise of a specific *genus* was due to stochastic processes, although extinctions of its *species* were due to widely varying and independent causes.

that it may be a caused event (albeit the result of a number of independent causes), proponents of the stochastic account are looking at the phenomena at a certain level of explanation. At the level of the individual occurrences of origination and extinction, we can talk about the individual causal histories leading to the origination of a new taxa or the extinction of a taxa. However, when viewed from higher taxonomic levels (where the process in question is occurring), the phenomena behave stochastically, because the causes of the higher level event are unrelated. The higher taxonomic event should not be given a deterministic explanation, proponents of the stochastic account argue, because there is no one cause of the *higher level* event. From the perspective of the higher taxonomic level, extinctions and originations are distributed randomly; the higher level event has occurred by a chance coincidence of events. As Schopf notes: "Determinism at the individual level fails when the questions asked are of a populational nature" (Schopf 1979, 342).

Proponents of the stochastic account suppose that if it is truly the case that the causes of origination and extinction are many, varied, and unrelated, then, over a long period of time, all biologic groups will be equally liable to branch and go extinct. This supposition is manifested in the MBL program by making the probabilities of branching and extinction identical across taxa. Extinctions and originations are thus considered to be randomly distributed across taxa. In other words, the actions of independent causes lead to a result that would be expected from chance. However, as we shall see in the next section, this random distribution does not preclude the production of seemingly nonrandom patterns.

4.4. "Non-randomness" Out of Randomness

Deterministic explanations have traditionally been provided for seemingly non-random[39] phenomena such as patterns of taxonomic diversity (for example, extreme increases or decreases in the number of taxa), the simultaneous increase or decrease of taxonomic diversity in a number of unrelated taxa, and morphological trends (directional

[39] I have deliberately not provided an analysis of the term "nonrandom". This reflects the fact that in some cases, biologists have looked at patterns that *appear to them* to be nonrandom, and as a consequence attempted to provide deterministic explanations for the patterns (although, of course, in other cases, the terms are used more precisely).

evolutionary changes in physical traits that take place over a lengthy time period). Proponents of stochastic accounts claim that their models can explain phenomena that were previously thought to require deterministic explanations (such as the radiation in the anoline lizard genera discussed above). Thus, it is often the case that stochastic explanations seek not merely to explain where deterministic explanations have been found lacking, but rather to supplant them (although proponents of stochastic accounts admit that that not all deterministic explanations can or should be eliminated from evolutionary biology). Yet it seems puzzling that stochastic models would be able to account for the same sorts of phenomena that deterministic models can account for; intuitively, one would expect very different sorts of results from a stochastic model than from a deterministic model. For example, you might expect that a stochastic process would show fluctuations in taxonomic diversity over time, but you would not necessarily expect it to show a marked increase in the number of taxa. However, the particular kind of system embodied by the stochastic models – the Markov process – can often result in seemingly non-random results that we would not intuitively expect from a stochastic process.

In Chapter 3, we saw how random drift, a Markov process, is able to explain apparent directionality (the fixation of one type over another in a population). Once again, the situation for stochastic models of macroevolution is very similar to the microevolutionary case. Recall that a Markov process is a process in which "future development is completely determined by the present state and is independent of the way in which the present state has developed" (Feller 1968: 420). The stochastic models we have been discussing incorporate a kind of Markov process called a branching process (Feller 1968). The models can be seen as embodying a branching process because at each point in time, each "particle" (in this case, a taxon) has a certain probability of terminating (going extinct), branching into two "particles" (producing two new taxa), or continuing into the next time interval (persistence).[40] Because each future outcome (two new taxa, extinction of taxon, or persistence of taxon) is dependent on the present taxa without being dependent

[40] This is only one kind of branching process; there are others, but they are not relevant to this discussion. For example, a simpler version of a branching process permits only branching and extinction, but not persistence.

on any of the outcomes prior to the present taxa, the branching process of the stochastic model is Markovian. Similar to the random walk process we discussed in Chapter 3, branching processes also have the property that future outcomes are partly dependent on the present taxa and partly dependent on the transition probability to the future outcome. In the "freely floating" model described above, the transition probability from the present state to a branching event is 0.1, the transition probability from the present state to an extinction state is also 0.1, and the transition probability from the present state to a "persistence state" (i.e., the taxon simply persists in time without branching or terminating) is 0.8. In the damped-equilibrium model, transition probabilities change over the course of a simulation. In the beginning of the simulation, the probability of branching is greater than the probability of termination, but once the optimum diversity is reached, the probabilities of branching and termination remain relatively constant and equal.

Despite the fact that the basic model used in the MBL program is a branching process, the program also uses a random walk model by superimposing a random walk on top of the branching process. This is done by counting the number of coexisting lineages at each time interval, where lineages are produced by a branching process. Gould et al. 1977 use this method in the "freely floating" model described above. As with the random drift case, the random walk produced by this model differs from the random walk of an ideal coin tossing game in that the number of coexisting lineages isn't limited to incrementing or decrementing by one unit. For example, if there are four lineages at time t, and at time $t + 1$ three of them branch while one goes extinct, the number of coexisting lineages increases from four to six. Another way proponents of stochastic models utilize a random walk model is by incorporating morphology (physical or behavioral characteristics of organisms) into the branching process. Raup and Gould (1974), employing this method, describe the results of simulations performed when ten hypothetical characteristics (treated independently) are allowed to persist or to change by one unit in either direction with each branching event. (They use the damped-equilibrium model, and simply add morphological change to the model). For the purposes of this discussion, we will focus on the random walk of diversity values (which is common to both random walks).

In Chapter 3, we saw how, in a random walk, the dependency of a future state on the present state has a tendency to constrain the future state; values "walk" from the present state to the future state. Furthermore, there is no general tendency to return to an equalization point; if the values in the system have been increasing up to the present state, they are no more likely to decrease at the next state then they are to increase. This is the situation with the "freely floating" model. For example, if the number of taxa at time t - 2 is 10, the number of taxa at t - 1 is 12, and the number of taxa at t is 15, the number of taxa at t + 1 will be an increment or a decrement from the value of 15. The fact that the two previous time intervals represented an increase in diversity value does not make it any more likely that there will be a decrease at t + 1. The damped-equilibrium model behaves similarly, except that there *is* some tendency to return to the optimum value that the program has set. Note, however, that since we are dealing with probabilities, even when the diversity exceeds the optimum value, it may still continue to increase in the next time period. It is just less *likely* to do so (and the farther away from the optimum value a simulation goes, the less likely such a departure becomes, by the design of the program). The characteristics of these two models imply (as we saw in Chapter 3) that the values of these simulations may increase (or decrease) for long periods of time without returning to the equalization point. (Of course, this is less true for the damped-equilibrium model than it is for the "freely floating" model). As a consequence, when the simulations are run, the patterns of diversity exhibit considerable variation. Some show what you would intuitively expect from a stochastic model: fluctuations in diversity, with no marked increases or decreases. Others, however, exhibit patterns where there are marked increases in diversity, or marked decreases in diversity. (See Figure 4.1).

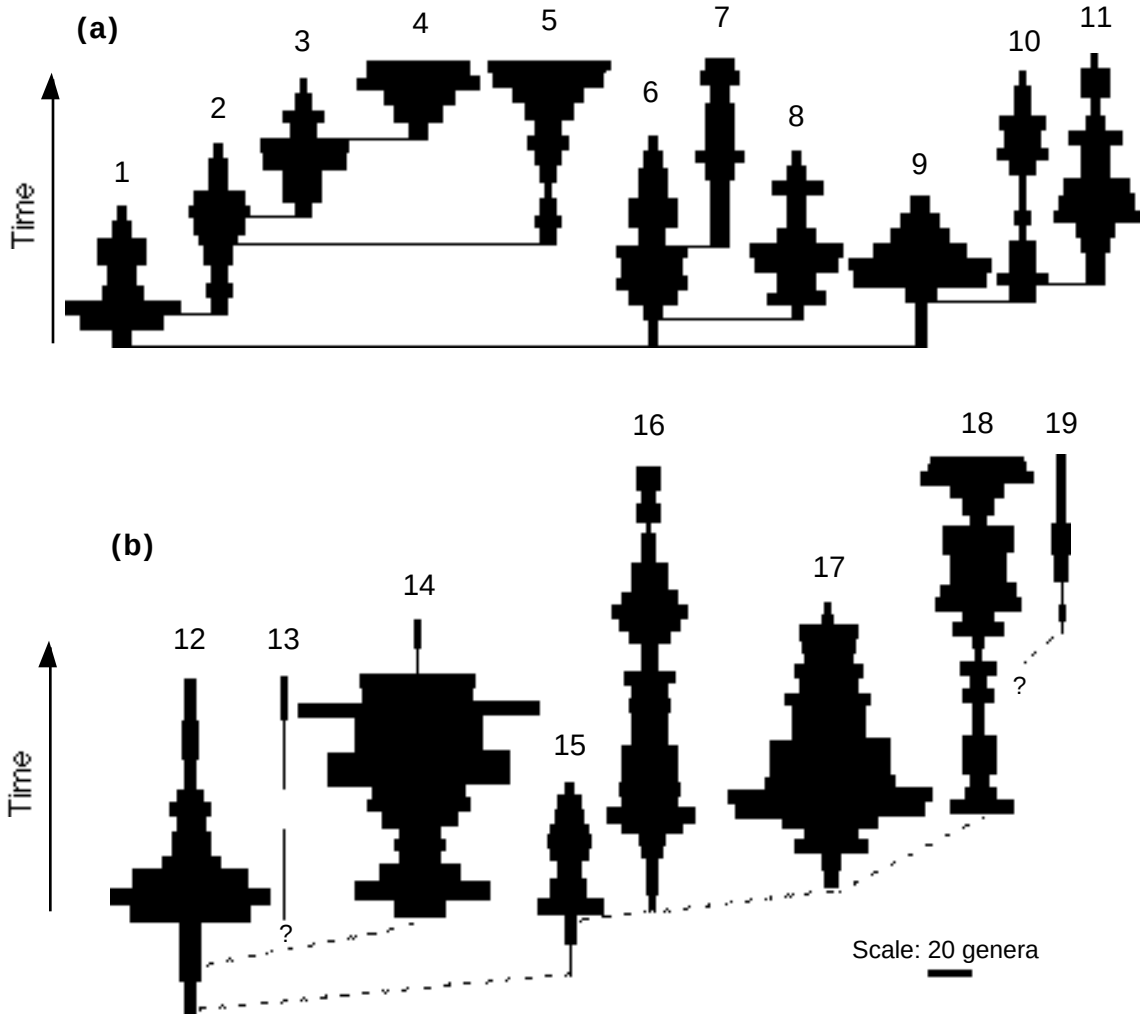

Figure 4.1: Comparison of stochastically generated diversity patterns with real diversity patterns. Numbers above diagrams are for ease of identification only. The width of a diagram at a given "point" in time represents the number of taxa at that time (more taxa = greater diversity). (a) Diversity diagrams for one run of the MBL program with probabilities of branching and extinction set to 0.1. (b) Diversity diagrams for genera within orders of Brachiopods. Diagrams 2 and 6 are what one might intuitively expect from a stochastic process: small amounts of changes in diversity over time. Notice, however (for example), the sudden increase in diversity in diagram 9, similar to that of diagram 12, and the long periods of stability in both diagram 7 and diagram 19 (From Gould et al. 1977).

We are now in a position to answer the question raised in the first paragraph of this section, namely, how is it that a stochastic model can account for the same phenomena that

were previously thought to require deterministic explanations? The Woods Hole Group charges that deterministic explanations are sometimes proposed merely because the phenomena being explained appears to be nonrandom (Gould et al. 1977). However, Markov processes can also (at least to some extent) produce these same sorts of results.[41] Thus, if seemingly non-random patterns are the *only* evidence, stochastic models can often provide explanations of the same phenomena. As D. C. Fisher notes, "An important lesson of Markovian models is that some degree of apparent order can arise by chance … As such, directional trends do not necessarily require explanations framed in terms of single causes acting throughout the duration of the trend" (Fisher 1986: 106). It follows that there is an underdetermination of patterns; you cannot determine whether a stochastic or a deterministic explanation is appropriate just from looking at the pattern. Thus, the empirical question remains as to how, in any given instance, one is to decide whether a deterministic or stochastic model provides the better explanation.[42] Our purposes here are simply to show how stochastic models provide an alternative to deterministic explanations. In the next section, we will show how these models can account for specific kinds of macroevolutionary phenomena (with a focus on one particular phenomenon).

One issue that this entire discussion presupposes is that the stochastic models are *explanatory*. In Chapter 3, I argued that random drift is explanatory under Humphreys's causal account of explanation. The above discussion (and the discussions to follow) suggest that if there are stochastic macroevolutionary processes operating in nature (as random drift is a process which operates in nature) that such processes would increase the probability of certain kinds of events. For example, it would increase the probability that diversity values would fluctuate over time. (Of course, as with the microevolutionary

[41] It should be noted that the "deterministic" process of unidirectional selection can be represented by a Markov process as well. However, the point of introducing the discussion of Markov processes is not to distinguish stochastic explanation from deterministic explanation (that is done with the criteria of "timeboundedness" and "taxonboundedness"), but to explain how it is that a stochastic model can account for patterns that one might intuitively think were produced by a deterministic process.

[42] Raup and Gould suggest that, with regard to morphological trends, "studies of the mechanics of form in relation to environment" might help to distinguish between deterministic and stochastic explanations. Although I did not pursue this empirical question with regard to empirically distinguishing natural selection and random drift either, my claim that natural selection and random drift are distinguished by causal relevance suggests that similar studies might be fruitful in the microevolutionary case as well.

case, selection in fluctuating environments could produce the same pattern. This makes it difficult to determine the cause in any particular situation, but it does not imply that a stochastic macroevolutionary process would be any less causal). Thus, under Humphreys' account of causal explanation, stochastic macroevolutionary processes would be causal, and therefore explanatory.

4.5. The Explanatory Use of Stochastic Macroevolutionary Models

Stochastic models have been proposed as explanations for a variety of macroevolutionary phenomena. There have been studies of phylogenetic patterns of diversity (Raup et al. 1973, Gould et al. 1977, Smith 1977); morphology (Raup and Gould 1974); rates of evolution (Schopf et al. 1975, Schopf 1979, Raup and Marshall 1980; Bookstein 1987); and faunal extinctions (Schopf 1974; Simberloff 1974). Each of these studies discusses phenomena to which stochastic models can potentially be applied. Thus, rather than attempting to be exhaustive, I will focus on one kind of phenomenon (from the study of phylogenetic patterns), and attempt to elucidate how stochastic models can provide an alternative explanation to deterministic accounts.

Real world phylogenetic patterns, as viewed in the paleontological record, exhibit a large amount of diversity. Some taxa experience a large amount of branching early in their history, followed by a long period where numbers taper off. Other taxa steadily increase to become quite diverse, and then quickly go extinct. And many taxa fit neither of these patterns, falling somewhere in between. These historical patterns have usually received deterministic explanations. In the first case, the high amount of branching is explained as an adaptive radiation; for example, it might be hypothesized that an adaptation allowed the exploitation of a new adaptive zone (group of related ecological niches) in the environment. In the second case, the sudden decrease in diversity is explained as a mass extinction, which could be due to such factors as a sudden and extreme change in the climate. In the third case, when faced with differences in patterns "paleontologists are inclined to suspect or even to postulate that the organisms involved are inherently different – that the various

taxonomic groups differ from one another in evolutionary potential because they differ in population structure, reproductive systems, and so on" (Raup et al. 1973, 534). When there are similarities in patterns, such as a number of taxa which go extinct during the same period of time, paleontologists are likely to propose singular causes to explain them.

