RANDOM DRIFT: CHANCE AND EXPLANATION IN EVOLUTIONARY BIOLOGY

by

Adam Goldstein

A dissertation submitted to The Johns Hopkins University in conformity with the

requirements for the degree of Doctor of Philosophy.

Baltimore, Maryland

March, 2006

Abstract

The central claim for which I argue in this dissertation is that there are important phenomena that occur by random drift that evolutionary biologists explain using a strategy I term "process explanation." This claim puts me at odds with an influential view about the nature of explanation that I term "Hempelianism." Hempelianism is the view that the scientific explanation of a particular event E requires (a) showing that E was to be expected, or indicating the degree to which it would have been rational to expect E's occurrence; and (b) laws of nature. My central claim entails that both (a) and (b) are false. A process explanation consists of a narrative describing events causally relevant to the event to be explained. These narratives need not contain laws, show that the event to explained ought to have been expected, or indicate the degree to which it would have been rational to expect the event. My position about random drift also puts me at odds with evolutionists who, influenced by Hempelianism, claim that only natural selection can explain evolution. In my argument, I articulate the strategy of process explanation and defend it against Hempelian critics; describe a mechanism of random drift known as "indiscriminate

sampling;" and describe process explanations of phenomena of drift that occur by indiscriminate sampling.

Advisor: Dr. Peter Achinstein.

Acknowledgements

I would like to warmly extend my thanks to my wife Abigail, who—in contrast with the chance fluctuations in the fortunes of this project, which she endured with grace and magnanimity—steadily provided me with encouragement and support. Additionally, I would like to acknowledge my graduate advisors, Karen Neander, Peter Achinstein, and Steven Stanley; Alexander Rosenberg of Duke University, who served as a committee member; my teachers at Johns Hopkins, Stephen Barker, George Wilson, Maura Tumulty, and Susan Wolf; the late David Sachs, for his endowment of a fellowship at Johns Hopkins, from which I benefited; former graduate students at Johns Hopkins Richard Richards, Kent Staley, and Chuck Ward; my undergraduate teachers David Hoy, W. E. Abraham, Léo Laporte, and Todd Newberry; and my colleagues at Drew University, Seung-Kee Lee, Erik Anderson, and Thomas Magnell. The Department of Philosophy at Johns Hopkins provided me with financial support and with a rich and challenging environment for research. The graduate school and university as a whole also deserve acknowledgement for these reasons. Much of this dissertation was written in the Rose Main Reading Room of the New York Public Library, and I would like to thank the staff there for research support, and for providing a quiet and dignified place in which to write. Sara Harrington of the Rutgers University

Art Library provided facilities for scanning and editing the images reproduced in Appendix A. The open-source computing community made a variety of first-class software packages freely available, including LaTeX, which I used, together with BibDesk and TeXShop for Mac OS X, to write and typeset the dissertation.

Contents

List of Tables

List of Figures

Chapter 1

Chance, Explanation, and

Narrative in Evolutionary Biology

This dissertation concerns a process of evolution known as random genetic drift, also known as "random drift" or just "drift." Drift is nonadaptive, and may be thought of as a process of evolution by chance or accident. In this regard, it differs from natural selection, which is adaptive, and may be thought of as purposive. The main claim for which I argue in this dissertation is that there are important phenomena that occur by drift that evolutionary biologists explain using a strategy that I term "process explanation."

By claiming that there are important phenomena that occur by drift that evolutionary biologists explain, I disagree with an influential and articulate group of evolutionists who adhere to what I term "the exclusivity thesis." The exclusivity thesis is the view that natural selection alone explains evolution—a view that entails that phenomena occurring by drift cannot be explained. The resolution of this disagreement does not depend on scientific

theory and observation alone; rather, it depends in part on the resolution of an important conceptual issue about the nature of explanation.

This conceptual issue is the nature and justification of process explanation, the strategy that (as I state above) I claim that evolutionary biologists use to explain phenomena that occur by drift. Process explanation, which I elaborate and defend in the dissertation, is incompatible with a widely-held view about explanation. I term this view "Hempelianism" to reflect its close association with the philosopher Carl Hempel. Hempelianism is the view that explaining a particular event E requires showing *why* E occurred by citing laws of nature. In contrast, a process explanation of a particular event E derives its force from a narrative describing E's causes; such a narrative describes *how* E occurred, and does not require any laws of nature.

This disagreement with the Hempelians reflects a deeper disagreement about the nature of explanation. Hempelianism entails what I term "universalism" about explanation. Universalism is the view that there is an unique set of criteria for explanation. For the universalist, the characteristics of the audience are irrelevant for judging the adequacy of an explanation. In contrast, my claim that there are process explanations reflects what I term "contextualism." Contextualism is the view that the characteristics of the audience are relevant for judging the adequacy of an explanation, which depends upon the explanation's context of utterance.

The goal of this chapter is to introduce the claims and arguments I have just mentioned, and to provide an overview of how I will develop them in the remainder of the dissertation. In section 1.1, I describe random drift; in section 1.2, I address issues concern-

ing Hempelianism, the exclusivity thesis, and my response to them; and in section 1.3, I summarize the chapter and provide a chapter-by-chapter overview of the remainder of the dissertation.

1.1 What Is Random Drift?

My view is that random drift is a kind of evolution by chance or accident, and my aim in this section is to indicate what I mean by this. I will proceed at a broad level, without providing the kind of detail that would be required for a philosophical analysis of the concept, a project I carry out, in part, in chapter 4. In section 1.1.1, I consider how scientists usually describe drift; In section 1.1.2, I describe the properties of an evolutionary process that an evolutionist would have in mind, if he or she believed that process to be one of random genetic drift; and finally, in section 1.1.3, I elaborate the sense in which, in contrast with natural selection, drift is a kind of evolution by accident.

1.1.1 A look at the literature

Sewall Wright, one of the inventors of the current notion of random drift, often describes random drift as a chance process, as random, or as a kind of accident. This can be seen in his 1931 "Evolution in Mendelian Populations" [156].[1]

---

[1]Here as in the remainder of the dissertation, citations appear in the text in square brackets. Each consists of at least two numbers. The first corresponds to one of the works referred to in the bibliography. In no instance do I cite more than one work within a single pair of square brackets. The second number refers to a page number, unless otherwise noted that it refers to a section, chapter, or some other element of the work cited. I sometimes cite additional pages, chapters, or sections. These appear, separated by commas, after the second number in brackets. So, to consider an example of one of the more complex references that I might make, I will refer to pages 22 through 24, and page 29, of work the work listed in the bibliography as work number 10, by the following: [10, 22 – 24, 29]. Other annotations may appear in the square brackets, as well: [10, 15, eqn. 3.1] indicates page 15, equation 3.1, in work 10, for instance.

The constancy of gene frequencies in the absence of selection, mutation, and migration cannot for example be expected to be absolute in populations of limited size. Merely by chance one of the allelomorphs may be expected to increase its frequency in a given generation and in time the proportions may drift a long way from the original values. The decrease in heterozygosis following inbreeding is a well known statistical consequence of such chance variation. In this case ... [there will be a significant change] merely as a result of random sampling among the gametes. [156, 107]

Wright makes similar remarks in a 1932 paper, "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution" [162].

[In addition to selection, mutation, and migration] another factor must be taken into account: the effects of accidents of sampling among those that survive and become parents in each generation and among the germ cells of these.... Gene frequency in a given generation is in general a little different one way or the other from that in the preceding, merely by chance. [162, 165]

Other evolutionists writing more recently have maintained Wright's view that drift is a kind of evolution by accident or chance. Consider the following passage from Gillespie's *Population Genetics: A Concise Introduction* [56]. Besides his explicit mention of chance and randomness, Gillespie describes drift as trendless, supporting my claim that scientists view drift to be a chance process, or a kind of evolution by accident: one would expect just such trendless, directionless evolution from such a process.

[R]andom changes in allele frequencies result from [two mechanisms of random drift:] variation in the number of offspring between individuals and, if the species is diploid and sexual, from Mendel's law of segregation.... One [important feature of random drift] of course, is that genetic drift causes random changes in allele frequencies.... [T]he direction of the random changes is neutral. There is no systematic tendency for the frequency of alleles to move up or down. [56, 21]

Jonathan Roughgarden, in another well-known introduction to population genetics, says that drift provides "our entrée into a stochastic theory of evolution" [124, 57]. Regarding stochastic theories, Roughgarden explains as follows.

> [I]n a stochastic theory we assume that it is possible for the system under study to be in many states, and we develop a theory to predict the probability of the system's being in each possible state as a function of time. [124, 57]

He elaborates, describing the role of random drift in evolution.

> Evolution is an affair of chance. Some gametes by chance are drawn from the gamete pool and incorporated into zygotes while others are washed to sea; some types of individuals may, on the average, leave more offspring than others but what actually happens depends in part on chance. [124, 57]

Reading these passages as follows, I interpret Roughgarden's view to be that drift is a chance process. A stochastic theory can indicate the probability that a population will have a certain allele frequency after a given amount of time, under certain conditions. However, the actual state of the population at the time in question is, at least in some degree, due to chance. This contrasts with what is often called a "deterministic" theory, which describes the time-course that the allele frequencies in a population will take, necessarily. Accordingly, Roughgarden suggests that when evolutionary processes described by stochastic theories occur—that is, when drift occurs—"what actually happens depends in part on chance."

Finally, I want to point out that Douglas Falconer, in his textbook on quantitative genetics [45], understands drift in much the same way as Wright and Gillespie.

> The random changes of gene frequency are called random drift. If the gene frequency in any one small population is followed, it may be seen to change in an erratic manner from generation to generation, with no tendency to revert to its original value. [45, 51]

Calling random drift "the dispersive process," he points to what he terms "the sampling of gametes," which is one of the biological mechanisms that can produce drift. The point is completely general, however, and applies to all processes of drift.

[The dispersive process] differs from the systematic processes in being random in direction.... In order to exclude this process from previous discussions we have postulated always a "large" population, and we have seen that in a large population the gene frequencies are inherently stable.... This property of stability does not hold in a small population, and the gene frequencies are subject to random fluctuations arising from the sampling of gametes.... This random change of gene frequency is the dispersive process. [45, 51]

1.1.2   The process of drift

So evolutionary biologists believe that random drift is a chance process, or a process of evolution by accident; what do they mean by this? My aim in this section is to answer this question by sketching an account of drift that highlights the connection between causal-mechanical and teleological terms used to describe drift. I proceed in further subdivisions of this section, the initial series of which is directed at developing a notion of biological purpose, which is essential for the account of drift I formulate in its own subsection. The concluding subsection describes an example that is illuminated by the account, catastrophic drift.