Proponents of the stochastic account generally admit that most taxa which are very diverse at the beginnings or endings of their histories should be given deterministic explanations (Raup et al. 1973, Gould et al. 1977). This is because these patterns involve sudden explosions of diversity (in the case of a radiation) or sudden, sharp decreases in diversity (in the case of a mass extinction). Although Markovian processes can produce patterns approaching these extremes, the more extreme the pattern, the less likely it is that a Markovian process will produce it. Indeed, these extreme patterns are not well-represented in the simulations produced by the MBL program (Raup et al. 1973, Gould et al. 1977). It is the rest of the cases, where the diversity is greatest somewhere in the middle of a taxa's history, over which there is contention. Here proponents claim that the stochastic model can account for phylogenetic patterns of diversity: "The results of the program and the comparison with reptiles demonstrate that an exceedingly simple stochastic model can produce branching and diversity patterns very like those described in the real world" (Raup et al., 1973, 539).

In some cases, similar patterns appear across unrelated (ecologically and/or phylogenetically) taxa during a particular time period. That is, a number of unrelated taxa will undergo increases and decreases in diversity at approximately the same points in time. Once again, such patterns are generally given deterministic explanations. For example, Flessa and Imbrie (1973) and Flessa and Levinton (1977) studied diversity patterns in 59 marine taxa and 20 terrestrial taxa of differing taxonomic levels (i.e., some of the taxa were on the family level, some were on the class level, etc.) during the time period known as the Phanerozoic. They sought to identify "groups of taxa whose rates of diversification, times of diversification, and times of decline are similar", calling such groups of taxa "diversity associations" (Flessa and Imbrie 1973: 253). According to their analyses, there were 10 diversity associations (that is, 10 groups of taxa sharing common patterns) accounting for

96% of the fluctuations in diversity of the marine taxa, and 4 diversity associations accounting for 91% of the fluctuations in diversity of the terrestrial taxa (Flessa and Imbrie 1973).[43] (See Figures 4.2 and 4.3). When many taxa change in concert, the fossil record shows periods of intense evolutionary change occurring during relatively short periods of time, followed by periods of less intense change (Smith 1977). The periods of intense change are referred to as "pulses". Flessa and Imbrie provide the following explanation for the pulsations that they argue are exhibited by the fossil record:

> That climatic regimes, global temperature gradients, oceanic circulation and perhaps environmental stability are due in part to the topological configuration of continental-ocean relationships seems clear. That these and other environmental factors regulate the distribution and diversity of taxa is also well known. It would therefore be logical to conclude that changes in continenent-ocean relationships due to plate motions would indirectly affect global diversity…Despite the difficulties inherent in both our analysis and the problems in reconstructing ancient plate positions, a remarkable correspondence between major biotic changes and major events in tectonic history of the earth is evident…the independent evidence presented in this paper strongly supports the hypothesis that changes in the continent-ocean configuration of the earth have exerted a significant influence on the evolutionary history of life (Flessa and Imbrie 1973: 276).

In other words, Flessa and Imbrie hypothesize that changing content-ocean configurations due to continental drift caused global environmental changes, which in turn caused diversity fluctuations to occur simultaneously across different taxa. These simultaneous fluctuations appear as pulses in the fossil record. Since this hypothesis attributes one specific cause to the pattern of diversity explained, it is clearly a deterministic explanation. Moreover, it is a *timebounded* explanation; certain times were "good" or "bad" for various taxa, depending upon global environmental changes. Flessa and Imbrie note that in spite of the fact that "[i]t is perhaps naive to search for *the* cause of the Phanerozoic changes in diversity," their "evidence indicates that major changes in continent-ocean relationships, due to plate-movements, appear to have influenced environmental stability and habitat diversity, and thus may be the major factor determining the patterns of Phanerozoic diversity" (Flessa and Imbrie 1973: 282). Figure 4.4 illustrates the major changes in

[43] Terrestrial taxa and marine taxa were not compared against one another.

diversity (“pulsations”) for the Phanerozoic as a whole and the correlation of those changes with major plate tectonic events.

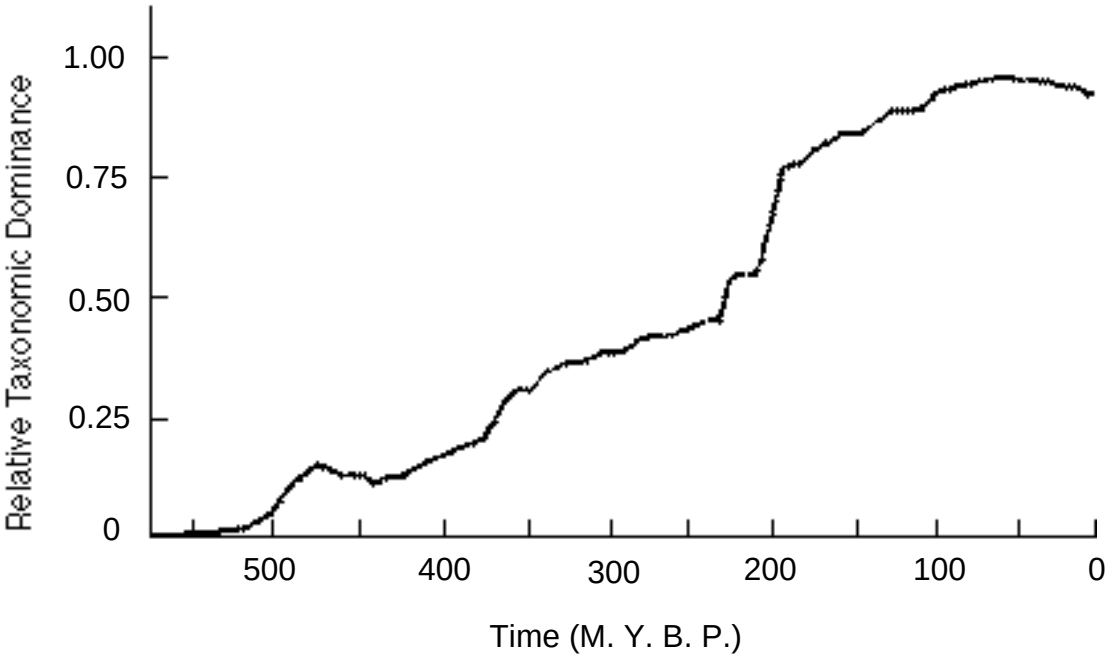

Taxon	Factor Score	Taxon	Factor Score
Regularia	2.0	Heterodonta	1.4
Scleractinia	1.8	Irregularia	1.4
Cyclostomata	1.8	Chlorophyceae	1.4
Holothuroidea	1.7	Chrysophyceae	1.3
Textulariina	1.6	Rhodophyceae	1.3
Rotaliina	1.6	Anomalodesmata	1.3
Stelleroidea	1.6	Dinophyceae	1.2
Miliolina	1.5	Palaeotaxodonta	1.2
Pteriomorphia	1.5	Gastropoda	1.2
Selachii	1.5	Ostracoda	1.2
Porifera	1.5	Cheilostomata	1.1
Terebratula	1.5	Actinopoda	1.0

Figure 4.2 : Marine diversity assocation. This marine diversity association explains the greatest percentage of the total amount of marine information (44.5%) compared to the other marine diversity associations. The higher a taxon’s factor score, the closer the pattern of diversification for the taxon matches the pattern of the shown in the graph. Values less than zero would indicate an inverse relationship to the graph shown (From Flessa and Imbrie 1973).

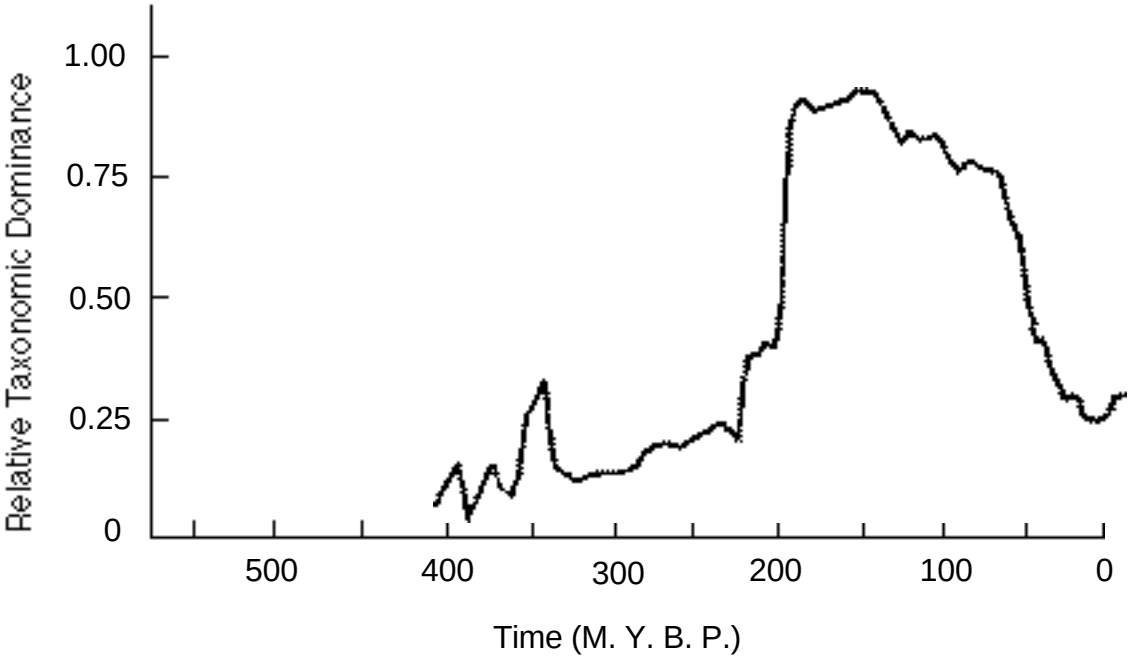

Taxon	Factor Score
Cycadopsida, Gnetopsida	2.7
Coniferopsida	1.9
Bryophyta	1.6
Archosauria	1.5
Charophyta	1.1

Figure 4.3: Terrestrial diversity assocation. This terrestrial diversity association explains the greatest percentage of the total amount of terrestrial information (33.5%) compared to the other terrestrial diversity assocations. See Figure 4.2 for explanation (From Flessa and Imbrie 1973).

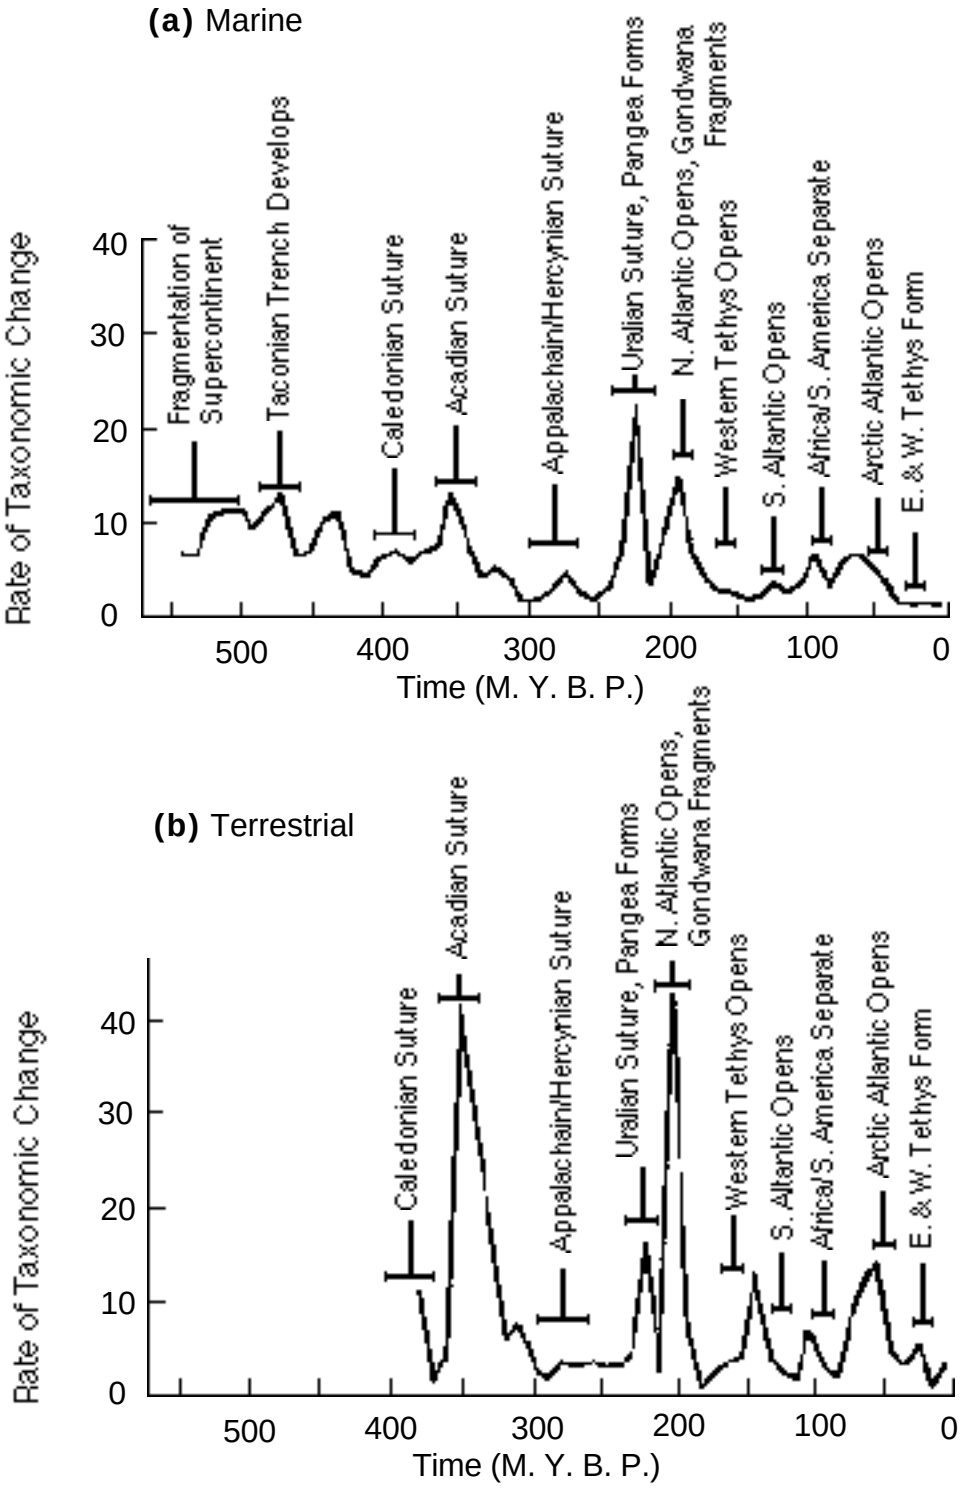

Figure 4.4: Comparison between rates of taxonomic change and major plate tectonic events through the Phanerozoic. Width of bar indicates approximately the time interval during which the plate techtonic events took place (From Flessa and Imbrie 1973).

Flessa and Levinton (1975) agree with the conclusions of Flessa and Imbrie (1973), and seek to show that the stochastic model of Raup et al. (1973) would not be very successful at explaining the patterns of diversity change discussed in the earlier (Flessa and Imbrie) paper. However, they do not offer any analysis to support this conclusion. Instead, they simply assert: "That different taxa do not diversify or decline at random, but tend to do so in concert" has been demonstrated, and that "[i]nsofar as a purely stochastic explanation of the patterns of diversity change would predict no correspondence among clades, covariance of diversity change would indicate further nonrandom evolutionary elements" (Flessa and Levinton 1975: 242). In other words, they argue that the similarity of the diversity patterns is a nonrandom pattern, and since the stochastic model produces random patterns, it cannot explain the patterns of diversity; a nonrandom (deterministic) evolutionary explanation is called for. Furthermore, Flessa and Levinton claim that they would not have been able to ascertain the diversity associations discussed above if the diversity patterns were stochastically produced, since (they claim) under a stochastic model, each of the 79 taxa studied would show a distinctively different pattern of diversification and decline. However, these arguments seem to deny the very claims made by proponents of the stochastic model (claims supported by the analyses of Markov processes) that apparently nonrandom *patterns* can be produced by random *processes.* In other words, Flessa and Levinton seem to miss the major thrust of the Woods Hole Group's arguments, and in so doing fail to convincingly show that the stochastic model would be unable to account for the changes in diversity.

The very analysis that Flessa and Levinton would need to be convincing is supplied by Smith (1977). Unfortunately for Flessa and Levinton, he reaches conclusions in opposition to theirs. Smith analyzes diversity patterns generated from the simulations reported in Raup et al. (1977), using the same method of analysis used by Flessa and Imbrie (1973).[44] In other words, he analyzes the patterns of diversity from the simulation data to see if diversity associations are found, similar to the diversity associations found by

[44] A Q-mode factor analysis, "one of several multivariate analytic techniques which aim to reduce the complexity of a large data matrix in order to reveal a simple structure" (Smith 1977: 43).

Flessa and Imbrie (1973). Smith found that 8 diversity associations accounted for 90% of the data variance in the stochastic simulations, arguing that: "There are many similarities between the present results and those of Flessa and Imbrie (1973)" (Smith 1977: 44). Specifically, Smith found that in both the stochastic simulation data and the Phanerozoic data: 1) a relatively small number of diversity associations can account for over 90% of the diversity patterns, 2) the common diversity pattern of the diversity association is reflected strongly for only a small percentage of the total time period under consideration (that is, the patterns in both the simulations and the Phanerozoic data are pulsational), 3) when one common diversity pattern corresponds strongly, the other common patterns do not, and 4) the combined time spans accounted for by all the common diversity patterns essentially account for the entire time range. Furthermore, Smith argues, the rates of taxonomic change produced by the stochastic simulation show pulsations similar to that of the Phanerozoic data (see Figure 4.5). Thus Smith concludes that:

> …stochastically generated phylogenies can have a pattern of [common diversity patterns] and turnovers similar to those observed in the fossil record. There are evolutionary pulsations in these phylogenies despite the fact that there are no environmental changes to precipitate their occurrence. Thus pulsations do not necessarily indicate the action of major environmental changes since pulsations may be caused by small scale random fluctuations in an equilibrium condition (Smith 1977: 47).

In other words, if Smith's analyses are correct, there is no more reason to accept a deterministic explanation for the common patterns of diversity change across disparate taxa than there is to accept a stochastic explanation. The patterns of diversity change may be the result of an untimebounded, stochastic process.

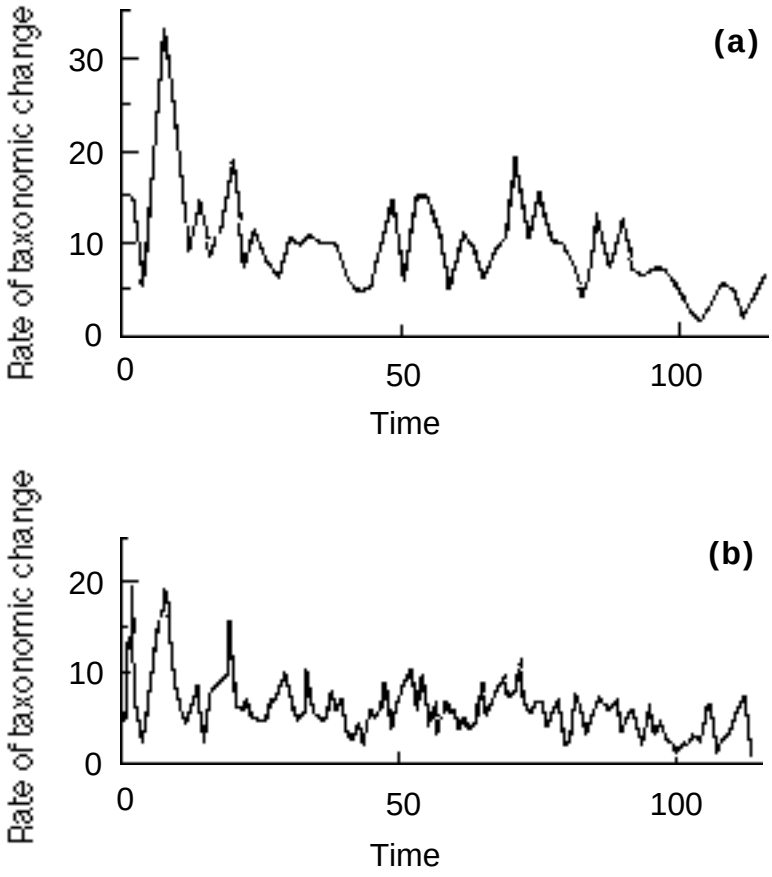

Figure 4.5: Rates of taxonomic change through time using data from stochastic simulations. In (a), every other sample is skipped, producing data more similar to that of Figure 4.4. According to Smith, Flessa and Imbrie skipped samples in order to have all intervals be equal. The difference between(a) and (b) shows that skipping samples filters out some of the "noise" and increases the maximum rate of taxonomic change (From Smith 1977).

The question is, how is it that the stochastic model can account for these patterns? The simple answer is, of course, that the patterns produced by the simulations are comparable to those found in the real world. But *why* can the model yield these patterns? The model used in this case is the damped-equilibrium model. As we saw above, because the damped-equilibrium model is a Markov process, we can expect it to produce a great variety of patterns; within the history of a lineage, it may undergo periods of increasing or decreasing diversity. We also saw that with Markov processes, the production of directionality was not unusual, but rather quite common. Indeed, lineages may contract and

expand at similar times without there being a common cause for the changes in diversity. Thus, the fact that three lineages undergo the same fate during the same time period might not be the result of a singular cause, but instead the result of chance: causes acting independently, coincident in time. Markov processes are processes which are constrained by the past; this leads them to produce orderly results regularly. In this case, proponents of the stochastic model argue, it leads them to produce diversity patterns with a seemingly nonrandom appearance.

4.6. Conclusion

The examination of stochastic models in macroevolution provides us with additional insight into the way chance is used in evolutionary biology. Through the development of untaxonbounded and untimebounded models, the Woods Hole Group brings to light a conception of chance at the macroevolutionary level. At a lower level, there may be causal processes operating, but the causes are numerous and independent. From the viewpoint of the higher level, there is no one cause of the event; lower level causes, coincident in time, "just happen" to produce the higher level event. The process, occurring at the higher level, is inherently probabilistic. This stochasticity, which can be modeled as a Markov process, can provide explanations of seemingly nonrandom patterns previously thought to require deterministic explanation. Stochastic models thus provide us with an additional explanatory tool, and add to the account of the explanatory nature of chance in evolutionary biology, developed in the previous chapter. It also raises the intriguing possibility that the evolution of life is stochastic on a grand scale – that major events in our earth's biological history are due to chance.

In the course of this chapter, we saw many parallels between macroevolutionary stochastic models (and the "opposing" deterministic style of explanation) and random drift (and the "opposing" process of natural selection). The scare quotes here indicate that in both the macroevolutionary and the microevolutionary cases, the two forms of explanation are not really opposing; they may 1) simultaneously explain different aspects of the same

phenomena, 2) be applied to different phenomena by different people, or even 3) be applied to different phenomena by the same people. Nonetheless, there is a sense in which the random models have (historically) provided a challenge to more traditional, Darwinian, selectionist explanations, and this aspect cannot be overlooked. Of course, this chapter saw many more parallels between the microevolutionary and the macroevolutionary cases, which I will summarize here. In both the macroevolutionary and the microevolutionary cases: 1) Nonrandom and random processes can account for the same phenomena, causing there to be 2) Similar debates in both areas, occurring on both general (in the real world as a whole) and specific (concerning particular phenomena) levels, concerning the extent of random versus nonrandom processes; 3) The random process is defined by physical differences being causally irrelevant to differences in "success" (the difference between the two processes being primarily the level at which they occur), implying that 4) A study of functional morphology might profitably serve in distinguishing between random and nonrandom processes; 5) The population-level nature of the random process implies that the process is inherently probabilistic; 6) The random process is Markovian and can be represented as a random walk, implying that 7) Seemingly nonrandom patterns can be produced, even though the processes themselves are random; and 8) The random process is causal, and therefore explanatory. Given the multitude of similarities, it does not seem very farfetched to suggest that the parallels between the microevolutionary process of random drift and the stochastic macroevolutionary models are deliberate, although the Woods Hole Group pointed more towards ecology than population genetics as their inspiration. In any case, it seems clear that, taken together, these two processes imply that chance can (at least potentially) play a very large role in evolution.

Chapter 5: Random Mutation and Directed Mutation

5.1. Introduction

Darwin's theory of natural selection is often seen as having opposed the earlier theories of Jean-Baptiste Lamarck, even though Darwin incorporated many of Lamarck's ideas (most notably in the later editions of the *Origin of Species*).[45] On Darwin's account, evolution by natural selection is a two-step process. The first step is the *production* of chance variations, where chance serves "to acknowledge plainly our ignorance of the cause of each particular variation" (Darwin [1859] 1964: 131). More than this, however, for Darwin a chance variation is a variation that is not "designed"; it could be beneficial to the organism, or harmful, or neither (Darwin [1859] 1964: 81). The second step is the *perpetuation* of the variations, and this, according to Darwin, takes place by the process of natural selection: variations that are beneficial have the best chance of being preserved and passed along to offspring (Darwin [1859] 1964: 127). Furthermore, the frequency of beneficial variations will tend to increase from generation to generation. Darwin's view of evolution as a two-step process has remained through the present day, as Mayr notes: "Evolution through natural selection is...a two-step process. The first step is the production (through recombination, mutation, and chance events) of genetic variability; the second is the ordering of that variability by selection" (Mayr 1978: 52).

However, on Lamarck's account, the distinction between the *production* of variation and the *perpetuation* of variation is blurred. According to Lamarck's "First Law", the use of a physiological character strengthens and enlarges that character, whereas the permanent disuse of the character weakens and deteriorates the character, until it disappears altogether (Lamarck [1809] 1914: 113). So, for example, according to Lamarck a giraffe that continually stretched its neck to reach the highest leaves of a tree would develop a longer neck. Lamarck's "Second Law" claims that these "acquired characters" can be inherited by

[45] To simplify the discussion here and below, I will ignore Darwin's use of Lamarck's theory. For Darwin, evolution is a two-step process, with natural selection as the second step, even if he allowed that sometimes the first step could be variation produced by use and disuse.

an organism's offspring (Lamarck [1809] 1914: 113). Thus, the progeny of two giraffes who developed long necks through stretching would have long necks as well. On Darwin's view, variations arise by chance (they may be beneficial, harmful, or neither); *then* change in a population occurs as the beneficial variations tend to become increasingly prevalent in each succeeding generation (through the action of natural selection). However, on Lamarck's view change occurs by the action of individual organisms in the population (the use of their physiological characters). If all (or most) of the organisms in a population are exposed to the same environmental stimuli, then we would expect that all (or most) of the organisms would acquire the same enhanced characters. These acquired characters are simply inherited by the offspring. There is no need to offer a separate explanation for how variations are perpetuated from how they are produced, as with the Darwinian account; both are accounted for by Lamarck's theory of the inheritance of characteristics acquired through use and disuse. In a sense, Lamarck's views render natural selection unnecessary. If Lamarckism were ever to raise a serious challenge to Darwinism, it would have the potential to undermine present-day evolutionary theory as we know it. We will revisit this issue at a later point in the chapter.

We have seen that for Darwin, evolution is a two-step process, whereas for Lamarck, it is only a one-step process. However, this is not the only element that distinguishes the two theories. On Lamarck's account, organisms will always get the characters that they need; acquired characters are beneficial because they are developed in direct response to environmental conditions. For Darwin, however, variations produced in a population may be beneficial, harmful, or neither. This difference is often captured by saying that on a Lamarckian view, variations are *directed*, while on the Darwin view, variations are said to be *chance* or *random*. But what does it mean to say that variations are random? As we shall see, this question is intertwined with the question of whether there is a distinction between the production and the perpetuation of variation.

It is a major tenet of present-day neo-Darwinian evolutionary theory that mutations, which lead to variation within a population, occur at random. Yet there are many senses in which mutations are nonrandom. According to Hartl and Clark, mutations are nonrandom

with respect to the distribution and frequency of mutational sites, as well as being nonrandom with respect to the kind of mutation (Hartl and Clark, 1989: 114-118). Elliot Sober argues that mutations are misconceived as being random in three ways: 1) that a given gene is equally likely to undergo mutation or not; 2) that mutations are unpredictable; and 3) that it is equally likely that a mutation will be harmful as advantageous (Sober 1984: 104). Jablonka and Lamb point out that the claim that mutations are random does *not* mean that: genes are equally likely to undergo mutation, backward and forward mutations of a gene are equally likely, mutation rates are uninfluenced by mutagens, or that all genes are affected the same (Jablonka and Lamb 1995: 56). In a recent paper, Sniegowski and Lenski assert: "By arguing that mutation is random, we have not meant to imply that mutation occurs at equal rates at all loci or in all environments, or that mutations do not have definable, proximate causes" (Sniegowski and Lenski 1995: 572). Clearly there is some overlap in this list, yet it serves to illustrate the number of ways in which mutation *could* be random, but is not.

In what sense, then, are mutations random? According to Sober, the "defensible idea in the claim that mutation is random is simply that mutations do not occur *because* they would be beneficial" (Sober 1984: 105). However, Sober claims, this idea does not bear on the stochasticity of mutation, since mutation would be conceived as a deterministic process within population genetics whether mutations are directed or not (Sober 1984: 109). Therefore, Sober maintains, it is misleading to assert that "mutation is random"; rather, mutation "stands in contradistinction" to chance (Sober 1984: 110). Thus, on Sober's account, mutations are not random in any important sense. Alexander Rosenberg seems to share Sober's view that mutations are not random in any important sense. According to Rosenberg, the important question is whether mutations are "radically indeterministic" in light of quantum mechanical considerations. Even if they are, claims Rosenberg: "Mutation is but one among several sources of evolutionary change, and certainly not the most important one" (Rosenberg 1994: 60).

However, in claiming that mutations are not random in any important sense, both Sober and Rosenberg are examining mutation in the context of the *perpetuation* of

variation. In population genetics, mutation is, as Rosenberg notes, “one among several sources of evolutionary change,” along with natural selection, random drift, and migration. Mutation does change the composition of a population, albeit very slowly (which is why it is not generally considered to be an important source of evolutionary change). And, as Sober notes, in this context, mutation is described deterministically. This is because we can estimate the probabilities of mutation within a given population and predict the effect that mutation will have on gene frequencies in subsequent populations. Population genetics considers the effect of the mutation rate over the course of generations, and in this light treats mutation as a deterministic force.