"Trait" and "variant"

I understand the category indicated by "trait" to be quite broad, including behavior, morphological structures, and physiological processes. Although I believe that it is not clear in general what ought to count as a trait, the account of drift I elaborate below does not depend on the resolution of this issue, and I will not examine it any further here. A given trait can take various forms, each of which I term an "alternative variant" or a "variant" of the trait. So, for instance, eye color is a trait; brown and blue are variants of this trait. Furthermore, it may be assumed that all of the traits and variants I consider

below are heritable.

Sex, diploidy, and the life cycle

In general, evolutionary success for a variant—or, more precisely, for the genes of a variant—consists in being passed on to the next generation. This means that organisms bearing the variant must succeed at all stages of the life cycle, each of which poses different challenges for organisms and their genes to meet. Among these challenges, there are the following: coming to term and being born, after having been created by fertilization; surviving to adulthood; and finding a mate and successfully contributing genes in sex with that mate. In order to keep the formulation of my account of drift as simple as possible, I will only consider drift as it occurs in certain early stages of the life cycle of sexually reproducing, diploid populations. In these early stages, survival, and not success in sexual reproduction, is the critical evolutionary challenge.

Populations with a "main variant"

A population with a main variant has one variant of the trait of interest that predominates, in the sense that it is the most frequent variant in the population. Because of its predominance, I term it the "main variant." The idea is that the main variant is the one that would be represented in a zoology textbook, for instance, as typical or normal for animals of a given species. Furthermore, for a variant to be the main variant, it must have become the most frequent in the population by natural selection, in the current environment.

Organic purposes

My account of organic purposes is as follows.

Analysis 1.1 *A trait T in a sexually reproducing, diploid species S (or in some sub-population of S, S^*) has the purpose P if and only if P is what the main variant V of T did better than alternative variants that caused V to become statistically dominant in S (or S^*) by natural selection.*[2]

To be clear, the idea is that the purpose of variants other than the main variant is indexed to the main variant, so that the main variant establishes the purpose for the trait and all its variants. While other variants do not carry out their purpose as well as the main variant, they have the same purpose that it does. Note that this notion of purpose is explicitly historical. The purpose of a trait is indexed to something that past instances of one of its variants did, that caused present instances of that variant to be the most frequent in the population, that is, to become the main variant.[3]

Let me illustrate my account of purposes with a hypothetical example. Suppose that there are two variants of a certain trait in a population. Suppose that the trait is the

---

[2]Throughout the dissertation, I will introduce statements and questions to which I want to call special attention.

Analysis A statement of the form "— if and only if —."

Explanation A set of statements answering an explanation-seeking question.

Question An interrogative, usually to request an explanation.

Statement Statements not of the form "— if and only if —" that I intend to address or argue for, or that are proposed necessary or sufficient conditions for a concept under discussion.

Each of the above are indexed to the chapter in which they appear: statement $x.y$ is the *yth* statement in chapter x. Analyses have an index that is independent of that for statements, e.g., a new statement does not increase the index number y for analyses in the same chapter. The same is true for the rest of the elements described above.

[3]My account of purposes converges with the so-called etiological theory of biological function [5]. I was not primarily motivated by the latter as I developed my notion of purpose, although it did provide a guide for me as I was doing so.

coat color, say, of some squirrels. The variants are different colors of brown. Darker brown squirrels, it may be supposed, have a higher probability of survival than do lighter brown squirrels.

Suppose that the reason for this is that the darker brown squirrels blend into the background better than lighter brown squirrels, and so, are harder for their predators to spot. This means that the darker brown coat color confers a greater advantage for survival on its bearers than the lighter brown coat color. I will also suppose that the darker coat color has become better represented in the population because of the advantage it conferred on its bearers: due to that advantage, they succeeded in natural selection over their lighter-colored conspecifics.

This means that the darker coat color is the main variant, and that, accordingly, the purpose of the coat color trait is to camouflage the squirrels. This is what the main variant (the darker-color coats) did that caused bearers of that trait to succeed in natural selection. The purpose of the lighter-colored coats, which are far less frequent than darker colored coats, is also to camouflage their bearers. This is because their purpose is indexed to that of the main variant.

Random Drift

Broadly speaking, I understand random drift in diploid, sexually reproducing populations, occurring due to differences in survival rates of variants, as follows.

Analysis 1.2 (Random drift) *In a sexually reproducing, diploid population P with trait T that has variants V_1 and V_2, one of which is the main variant of T, random drift of T*

occurs in the degree to which cross-generational changes in the relative frequencies of V_1 and V_2 are due to differences in rates of survival of organisms with alternative variants that are not due to differences in the abilities of those variants to carry out their purposes, ceteris paribus.

The idea is that, when populations of the sort that I have been discussing evolve by drift, the superior advantage for survival that the main variant confers on its bearers does not result in a commensurately greater rate of survivorship among organisms with the main variant. This is possible because, in general, having a variant that confers a greater ability to survive does not guarantee survival. Indeed, it is to be expected that some organisms with the main variant will die each generation. In general, what is expected is that the relative number of organisms with the main variant that succeed in carrying out their purposes and whose bearers survive should be roughly proportional to the relative advantage it confers on its bearers, in contrast with other variants.

Nonetheless, it may happen that, for whatever reason, there is a disproportionately large number of organisms with the main variant whose copies of the variant do not successfully carry out their purposes, and whose bearers die because of this. This is drift, in the form of what might be termed "survival of the unfit." Likewise, there might be a disproportionately large number of organisms with the main variant that carry out their purposes successfully, and whose bearers survive because of this. This, too, is drift, in the form of a chance flush of fitter organisms.

The *ceteris paribus* clause I have appended to my account of drift above is intended to indicate that drift occurs in the degree to which relative frequencies of the variants in

question *would have changed* if there were no countervailing evolutionary processes. For instance, suppose that there is some population in which drift of some trait occurs in some degree, but in which natural selection occurs in an equal and opposite degree. Thus, there is no net change in the relative frequencies of the variants in question. However, this does not mean that drift and natural selection did not occur: they occurred, but "canceled out" one another.

Additionally, note that the account requires that changes in the relative frequency of the variants be cross-generational. This reflects the generally accepted definition of "evolution," according to which evolution of a trait T occurs only if there are cross-generational changes in the relative frequency of T's variants. Also, note that I state the account for traits with two variants only, although it is perfectly general, and—with some new notation, which is less illuminating—can be extended to any number of variants.

Let me illustrate my account of drift with a numerical example. Suppose that the main variant gives its bearers, say, twice as much of an advantage for survival as another variant does. If survivorship among these organisms were due to differences in the abilities of the alternative variants each bears to carry our their purposes and nothing else, it would be the case that, for every two organisms with the main variant that survives, there would be only one survivor with the other variant.

Nonetheless, such proportionality may not obtain. Suppose that a particularly large number of organisms that have the better variant do not survive, because their copies of the variant in question failed to carry out their purposes: although the better variant will in general perform twice as reliably as the other variant in the long run, it fails to do

so in the case at hand. This would be drift. Though there is a difference in the abilities of alternative variants to carry out their purposes, this is not reflected in the relative rates of survival of organisms with each variant.

Catastrophic drift

In this section, I want to provide an example of random drift: what Steven Stanley (pers. comm.) calls "catastrophic drift." I do not have any actual case of catastrophic drift in mind, so I will just outline how a case of it would look, if it occurred. Nevertheless, Stanley (pers. comm.) believes that catastrophic drift is very frequent in the history of life, and so I think that I am not considering a mere possibility.

Catastrophic drift proceeds as follows. Suppose that, every generation, there is a small chance that some catastrophe such as a forest fire, storm, or other disaster occurs. Considering factors other than the catastrophe, it might be the case that the population has variants that make their bearers better able to find food, survive disease, drought, weather, predation, and so on. Suppose that some organisms have a variant that gives them a five-fold advantage for surviving predation, and that this is its purpose. However, considering the disaster, suppose that all organisms are equal, in the sense that no organism possesses any advantage for survival with regard to the storm. The idea is just that, even though the storm affects each organism's probability of survival, it does so in an equal degree for each organism, and so does not result in any changes in their relative fitness.

Now, suppose the storm occurs, reducing the population to a few individuals. It might be that the remaining individuals all have the variant that confers less of an advantage for survival with regard to predators. Or, rather, it might be the case that each

of the remaining individuals has the variant that confers more of an advantage for survival with regard to predators. Either way, the proportion of survivors with each type of variant is not due to differences in their ability to carry out their purposes. As I suggested above, no variants differ in their ability to survive the catastrophe. Because of this, evolution occurring due to it is drift.

1.1.3 Purpose and accident in evolution

In the last section, my aim was to explain the meaning of the claim that drift is a kind of evolution by accident. In this section, I would like to establish *why* random drift should be understood as a conception of chance in evolutionary biology. My position is that random drift is a conception of chance because of its role in teleological interpretations of evolution used by evolutionists. I want to begin my argument for this claim by introducing a contrasting conception of chance that I believe is weaker than the concept of chance that random drift represents.

This weaker conception of chance may be illustrated by a coin tossing example. Suppose a fair coin is tossed 1,000 times, and the resulting sequence of heads and tails has 530 heads and 470 tails. Although every possible sequence of heads and tails has an equal probability of occurring, it is not particularly improbable, in a sequence of 1,000 tosses of a fair coin, for some sequence to occur that has somewhere between, say, 450 and 550 tails. Usually, these deviations from a fifty-fifty distribution of heads and tails would be ascribed to chance. The idea seems to be that there is a main factor—the symmetry of the coin, that causes it to be fair, in this case—and this factor is responsible for a trend, or, as a statistician would say, "central tendency." Any results deviating from this central tendency

13

are ascribed to chance.

I want to be clear that I agree that random drift may be understood to be something like deviation from a central tendency. Differences in how well variants of a trait carry out their purposes establish the "central tendency;" deviations from this central tendency constitute the influence of chance. However, I do not think that these patterns of deviation from a central tendency are the characteristic of random drift that is most responsible for its being considered a conception of chance. Rather, my view is that random drift is a conception of chance because of the interpretation of the kinds of deviations from a central tendency that random drift creates. This interpretation takes shape in contrast with natural selection. Let me elaborate.