Yet we saw above that question of the randomness of mutation arises not within the context of the *perpetuation* of variation, but within the context of the *production* of variation. Mutation, along with recombination, provides the variation upon which the other evolutionary processes (such as natural selection) operate. Sober’s “defensible idea” of random mutation is that, “mutations do not occur *because* they would be beneficial." But the occurrence of mutation is the production of variation, rather than its perpetuation. Yet Sober’s arguments for the determinism of mutation concern the effect of mutation on subsequent generations, or the perpetuation of variation. In arguing that mutation isn’t random in any important sense, Sober has not only shifted the context, he has shifted to the *wrong* context. Likewise, in labeling mutation as an unimportant source of evolutionary change, and therefore an unimportant source of randomness, Rosenberg is considering mutation in the wrong context. It is in the context of the production of variation that mutations can be considered to be random, if at all. Thus, in arguing that mutation is not random in any important sense, both Sober and Rosenberg shift from a context of the production of variation (where mutation is generally considered to be random) to a context of the perpetuation of variation (where mutation is generally considered to be nonrandom).

The arguments of Sober and Rosenberg, taken with the multitude of ways in which mutation is considered *not* to be random (as discussed above), highlight the importance of being very clear about what is meant by the claim that mutations are random. Yet even

Sober's "defensible idea," that mutations are random in the sense that they do not occur *because* they would be beneficial, needs further clarification. As Sober notes, we don't say that weather occurs at random because rain doesn't fall when it would be beneficial (Sober 1984: 105). Why should we say that mutation is random simply because organisms don't get the mutations they need because they need them? Furthermore, the "defensible idea" is defined in the negative: random mutations are mutations that do *not* occur because they would be beneficial; using the distinction introduced above, they are *un*directed. This means that in order to fully characterize what a random mutation is, one must first characterize what its opposite, a directed mutation is. Yet we shall see that 'directed mutation' is as difficult to characterize as 'random mutation'. But if we can't get clear on what a 'directed mutation' is, then we shall have no hope of ever characterizing a 'random mutation'.

Thus, while my primary focus here is on random mutation, I will spend some time discussing directed mutations in order to clarify the distinction. In discussing what it is to be a random mutation versus what it is to be a directed mutation, I will be distinguishing between the *concepts* of random mutation and directed mutation. However, there is also a question as to how one distinguishes between random mutations and directed mutations *empirically*. From the time of Luria and Delbrück's 1943 paper (which claimed to experimentally demonstrate that directed mutations do not occur) through 1988, few biologists doubted the claim that all mutations are random. However, in 1988, Cairns, Overbaugh, and Miller claimed to have demonstrated the occurrence of directed mutation in bacteria, setting off a series of experiments and papers, some in an attempt to support their claims and others to refute them. While it is probably still true that most biologists believe that all mutations are random, it is perhaps less true than it once was; many respected biologists now take seriously the idea that some mutations might be directed, and some would say the issue is as yet unresolved. I mention these experiments not because it is the job of a philosopher to settle the issue – that is clearly the job of the biologist – but because the recent debate has forced biologists to further refine their definitions of directed and random mutations, and the ways to empirically distinguish between them. An examination

of the recent literature therefore provides an opportunity for philosophers to assist biologists in clarifying these conceptual and empirical distinctions.

The structure of the chapter is as follows: First, I provide a short history of the directed mutation controversy. Second, I attempt to clarify and distinguish between the concepts of 'directed mutation' and 'random mutation', by examining three candidate sets of definitions, one from Sarkar, one from Lenski and Mittler, and one from Jablonka and Lamb. Last, I describe the nature of the directed mutation controversy and suggest a way that the opposing hypotheses can be distinguished empirically.

5.2. Background

Before we begin it will be helpful to have some background on mutations and on the directed mutation controversy. However, my remarks here will be rather brief. For recent and more complete accounts of the directed mutation controversy, see Jablonka and Lamb (1995), and Sniegowski and Lenski (1995).

A *mutation* is said to have occurred when the genetic material (usually DNA) of an organism does not replicate faithfully; in other words, when there is a sudden, heritable change or "mistake" in the genetic material (Gardner et al. 1991: 289). Mutation allows parents to transmit their genetic material to their offspring in a changed form, leading to new genes. There are a number of different kinds of mutation, including base substitutions, frame-shift mutations, inversions, translocations, and gene conversions (Hartl and Clark 1989).

Lamarck is the person usually credited with the idea that the inheritance of acquired characters is the source of evolutionary change, based on a theory of "use and disuse". Although Lamarckism is often contrasted with Darwinism, as discussed above, Darwin agreed with Lamarck that organisms could acquire new characters through use and disuse and pass them along to their offspring.[46] Instead, it was the work of August Weismann in the 1880's that formed the primary challenge to Lamarck's ideas. In essence, Weismann

[46] Again, to simplify the discussion, we will ignore Darwin's Lamarckian views (see note 1).

claimed that while changes in the germ line produced changes in the soma, there was no mechanism by which changes in the soma could produce changes in the germ line (Jablonka and Lamb 1995: 8). In other words, there was no mechanism by which changes an organism acquired in its lifetime, such as a longer neck from stretching or strong muscles from running, could affect the organism's genetic material. Therefore, there was no mechanism by which acquired characteristics could be inherited by an organism's offspring.

By the time that Luria and Delbrück were performing experiments in the early 1940's, Weismann's ideas had gained general acceptance, and it was believed that new characters arose randomly, without respect to the use or benefit of the organism. However, the case was still somewhat unclear for bacteria, for a number of reasons. First, "it was impossible to observe the origin of an individual bacterial variant in circumstances in which it was disfavored; the only way to isolate a specific bacterial variant was by altering the environment so as to favor its phenotype" (Sniegowski and Lenski 1995). Second, bacteria lack the separation of germ line and soma found in most other, "higher" organisms. And third, it was well known that bacteria had the ability to adapt quickly to adverse conditions, causing some to speculate that bacteria were producing directed mutations – mutations developed in response to the environmental stress that would allow them to survive in the new environment – although others maintained that what was observed was the proliferation of mutations that had arisen randomly.

These proponents of directed mutations were considered to be "neo-Lamarckians", because they held that the beneficial characteristics acquired by the bacteria would be inherited by subsequent generations of bacteria. However, there is some question as to whether directed mutations may legitimately be considered "Lamarckian", because they do not involve the principle of use and disuse. Instead, the mutations were held to be directly induced by the environment, a mechanism which Lamarck explicitly rejected (Jablonka and Lamb 1995: 6). However, just as neo-Darwinian ideas stem from, but are not identical to, Darwin's own ideas, neo-Lamarckian ideas stem from, but are not identical to, Lamarck's ideas. I will thus continue the common practice of referring to proponents of directed

mutations as “neo-Lamarckian” and critics of directed mutation as “neo-Darwinian” – in spite of the differences between Darwin’s and Lamarck’s views and those of their present-day advocates.

For many biologists, Luria and Delbrück’s 1943 paper, and the papers that confirmed its results (e.g., Newcombe 1949, Lederberg and Lederberg 1952, Cavalli-Sforza and Lederberg 1956), put to rest the idea that directed mutations occurred in bacteria. Luria and Delbrück performed a “fluctuation analysis” in an attempt to decide between what they termed the “acquired hereditary immunity” hypothesis and the “mutation” hypothesis (Luria and Delbrück 1943). In present-day terms, the “acquired hereditary immunity” hypothesis corresponds to the hypothesis that mutations are directed; the “mutation” hypothesis corresponds to the hypothesis that there are random mutations. Luria and Delbrück exposed a number of pure bacterial cultures to a virus, and speculated thus: random mutations could have occurred any time prior to the application of the virus, but they would have occurred at different points in time. This would mean that the clones of resistant bacteria would be of different sizes (Luria and Delbrück 1943: 493), since the earlier mutants would have more time to replicate, and so would be of a greater number than later mutants. This kind of distribution has been referred to as a “jackpot” distribution – the early mutants “won” early, enabling them to acquire a “jackpot”. “Jackpot” distributions have a high variance-to-mean ratio; the distribution would show a high variance, due to the differences in sizes of the jackpots. However, directed mutations would of course occur only subsequent to the application of the virus. So, Luria and Delbrück theorized that directed mutations, if they occurred, would be distributed randomly among cultures (Luria and Delbrück 1943: 493), since they would presumably be equally likely to occur in each culture. More specifically, they would exhibit a Poisson distribution (Luria and Delbrück 1943: 494), which is a special case of the well-known binomial distribution for very rare events (such as mutation). A Poisson distribution has the special property that the variance-to-mean ratio is equal to one. Thus, according to Luria and Delbrück, if the distribution had a high variance-to-mean ratio, the mutation hypothesis (random mutation) would be supported; if the distribution had a variance-to-

mean ratio of one, the acquired hereditary immunity hypothesis (directed mutation hypothesis) would be supported. A high variance-to-mean ratio was found (Luria and Delbrück 1943: 503-7); thus, for Luria and Delbrück (and subsequent generations of biologists), mutations in bacteria were considered only to occur randomly. Neo-Lamarckism was seen as having been refuted.

It wasn't until the experiments of Cairns, Overbaugh, and Miller (1988) that directed mutation again began to be taken seriously, at least by some biologists. Other biologists before Cairns, Overbaugh, and Miller had attempted to show the occurrence of directed mutation, such as Shapiro (1984), but Cairns, Overbaugh, and Miller were more influential, perhaps because the reputation of the authors, or perhaps because of their controversial language: "cells may have mechanisms for choosing which mutations will occur" (Cairns et al. 1988: 142). Cairns, Overbaugh, and Miller criticized the experiments of Luria and Delbrück on two grounds: 1) The hypotheses Luria and Delbrück considered were "extreme": 'all mutations are random' versus 'all mutations are directed'. Cairns, Overbaugh, and Miller proposed testing two less extreme hypotheses: 'all mutations are random' versus 'not all mutations are random (some are directed)'. 2) Directed mutations would not have had time to occur in Luria and Delbrück's experiments, since the bacteria would be killed before they had time to develop the mutations that would provide them resistance to the virus. Cairns, Overbaugh, and Miller thus used non-lethal "selective" environments for their bacterial experiments. Using Luria and Delbrück's fluctuation analysis, Cairns, Overbaugh, and Miller found that the distribution of mutants was neither a "jackpot" distribution (which would indicate that all the mutations were random) nor a Poisson distribution (which would indicate that all the mutations were directed), but somewhere in between. On this basis, they concluded that both directed mutation and random mutation were occurring in the population under study (Cairns et al. 1988). They also performed a controlled experiment (discussed below) and attempted to confirm the results of Shapiro (1984).

Since the work of Cairns, Overbaugh, and Miller was published, there has been a flurry of papers, some in support of Cairns, Overbaugh, and Miller's claims, and some

attempting to refute them (see Jablonka and Lamb 1995 and Sniegowski and Lenski 1995 for references). Some would say the issue is not yet resolved, although many (perhaps most) biologists continue to hold the belief that directed mutations do not occur. Most of the work has been done on bacteria, although Jablonka and Lamb (1995) discuss experiments that attempt to show directed mutation in eukaryotes. My goal here is to examine these findings for what they can tell us about the concepts of directed and random mutation generally (as they apply to all organisms, not just bacteria) as well as what they can tell us about the nature of the empirical controversy.

5.3. The Conceptual Distinction Between Random Mutations and Directed Mutations

As I suggested earlier, the concepts of directed and random mutation are quite muddled. The term 'random mutation' is defined in negative terms, but it is not clear what it is being contrasted with. A random mutation is an undirected mutation, but what is a directed mutation? If, as Sober suggests, random mutations are mutations that do *not* occur because they would be beneficial, what does it mean for a mutation to arise because it is beneficial? In order to answer these questions and further refine our definitions of random and directed mutations, we will examine three pairs of candidate definitions: one pair from a philosopher and two pairs from biologists (one pair neo-Darwinian, the other pair neo-Lamarckian). We shall see that none of these definitions is entirely appropriate; however, our examination of them will point us in the direction of a suitable definition.

In the following discussion, I shall assume that the goal is to define random and directed mutation in such a way so that they are mutually exclusive and exhaustive. In addition, we shall consider the following factors: do the definitions reflect the terms of the current controversy? Do they reflect the terms of historical debates? How do they fit in with other concepts in current evolutionary theory? These criteria shall serve to help us determine whether a purported set of definitions for random and directed mutation is adequate.

5.3.1. Sarkar: A Philosopher's Definitions

According to Sahotra Sarkar, a mutation is random "if and only if the probability of its occurrence in an environment has no correlation with the fitness of the phenotype induced by it in that environment" (Sarkar 1991: 237), whereas directed mutations are mutations for which there is "a positive correlation between the mutation rate and the fitness of the associated phenotype" (Sarkar 1991: 239). (Let us call this definition **S1**). This definition of directed mutations is misleading; it seems to imply that the fitter a mutation will make an organism in its environment, the higher the organism's mutation rate will be. If this is what Sarkar means, then his definition **S1** of directed mutation is too strict. One phenotype can be fitter than another, yet the rate of mutation to that phenotype may be lower. The important question is whether the mutation in question is fitness enhancing at all; the degree to which it produces fitness, and the correspondence to the mutation rate, is really besides the point. By requiring that "fitter" mutations have higher mutation rates, Sarkar potentially excludes directed mutations in his definition **S1**. Thus, the definition fails because it fails to be exhaustive.

Somewhat more promising is Sarkar's restatement of the definition of directed mutation. Sarkar states:

> A strong definition [of directed mutation] would be one that requires that a mutation be considered directed *if and only if* it occurs (or occurs more frequently) in an environment where its associated phenotype has an enhanced fitness. A weaker definition would be one that dropped the "only if" clause: a mutation is then considered directed if it occurs (or occurs more frequently) in the fitness-enhancing or "selective" environment as characterized above… Only the latter will be used here (Sarkar 1991: 239; emphasis in original).

Thus, the "weak definition" of directed mutation, which Sarkar adopts, suggests that a mutation is directed if it occurs (or occurs more frequently) in an environment where its associated phenotype has an enhanced fitness. (Let us call this definition **S2**). However, the parenthetical clause "occurs more frequently" is ambiguous. Does Sarkar mean that a

mutation is directed if it occurs more frequently than mutations that are less fit in the given environment? Or does he mean that a mutation is directed if it occurs more frequently in environments where it is advantageous than in environments where it is less advantageous? If it is the former, this is the interpretation we gave to **S1** above; we saw that unduly required directed mutations to be prolific. If it is the latter, we will have more to say about this definition below, in discussing the views of Lenski and Mittler.

No matter which interpretation you take, however, definition **S2** as it stands is problematic, because of the possibility of satisfying the definition without the parenthetical clause. In other words, definition **S2** leaves open the possibility that a mutation that occurs in an environment where its associated phenotype has increased fitness is a directed mutation. Yet we saw above that the two-step Darwinian process requires that a variation be either harmful, or neutral, *or* beneficial; natural selection then tends to increase the proportion of beneficial variations within the population. If variations were never beneficial, then evolution by natural selection would never be able to explain the adaptation of organisms to their environment. Assuming populations had sufficient traits to survive in their environments, without beneficial variations neutral variations would tend to become more prevalent than harmful ones, and the less harmful variations would tend to become more prevalent than the more harmful ones. But it is the *beneficial* variations that lead to an organism's "fit" with the environment. Yet, as we saw, Darwinian variation is not directed variation; it is explicitly chance variation (or, in present-day terms, random mutation). So to define a directed mutation as a beneficial mutation, as **S2** does, is to include under the umbrella of 'directed mutation' mutations that perhaps should properly be considered to be random mutations. Clearly some beneficial mutations should be directed mutations (definitionally speaking), but not all. Sarkar's definition **S2**, by including what might be considered to be random mutations under the definition of directed mutations, thus fails to produce a pair of definitions that are mutually exclusive. He therefore fails to distinguish between the concepts of directed mutation and random mutation.

As stated in the quotation above, Sarkar makes a distinction between "strong" and

"weak" definitions of directed mutation, the difference being that the "weak" definition lacks an "only if" clause. Sarkar is a proponent of the "weak" definition. That is, he claims that a mutation is directed *if* it occurs (or occurs more frequently) in an environment where its associated phenotype has an enhanced fitness, but he does not claim that a mutation is directed *if and only if* it occurs (or occurs more frequently) in an environment where its associated phenotype has an enhanced fitness. Sarkar maintains that the "strong", "if and only if" definition is propounded by Richard Lenski (1989), and is *too* strong: "it excludes any random (or selectively neutral) occurrence of these mutations" (Sarkar 1991: 239; see also n. 17). However, Sarkar misreads Lenski on this issue; the directed mutation hypothesis that Lenski (a neo-Darwinian) disagrees with is that "certain mutations in bacteria occur more frequently in the presence of environmental agents that favor the resulting mutants than in the absence of these agents" (Lenski 1989: 148). The key difference is that Sarkar claims that Lenski says, "occur (or occur more frequently)" whereas Lenski says only, "occur more frequently." To say that a mutation is directed if and only if it *occurs* in an environment where it would be beneficial *is* to exclude the possibility that the same mutation could occur randomly in a neutral or disadvantageous environment. But to say 'a mutation is directed if and only if it *occurs more frequently* in environments where it is advantageous than in environments where it is less advantageous' does not exclude the possibility of it occurring randomly in neutral or disadvantageous environments; it suggests only that there will be more mutations *overall* in the favorable environment. This is quite reasonable, since if random mutations are occurring in all the environments, but directed mutations are occurring only in the environments where they would be advantageous, we would expect there to be more of them in that environment. Thus, Sarkar's criticism of Lenski's definition does not hold; it is based on a misreading. Let us then turn to a version of Lenski's definition, as expounded in Lenski and Mittler

(1993) and see if it fares any better than Sarkar's definition.[47]

5.3.2. Lenski and Mittler: Neo-Darwinian Definitions

According to Richard Lenski and John Mittler, persistent critics of the directed mutation hypothesis, random mutations are random in the sense that "the likelihood of any particular mutational event is independent of its specific value to the organism" (Lenski and Mittler 1993: 188). A directed mutation, on the other hand, would be a mutation which "occurs at a higher rate specifically when (and even because) it is advantageous to the organism, whereas comparable increases in rate do not occur i) in the same environment for similar mutations that are not advantageous, and ii) for the same mutation in similar environments where it is not advantageous" (Lenski and Mittler 1993: 188). In other words, on this definition directed mutations are mutations that are specifically induced and which occur more often when they would be beneficial to the organism than when they would not.

One thing to notice about Lenski and Mittler's definition is that while it hints at causality (in the phrase, "and even because"), the causality is not explicit or required. That is, it is unclear what sort of causal connection (if any) is to obtain between any of the factors involved in the definition in order for a mutation to be considered directed. This is similar to Sober's definition, which suggested that a directed mutation is one that occurs *because* it is beneficial. On Sober's definition, causality is implied, but it is unclear exactly how being beneficial is a causal factor in the production of a directed mutation. Likewise, the role of causality in Lenski and Miller's definition is unclear. Why is it that mutations are occurring more often in certain environments? Are environmental stresses the key causal factor? The definition suggests this interpretation, but it is not made explicit,

[47] Lenski and Mittler's 1993 definition is essentially the same as Lenski's 1989; I use the 1993 version below because it is more clearly articulated than the 1989 version. In fairness to Lenski and Mittler, I should point out that there are numerous variations on their definitions of random and directed mutations in their extensive work on the subject, some of which will not be subject to the criticisms I will levy against them. I consider the 1993 definition not only because of its clarity, but because it is illustrative as to where it goes right and where it goes wrong.

leaving the definition open to other possible interpretations.

One problem with the lack of explicit causality in Lenski and Mittler's definition is that causality, and in particular the causality of the environment, has been an important element of the debate between proponents and critics of directed mutations since before Luria and Delbrück. Luria and Delbrück note that while "many… investigators believed that the virus by direct action induced the resistant variants," other investigators "believed that the resistant bacterial variants are produced by mutation in the culture prior to the addition of the virus" (Luria and Delbrück 1943: 491). In other words, the proponents of directed mutation believed that the virus (part of the bacteria's environment) caused the mutation that allowed the bacteria to survive the virus. The critics of directed mutation, on the other hand, believed that the mutations occurred prior to the introduction of the virus, making it impossible for the virus to play a causal role. Luria and Delbrück, in attempting to experimentally decide between these two hypotheses, are attempting to ascertain whether the environment (in particular, the virus) plays a causal role in the production of the mutations or not.

The importance of causality in the debate over directed mutation has continued into the present day. This is reflected in the language used by biologists. Thus, Cairns, Overbaugh, and Miller declare: "Our problem is to determine how many of these variants are arising as a *direct and specific response* to the selection pressure (would not have occurred in its absence) and how many are 'spontaneous' (would have arisen even in the absence of selection)" (Cairns et al. 1988: 142; emphasis added). Foster refers to directed mutations as mutations that arise "in response to the selective agent" (Foster 1991: 215). Even Lenski himself uses this language in an earlier paper: "According to the hypothesis of induced or directed mutation … the mutant genotype appears, with some low probability, as *the direct result of its exposure to the selective environment*" (Lenski 1989: 148; emphasis added); why is this causal language not part of his and Mittler's 1993 definition?

Thus, the problem with Lenski and Mittler's definition is that it does not explicitly exclude other causal factors, leaving open the possibility that other causal factors besides environmental stresses are what really is at stake. For example, the death of mutants in

disadvantageous environments will produce results such that the mutation rate appears to be higher in favorable environments than in unfavorable ones. Lenski and Mittler's definition needs to rule out such causal factors. Moreover, in failing to explicitly incorporate this causal element in their definitions, Lenski and Mittler fail to delineate a key element of the debate between proponents and critics of directed mutation, as reflected in their language.

One point has been glossed over in the above discussion, and that is the issue of *specificity*. Part (i) of Lenski and Mittler's definition, quoted above, points out that a mutation should only be considered directed if the increase in mutation rate is not part of a general increase in the mutation rate for *all* mutations. In other words, the increase in mutation rate should be *specific*; the environmental stress should produce an increase in the mutation rate only for those mutations that are an adaptive response to the stress, not for all mutations in general. The reason that Lenski and Mittler require that the increase in mutation rate be specific to the beneficial mutation(s) in question is that if a particular environmental challenge increases the mutation rate for advantageous, neutral, and disadvantageous mutations alike, then the environmental challenge has produced a general increase in the mutation rate, but it has not produced a directed mutation. Such conditions are well known and uncontroversial. For example, it has long been known that radiation produces a general increase in the mutation rate. Some of these mutations might be beneficial (although most probably will not be), yet they should not be considered to be *directed* mutations. Only mutations that undergo an increase in mutation rate that is specific to that mutation, and not part of a general increase in mutation rate, should be considered to be directed mutations. Thus it is important that definitions of directed and random mutation include specificity as well as causality.

5.3.3. Jablonka and Lamb: Neo-Lamarckian Definitions

Eva Jablonka and Marion Lamb, in their recent book, *Epigenetic Inheritance and Evolution: The Lamarckian Dimension*, propose that neo-Lamarckian systems – defined as

the transmission of acquired characters – play a much larger role in evolution than is generally thought. Although their research concerns neo-Lamarckian inheritance in broad terms (including, for example, non-DNA cellular transmission of hereditary information), our focus here will be only on their neo-Lamarckian views regarding mutation.

Before beginning, it is important to understand exactly what kind of challenge (if any) Jablonka and Lamb pose to the neo-Darwinian paradigm. In the Introduction to this chapter, we briefly discussed the way in which the Lamarckian theory of use and disuse seems to make natural selection unnecessary. If organisms get the beneficial variations that they need and transmit them to their offspring, there seems to be no need for natural selection – no need to explain how beneficial variations are perpetuated in a population. Yet Jablonka and Lamb steadfastly maintain that their intention is not to overturn neo-Darwinism, but to enlarge it. Similar to Darwin himself, they contend that not all inherited variations are random; some are directed (Jablonka and Lamb 1995: 1).

Furthermore, Jablonka and Lamb argue that not only is neo-Lamarckism *compatible* with neo-Darwinism, it *requires* it. Lamarckism has often been criticized because it depends upon, but fails to explain, the physiological adaptability of organisms. In the giraffe example we discussed previously, giraffes stretch their necks to reach the highest leaves, lengthening their necks in the process. The Lamarckian theory holds that these lengthened necks can be inherited by the giraffes' offspring. But why do giraffes have the ability to lengthen their necks through stretching in the first place? What is the explanation for this adaptability? Lamarck himself offered no explanation; "for him, it was inherent in the very essence and definition of life" (Jablonka and Lamb 1995: 6). However, Jablonka and Lamb argue that the explanation for the physiological adaptability of organisms is natural selection itself; they "do not doubt that the basic mechanisms underlying the inheritance of acquired variations evolved in a Darwinian fashion by the selection of accidental variations" (Jablonka and Lamb 1995: 25). Thus, while it may be the case that the definitions of random and directed mutation that Jablonka and Lamb propose are influenced by their neo-Lamarckian beliefs (as we shall discuss below), this influence should not be interpreted as anti-Darwinian. Jablonka and Lamb's beliefs both incorporate

and depend upon the principles of Darwinian natural selection.

Jablonka and Lamb (1995) offer definitions of directed mutation and random mutation which incorporate the properties of causality and specificity discussed above, but they add an additional twist to the debate. According to Jablonka and Lamb, the term 'directed mutation' has itself been used in two different ways:

> **JL1**: A mutation is directed if and only if a "particular environmental challenge produces specific and repeatable changes in the hereditary material. The changes can be advantageous, detrimental, or neutral" (Jablonka and Lamb 1995: 57).
>
> **JL2**: A mutation is directed if and only if "new variation is directed towards an adaptive end: the mutations produced in response to an environmental challenge are both specific and selectively advantageous" (Jablonka and Lamb 1995: 57).

What the two definitions have in common is that the directed mutations are produced in response to environmental challenge (in other words, environmental stress is a causal factor in the production of the mutations); furthermore, the changes produced by environment stress are *specific* (not part of a general increase in the mutation rate). Thus, both definitions satisfy the conditions outlined above. They differ in regard to adaptiveness, however. According **JL1**, the changes can be beneficial, harmful, or neutral (the mutations are *directed* in the sense that the environment determines the nature and direction of the mutation). According **JL2**, the changes must be beneficial to be considered directed (again, the environment determines the nature and direction of the mutation, but in this case the nature and direction are adaptive). In other words, the decision between the two definitions of directed mutation boils down to this: does a mutation have to be adaptive (beneficial) in order to be considered directed, or not?

Corresponding to each definition of directed mutation is a definition of random mutation, since random mutation is defined in opposition to directed mutation. However, as Jablonka and Lamb note, there are really three cases to be to be described, not two:[48]

[48] Evelyn Fox Keller makes the alternative suggestion that we see not neo-Darwinian and neo-Lamarckian alternatives, but a continuum of phenomena between the two (Keller 1992). I find this viewpoint intriguing; however, it is an issue that will only be resolved empirically, not philosophically. Also, I think it has a tendency to downplay the very real differences that separate participants in this debate.

> **Case A**: Mutations occur that are either non-specific (the environment produces a general increase in the mutation rate) or not caused by the external environment. These mutations may be beneficial, harmful, or neutral.
>
> **Case B**: Mutations are produced by environmental stress and are specific (only these mutations experience an increase in mutation rate, or these mutations experience an increase in the mutation rate to a greater degree than other mutations). Some of the mutations may be beneficial; however, some of them are harmful and/or neutral as well.
>
> **Case C**: Mutations are produced by environmental stress and are specific. Only beneficial mutations are produced.[49]

On both **JL1** and **JL2**, **case A** mutations are random mutations and **case C** mutations are directed mutations. The question is, what are **case B** mutations? On **JL1**, **case B** mutations are directed mutations because they are specific. However, according to **JL2**, **case B** mutations would be considered random mutations, because in order to be a directed mutation the mutation must be beneficial, and in **case B** not all (or perhaps none of) the mutations produced are beneficial.

What is the best way to solve this puzzle? According to Jablonka and Lamb, **JL2** is "very narrow", "not what most Darwinians have in mind", and "poses terminological problems" (Jablonka and Lamb 1995: 56). They give the example of a mutagen that induces a high rate of mutation in just one or two genes. According to Jablonka and Lamb, even if the mutations are not adaptive, most biologists would not consider them to be random mutations. Thus, Jablonka and Lamb maintain, we should use **JL1**, which corresponds to the more "common usage" of the terms (Jablonka and Lamb 1995: 57).

However, I do not wish to enter into a discussion of which is the more common usage. Both usages appear in the literature. Even if we were to make a count of those usages, and find that one or the other definition was more prevalent, it is not clear what that would prove. Instead, we will see if there are other grounds on which to decide the issue. Let us briefly reconsider neo-Darwinian biologists, Lenski and Mittler. As we saw, Lenski and Mittler's definition of directed mutation corresponds more closely with

[49] Note that if a mutation is caused by environmental stress, specific, and beneficial, and there are no non-beneficial mutations caused by the environmental stress, then it is a **case C** mutation, not a **case B** mutation. This is because **case B** mutations require that at least one of the mutations produced by the environmental stress be either neutral or deleterious. Thus, there is no overlap between **case B** and **case C** mutations.

JL2, in that it insists that directed mutations be adaptive. More tellingly, they do not deny that the occurrence of **case B** mutations has been demonstrated; they maintain only that **case C** mutations (environmentally induced, specific, adaptive mutations) have yet to be demonstrated (see, e.g., Mittler and Lenski 1990). It may be that the choices of definitions on the parts of Lenski and Mittler on the one hand, and Jablonka and Lamb on the other hand, are not accidental. The occurrence of **case B** mutations is relatively uncontested. Thus, by calling such mutation directed, Jablonka and Lamb can point to more experiments as being in support of their neo-Lamarckian claims. Similarly, Lenski and Mittler can insist that **case B** mutations are random, leaving neo-Darwinism intact. In other words, the choice of definition may have more to do with doctrinal allegiances than with biological or historical considerations. Both sides may to some extent be using rhetoric to bolster their respective positions.

Speculation aside, what this discussion illustrates is that there *is* a point of controversy here, but it is not over the existence of **case B** mutations, only over **case C** mutations. Given this, one could argue that the distinction between 'random' and 'directed' should reflect the controversy. On **JL1**, all random mutations are uncontroversial, whereas some (but not all) directed mutations are controversial. On **JL2**, all random mutations are uncontroversial, and all directed mutations are controversial. **JL2** thus better reflects the parameters of the current debate about directed mutation.