Natural selection, as a general rule, results in adaptive trends. This property of natural selection is embodied in an observation that was made by both R. A. Fisher and Sewall Wright [56, 59]: as a general rule, natural selection will increase the mean fitness of a population. What I want to indicate about this kind of trend is how it is interpreted by evolutionary biologists. Adaptive trends in a population are accompanied by an important qualitative change in that population: they create what appears to be intelligent design, and, across generations, promote the spread of apparently better-designed variants of a trait, removing the apparently less-well-designed traits from the population.

This is reflected in the account of natural purposes that I elaborated in the previous section. Recall that I identified the purpose of a trait as what the more successful of its ancestors did to promote its bearers' success in natural selection. The term "purpose" is not meant to be merely figurative or evocative. Rather, it is intended to reflect evolu-

tionists' practice of recognizing that products of natural selection are apparently products of intelligent design, just as artifacts are.

I think it would be easy to underestimate how forcefully contemporary evolutionists are struck by the impression that natural selection produces design. This would be a serious mistake. Consider the following passage, from a paper entitled "The Problem of Plan and Purpose in Nature" by paleontologist George Gaylord Simpson, one of the inventors of evolutionary theory as it stands today.

> We feel, almost instinctively, that there is a pattern.... There is, or seems to be, an essential order or plan among the forms of life in spite of their great multiplicity. There seems, moreover, to be purpose in this plan.... It is a habit of speech and thought to say that fishes have gills in order to breathe water, that birds have wings in order to fly, and that men have brains in order to think.... This appearance of purposefulness is pervasive in nature.... Accounting for this apparent purposefulness is a basic problem for any system of philosophy or of science. [133, 481]

Simpson adds that "adaptation does exist and so does purpose in nature.... Denial of this does violence to the most elementary principles of rational thought" [133, 489]. He concludes that "adaptation is real, and it is achieved by a progressive and directed process. This process is natural, and it is wholly mechanistic in its operation. This natural process achieves the aspect of purpose, without the intervention of a purposer" [133, 495]. Of course, the natural, mechanistic process he is referring to is natural selection.

The views expressed by Simpson are held by a near-total majority of contemporary evolutionists, including the most influential of them. George C. Williams expresses the view in *The Pony Fish's Glow: Clues to Plan and Purpose in Nature* [147]. Maynard Smith, cited approvingly by Dawkins [71, 16] remarks that "the main task of any theory of evolution is to explain adaptive complexity, i.e., to explain the same facts that Paley used as evidence of a

Creator." Michod [103, 163] expresses the view, arguing—in a section suggestively entitled "What Makes Biology Different?"—that biology is unique among natural sciences because it is concerned with the study of purposive objects, i.e., organisms and their traits. Ayala [7] also clearly expresses a teleological interpretation of evolution, as does Rudwick [125]. Williams and the others, of course, believe that natural selection is the process responsible for organic design, equaling or exceeding what a human designer would be able to achieve; none believe that a conscious designer acts in evolution.[4]

Now, having clarified the sense in which natural selection is identified by evolutionary biologists as a process of design, let me indicate the sense in which, by contrast, random drift is a process of evolution by chance or accident. There are two senses in which drift is a conception of a chance process or a process of accidental evolution: a causal sense and a statistical sense. To understand the causal sense in which drift is a kind of evolution by accident, recall my suggestion in analysis 1.2 above (pages 9 – 10) that if drift occurs there is some amount of evolution that is not caused by differences in the ability of alternative variants to carry out their purposes.

In purely causal terms, this means that there is some amount of evolution that is not due to differences in the ability of alternative variants to do what the main variant of the population was selected by natural selection to do. Interpreted teleologically, the idea is that there is a disproportionate number of organisms whose copies of a given variant do not operate in a manner that accords with what seems to be the intentions of their designer. Since the main variant has been selected by natural selection over the generations, it will

---

[4]Lewens [88] provides a particularly lucid description of what he terms the "artifact model," one form of reasoning according to which organisms are viewed as products of design.

have the look of having been designed. When drift occurs, the main variant will seem to have failed to perform as reliably as its designer intended, as though plagued by accidents.

The statistical sense in which random drift is a conception of chance has to do with the kinds of trends—or, more precisely, the lack thereof—that result from random drift. Consider an example. Suppose that, based on the differences in their ability to carry out their purposes, it would be expected that three organisms that have variant V_1 survive for every two organisms that have variant V_2 that do. Natural selection, if it acts without interference by drift or any other process, would result in an increase in the proportion of organisms with V_1, given many generations. Interpreted teleologically, this would be understood by biologists as an improvement in the population, by which the purpose of the trait at issue is refined. In contrast, drift will result in the frequencies of the two variants fluctuating back and forth across the mean value, "wandering" or "drifting."

This is a well-known consequence of drift, reflected in many of the statistical models of it, for instance, those described by Gillespie [56, ch. 2]. These models show the following. Suppose that, at a time T, variant V has a frequency F_T. Next, suppose that there is a frequency F_{T+t} that variant V might have at some later point in time, $T + t$. Suppose, furthermore, that F_{T+t} is greater than F_T by the amount $\triangle F$. Additionally, consider some alternative frequency F'_{T+t}, so that F'_{T+t} is less than F_T in the amount $\triangle F$. The important result of the mathematical models of drift that I want to call attention to is that, when drift acts alone, the variant V has an equal chance of having the frequency F_{T+t} as it does of having the frequency F'_{T+t}, at time $T + t$.

The idea is that the probability that drift will cause a variant to increase in fre-

quency by a given amount is the same as the probability that drift will cause that variant to decrease in frequency by that same amount. So, over many generations, drift introduces no trend into a population; the variant's frequency will fluctuate across the mean in an erratic manner. Interpreted teleologically, this means that there will be no pattern of improvement in the trait's design.

In closing, I want to contrast the sense in which I have claimed that drift is a kind of chance process with the weaker sense in which the results of coin tossing may be understood as a kind of chance process. Although the sequence of heads and tails produced by tossing a fair coin is due to chance, it would seem strange to say that the coin failed to operate in accord with its purpose because it failed to produce precisely equal proportions of heads and tails. Perhaps coin tossing is not a good example, because someone might think that a fair coin is designed to produce sequences with a fifty-fifty proportion of heads and tails. I think that the coin is best thought of as being designed to have an equal chance of coming up heads or tails, but in any case, to consider another example that may be clearer, suppose that the unusually high incidence of thunderstorms in some region in some period of time is ascribed to chance.

I think that only those who believe that the weather is controlled by a deity would think that the weather is not behaving as it was designed. Others ascribing the deviation to chance may imply that there is no systematic explanation for it; however, they do not imply that the weather system is not carrying out its purpose. While a purely statistical and causal-mechanical conception of chance can be applied to coins and the weather, a teleological conception of chance cannot be. This is disanalogous with drift, to which both

conceptions of chance apply, with the teleological conception doing a good deal of the work: drift encompasses the accidents in evolution.

1.2 Process Explanation in Evolutionary Biology

My aim in this section is to provide an overview of the controversy between Hempelians and myself concerning whether random drift has any explanatory power. Section 1.2.1 concerns a group of evolutionists who affirm what I term "the exclusivity thesis." The exclusivity thesis is the view that if and only if an evolutionary event can be explained, natural selection explains it. The evolutionists who hold the exclusivity thesis believe that it follows from Hempelianism; accordingly, I term these evolutionists "the Hempelian evolutionists." In section 1.2.2, I describe process explanation and contextualism, and how they inform the view—contrary to the exclusivity thesis—that there are events in evolution that can be explained by random drift.

1.2.1 Hempelianism and the Hempelian evolutionists

I divide my account of the Hempelian evolutionists' position into two further subsections: first, I describe Hempelianism and universalism; second, I outline the argument, informed by Hempelianism, for the exclusivity thesis.

Hempelianism and universalism

Consider the following passage from Mill's *Science of Logic*:

[E]xplaining, in the scientific sense, means resolving an uniformity which is not a law of causation from which it [the uniformity] results, or a complex law of

causation into simpler and more general ones from which it is capable of being deductively inferred.... [3, 216]

On the one hand, the account of explanation Mill articulates in this passage is not precisely coextensive with Hempelianism: Hempelianism does not apply only to causal explanation; it does not require deduction; and—although there may be a more general formulation of it—Hempelianism, as I construe it, is a doctrine about the explanation of particular events, not generalizations ("uniformities"). On the other hand, with characteristic pithiness, Mill captures the essential motivation for Hempelianism: to explain a particular event E, scientists adduce at least one law of nature for E's occurrence.

Carl Hempel's [64] famous "deductive-nomological" ("D-N") model of explanation exemplifies the main principles of Hempelianism. Roughly speaking, a D-N explanation of a particular event consists of a deductive argument whose premises contain at least one law of nature, and a description of conditions that obtained prior to the occurrence of the event to be explained, and whose conclusion is a statement to the effect that the event to be explained occurred. Deducing a claim describing the event to be explained from natural laws and local conditions, according to Hempel, shows why an event occurred by showing that it was to be expected as a matter of natural necessity.

To see how this works, consider a specimen D-N explanation of the event described by the following statement.

Statement 1.1 *The gravitational force between the Earth and Sun at time t is $F_G(E, S, t)$.*

The argument that explains the event described by statement 1.1 is as follows.

Explanation 1.1 *At any time T, for any objects O_1 and O_2 separated from one another*

20

by a distance R, each having a mass of M_1 and M_2 respectively, and where $F_G(O_1, O_2, T)$ is the force exerted on one body by the other due to gravity at time T, and G is a gravitational constant:

$$F_G(O_1, O_2, T) = G\frac{M_1 M_2}{R^2}.$$

At time t, the Sun has a mass of M_S, the Earth has a mass of M_E, and the two are separated by a distance of $R_{E \to S}$.

Therefore, at time t, the Earth exerts the following force on the Sun, due to gravity:

$$F_G(E, S, t) = G\frac{M_E M_S}{R^2_{E \to S}}.$$

The law of nature required by Hempelianism for the explanation of particular events appears in the first premise, which states Newton's law of gravitation; the second premise completes the argument by describing conditions local to the Earth and Sun. As the argument is deductively valid, it meets the Hempelian requirement that an explanation show why the event to be explained occurred: given relevant laws of nature, and given the circumstances antecedent to that event, it had to occur, as a matter of natural necessity.