In addition to this consideration, which has to do with the current biological practice, we can consider the role that the terms play within evolutionary theory. Jablonka and Lamb note: "If some new variation is directed, then even if it is not adaptive, it could affect the direction of evolutionary change" (Jablonka and Lamb 1995: 75). This, of course, is why evolutionary theory should pay attention to specific, non-adaptive mutations. Furthermore, Jablonka and Lamb are right in that there is a sense in which **case B** mutations are directed, or at least *directional*. A hypothetical example will serve to illustrate this point. Suppose we produce a new chemical that increases the rate of mutation to albinism in humans (but does not cause an increase in the rate of any other mutation). Since humans can spend much of their lives indoors, albinism can be

considered a neutral, or at least mildly deleterious trait. It is therefore a **case B** mutation. Currently, albinism is extremely rare among humans. However, if many people were to be exposed to this chemical, the number of Albino individuals in the population would increase, depending upon the number of people who were exposed to the chemical. Thus, **case B** mutations, which are specific but not beneficial, could lead the course of human evolution in a direction that it might not have gone in otherwise. This is the sense in which **case B** mutations are directed.

It is also interesting to note that this sense of directedness involves the dissemination of **case B** mutations throughout a population: what the evolutionary effects of **case B** mutations are, how they change the composition of a population in a certain manner. So this sense of directedness has to do with the perpetuation of variation. Yet, it also concerns the production of variation: what kinds of variation are produced in the population (specific; beneficial, harmful, or neutral). Thus, Jablonka and Lamb's definition **JL2**, in considering **case B** mutations as directed mutations, does what Lamarck's theory did; it blurs the distinction between the production of variation and the perpetuation of variation. In this way, Jablonka and Lamb are being consistent Lamarckians. Yet is this the sense of directed which was most important to Lamarck? Below, I will argue that it is not.

Even though **case B** mutations can affect the direction of evolutionary change, there is an important sense in which **case B** mutations are *not* directed. The *general* direction these mutations, taken as a whole, will push evolution in is indeterminate; it could be in adaptive, harmful, or neutral directions. Contrast **case B** mutations with natural selection, which (in the long run and for the most part) we expect will lead evolution in adaptive directions. If, over the long run, beneficial variations are preserved and inherited by offspring, the course of evolution will follow an adaptive path. Here we should keep in mind that "path" is a metaphorical term; changing environments and other factors may intervene to change the course of evolution. Natural selection is itself not going in any one *particular* direction. The process of natural selection is directed in the same way the instruction "head south" is directed; natural selection is not like someone who tells you,

"Well, you take I280 south for about 5 miles, exit onto route 17, take the first exit on your right..." **Case B** mutations, on the other hand, say, "Go wherever you want, and take whatever path you want to get there." Thus, compared to natural selection, **case B** mutations are *not* directed, though neither of them is going in any one particular direction.

But why should we compare **case B** mutations to natural selection? Here we must turn again to the history of these issues, to the questions Darwin was trying to answer. Darwin was trying to explain what to him was one of the most striking feature of our natural environment: the adaptation of organisms to their environment. As Darwin asks, "How have all those exquisite adaptations of one part of the organisation to another part, and to the conditions of life, and of one distinct organic being to another, been perfected?" (Darwin [1859] 1964: 60). The production of variation was, as we have seen, part of Darwin's answer to that question. But ultimately, it is not what leads to adaptation; for that, according to Darwin, you need natural selection. So if the question is "what has produced adaptation?", Darwin's answer is natural selection. We understand other evolutionary processes in contrast to natural selection, because we are trying to understand exactly how the different processes lead to adaptation, and we should be clear on which lead directly to it, and which only indirectly lead to it, or lead away from it. Mutations that are specific, but do not have a general direction, are part of the evolutionary process, in that they produce variation, but they do not necessarily lead to adaptation. So we contrast **case B** mutations with natural selection, that which does lead to adaptation, and that which does ultimately answer the question we are asking.

In current evolutionary theory, it is still the case that other evolutionary processes are described in contrast to natural selection. As we have discussed in earlier chapters, random drift is described in this way: random drift is a process in which physical differences between organisms are causally irrelevant to their survival and reproductive success, in contrast to natural selection, a process where physical differences between organisms *are* causally relevant to their survival and reproductive success. Moreover, random drift *is* and *is not* directed in exactly the same way as **case B** mutations are and are not directed. Random drift, in small populations, can lead to a trait becoming ubiquitous

within a population (as we discussed in Chapter 3), and so, like **case B** mutations, are directed in that they can change the course of evolution in a particular way. However, random drift is undirected in the same way: it likewise can push evolution in adaptive, harmful, or neutral directions. Thus, random drift draws the same contrast to natural selection that **case B** mutations do. If we are going to be consistent about the contrasts we draw and our uses of the term 'random', then, it seems to make more sense to classify **case B** mutations with random mutations than with directed mutations.

At this point the reader may very well question the fairness of trying to fit a neo-Lamarckian definition within Darwinian and neo-Darwinian evolutionary theory. In one obvious sense, it *is* fair, because that is exactly what Jablonka and Lamb are trying to do: find room for neo-Lamarckism within neo-Darwinian theory. However, there are other considerations that make it an appropriate line of reasoning. Adaptation was the principle question for Lamarck as well as Darwin. His theory of use and disuse also sought to explain the adaptation of organisms in their natural environment; continuous use strengthened and enlarged physical characteristics, whereas permanent disuse weakened and deteriorated them. The result would be a population of organisms that were well adapted to handle the challenges the environment posed. "Directed" variations, for Lamarck, were variations in a specifically *adaptive* direction. In this way, **JL1**, in considering non-adaptive mutations as directed, is no more consistent with Lamarck's views than it is with Darwin's. Jablonka and Lamb have changed the terms of the debate, but in a way that serves more to blur the important issues than to clarify them.

From three points of view – the point of view of the issues at stake in the present controversy, from the point of view of the historical issues, and from the point of view of present-day evolutionary theory – it seems to make more sense to classify **case B** mutations as random mutations, rather than directed mutations. On these grounds, it is **JL2** that is the more appropriate.

Of course, there is an easier way out of the problem of whether to classify **case B** mutations as directed or random, and that is to have three terms rather than two. This is perhaps also the clearest and least confusing way, assuming the terms for each case were

chosen carefully. However, to do this would be to ignore a large body of existing literature and the history of the debate, and such attempts are usually ill-fated from the start. Thus, I would argue that **JL2** is the definition that we should adopt. In that case, we would say that a mutation is directed if and only if it is specifically caused by environmental stress in an exclusively adaptive manner. Otherwise (if the mutation is non-specific, or specific but not exclusively adaptive, or not caused by environmental stress) it is a random mutation.

Is this definition of 'random mutation' "misleading" in its use of the term 'random', as Sober claims? I maintain that it is not – that the notion of randomness employed by the definition of 'random mutation' is a legitimate one. One commonly accepted definition of 'random' or 'chance' is 'not caused'. This definition implies *physical* randomness, or indeterminism. While it may turn out that mutations are random in this sense (at the quantum mechanical level), the notion of 'random mutation' that we have developed here does not depend on any such finding. Rather than meaning 'not caused *simpliciter*', random mutation implies 'not caused by environmental stress in a specific, adaptive manner'. Or, to put the point another way, random mutations are random with respect to their advantageousness, since a favorable environment is not the cause of their existence. Notice that this sense of randomness is analogous to the sense in which 'random drift' is random; random drift means that the survival and reproductive success of organism is not caused by the fitness of the organism, such that the survival and reproductive success of the organisms is random with respect to their fitness. These biological conceptions of 'random' thus extend the common conception of 'random' as 'not caused' to include 'not caused by X', such that the randomness is not general, but with respect to a particular parameter.

5.4. Distinguishing Empirically Between Random Mutations and Directed Mutations

Our discussion up to this point has focused on the conceptual distinction between random mutation and directed mutation: how the terms are defined and how they are

distinct from one another. Distinguishing between the concepts of random and directed mutation is related to, but distinct from, another important task: how we distinguish between random mutation and directed mutation *empirically*, i.e., what sort of evidence would decide the debate between proponents and critics of directed mutation.

In previous chapters, we discussed how it is that random drift could be explanatory (Chapter 3), and how it is that stochastic macroevolutionary models could be explanatory (Chapter 4). In both of those cases, the chance element served as an alternative explanation for the non-chance (or less chancy) element (with respect to random drift, this is natural selection; with respect to macroevolutionary stochastic models, this is deterministic explanation). In the case of random mutation, the important issues are somewhat different. On the one hand, there is no puzzle as to how random mutation can explain the production of variation within populations; most variations appear to be non-specific, or specific but not exclusively adaptive, or not caused by environmental stress. On the other hand, there is a real question as to whether directed mutation occurs *at all*, making it somewhat premature to refer to random mutations as an alternative to "directed mutation explanations" (even though they are defined in this manner). Thus, rather than focusing on the issue of explanation, we will examine the more fundamental question of what sort of evidence would be sufficient to demonstrate the existence of directed mutations.

However, before we begin that task, we must first clarify exactly what is at stake in this debate between the proponents and the critics of directed mutation. We will then examine Sahotra Sarkar's claims concerning the empirical distinction. I will argue that Sarkar's arguments go too far, and offer a more temperate solution.

5.4.1. Nature of the Directed Mutation Controversy

The most important thing to be clear on in this debate is that proponents of directed mutation ("neo-Lamarckians") are not claiming that all mutations are directed, or even that most mutations are directed; they readily acknowledge the occurrence of random

mutations. They maintain only that directed mutations occur *sometimes*, for certain organisms under certain conditions (see, e.g., Cairns et al. 1988, Jablonka and Lamb 1995). Critics of directed mutation ("neo-Darwinians"), however, deny that directed mutations ever occur, although some critics admit that "the possibility can never be excluded that some arbitrarily small fraction of mutants is the result of a directed process" (Mittler and Lenski 1990: 175). This admission is presumably because of the logical difficulty of the critic's position; in order to show that proponents of directed mutation are wrong, they would have to show that directed mutations *never* occur, or else demonstrate that they are impossible. Proponents of directed mutation have much easier time of it, logically speaking; they need only show that directed mutation occurs in *one* case to support their position. Thus most of the discussion takes place on the level of particular cases: this strain of bacteria, this mutation, these conditions. The question under discussion is usually whether directed mutation has been demonstrated in a particular instance, with critics arguing that the experimenters have failed to experimentally make their case.

In the sense that the debate between proponents and critics of directed mutation is not strictly an either/or question – the competing hypotheses are not 'all mutations are random' vs. 'all mutations are directed' – it parallels the debate between proponents and critics of random drift. In a similar fashion, the participants in that debate do not claim that 'all evolution is due to natural selection' or 'all evolution is due to random drift' (Beatty 1984). However, critics of directed mutation are much more likely to deny directed mutation outright than participants in the natural selection/random drift debate would be to deny their opponents position outright, so the debates are not strictly analogous. (However, as with the directed mutation controversy, the natural selection/random drift controversy more often takes place at a particular, rather than general, level, although participants in both debates will sometimes try to draw general conclusions).

There is another way in which the directed mutation controversy is analogous to the random drift/natural selection controversy, and that is that, for the most part, it is not the data that are in question, but the interpretation of the data. Thus, when critics of directed

mutation argue against proponents, they generally do not argue with the results of the experiments; rather, they offer alternative explanations for the results (see, e.g., Charlesworth et al. 1988, Lenski 1989, Lenski and Mittler 1993, Mittler and Lenski 1990, Symonds 1989; Keller 1992 and Sarkar 1991 make this point as well). The collection of letters written in 1988 (in the wake of Cairns, Overbaugh, and Miller's 1988 paper) by D. Charlesworth, B. Charlesworth, Bull, Grafen, Holliday, Rosenberger, Van Valen, Danchin, and Tessman is a particularly striking example; it shows biologist after biologist offering alternative explanations for the results of Cairns, Overbaugh, and Miller). Likewise, in the natural selection/random drift controversy, participants will often contemplate the same data, yet offer differing explanations to account for it.[50]

5.4.2. Sarkar's Reductionist Account

In two recent papers, Sarkar claims that the only way to empirically distinguish between the claims of the two camps is by discovering the mechanisms controlling the mutations for the systems in question (Sarkar 1990, 1991; the argument is virtually identical in both papers). However, I will argue that Sarkar has overlooked an important class of experiments involved in the debate, namely controlled experiments, and thus goes too far in his conclusions. Finding the mechanisms responsible for purported instances of directed mutation would certainly settle the controversy, but resolving the controversy need not rest on finding such mechanisms. Or so I shall argue.

Sarkar's argument rests on findings from Stewart et al. (1990), who demonstrate that in a fluctuation analysis (such as those performed by Luria and Delbrück and Cairns, Overbaugh, and Miller, discussed above) a number of factors may shift the distribution of mutants towards a Poisson distribution and away from a "jackpot" distribution. Thus, what appears to be directed mutation may not be; the "Poisson-like" distribution is not the result of a random distribution of a rare event, but other factors such as "phenotypic lag,

[50] These parallels between the directed mutation controversy and the natural selection/random drift debate can also be drawn with the debate over stochastic versus deterministic accounts of macroevolutionary phenomena.

the differential fitness of mutant and non-mutants, a mutation rate dependent on the nutritional state of the cells and the possibility that plating efficiency is less than 1" (Sarkar 1990: 117). Lenski and Mittler (1993) explain one way such a shift could occur. Some mutants favored under "selective" conditions may grow more slowly than non-mutants in "non-selective" conditions (Lenski and Mittler 1993: 190). The result is that any random mutations, produced prior to exposure to "selective" conditions, will leave relatively fewer progeny. The "jackpot" cultures are therefore smaller than you would expect if the growth rate of mutants was thought to be equal to non-mutants, and thus there appears to be a shift in the direction of the Poisson, because there is less variance in the distribution.

On the basis of these findings, Sarkar claims that "fluctuation analysis alone cannot resolve the dispute over whether some directed mutations occur in bacteria" because "[a]ny observation of a deviation from the Luria-Delbrück [jackpot] distribution can be explained by invoking subsidiary interactions" such as those I have mentioned above (Sarkar 1991: 255). Sarkar admits that (of course) anyone who puts forth such explanations must in turn provide experimental evidence for these explanations, and that proponents of directed mutation can attempt to "experimentally rule out" these factors. Furthermore, Sarkar acknowledges that if proponents of directed mutation were successful in ruling out all the factors that critics have proposed and were still able to produce evidence suggesting directed mutation, "then this evidence can be taken to suggest the plausibility of the neo-Lamarckian position" (Sarkar 1991: 255). However, Sarkar asserts that given the number of alternative explanations that critics have produced, even this evidence would not be convincing to critics of directed mutation. To "convince the skeptic," Sarkar claims, one must find the *mechanisms* for directed mutations (Sarkar 1991: 255).

It is in trying to "convince the skeptic" that I think Sarkar goes too far. Sarkar is right to demand that other factors be ruled out before any conclusions are drawn solely on the basis of fluctuation analysis. However, even if proponents of directed mutation are successful in doing this, should we still disbelieve their claims? Is it reasonable to try answer the skeptic, who continues to make demands in the face of contrary evidence? At some point, shouldn't evidence cause us to change our beliefs, no matter how seemingly

implausible the hypothesis supported by the evidence is?