Because Hempel believes that showing why an event occurred means showing that it was to have been expected, he rejects the idea that highly improbable events can be explained at all: if an event is highly improbable, it cannot be shown to have been expected. In contrast, some Hempelians believe that it is possible to explain highly improbable events because they believe that showing why an event occurred does not mean showing that it was to be expected. Rather, these Hempelians believe that laws of nature are invoked to show that there is some *degree of expectation* that it would have been rational to have that the event occur.

To see how this works, consider the following probabilistic law describing the radioactive decay of Uranium 238 (U^{238}), derived from an example proposed by Peter Railton [119, 125], a notable Hempelian.

Statement 1.2 (Law of decay of U^{238}) *All nuclei of U^{238} have a probability p of emitting an alpha-particle during any interval of length θ, unless subjected to environmental radiation.*

Supposing the time-interval θ to be relatively small, Railton [119, 124] points out that "since the mean half-life of U^{238} is 6.5×10^9 years, the probability p of observing a decay ... during this interval is exceedingly small" as long as no environmental radiation is applied. Thus, it would not be correct to expect that any decay be observed during an interval θ units of time in length; rather, it would be correct to expect that no decay be observed during that time. Despite this, Railton's view is that any decay that is in fact observed can be explained because the probabilistic law expressed by statement 1.2, by describing the probability of decay p, indicates the degree in which would have been rational to expect that the decay would occur. Railton believes that, even though p is diminishingly small, citing it still provides a reason why the decay occurred, thus answering the explanation-seeking question about particular events that Hempelians believe to be canonical, "Why did the event to be explained occur?"

As I have mentioned previously, Hempelianism is a variety of what may be termed "universalism" about scientific explanation. The central principle of universalism is that the criteria for a good explanation do not include reference to the audience's intentions and cognitive states. Hempel [64, 425 - 428] expresses this view using an illuminating analogy:

a good explanation, he claims, is like a good mathematical proof.

If a mathematical proof is a good one, its conclusion follows from its premises, a matter independent of the intentions or cognitive states of its intended audience. Surely, Hempel would point out, the fact that only a select group of mathematicians can follow each step of Andrew Wiles's proof of Fermat's last theorem [135] does not make it inadequate. Extending the analogy, the idea is that the gravitation explanation above is good if and only if it meets the criteria for a D-N explanation; these requirements do not include any reference to an audience. Rather, they describe certain kinds of deductive arguments. Other Hempelians propose similar theories.

The universalist does not deny that there are pragmatic issues associated with explanation. Hempel would agree, of course, that an explanation stated in terms that its audience cannot understand is of no use to that audience as an explanation of anything. The Hempelian recognizes a clear distinction between a *good explanation* and a *good explanation for so-and-so*: a good explanation may not be good for a particular audience, in the sense that it may not answer a question posed by that audience, or be understandable by it. Nonetheless, it does not fail to be a good explanation *simpliciter* because of this.

The Hempelian evolutionists

I will begin my account of the Hempelian evolutionists and their position with an account of the exclusivity thesis. Before stating the thesis, there are two sets of issues that I want to clarify: I want to describe how I will be using certain terms that appear in the thesis, and I want to describe the Hempelian evolutionists' attitude toward the principle of natural selection.

The first term I want to clarify is "evolutionary event." By this, I mean to indicate the occurrence of any event that constitutes the occurrence of evolution. Although I will not define "evolution" here, the central idea is as follows. An evolutionary event requires some form of heritable change in some biological aspect of a population. This includes events such as the following: the origin of a new allele; cross-generational changes in allele frequencies; cross-generational changes in the relative frequency of a variant; and the origin or extinction of a species. It does *not* include the following: the birth or death of an organism, unless it amounts to the origin or extinction of a species; an increase in population number within a generation; an instance of a biochemical mechanism operating in some organism; and sex among organisms.

The second term I want to clarify is "adaptation." By this, I mean to indicate the following. Consider a population with variants V_1 and V_2 of trait T.

Analysis 1.3 V_1 *is an adaptation for carrying out some activity A if and only if, by carrying out A, it makes a greater contribution to the relative fitness of its bearers than does* V_2.

Note that V_1 need not have any history of selection for A; it need only carry out A in the current environment, and satisfy the fitness requirement described in analysis 1.3.[5]

Next, I want to consider the Hempelian evolutionists' attitude toward the principle of natural selection, one formulation of which is as follows:

---

[5]I owe this conception of adaptation to Reeve and Sherman [120]. It is at odds with the conception of adaptation popular in the philosophical literature, described, for instance, by Sober [139, ch. 6]. I do not want to discuss why I think that the Reeve and Sherman's account should be adopted; suffice it to say that I think that the arguments that Reeve and Sherman make against Sober's conception of adaptation (as well as a variety of others) warrant abandoning Sober's in favor of theirs.

Statement 1.3 (Principle of Natural Selection) *If variant V_1 is an adaptation for activity A, then, as a result of doing A better than other variants, organisms with V_1 will survive in greater proportion than organisms with those other variants,* ceteris paribus.

As I indicate above, a variant is an adaptation for an activity A if and only if, by doing A, it makes a greater contribution to the fitness of its bearers than do other variants; so statement 1.3 amounts to the claim that organisms that bear a variant that confers greater adaptedness on its bearers will survive in greater proportion than those that bear a variant that confers lesser adaptedness on its bearers, *ceteris paribus.*[6] Furthermore, statement 1.3 states that the reason for the difference in rates of survivorship among organisms with alternative variants of the trait in question is that V_1 is better able to carry out A than its competitors.

The attitude that Hempelian evolutionists take toward statement 1.3 is as follows: they believe that statement 1.3 is the *only* lawlike statement that can be made about evolution. The idea is that statement 1.3 or similar statements of the principle of natural selection establish a lawlike relationship between organisms' bearing traits that make them better adapted, and their rates of survivorship; and there is no other evolutionary process that may be described in a similar manner, that is, in terms of a relationship of natural necessity between some one property of organisms (or other biological entities) and the evolutionary fate of variants, genes, or any other biological entity.

Having made these points of clarification, I am now in a position to state the

---

[6]The *ceteris paribus* clause protects E from spurious correlations: suppose that organisms bearing V_1 also happen to have a variant of another trait that detracts from their fitness. The idea is that, discounting such effects, V_1 makes a greater contribution to fitness than its alternative, even though it is masked by a correlation with another trait; in the long run, such effects can be expected to "cancel out."

exclusivity thesis, and to explain how Hempelianism motivates it. The thesis is as follows:

Analysis 1.4 (The exclusivity thesis) *Statement S expresses a proposition explaining an evolutionary event if and only if (1) S describes the process of natural selection responsible for adaptation A's having spread in P to the extent that it did between the time of its initial appearance and time T, and (2) S answers the question, "Why, between the time of its initial appearance in population P and time T, did adaptation A spread to the extent that it did, in P?"*

(The thesis need not be stated in terms of propositions, and I invite those with strong opinions on such matters to restate it as they please.)

The exclusivity thesis is motivated by a combination of Hempelianism and the attitude toward statement 1.3 that I attribute to the Hempelian evolutionists above—the belief that it is the only correct lawlike statement about evolution. As I emphasized previously (pages 19 – 20), Hempelians believe that scientific explanations of particular events cite at least one law of nature. If someone believes that the principle of natural selection is the only relevant law of nature, then—on the Hempelian principle I have just enunciated—he or she must also believe that the explanation of an evolutionary event must make reference to that principle. This is reflected in the first provision of the thesis. Additionally, Hempelians believe that the canonical form of an explanation-seeking question in science about a particular event E is, "Why did E occur?" This is reflected in the second provision of the thesis.

To conclude my account of the Hempelian evolutionists, I would like to provide a more precise statement of their position. I will use "Hempelian evolutionist" to indicate

26

someone who affirms the exclusivity thesis because he or she is also a Hempelian.

Analysis 1.5 *A person is a Hempelian evolutionist if and only if he or she affirms both the exclusivity thesis and Hempelianism, and affirms the former because he or she affirms the latter.*

1.2.2 Narrative explanations of evolution

In this section, I articulate the main principles of my opposition to the Hempelian evolutionists. I proceed in two further subsections. In the first, I describe, in outline, the nature of process explanation and the contextualism that it presupposes. In the second, I outline my argument that process explanation is used to explain evolutionary events that occur by random drift.

Process explanation and contextualism

Process explanations fail to meet the two criteria for good explanations that Hempelians claim are essential: they need not contain laws, and they do not aim to answer the question, "Why did E, the event to be explained, occur?" Rather, these explanations aim at answering the question "How did E occur," and they do so in a narrative fashion, describing a sequence of events that is causally relevant to the occurrence of the event to be explained. To see how this works, let me provide an example of a process explanation, adapted from Jeffrey [75, 25].

Consider a process explanation of the following statement, which concerns the sex determination of a human baby.

Statement 1.4 *Mr. and Mrs. Z.'s baby, Zoe, is a girl.*

A process explanation of the event described by statement 1.4 goes as follows.

Explanation 1.2 *During Zoe's conception, a gamete of Mr. Z.'s that carried an X chromosome fertilized a gamete of Mrs. Z.'s, creating a zygote with the "XX" genotype. As a result of a process of normal development—culminating in the birth of a child—this zygote matured into the baby known as "Zoe," who was born with the normal characteristics of a human female.*

Explanation 1.2 exemplifies the important characteristics of process explanation. First, it does not make explicit mention of any laws. Note that explanation 1.2 may presuppose laws, or that it may be possible to supplement it by reference to laws; however, it is explanatory notwithstanding their failure to appear in it. Second, it contains no claims to the effect that it ought to have been expected in any degree that the child of Mr. and Mrs. Z. would be a girl. Supposing that Mr. and Mrs. Z. are normal *Homo sapiens*—given only the information that they would have sex—it would be incorrect, in fact, to have any expectation about the sex of the child. As is well known, there is a chance of approximately 50% that a normal such couple produce a girl, and a chance of approximately 50% that such a couple produce a boy.

Accordingly, explanation 1.2 does not provide any insight into *why* Mr. and Mrs. Z.'s child turned out to be a girl. Indeed, perhaps there is no reason at all why the child is a girl. Rather than provide a reason why Mr. and Mrs. Z.'s child ended up female, explanation 1.2 provides an answer to the question, "How did that event occur?" The means

by which explanation 1.2 does this—the means by which which it bears explanatory weight—is by describing the sequence of steps in the causation of Zoe's sex: a paternal gamete bearing an X chromosome fertilized a maternal gamete also bearing an X chromosome, and by the normal processes of development, a female baby was born.

Explanation 1.2 also exemplifies the contextualism about explanation that I believe is required by process explanation. As I indicate above (page 19), Hempelians believe that (1) there is a single canonical kind of explanation-seeking question ("Why did E occur?") and (2) a single canonical kind of answer to that kind of question (citing laws to show that the event in question was to be expected, or expected in some degree). In contrast, I believe that explanation is a contextual matter because audience demands determine whether an explanation is a good one. This puts me at odds with Hempelians concerning both of the reasons that I have just suggested, by (1) and (2) above, that they have for believing that audience demands are irrelevant. Let me elaborate.

Explanation-seeking questions vary across contexts. "Where were the effects of the atomic bombs dropped on Hiroshima and Nagasaki most severe?" "What are the the effects of ionizing radiation on organisms?" "What is the function of the heart?" "When has a patient made a full recovery from bypass surgery?" On encountering a particularly puzzling or confusing state of affairs, it is common to ask simply, "What happened?" Requests for directions or instructions constitute a further type of explanation-seeking question: "How do we get to the museum?" Although such questions can simply constitute requests for information, they often demand the depth required of explanation, which aims to bring about understanding. None of these questions can be reformulated into the kind of why-

question that Hempel believes to be canonical.

The content required of responses to explanation-seeking questions is also determined by the audience's intentions and cognitive states. Consider statement 1.4, "Mr. and Mrs. Z.'s baby, Zoe, is a girl." The information that an embryologist would require for a process explanation of the event described by statement 1.4 would look different from the information that the child herself would require at age six. The embryologist would be interested, most likely, in how the Z.'s child developed into a girl, as opposed to a boy: given that it is usually the case that neither outcome is to be expected, what mechanisms operated to decide the case in this instance? This would require more detail about the process of sex than is provided by explanation 1.2. In contrast, explanation 1.2 is clearly inappropriate for the six year old Zoe, who may need to be told, for instance, that neither her mother's steady diet of sugar and spice nor the actions of a stork made the difference.

Process explanation and drift

The position I take up against the Hempelian evolutionists is as follows. Universalism is false, but contextualism is not; and there is no reason to think that explanations necessarily contain laws, or that they aim at explaining why the event to be explained occurred. On the contrary, there is ample reason to believe that the widespread use of the strategy of process explanation in science is well warranted; and evolutionary biologists use just such a strategy to explain evolutionary events occurring by random drift.

There are five such kinds of events that I believe to be of particular importance. In the past, these kinds of evolutionary events have been topics of the first importance across a wide range of evolutionary disciplines; in the present, each continues to be the

focus of active research, and each remains fundamental to evolutionary biology. I divide these events into two classes, which I designate using "N" to indicate population size, the notation commonly used for it by population geneticists. The two classes are as follows: events whose process explanation refers to N, and events whose process explanation does not refer to N.

I draw this distinction for the following reason. Drift in small populations generates a striking pattern, causing evolution of great magnitude in rapidly shifting directions. Accordingly, population size N plays a central role in the process explanation of evolution due to drift in such populations. In contrast, no such distinctive pattern is characteristic of drift in large populations, and N does not have any special relevance to process explanation of the evolution of them.

The events occurring by drift that are explained by process explanation in which N need not be mentioned are as follows.

The chance elimination of a rare but favorable allele A rare, favorable allele can be lost by drift, despite its advantage in fitness. For instance, suppose that there is only a single copy of the favorable mutation. There is a good chance that its bearer will die for reasons unconnected with its purpose, or that it will not be passed on during Mendelian reproduction.

Molecular evolution In the late 1960's new findings suggested that the evolution of DNA and its immediate protein products is mostly due to mutation and drift, rather than natural selection. This view still has currency with many, and in the debates around it, the burden of proof rests with "selectionists" rather than "neutralists."

Events occurring by drift that are explained by process explanations in which N plays a central role are as follows.

The shifting balance process Sewall Wright is famous for what he termed the "shifting balance theory," which he believed described the conditions under which adaptive evolution is most probable. The shifting balance process occurs in a large population that is subdivided into smaller local populations, between which there is a small but steady flow of migrants. Drift plays an important role in the local populations. The theory continues to be debated, having a good number of vocal supporters as well as detractors, and with many of its key ideas spreading to related areas of study.

The origin of species According to the "biological species concept," a group of organisms is a species if and only if it is genetically isolated from others, that is, cannot interbreed with them. Speciation, or branching, is the process by which a population that will go on to form a new species attains genetic isolation. According to a number of theories intended to apply to the biological species concept, drift plays an important role in bringing about speciation.

Punctuated equilibria and the shape of phylogeny On the theory known as "punctuated equilibrium," daughter species branch away from their parent species rapidly in geological time. Furthermore, on the punctuational view, once a species is formed, it does not exhibit a great deal of morphological change over its lifetime. Drift has been cited to explain the rapid branching of daughter species from their parents that generates the punctuational pattern.

1.3 Chapter Summary and Overview of Chapters

I conclude the chapter with a look back at some of the main claims that I have made in it (section 1.3.1) and a look ahead at what I intend to accomplish in each of the remaining chapters of the dissertation (section 1.3.2).

1.3.1 Chapter summary

As I stated in the opening passages of this chapter, the central claim for which I argue in this dissertation is that there are important phenomena that occur by random drift that are explained by evolutionary biologists using a strategy of explanation I term "process explanation." This claim is at odds with the views of evolutionists and philosophers that hold what I term the "exclusivity thesis," which is the view that only evolution occurring by natural selection can be explained. Our disagreement rests on more than just whether the science is adequate to support one view or the other: it also rests on a methodological dispute.

The methodological dispute concerns the nature and justification of process explanation. My view is that certain narratives that carry explanatory weight—process explanations—do not contain laws, and have the aim of showing *how* (rather than why) the event in question occurred. Those who affirm the exclusivity thesis do so because they believe that such narratives do not carry any explanatory weight. This follows from their Hempelianism, according to which explaining a particular event E requires citing laws of nature that show why E occurred.

The issue of whether process explanations are adequate turns on a further issue

about the nature of explanation. While Hempelians are what I term "universalists" about explanation, I am a "contextualist." The universalist view is that the adequacy of an explanation depends only upon whether it meets criteria that are invariant across contexts of utterance, the latter being irrelevant to whether the explanation is a good one. Contextualism is the view that the context of an explanation's utterance is relevant to its evaluation: what makes an explanation a good one is that it meets the requirements of the audience.

In the context of the issue of whether the exclusivity thesis is correct, these debates about process explanation, contextualism, and universalism play out as follows. Proponents of the exclusivity thesis affirm it because they view natural selection as the only lawlike process of evolution, and—informed by their Hempelianism—view the principle of natural selection as the only principle of evolution that has any explanatory power. I take a different view, because I reject Hempelianism and the universalism that it entails. Process explanations function without laws, explaining how (rather than *why*) the event to be explained occurred by describing the sequence of events that caused it. It is just this strategy of explanation that evolutionary biologists use to explain phenomena that occur by random drift.

1.3.2   Overview of chapters

I conclude this introductory chapter by indicating how each subsequent chapter will contribute to my argument.

Chapter two I describe Hempelianism, explaining the sense in which it is a universalist view, and I argue that the view is held by Hempel, Salmon, and Railton, and I

34

articulate contextualism about explanation and process explanation.

Chapter three I consider Hempelian objections to process explanation and contextualism, and I respond to those objections.

Chapter four I state a probabilistic theory of a mechanism of random drift known as "indiscriminate sampling."

Chapter five I argue that Daniel Dennett, Richard Dawkins, and participants in what I term the "theory of forces dispute" are Hempelian evolutionists. Referring to the five kinds of phenomena I describe above (pages 30 – 32), I make the main argument of the dissertation: evolutionary biologists use process explanation to explain important phenomena resulting from drift. I use the account of indiscriminate sampling I develop in chapter 4 to describe each phenomenon, and to describe the explanation of each by process explanation.

Chapter six In this concluding chapter, I review the argument of the dissertation.

Chapter 2

The Nature of Process Explanation

My aim in this chapter is to elucidate the strategy of explanation that I term "process explanation." Broadly speaking, process explanations generate understanding of an event to be explained E by providing a narrative of events, terminating with E, that are causally relevant to E's occurrence. My view that there are such explanations places me squarely within the camp of a group of philosophers that may be called "contextualists" about scientific explanation. I distinguish contextualists from what may be termed "universalists" about scientific explanation:[1]

Universalism The conditions for evaluating scientific explanations are invariant across

contexts—there is one and only one set of criteria for evaluating scientific explanations, and those criteria apply regardless of the intentions or cognitive states of the audience of the explanation, or of its producers. Prominent universalists include Carl Hempel, Wesley Salmon, Peter Railton, Philip Kitcher, and Michael Friedman.

---

[1] I have adapted "contextualist" and "universalist" from Achinstein [1, 119].

Contextualism The criteria to be used for evaluating explanations differ from one context of utterance to the next; there is no single set of conditions that a successful explanation meets. Prominent contextualists include Peter Achinstein, Micheal Scriven, William Dray, Sylvain Bromberger, and, on some interpretations, Bas van Fraassen.

I take issue with one particularly important group of universalists, whose view I term "Hempelianism." Although the view is widely held among philosophers, I have chosen to name it "Hempelianism" because Carl Hempel is its best known, most forceful, and most consistent advocate, as well as being one of its earliest. The central tenet of Hempelianism is the well known "covering law" requirement embodied in Hempel's famous models of explanation: roughly, Hempelianism is the view that scientists explain particular events by subsuming them under laws of nature.

Different Hempelians incorporate this requirement into conditions they believe explanations must meet, supplementing it with others. Nevertheless, all Hempelians are universalists because they each believe that there is some set of conditions—conditions incorporating subsumption under law—that, across any and all contexts of utterance, provide criteria for a good scientific explanation of a particular event. My proposal that there are process explanations is incompatible with the Hempelians' universalism. This is because I claim that there are some contexts in which it is inappropriate to explain an event by subsuming it under laws of nature: in these contexts, requests for explanation require describing a sequence of events that are causally relevant to the event to be explained E. The descriptions of such sequences do not require mention of laws.