In general, I think it is unreasonable to bow to the demands of this kind of skeptic. Yet in this case there *are* reasons to doubt claims based on fluctuation analysis alone, reasons that go beyond unrelenting skepticism. Sarkar claims that the debate over directed mutations occurs at the level of sub-cellular processes, and thus that the failure of fluctuation analyses (which take place at the cellular level) is a vindication for reductionism (Sarkar 1991: 261). However, I maintain that the debate is *not* over processes at the sub-cellular level, at least not exclusively; primarily, the debate is over whether directed mutations occur, i.e., whether the environment causes specific, adaptive mutations to occur in some organisms. Certainly part of the neo-Darwinian skepticism stems from an inability to imagine any kind of mechanism that could produce directed mutations. However, what is going on at the sub-cellular level, i.e., what mechanisms produce this directed mutation, are in some sense irrelevant. If such mechanisms were discovered, it would certainly help settle the debate. But the debate is not *about* the mechanisms, it is about whether directed mutation occurs at all, whatever the underlying mechanism might be. Consider, for example, a case where a proponent of directed mutation proposed a particular mechanism, which was subsequently shown not to exist. Would the directed mutation hypothesis thereby have been "falsified" for all time? No, because it is possible that *other* mechanisms could be producing directed mutation, and directed mutation requires only that environmental stress cause the mutation in a specific, adaptive manner.

It is this latter requirement, derived from the conceptual distinction between directed mutation and random mutation outlined above, which makes fluctuation analyses inadequate. While the alternative hypotheses involved in the fluctuation analysis make causal claims, the fluctuation analysis cannot really provide complete evidence for them. As Lenski and Mittler put the point:

> If care is taken to ensure that the assumptions of [the fluctuation analysis] have been met (or if any violations are taken into account quantitatively), then deviations may be used to show that mutations occur after the imposition of selection. Even so, it does not follow that any such post-selection mutations are caused by the selection conditions. Rather, to establish that some post-selection mutations are directed would require a

> demonstration that the mutation rate is higher under selective than nonselective conditions, all else being equal (Lenski and Mittler 1993: 190).

In other words, fluctuation analyses do not provide evidence for the causal role played by the environment, because they do not compare mutation rates for environments where the mutations are adaptive with mutation rates in environments where the mutations are not adaptive. And if they do not provide causal evidence, they cannot distinguish between directed mutation and random mutation, because these concepts are distinguished by the causal role played by the environment.

However, there *is* a kind of experiment designed specifically to ascertain this causal role, namely, the controlled experiment. Consider one of the experiments which Cairns, Overbaugh, and Miller describe. Cairns, Overbaugh, and Miller examined replicate cultures of Lac^- *Escherichia coli* (bacteria unable to utilize lactose) in environments containing lactose and in environments without lactose. They found that Lac^+ mutants (bacteria able to utilize lactose) accumulated over time in the presence of the lactose, but failed to accumulate in the absence of the lactose (Cairns et al. 1988). Biologists argue over the meaning of these results – critics have pointed out, for example, that Cairns, Overbaugh, and Miller fail to fully demonstrate that the accumulation of Lac^+ mutants is a *specific* response to the lactose and not simply a general increase in the mutation rate – but what it is that they are attempting to prove is not in question. By considering replicate cultures where all factors are identical (to the best of the investigator's abilities) except for the one crucial factor, the lactose, Cairns, Overbaugh, and Miller try to show that it is the lactose that is the causal factor at work. In other words, the setup which Cairns, Overbaugh, and Miller use fits the model of a controlled experiment, where the goal of the experiment is, as Nancy Cartwright describes, "to find out if our causal stories are right or wrong" (Cartwright 1980). More specifically, we can say that the controlled experiment which Cairns, Overbaugh, and Miller perform is an attempt to demonstrate that environmental stress is a causal factor in the production of the Lac^+ mutants – it is an

attempt to demonstrate the occurrence of directed mutation.

For some reason, Sarkar glosses over the controlled experiments that are described in Cairns et al. (1988). After asserting that Cairns, Overbaugh, and Miller's conclusions rest largely on fluctuation analysis, Sarkar remarks in a footnote that "they carry out, of course, a series of more pointed experiments which are also generally based on statistical analysis" (Sarkar 1990: 120 n. 5). In so doing, Sarkar lumps controlled experiments together with fluctuation analysis under the heading of "statistical analysis". And according to Sarkar, "mere statistical analysis of the mutants, by itself, is unlikely to be able to settle the controversy regarding their origin" (Sarkar 1990: 112).

But why should controlled experiments be unlikely to settle the controversy, to settle the empirical question between random mutation and directed mutation? It is true that controlled experiments must rule out the possibility that other factors besides the ones in question have produced the results, and Cairns et al. (1988) have been criticized on those grounds (see, e.g., Charlesworth, Charlesworth, and Bull 1988). However, if someone *were* to perform an experiment in which the mutation rate was shown to be higher in advantageous environments than in non-advantageous environments, where there was not a likewise increase in the mutation rate for other mutations of the same kind (specificity requirement), and where all other factors that critics could come up with were controlled for, why should we not acknowledge the existence of directed mutations? Of course, there may still be factors unaccounted for, factors that critics are unaware of. But, as Cartwright notes, it can happen in any kind of experiment, not just controlled experiments, that we do not have enough of the knowledge that we need (Cartwright 1989: 69). Controlled experiments are a well-accepted methodology in other sciences besides biology, such as physics and chemistry. There is no more reason to doubt their results in biology than would be the case in other sciences, and thus, no reason to doubt that a controlled experiment can settle the empirical question of how to decide whether directed mutations or random mutations are occurring.

5.5. Conclusion

In discussing any controversial issue, it is possible to gloss over the historical and contemporary issues at stake in the debate. However, to do so risks spreading confusion instead of clarity, and it will ultimately fail to serve your purposes; the opposition can too easily claim that you have not proved anything, only changed the terms of the debate. For this reason, the person that ignores what is a really at stake in a controversial debate does so at his or her peril.

Recent work in biology has produced a resurgence of interest in the question of directed mutations. Along with that interest has come a desire, and a need, to specify exactly what a directed mutation is and what a random mutation is, in order to clarify the terms of the debate. Sarkar provides us with two versions of a pair of definitions for directed and random mutation. However, as we saw, the first fails to include mutations that should properly be considered directed, and the second classifies what should be random mutations (in order for Darwinian evolutionary theory to function) as directed mutations. Thus, his first pair of definitions is inadequate because it failed to be exhaustive; the second pair fails because the definitions are not mutually exclusive.

The definitions provided by Lenski and Mittler include the important characteristic of specificity, but they lack the concept of causality, a concept that has been a key issue of the debate from the time of Luria and Delbrück's experiments to the present day. We need to be clear about the fact that environmental stress is the causal factor in question in determining whether a mutation is directed or random. Not only does an explicit incorporation of causality in our definitions allow us to capture the terms of the debate, but it provides us with guidance concerning how to distinguish between the directed mutation hypothesis and the random mutation hypothesis empirically. The goal of a controlled experiment is to isolate the causal factors at work in a particular system. It is possible that a well-designed controlled experiment could settle the debate between neo-Lamarckians and neo-Darwinians. Settling the debate need not wait until a mechanism for directed mutation is found, as Sarkar argues.

Jablonka and Lamb provide us with two sets of definitions, both of which incorporate specificity and causality, but they prefer a definition of directed mutation that does not include adaptiveness. However, as we saw, in order to reflect the terms of the historical debate between Darwin and Lamarck and the contemporary debate between neo-Darwinians and neo-Lamarckians, the definition of directed mutation must incorporate adaptiveness. Jablonka and Lamb's neo-Lamarckism is not a threat to neo-Darwinism in the same way that Lamarck's views were a threat to Darwin's, but they fail to respect the terms of the debate. In blurring the distinction between the production of variation and the perpetuation of variation, they blur it even farther than Lamarck did himself. In evolutionary theory, we often distinguish between processes that lead to adaptation and processes that do not. Jablonka and Lamb's preferred definition blurs that distinction, and that can only serve to make our explanations more convoluted and confusing.

Seeing where the above proposed definitions fail and where they succeed points us in the direction of a new pair of definitions. I argue that we should adopt the following definitions: a mutation is directed if and only if it is specifically caused by environmental stress in an adaptive manner. Otherwise (if the mutation is non-specific, or specific but not exclusively adaptive, or not caused by environmental stress) it is a random mutation. These definitions suggest that random mutations are random with respect to their advantageousness. In other words, random mutations are mutations in which environmental stress is causally irrelevant to the type of mutation produced: advantageous, disadvantageous, or neutral. As I noted above, this sense of randomness is analogous to the sense in which 'random drift' is random. Random drift is a process in which the survival and reproductive success of organisms is random with respect to their fitness, i.e., random drift is a process in which fitness differences between organisms are causally irrelevant to their survival and reproductive success.

The definitions of directed mutation and random mutation that I have argued for here are consistent with current evolutionary theory and respect the terms of the historical and contemporary debates. Causality and specificity are, and have been, key issues in these debates. Disputants disagree over whether environmental stress can specifically cause

beneficial mutations, a dispute that primarily concerns the production of variation, not the perpetuation of variation. In the context of the production of variation, the key issue that separated Darwin and Lamarck concerned the adaptedness of the variations produced, and this is the issue which remains in contention today. The proposed definitions respect these considerations, and thus provide a vehicle for resolution of the issues at stake, without allowing disputants to talk past one another.

Conclusion

Random drift, stochastic macroevolutionary models, and random mutations are three distinct areas of evolutionary theory where chance plays a role. Random mutations concern the *production* of variation within populations. Random drift, a microevolutionary process, and stochastic macroevolutionary models, both concern the *perpetuation* of variation within populations. In each area, chance is defined slightly differently (as one would expect, given that they pertain to different areas), yet there are striking similarities between the different areas.[51] Random drift is a process where heritable fitness differences (or heritable physical differences, in the absence of fitness differences) are causally irrelevant to differences in reproductive success. That is, random drift is a process where the reproductive success of organisms occurs randomly with respect to their physical traits. Stochastic macroevolutionary models are untimebounded and untaxonbounded; branching and extinction occur randomly with respect to time interval and taxon. As we discussed in Chapter 4, there is no corresponding requirement of "untimeboundedness" for random drift, or random mutation for that matter.[52] An untaxonbounded model, on the other hand, claims that differences between taxa are causally irrelevant to differences in rates of branching and extinction within the taxa. Random mutations are mutations that are not caused by environmental stress in a specific, adaptive manner. Or, to put the point another way, random mutations are random with respect to their advantageousness, since a favorable environment is not the cause of their existence.

If we examine these definitions, we find the following similarities between the different conceptions of chance:

1. All three conceptions are defined in contrast to another process

[51] There are more similarities between the two processes which concern the perpetuation of variation than there are between either of these processes and random mutation, which is not surprising given that they refer to different aspects of the evolutionary process. In addition, both random drift and the stochastic macroevolutionary model are random walk processes; random mutation is not.

[52] Indeed, one could develop a stochastic macroevolutionary model that was timebounded, yet untaxonbounded. I speculate that the Woods Hole Group did not do so because of their desire to develop a law-like model, similar to those of other sciences (as was discussed in Chapter 4).

occurring in the same evolutionary sphere.

Random drift is defined in contrast to natural selection. Stochastic macroevolutionary models are defined in contrast to "deterministic" modes of explanation.[53] Random mutations are defined in contrast to directed mutations.

2. All three conceptions involve the negation of a specific causal process.

The term 'chance' often means 'not caused'. These biological conceptions are a modification of this common usage; rather than meaning 'not caused *simpliciter*', they involve processes whose outcomes are not caused *in a specific manner*. Random drift is a process in which physical differences do not cause differences in reproductive success. Stochastic macroevolutionary processes (should they occur in nature) are ones in which differences between taxa do not cause differences in rates of branching and extinction within the taxa. Random mutations are mutations that are not caused by environmental stress in a specific, adaptive manner.

3. All three conceptions capture processes which are random with respect to the general course of natural selection (i.e., adaptation).

Natural selection is a process in which fitness differences between organisms *are* causally relevant to differences in reproductive success. That is, the fitter organisms are more likely to have greater reproductive success than less fit organisms. As natural selection occurs over the course of many generations, in a relatively constant environment we would expect organisms to be increasingly better adapted to that environment. With random drift, on the other hand, differences between organisms are *not* causally relevant. Thus, random drift may

[53] As we discussed in Chapter 4, "determinism" here does not refer to Laplacean determinism, but to macroevolutionary explanations which invoke a specific cause. I will follow the practice of that chapter by explicitly using the term "Laplacean determinism" when necessary.

> lead to increasing adaptation – but it may also lead to decreasing adaptation, or simply the perpetuation of neutral traits from generation to generation. Similarly, stochastic macroevolutionary processes and random mutations may be adaptive, maladaptive, or neutral. At the level at which each of the processes operate, it is matter of chance which direction the process will proceed in (although the processes may be Laplacean deterministic at a lower level).[54]

Taken together, these three properties capture a distinctly biological notion of chance, one that is not (as we discussed in earlier chapters) dependent on whether evolution is Laplacean deterministic or indeterministic (thus, it is not dependent on conclusions drawn from quantum mechanics). Random drift and stochastic macroevolutionary processes are population-level processes; perhaps Laplacean deterministic stories can be told at a lower-level, but those stories fail to fully capture the changes that are occurring at the population-level. Both of these processes require population-level explanations and are thus inherently probabilistic; they abstract away from individual causal differences, leaving us with probabilistic generalizations. It is a little harder to see the argument in the case of random mutation. Of the three processes, mutation is the most likely to be genuinely indeterministic, since (in the terminology of Chapter 3) the quantum-level probabilities do not have to "percolate up" to as high of a level. Nonetheless, this is a distinct issue from the concept of random mutation that we have been discussing, which has to do with the kind of mutations (beneficial, harmful, or neutral) that organisms get. Yet mutation, in the context of the production of variation, is not a population-level process. Nonetheless (and this is more of a suggestion than an argument) it may be possible to claim that the level of causality at which the environment acts – or rather doesn't act, in the case of non-directed, random mutation – is not the same level at which the mutation occurs. That is, the

[54] It is interesting to speculate as to why all three conceptions have this character. The reasons may be historical; the concepts are defined in contrast to Darwin's theory of natural selection (although, as we have discussed, our conception of random mutation derives from Darwin's). The reasons may be pragmatic; human beings are interested in adaptation, and we thus divide evolutionary processes into those that are in an adaptive direction and those that do not have a specific direction. The reasons may be biological; perhaps the concepts capture genuinely different kinds of processes in the world. Or perhaps some combination of the three kinds of reasons provides the explanation for the commonalities.

environment impinges upon the organism, but the mutation occurs in the DNA. So, perhaps similar arguments can be made for the mutation case as were made in the random drift and stochastic macroevolutionary cases.

These conceptions of chance play a distinct role in evolutionary theory: an explanatory role. As we discussed in Chapter 5, there is no puzzle as to the way in which random mutations explain the production of variation in a population. Indeed, both proponents and critics of directed mutation accept that random mutations occur. Nor is the issue (at this point in time) really one of contrasting explanations. For most biologists, random mutation is the *only* explanation for the mutations which occur in natural and experimental populations. (Thus, in Chapter 5 we explored the more interesting question of how to determine empirically whether directed mutations occur *at all*). For random drift and stochastic macroevolutionary models, the situation is somewhat different. In these cases, the alternative form of explanation (natural selection explanations and deterministic explanations, respectively) is the more traditional, accepted explanation. In addition, there is a real question as to how these random processes can explain apparently non-random phenomena. However, once we understand that random drift and stochastic macroevolutionary processes are *causal* processes, we can see how, under a causal theory of explanation, they can be explanatory. Furthermore, once we understand the nature of the Markov process, it is no longer puzzling as to how random drift and stochastic macroevolutionary models can account for apparently non-random phenomena, such as the fixation of a trait within a population (for random drift) or the long periods of increase in taxonomic diversity (for stochastic macroevolutionary models). These causal and Markovian elements endow random drift and stochastic macroevolutionary models with the ability to serve as explanatory alternatives to their more traditional counterparts. In sum, all three concepts of chance play an explanatory role within evolutionary theory: random mutation in the production of variation, random drift in the perpetuation of variation at the microevolutionary level, and stochastic macroevolutionary models in the perpetuation of variation at the macroevolutionary level.