My argument against Hempelianism plays a central role in the overall project of this dissertation, that is, in my argument that random genetic drift plays an important explanatory role in evolutionary biology. The significance of my argument against Hempelianism for my claims about evolutionary biology is that it opens the way for the argument that random drift possesses considerable explanatory power. This is because Hempelianism is not restricted to the literature on the nature of explanation. There are prominent evolutionists whose Hempelianism drives them to the view that random drift cannot explain evolution. As well, like my argument against Hempelianism, my elaboration and defense of process explanation plays a central role in the overall project of this dissertation. My view is that drift-explanations proceed by the strategy of process explanation. I postpone my more detailed account of these issues until chapter 5.

This chapter has three main sections. In section 2.1, I describe Hempelianism, and I argue that it has wide influence among philosophers, including Hempel's critics. In section 2.2, I describe the alternative to Hempelianism that I propose and that figures so prominently in the central argument of the dissertation: process explanation. I offer a brief concluding summary and overview of the chapter in section 2.3.

2.1   Hempelianism and the Hempelians

My aim in this section is to describe the universalist view that I term "Hempelianism" and to argue that some important philosophers hold the view. I proceed in three sections. In section 2.1.1, I state the Hempelian position, and indicate the origin of it in Hempel's own works. In sections 2.1.2 and 2.1.3, respectively, I argue that Wesley Salmon

and Peter Railton, two prominent philosophers of science, are Hempelians.[2]

2.1.1 Hempelianism

In this section, I want to state and explain Hempelianism, and also, to indicate its origins in the works of Hempel himself. I begin with Hempelianism, which is a two-part claim about scientific explanations: it is a claim about the kinds of explanation-seeking questions that scientists ask, together with a claim about what is required to answer those questions.

Statement 2.1 (Hempelianism) *Any scientific explanation of a particular event E answers the question "Why did event E occur?" by subsuming E under laws of nature.*

Regarding the idea that all scientific explanations answer the question "Why did event E occur?" consider the following. The Hempelian does not think that the context of utterance of an explanation-seeking question makes any difference to what that question implies about how the explanation offered in response ought to be evaluated; so the Hempelian imposes a uniform interpretation on why-questions, taking them to provide the canonical form for explanation-seeking questions in science. This interpretation is as follows: to ask why a particular event occurs is to ask for reasons that the event's occurrence should have been expected, or that someone ought to have had a certain degree of expectation regarding

---

[2]Although I will only discuss Hempel, Salmon, and Railton in depth in this chapter, I believe that there are many more Hempelians. Popper [116] is one notable Hempelian. I would also include Fetzer [46] and Humphreys [72] among the Hempelians, as well as Hull [70]; and I would argue that, while not a perfect fit with Hempelianism, the so-called "unification" approach to explanation advocated by Friedman [51] and Kitcher [82] exemplifies central elements of it, as does the mechanism-based approach favored by Glennan, ([59] and [58]), Machamer and his co-authors [91], and Craver [26]. On both the unification and mechanism approaches to explanation, scientific explanations of particular events are understood to require the use of generalizations, which, as I have mentioned above in several places already, is a central tenet of Hempelianism.

its occurrence. As will be seen, different Hempelians understand what is required to meet this requirement in slightly different ways, although all emphasize the notions of expectation and necessitation.

I do not think that Hempelianism commits one to any particular view about laws of nature. As well, I have used the term "subsumption" to indicate a broad set of relationships between laws of nature and events to be explained. As will be seen, Hempel believes that, in scientific explanations, a statement describing an event to be explained is the conclusion of an argument whose premises consist of a law and the description of facts obtaining prior to the occurrence of the event to be explained. Salmon, in contrast, does not believe that explanations are arguments, and, accordingly, has a different account of the relationship between laws and the event to be explained. Nevertheless, all Hempelians share the belief that laws are essential for explanation.

The Hempelian attitude toward pragmatic matters is exemplified by Hempel's [64, 425 - 428] view, which I mentioned in chapter 1 (pages 22 - 23), that explanations can be likened to mathematical proofs. Some are easier to understand than others; some are more elegant than others; some are more significant than others. These are pragmatic matters that depend on the interests and cognitive states of the audience of a proof. Nonetheless, whether the theorem in question does in fact follow from the premises advanced in the proof is a matter of logic; this does not depend on the beliefs or interests of any particular person. Explanations, similarly, can be judged according to a single set of criteria that apply across contexts.

Turning to Hempel's own views, I now want to argue that they conform to the first of the two claims of Hempelianism, that is, the claim that a scientific explanation of an event E answers the question "Why did E occur?" This is not a difficult argument to make, because Hempel states it explicitly, as well as providing examples that clearly suggest that this is his view. To see this, consider the following.

He states that explanation-seeking questions in science can "be expressed in the form 'Why is it the case that p?', where the place of 'p' is occupied by an empirical statement specifying the ... [phenomenon to be explained]" [64, 334]. Referring to one of his models of explanation, the "deductive-nomological" ("D-N") model, Hempel claims that "a D-N explanation answers the question '*Why* did the ... [phenomenon to be explained] occur'," adding that a D-N explanation "enables us to *understand why* the phenomenon [to be explained] occurred" [64, 337]. Hempel's many examples of explanation-seeking questions answered by D-N explanations include "Why did Hitler go to war against Russia?" [64, 334] and "Why did the television apparatus on Ranger VI fail?" [64, 334].

Hempel does not explicitly return to this point in his discussion of his inductive models of explanation, the "inductive-statistical" ("I-S") model, in particular. However, I take it that his remarks about the kinds of questions that D-N explanations are intended to answer also apply to I-S explanations. As an example of an explanation-seeking question to be answered by an I-S explanation, Hempel cites "Why [did] patient John Jones ... [recover] from a streptococcus infection?" [64, 381]; many similar examples may also be easily found ([68, 245], [66, 232 – 3], and [67, 298; 304 – 5]).

As in the case of the claim that scientific explanations answer why-questions, Hempel is quite clear in his explicit remarks that he believes that laws are required for explanation. For instance, regarding D-N explanations, his remark that "reliance on general laws is essential" [64, 337] is typical, and similar remarks are readily found in many places in Hempel's work ([66, 231], [68, 246], and [67, 298 - 303]).

The claim that laws are required for explanation is also reflected in the D-N and I-S models. According to Hempel [64, sec. 2], a D-N explanation of why a particular event occurred is a deductive argument that consists of three kinds of statements.

1. One of the premises of the argument contains one or more natural laws, that is, lawlike statements of the form "All F are G."

2. The other premise of the argument is a statement of conditions obtaining at a time before the event to be explained occurred.

3. The conclusion of the argument describes the event to be explained.

The idea is just that, when scientists explain why an event occurred, they do so by deducing a statement describing the event from a law or set of laws and a statement describing conditions that obtained prior to the event's occurrence. For instance, suppose that a member of an artillery regiment wanted to answer the explanation-seeking question, "Why did the mortar land in location L?" He or she could give a simplified D-N explanation of this by stating the relevant laws of mechanics, describing the initial impulse given to the mortar and the angle at which it was launched, and deducing from these statements a statement describing the mortar's landing in location L. According to Hempel, this explains the event.

Hempel's I-S model works much in the same way that his D-N model does [64, sec. 3.3]. As in the D-N model, Hempel understands I-S explanations to be arguments containing one or more laws as premises. Despite this, unlike in the D-N model, Hempel believes that these laws are probabilistic, having the form "All F have a probability p of having G." Hempel believes that the probability p must be high, in particular, that it must provide what Hempel terms "practical certainty" that an individual that is F is also G. This is an important point, because many other Hempelians, including Salmon and Railton, whose views I will discuss below, do not believe this.

The other premise is a statement to the effect that some individual f has the property F. This is analogous to the premise in the D-N argument in which conditions obtaining before the event to be explained are described. The conclusion of the argument describes the phenomenon to be explained, that is, individual f's having property G. Hempel believes that these arguments explain why the individual f has property G, this being construed as a particular event occurring at a particular time.

As I suggest above, Hempel believes that the explanatory force of the scientific explanations of particular events stems from their capacity to provide a reason that the event in question was to be expected. This is embodied in his famous "symmetry thesis" [64, 366]. Hempel believes that, if a statement S is an explanation of an event to be explained E, then it could have served, before the fact, as a prediction of E; and he also believes that, if S can be used to predict E, then it can explain E after the fact. This establishes the identity of explanation and rational expectability, because predicting an event requires showing that it ought to be expected; and this, as I have been describing, also amounts to explaining it,

on Hempel's view.

It should be clear how this works in the case of the D-N model. If the occurrence of an event can be deduced, then, clearly, it ought to be expected. The requirement of expectability is met in the case of I-S explanations by Hempel's high-probability requirement. As I mention above, Hempel believes that an I-S explanation must show that the event to be explained has a high probability of occurring, that is, that it be "practically certain" that it occur: such events ought to be expected, as opposed to events with a low probability of occurring, which should not be.

2.1.2 Salmon

What I want to do in this section is argue that another prominent philosopher of science, Wesley Salmon, is a Hempelian—accepting the basic Hempelian notion that explaining particular events requires laws, and also, Hempel's universalism. To begin with, let me consider some explicit remarks of Salmon's from *Statistical Explanation* [126] that indicate that he is committed to Hempelianism about laws.[3] Salmon claims that the crash of a small plane "is explained by invoking *general laws* ... under which the ... [plane crash] can be subsumed" [126, 4], and he emphasizes that, for any explanation, some of "the explanatory facts ... thus consist of ... general facts—embodied in ... general laws" [126, 4]. As well, he states that "it is evident that explanations ... are nomological.... [E]very explanation must contain at least one such [statistical] generalization.... [A]n explanation

---

[3]I am aware that Salmon has updated his analysis of explanation [128]; however, I believe that he views his later theory as a supplement to his earlier theory, intended to contribute to an account of causal explanation, and that he remains committed to many of his earlier views on statistical relevance, which reflect a commitment to Hempelianism. As his earlier views are simpler, I think it is more appropriate to consider them here, as opposed to his later views.

essentially consists of a set of statistical generalizations..." [126, 78].[4]

Regarding explanation-seeking questions, he states (referring to the plane crash, mentioned above) that "we ask why the crash occurred" [126, 3]. Other examples—explaining why a sample of salt dissolves in water [126, 33], or why Times Square has no tigers in it [126, 34]—substantiate the view that Salmon is a Hempelian about explanation-seeking questions. I think that examples concerning low-probability events are particularly telling in this regard.[5] Let me elaborate.

On Hempel's I-S model of explanation, improbable events cannot be explained; Hempel does not believe that there are any reasons why such events occur. Salmon disagrees, as indicated by his "statistical relevance" ("S-R") model of explanation. This model informs a view about explaining why an event occurs that is incompatible with Hempel's: Salmon believes that to explain why an event E occurs, it is enough to indicate the *degree* to which E was to be expected, rather than *that* E was to be expected [126, 78]. The point I want to make here is that, by engaging with Hempel on this issue, Salmon shows that he sees the philosophical problem in the same way that Hempel does—to account for how scientists answer "Why did E, the event to be explained, occur?"

Having cited some remarks that clearly point to Salmon's Hempelianism, I would now like to extend my argument by considering how his Hempelianism is exemplified by the S-R model. To begin with, consider the kinds of questions answered by S-R explanations, which have the form "Why does individual x, that has the property A, also have the property B?" Suppose that an individual's having (or coming to have) a property is understood as

---

[4]I have heavily edited these quotations to improve their readability.
[5]He provides a quantum-mechanical example [126, 9], and treats the low-probability issue in depth [126, secs. 10 and 11].

a kind of event. Then, the question "Why does individual x, that has the property A, also have the property B?" is indeed the kind of explanation-seeking question that Hempelians claim are typical of scientific explanations of particular events. Now, in order to see how laws figure into answering these kinds of questions on the S-R account, I need to provide a brief account of Salmon's [126, secs. 3 and 4] views on probability. I begin by introducing some terminology.

Let B represent what is termed the "attribute class," and let A represent what is termed the "reference class." According to Salmon, who is a frequentist about probability, the probability of an event of kind B in a reference class A is the relative frequency of events of kind B that are also events of kind A, in an infinitely long sequence of events of kind A. I will refer to these infinitely long sequences of events of one kind (i.e., kind A, here) as "long-run sequences." I will not consider any issues concerning the nature of long-run sequences, intriguing as they are; Salmon, for his part, thinks that ascriptions of probability entail that it is true that such sequences exist.

The frequentist account of probability only applies to kinds of events, as I have indicated in my description of it, above; so, there is a question about how it applies to single cases. The problem is to choose an appropriate reference class, out of the infinity of such classes that a given single event might fall into. For instance, suppose that I am tossing a coin C. What is the probability that C lands heads on the next toss? Is it the relative frequency of the heads in every toss of the coin, past and future? Suppose I am tossing the coin on a Wednesday. Is the appropriate reference class the class of Wednesday tosses? Each of these descriptions might apply equally well to coin C, and some guidance is needed

to determine which is appropriate for determining the probability of the coin's coming up heads the next time I toss it.[6]

Salmon's [126, 43] proposal is that the appropriate reference class for determining the probability of any given particular event is the widest possible *homogeneous reference class* into which that event can be placed. Salmon understands this notion as follows.

Analysis 2.1 (Homogeneity) *A reference class A is homogeneous with regard to an attribute class B if and only if there is no partition P of A of which the following obtains:*

$$p(B|A\&P) \neq p(B|A).$$

The idea is just that a reference class is homogeneous if and only if there is no further subdivision of that class in which the probability of an individual's having B differs from what it is in the class as a whole.

Statistical relevance, which Salmon [126, 42] understands as follows, is the other side of the coin, so to speak, from homogeneity.

Analysis 2.2 (Statistical relevance) *A partition P of a reference class A is a statistically relevant partition with regard to attribute class B if and only if the following obtains:*

$$p(B|A\&P) \neq p(B|A).$$

Here, the idea is that a reference class has a statistically relevant partition if and only if there is some further subdivision of the reference class in which the probability of an individual in that subdivision being in the attribute class B differs relative to the reference class as a whole.

---

[6]Strictly speaking, "weights," not probabilities, apply to single cases. I will use "probabilities" in the interest of economy; no confusion should result, and this accords with Salmon's usage.

For instance, suppose that the long-run relative frequency of heads among Wednesday tosses of coin C is 99%, while the long-run relative frequency of heads among all tosses of the coin is 50%. The class of all tosses is not an appropriate reference class, on Salmon's account, because it is not homogeneous: it is possible to isolate certain tosses, the Wednesday tosses, whose probability of heads differs systematically from the rest of the tosses. Accordingly, whether it is Wednesday is statistically relevant to the result of the tosses of coin C. Suppose that T_C means "Coin C is tossed;" C_H means "Coin C lands heads;" and W means "It is a Wednesday."[7] Then, the relationship between Wednesdays and the results of tossing the coin may be represented as follows:

$$p(C_H|T_C \& W) \neq p(C_H|T_C).$$

At this point, I am in a position to indicate the role of laws in Salmon's S-R analysis, by way of stating that analysis [126, sec. 13]. S-R explanations contain laws in the form of probability statements, which Salmon understands as statistical generalizations. This can be seen by considering what an S-R explanation that answers the question "Why does individual x have the property B?" looks like. First, there are two statements:

1. $x \in A \& C_i$; and

2. The C_i's partition A.

Next, forming the core of the explanation, there are the following probability

---

[7] These statements may be construed to refer to either types or tokens. If they are interpreted to be tokens, the probability statements I use below must be understood as "weights."

statements, i.e, statistical laws:

$$p(B|A\&C_1) = p_1$$

$$p(B|A\&C_2) = p_2$$

$$\vdots$$

$$p(B|A\&C_n) = p_n,$$

so that $A\&C_1, A\&C_2, \ldots, A\&C_n$ are homogeneous with respect to B, and $p_i = p_j$ only if $i = j$.

The idea is as follows. Explaining an event E requires indicating all factors statistically relevant to E's occurrence, and the degree to which each is statistically relevant to E's occurrence—this is the role of the probabilistic laws. Additionally, explaining E requires indicating which of those factors was at work in the case of E—this is the role of the statement that the individual x is an element of some one of the classes $A\&C_1, A\&C_2, \ldots, A\&C_n$.

Salmon's [126, 78] view is that this explains an event because it provides enough information so that, whatever happens, the degree of rational expectation that one should have regarding the event's occurrence can be determined. For instance, to explain why coin C landed heads the last time it was tossed, someone would state that the last time it was tossed, it was a Wednesday, and indicate the probability of heads on Wednesdays; and state the probability of heads that would obtain if the coin possessed another statistically relevant property, providing a comprehensive list of such properties and the corresponding probabilities. Certainly, with this information, and with adequate information about the coin, someone could determine the amount that it would be reasonable to bet on heads across the range of possible circumstances.

Salmon's view that particular events are explained by statistical relevance relation-

ships is analogous to Hempel's view that explanations of particular events work by showing that those events ought to have been expected. Salmon's account is that, rather than show *that* a particular event to be explained ought to have been expected, explanations function by showing the *degree* of expectation that one ought to have had concerning their occurrence. This, in turn, reflects Salmon's commitment to universalism, which he shares with Hempel: statistical relevance relationships are always explanatory, on Salmon's account, with regard to particular events—regardless of the beliefs or interests of the audience.

In conclusion to my discussion of Salmon, I want to note that his analysis of explanation differs from Hempel's on two key points.

1. Salmon [126, 11] does not think that explanations are arguments. Rather, as I indicated above, they are lists of probability statements, with the addendum that the individual in question has a certain property, and that certain classes partition the main class of interest.

2. Salmon [126, sec. 10] does not believe that high probability is required for explanation. All that is required is to indicate which statistically relevant partition of a broader reference class that the individual in question is a member of, regardless of the *degree* of relevance.

This is important because it indicates that Hempelianism is a broad view, representing a fundamental set of commitments about the goals and strategies of explanation: those who affirm it also affirm views about the nature of explanation that are strictly at odds with one another. Likewise, it shows that the universalist view is quite liberal, encompassing accounts of explanation that differ on fundamental points.

2.1.3 Railton

The third and final Hempelian philosopher that I will consider is Peter Railton. Like Hempel and Salmon, he is both a Hempelian and universalist. Railton formulates a model of explanation that he calls the "deductive-nomological-probabilistic" model, or, for short, the "D-N-P" model of explanation. Let me begin my account of Railton's Hempelianism by arguing that Railton, like Hempel and Salmon, believes that scientific explanations aim at answering the question, "Why did E, the event to be explained, occur?"

Railton [119] does not make any explicit remarks to this effect, as Hempel and Salmon do. My view is that this is because Railton takes it to be such a fundamental part of the background to the discussion of explanation that it is not in need of articulation, let alone defense. Furthermore, Railton seems to believe that the main problem facing philosophers working on explanation is to rehabilitate Hempel's models of explanation in response to criticisms of it; the framework of ideas concerning explanation developed by Hempel is essentially correct, even if it needs some minor adjustment. If this is so, the kinds of explanations at issue are those that answer the question, "Why did E, the event to be explained, occur?" As I have just argued above, these are just the kinds of explanations that Hempel concerns himself with.

Railton's examples bear this out, particularly his quantum-mechanical example [119, 124 - 128]. In this example, the phenomenon to be explained is the emission of a radioactive particle during a certain period of time. If Hempel's I-S account is accepted, this event cannot be explained. As I remarked above in my discussion of the I-S model, Hempel believes that explaining such an event requires an inductive argument to the effect

that there is a practical certainty that the event to be explained occur. The decay of a radioactive nucleus is very improbable, and so no such argument is possible. So, I take it that, by using this example, Railton believes that he can beat Hempel at his own game, that is, showing that there is a reason why certain events occur.[8]

Also telling is Railton's [119, 123] denial of Jeffrey's claim that there is no reason why improbable events occur: Railton believes that there is, and that his D-N-P model provides the strategy for such explanations. This accords with the view that Railton believes that the aim of scientific explanations of particular events is to answer explanation-seeking questions about why those events occurred.

Now what I would like to argue is that Railton adheres to the second claim embodied in Hempelianism, that is, the claim that laws are necessary for explanation. As well as his explicit statements to the effect that he believes that explanation requires laws [119, 119, 121, 124], Railton's D-N-P model exemplifies this belief. He believes that these explanations contain four elements [119, 127 − 128]. In my description of the D-N-P model, I will suppose that the question that is answered by a D-N-P argument is, "Why does the object f, that is of kind F, have the property G?" The four elements of D-N-P explanations are as follows.