Of course, it is not the case that in any of the three areas, all biologists see the

alternative explanations as *equal* alternatives. As we have discussed, differing explanations are often proposed for the same data in each of the three areas. Numerous biologists have offered alternative explanations (all involving random mutation) for the experimental results of Cairns, Overbaugh, and Miller, who maintained that they had demonstrated directed mutation. The fluctuation of gene frequencies from generation to generation in a particular population is explained by some biologists as being due to natural selection in a fluctuating environment, and by others as the results of random drift. Common patterns of diversity across certain taxa are given a deterministic explanation by some, and a stochastic explanation by others. Thus, biologists often disagree over which explanation is the more appropriate in a *specific* instance (sometimes by granting a small role for the less favored explanation). At other times, biologists in all three areas disagree over the prevalence of the alternative processes *in general*. That is, biologists disagree over whether all mutations are random, or whether some mutations are directed. They disagree over whether microevolution is largely the result of natural selection, or largely the result of random drift. And they disagree over whether macroevolutionary patterns of diversity change are largely due to deterministic or stochastic processes. In the directed/random mutation controversy, most critics of directed mutation deny the soundness of directed mutation explanations outright. In the other two controversies, disputants are more likely to argue that their favored explanation is to be preferred over the other, without disallowing a role for the other form of explanation entirely.

As Beatty has argued, it is difficult to see how one could precisely formulate such relative significance disputes, much less resolve them. Thus, it is difficult to determine the extent of the role of chance in the evolutionary *process.* In terms of evolutionary *theory,* random mutation clearly has a larger and more entrenched explanatory role than directed mutation. Of the other two areas, my general impression from the literature is that random drift is a more accepted form of explanation in microevolution than stochastic models are in macroevolution. Most evolutionary biology textbooks include large sections on both random drift and natural selection. Stochastic macroevolutionary models, on the other hand, are often given little or no coverage in textbooks. Moreover, Gould, a proponent of

stochastic macroevolutionary models, bemoans the prevalence of deterministic macroevolutionary explanations in the journal *Paleobiology* (Gould 1995).

Thus while it is clear that chance plays a fairly large role in evolutionary theory, it is hard to know the extent of its role in the evolutionary process. However, given the potential for probabilistic processes to leave their signature on the course of evolution (both on a small scale and on a large scale), it is seems fair to say that the evolution of life on our planet has at least in some part been due to chance.

References

Beatty, J. (1984), "Chance and Natural Selection", *Philosophy of Science 51*: pp. 183-211.

Beatty, J. (1992), "Random Drift", in E. F. Keller and E. A. Lloyd (eds.), *Keywords in Evolutionary Biology.* Cambridge, MA: Harvard University Press, pp. 273-281.

Bookstein, F. L. (1987), "Random Walk and the Existence of Evolutionary Rates", *Paleobiology 13*: 446-464.

Brandon, R. (1990), *Adaptation and Environment*. Princeton, NJ: Princeton University Press.

Brandon, R. and S. Carson (1996), "The Indeterministic Character of Evolutionary Theory: No 'No Hidden Variables' Proof But No Room for Determinism Either", *Philosophy of Science 63*: 315-337.

Buri, P. (1956), "Gene Frequency in Small Populations of Mutant *Drosophila*", *Evolution 10*: 367-402.

Cain, A. J. (1951), "So-Called Non-Adaptive or Neutral Characters in Evolution", *Nature 168*: 424.

Cain, A. J. and J. D. Currey (1963), "Area Effects in *Cepea*", *Philosophical Transactions of the Royal Society of London B 246*: 1-81.

Cairns, J., J. Overbaugh, et al. (1988), "The Origin of Mutations", *Nature 335*: 142-145.

Cartwright, N. (1980), "The Reality of Causes in a World of Instrumental Laws", in P. Asquith and R. Giere (eds.), *PSA 1980*. East Lansing, MI: Philosophy of Science Association, pp. 38-48.

Cartwright, N. (1989), *Nature's Capacities and Their Measurement*. Oxford: Clarendon Press.

Cavalli-Sforza, L. L. (1969), "Genetic Drift in an Italian Population", *Scientific American 223*: 30-37.

Cavalli-Sforza, L. L. and J. Lederberg (1956), "Isolation of Preadaptive Mutants in Bacteria by Sib Selection", *Genetics 41*: 367-81.

Charlesworth, D., B. Charlesworth, et al. (1988), "Scientific Correspondence: Origin of Mutants Disputed", *Nature 336*: 525-527.

Darwin, C. ([1859] 1964), *On the Origin of Species: A Facsimile of the First Edition*. Cambridge, Massachusetts: Harvard University Press.

Dobzhansky, T. (1959), "In Discussion of Lamotte, M.", *Cold Spring Harbor Symposia on Quantitative Biology 24*: 85-6.

Dodson, E. O. and P. Dodson (1985), *Evolution: Process and Product*. Boston: Prindle, Weber and Schmit.

Feller, W. (1968), *An Introduction to Probability Theory and Its Applications*. New York: John Wiley & Sons.

Fisher, D. C. (1986), "Progress in Organismal Design", in D. M. Raup and D. Jablonski (eds.), *Patterns and Processes in the History of Life*. Berlin: Springer-Verlag, pp. 99-117.

Fisher, R. A. (1953), *Design of Experiments*. London: Oliver & Boyd.

Flessa, K. W. and J. Imbrie (1973), "Evolutionary Pulsations: Evidence from Phanerozoic Diversity Patterns", in D. H. Tarling and S. K. Runcorn (eds.), *Implications of Continental Drift to the Earth Sciences*. London: Academic Press, pp. 247-285.

Flessa, K. W. and J. S. Levinton (1975), "Phanerozoic Diversity Patterns: Tests for Randomness", *Journal of Geology 83*:

Foster, P. L. (1991), "Directed Mutation in *Escherichia Coli:* Theory and Mechanisms", in A. I. Tauber (ed.), *Organism and the Origins of Self*. Dordrecht: Kluwer, pp. 213-234.

Futuyma, D. (1988), "Sturm Und Drang and the Evolutionary Synthesis", *Evolution 42*:

Gardner, E. J., M. J. Simmons, et al. (1991), *Principles of Genetics*. New York: John Wiley and Sons, Inc.

Giere, R. N. (1980), "Causal Systems and Statistical Hypotheses", in L. J. Cohen and M. Hesse (eds.), *Applications of Inductive Logic*. Oxford: Clarendon Press, pp. 251-270.

Giere, R. N. (1991), *Understanding Scientific Reasoning*. Fort Worth: Holt, Rinehart, and Winston, Inc.

Gigerenzer, G., Z. Swijtink, et al. (1989), *The Empire of Chance: How Probability Changed Science and Everyday Life*. Cambridge: Cambridge University Press.

Gould, S. J. (1995), "A Task for Paleobiology at the Threshold of Majority", *Paleobiology 21*: 1-14.

Gould, S. J., D. M. Raup, et al. (1977), "The Shape of Evolution: A Comparison of Real and Random Clades", *Paleobiology 3*: 23-40.

Guyer, C. and J. B. Slowinski (1993), "Adaptive Radiation and the Topology of Large Phylogenies", *Evolution 47*: 253-263.

Hartl, D. L. and A. G. Clark (1989), *Principles of Population Genetics*. Sunderland, MA: Sinauer Assocs.

Hodge, M. J. S. (1987), "Natural Selection as a Causal, Empirical, and Probabilistic Theory", in L. Krüger (ed.), *The Probabilistic Revolution*. Cambridge, MA: MIT Press, pp. 233-270.

Horan, B. L. (1994), "The Statistical Character of Evolutionary Theory", *Philosophy of Science 61*: 76-95.

Hume, D. ([1777] 1988), *Enquiries Concerning Human Understanding and Concerning the Principles of Morals*. New York: Oxford University Press.

Humphreys, P. (1989), *The Chances of Explanation: Causal Explanation in the Social, Medical, and Physical Sciences*. Princeton, NJ: Princeton University Press.

Jablonka, E. and M. J. Lamb (1995), *Epigenetic Inheritance and Evolution: The Lamarckian Dimension*. Oxford: Oxford University Press.

Keller, E. F. (1992), "Between Language and Science: The Question of Directed Mutation in Molecular Genetics", *Perspectives in Biology and Medicine 35*: 292-306.

Kimura, M. (1983), *The Neutral Theory of Molecular Evolution*. Cambridge: Cambridge University Press.

Kimura, M. (1991), "Recent Development of the Neutral Theory Viewed from the Wrightian Tradition of Theoretical Population Genetics", *Proceedings of the National Academy of Science 88*: 5969-5973.

Lamarck, J. B. ([1809] 1914), *Zoological Philosophy*. London: Macmillan.

Lederberg, J. and E. M. Lederberg (1952), "Replica Plating and Indirect Selection of Bacterial Mutants", *Journal of Bacteriology 63*: 399-406.

Lenski, R. E. (1989), "Are Some Mutations Directed?", *Trends in Ecology and Evolution 4*: 148-150.

Lenski, R. E. and J. E. Mittler (1993), "The Directed Mutation Controversy and Neo-Darwinism", *Science 259*: 188-194.

Luria, S. E. and M. Delbrück (1943), "Mutations of Bacteria From Virus Sensitivity to Virus Resistance", *Genetics 28*: 491-511.

Mayr, E. (1970), *Populations, Species, and Evolution*. Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

Mayr, E. (1978), "Evolution", *Scientific American 239*: 46-55.

Mayr, E. (1983), "How to Carry Out the Adaptationist Program?", *American Naturalist 121*: 324-334.

Mills, S. K. and J. Beatty (1979), "A Propensity Interpretation of Fitness", *Philosophy of Science 46*: 263-286.

Mittler, J. E. and R. E. Lenski (1990), "New Data on Excisions of Mu From *E. coli* MCS2 Cast Doubt on Directed Mutation Hypothesis", *Nature 344*: 173-175.

Newcombe, H. B. (1949), "Origin of Bacterial Mutants", *Nature 164*: 150-1.

Peterson, J. A. (1983), "The Evolution of the Subdigital Pad in *Anolis*. I. Comparisons Among the Anoline Genera", in G. J. Rhodin and K. Miyata (eds.), *Advances in Herpetology and Evolutionary Biology*. Cambridge, MA: Museum of Comparitive Zoology, Harvard University, pp. 245-283.

Provine, W. (1986), *Sewall Wright and Evolutionary Biology*. Chicago: University of Chicago.

Raup, D. M. (1977), "Stochastic Models in Evolutionary Paleontology", in A. Hallum (ed.), *Patterns of Evolution As Illustrated by the Fossil Record*. Amsterdam: Elsevier, pp. 59-78.

Raup, D. M. (1985), "Mathematical Models of Cladeogenesis", *Paleobiology 11*: 42-52.

Raup, D. M. and S. J. Gould (1974), "Stochastic Simulation and the Evolution of Morphology – Towards a Nomothetic Paleontology", *Systematic Zoology 23*: 305-322.

Raup, D. M., S. J. Gould, et al. (1973), "Stochastic Models of Phylogeny and the Evolution of Diversity", *Journal of Geology 81*: 525-542.

Raup, D. M. and L. G. Marshall (1980), "Variation Between Groups in Evolutionary Rates: a Statistical Test of Significance", *Paleobiology 6*: 9-23.

Ridley, M. (1993), *Evolution.* Boston: Blackwell Scientific Publications.

Rosenberg, A. (1994), *Instrumental Biology or the Disunity of Science*. Chicago and London: University of Chicago Press.

Roughgarden, J. (1996), *Theory of Population Genetics and Evolutionary Ecology: An Introduction*. Upper Saddle River, NJ: Prentice Hall.

Sarkar, S. (1990), "On the Possibility of Directed Mutations in Bacteria: Statistical Analyses and Reductionistic Strategies", in A. Fine, M. Forbes and L. Wessels (eds.), *PSA 1990*. East Lansing, Mich: Philosophy of Science Association, pp. 111-124.

Sarkar, S. (1991), "Lamarck *Contra* Darwin, Reduction Versus Statistics: Conceptual Issues in the Controversy Over Directed Mutagenesis in Bacteria", in A. I. Tauber (ed.), *Organism and the Origins of Self*. Dordrecht: Kluwer, pp. 235-271.

Schopf, T. J. M. (1974), "Permo-Triassic Extinctions: Relation to Sea-Floor Spreading.", *Journal of Geology 82*: 129-143.

Schopf, T. J. M. (1979), "Evolving Paleontological Views on Determinism and Stochastic Approaches", *Paleobiology 5*: 337-352.

Schopf, T. J. M., D. M. Raup, et al. (1975), “Genomic Versus Morphology Rates of Evolution: Influence of Morphologic Complexity”, *Paleobiology 1*: 63-70.

Shanahan, T. (1991), “Chance as an Explanatory Factor in Evolutionary Biology”, *History and Philosophy of the Life Sciences 13*: 249-269.

Shapiro, J. A. (1984), “Observations on the Formation of Clones Containing *araBlacZ* cistron fusions”, *Molecular and General Genetics 194*: 79-90.

Simberloff, D. S. (1974), “Permo-Triassic Extinctions: Effects of Area on Biotic Equilibrium”, *Journal of Geology 82*: 267-274.

Smith, C. A. F. (1977), “Diversity Associations as Stochastic Variables”, *Paleobiology 3*: 41-48.

Sniegowski, P. D. and R. E. Lenski (1995), “Mutation and Adaptation: The Directed Mutation Controversy in Evolutionary Perspective”, *Annual Review of Ecology and Systematics 26*: 553-78.

Sober, E. (1994), “Evolution, Population Thinking, and Essentialism”, in E. Sober (ed.), *Conceptual Issues in Evolutionary Biology*. Cambridge, MA: MIT Press, pp. 161-189.

Stanley, S. M., P. W. S. III, et al. (1981), “Natural Clades Differ From "Random" Clades: Simulations and Analyses”, *Paleobiology 7*: 115-127.

Stewart, F., D. Gordon, et al. (1990), “Fluctuation Analysis: The Probability Distribution of the Number of Mutants under Different Conditions”, *Genetics 124*: 175-185.

Symonds, N. (1989), “Anticipatory Mutagenesis?”, *Nature 337*: 119-120.

Thayer, C. W. (1979), “Biological Bulldozers and the Evolution of Marine Benthic Communities”, *Science 203*: 458-461.

Turner, J. R. G. (1987), “Random Genetic Drift, R. A. Fisher, and the Oxford School of Ecological Genetics”, in L. Krüger (ed.), *The Probabilistic Revolution*. Cambridge, MA: MIT Press, pp. 313-354.

Uhen, M. D. (1996), “An Evaluation of Clade-Shape Statistics Using Simulations and Extinct Families of Mammals”,

Wright, S. (1949), “Population Structure in Evolution”, *Proceedings of the American Philosophical Society 93*: 471-478.

Wright, S. (1955), “Classification of the Factors of Evolution”, *Cold Spring Harbor Symposium on Quantitative Biology 20*: 16-24.

Wright, S. (1964), “Biology and the Philosophy of Science”, *The Monist 48*: 265-290.