1. A law, "All F have a probability p of having G," derived from a fundamental theory.

2. A statement that f is of kind F.

3. A statement that there is a probability p that f will have property G. This statement is the conclusion of an argument that has the law in (1) and the statement in (2) as

---

[8]Other examples of low-probability events include a medical example [119, 121] and an unlikely result in roulette [119, 123].

premises. Note that, because Railton does not require high probability, the probability p may take any value.

4. A parenthetical addendum to the above statements that indicates whether, in fact, f has property G.

To see how this works, let me provide a sketch of an example offered by Railton in his initial exposition of the D-N-P model [119, 124 – 127]. Suppose that scientists are in possession of a sample of some radioactive element R; call the sample, "S_R." Suppose that S_R is observed for some period of time T, during which it emits a radioactive particle. The explanation-seeking question to be answered is, "Why did S_R emit a radioactive particle during T?" On Railton's account, the answer to this question is as follows.

Corresponding to (1) above, the explanation must include a probabilistic law, derived from some fundamental theory. In this case, the law—derived from atomic theory— would be something like the following: during a given length of time T, there is a probability p that a sample of radioactive element R will emit a radioactive particle. Second, there is the statement that the sample S_R is a sample of the radioactive element R in question. This is an instance of (2) above.

Third, there is an inference to the effect that, because S_R is a sample of R, it has a probability p of emitting a particle during a time period T units of time long. This probability statement is represented as the conclusion of an argument with the probabilistic law and the statement that S_R is a sample of R as premises. This is point (3) above. Finally, the D-N-P explanation requires a parenthetical addendum to the argument just described that indicates that S_R did, in fact, emit a radioactive particle during the time period in

question. This is point (4).

Note that, as I have suggested already, Railton does not require high probability for explanation. In this, he differs from Hempel, who does require it, and agrees with Salmon, who does not. In the example I am considering here, this means that, even though the sample of radioactive material under consideration may have a negligible probability of decaying in the time frame at issue, the fact that it did so is still explained, in part, by reference to that probability. Despite this difference with Hempel, Railton conforms to the broader tenets of Hempelianism: laws are required for scientific explanations of particular occurrences, and what is explained by them is why the event in question occurred. This reflects his commitment to universalism: he believes that any set of statements meeting the D-N-P requirements is explanatory, regardless of who is using them and what his or her purposes or beliefs are.

2.2   The Nature of Process Explanation

Now that I have described some Hempelian views about explanation in science—views that are also universalist in nature—I would like to elaborate an alternative that I term "process explanation." Process explanations do not provide reasons that an event ought to have been expected, or to have been expected in some degree; and they do not require laws. Rather, they are explanatory *narratives*. Process explanation is founded on contextualism about scientific explanation, and is incompatible with Hempelianism for that reason. The idea is that there are some contexts that demand an explanatory narrative, rather than the kind of explanatory power that derives from laws of nature.

In the remainder of this section, which has no further subsections, I proceed as follows. First, I sharpen my account of the differences between contextualism and universalism in order to formulate the philosophical problem of describing process explanation, a contextualist view, in a precise manner. Second, I engage in that philosophical problem, describing the logic of process explanation. Third and finally, I provide examples of process explanations from the sciences.

I would now like to begin the first task of this section, mentioned above, by considering the question, "What is the difference between universalism and contextualism?" The answer to this question is as follows. The point on which contextualists and universalists differ regards the criteria for evaluating answers to explanation-seeking questions. For S to constitute a good answer to an explanation-seeking question, according to the contextualist, it must meet some set of conditions that are implied by the question, together with the context in which the question is asked. The context includes both the beliefs and the intentions of the audience for the explanation—that is, what the audience already knows and believes, and what its purposes are in requesting an explanation. In contrast to contextualists, universalists believe that explanations should be evaluated independently of their context of utterance.

Allowing for context in the assessment of explanations generates a broadly pluralistic view: as there are many types of contexts in which explanation-seeking questions are asked, there is a corresponding variety of criteria for the adequacy of explanations. The idea is that the conditions for answering one kind of explanation-seeking question differ from the conditions for answering another, and so there are as many ways to produce a

good explanation as there are ways to ask an explanation-seeking question.

These views represent a sharp contrast to those of the universalist, who allows that there is one and only one set of criteria for evaluating explanations. As I indicate above (page 40), Hempel—a paradigm universalist—believes that explanations are analogous to mathematical proofs; whether they succeed is person- and context-independent. The universalists' belief that explanation can be characterized in an abstract manner accounts for their practice of inventing models of explanation, which describe relationships among statements that must obtain if an explanation is to be successful.

Let me provide a brief account of the way that conditions for the adequacy of an explanation can differ across contexts. Consider the following Hempelian question.

Question 2.1 *Why did person P die, i.e., why would a rational person have expected P's death?*

Suppose, furthermore, that question 2.1 is posed in two different contexts.

1. Question 2.1 is asked of a pathologist responsible for *P*'s autopsy, by the chair of a hospital review board.

2. Question 2.1 is asked of *P*'s family physician, by *P*'s spouse.

There are differences in what the two proposed audiences—spouse and review board chair—would expect from those to whom question 2.1 is posed—pathologist and family doctor. The pathologist would be expected to provide detailed information about *P*'s admission to the hospital, the manner of *P*'s diagnosis, the initial treatments proposed and how they were administered; and an account of the disease process as it overcame

P. This information is appropriate because the cognitive abilities of the review board chair include extensive knowledge of hospital policies, organization, and general knowledge of medical diagnosis and disease processes, and because his or her interest is to discover whether the hospital can improve its diagnostic services and treatment facilities.

In contrast, *P*'s family physician would be expected to answer question 2.1 by providing the name of the disease and some basic information about it, for instance, whether it was viral or bacterial; whether *P* had possibly inherited a predisposition for the disease; how *P* might have contracted it; appending, perhaps, that *P*'s death could not have been helped, given how far along the disease was before it was discovered. Although the family physician may be able to answer the same kinds of questions that the pathologist can, the audience does not demand the same level of detail about the biology of the disease, and is not interested in looking at *P*'s death as an opportunity to improve the hospital.

This brings me to the point at which I can state the philosophical problem that the contextualist must solve, if he or she is to advance the view against the universalist. The general problem is to develop counterexamples to the Hempelian-universalists' claim that there is only one way to evaluate an explanation. The challenge is as follows: identify contexts in which requests for explanation in science are made that are not requests for laws, laws that, together with initial conditions, necessitate the event to be explained *E*, or that entail that a rational person would have had some degree or other of rational expectation that *E* occur.

Having sharpened my account of the philosophical challenge that faces someone such as myself who wants to advance a theory of explanation that entails contextualism,

I now to turn meeting that challenge, taking up the second main task of this section: describing the logic of process explanation. Consider the following. Both inside and outside of science, there are contexts in which people want to know the events in a causal chain leading up to an event to be explained E, but do not want to know whether E was necessary, or the degree of rational expectation that someone ought to have had concerning E's occurrence. These are the contexts that call for process explanations, and in order to provide such explanations—and so, to satisfy the questioner's requirements for a good answer to his or her explanation-seeking question—what is required is a narrative that describes a sequence of events, terminating in E, that are causes of E's occurrence.

In claiming that narratives can be explanatory, I take a view much like that elaborated by William Dray in *Laws and Explanation in History* [38]. In that work, Dray provides an excellent example of process explanation, which I would like to cite here. The example concerns the explanation of the seizing-up of a car's engine.

> If I am to understand the seizure, I shall need to be told something about the functioning of an auto engine. . . . I need to be told, for instance, that what makes the engine go is the movement of the piston in the cylinder; that if no oil arrives the piston will not move because the walls are dry; that the oil is normally brought to the cylinder by a certain pipe from the pump, and ultimately from the reservoir; that the leak, being on the underside of the reservoir, allowed the oil to run out, and that no oil therefore reached the cylinder in this case. I now know the explanation of the engine stoppage. [38, 67 – 68]

Dray's view, which reflects my own, is that "my understanding of the engine seizure is very directly related to the fact that I can now *trace the course of events by which it came about*. . . . I can now envisage a *continuous series of happenings* between the leak and the engine seizure" [38, 68]. This does not require laws, which would be particularly out of place in this case. For instance, it would not help in this case to know that "whenever oil

reservoirs have leaks, the engine sooner or later seizes up" [38, 67]. This is so, even if "such a generalization [were arrived at] by the most careful inductive procedure" and if "there may never have been a contrary case in the records of this garage, or any other one I examine: whenever reservoirs are leaky, engines may have seized up" [38, 67]. The car owner wants to know the sequence of events that led to the engine's failure, not what circumstances existed such that the engine had to fail, or such that it would have been reasonable to have a certain degree of belief that it would do so.

Dray gives particularly clear expression to the contextualists' assessment of the role of context in fixing what is required to explain the engine's seizure. The adequacy of the explanation of the engine seizure that I have just cited depends "on who says what to whom ... it depends on what else is presupposed, or contextually supplied" [38, 67]. For instance, suppose that the mechanic had simply declared, when asked by the car's owner to explain the engine's seizure, that "it's due to a leak in the oil reservoir" [38, 67]. For those lacking knowledge of auto engines, such as the owner of the car, this is not an explanation. However, "to the assistant mechanic standing near by ... it may very well be an explanation" [38, 67].

An audience's intentions as well as its beliefs can have significant consequences for the content of an explanation. One can imagine a context in which a further explanation of the engine failure directed toward the assistant would emphasize the failure of a certain part P of the lubrication system. The garage has been keeping track of the rates of failure of different parts of the lubrication system, so whether part P failed is of particular interest to the assistant mechanic, who just replaced a faulty part P the day before. To the owner

of the car, this information would be useless, and would not be mentioned.

I would now like to address a question about the general characterization of process explanations. Is there a canonical form for them? Although I do not think that there is one, I do think that there are a few forms that they commonly take. Explanation-seeking questions demanding process explanations often take the form "How did E occur?" or, implicitly referring to E, "What happened?"

For instance, suppose that police come upon a crime scene, and find someone who was present when the crime was committed. They might ask, "How did the victim die?" or simply, "What happened?" In response, the police would expect a description of the events that led up to the victim's death. I think it is clear that this is a request for an explanation; and I think that it is equally clear that what is being requested is not a description of laws that, together with conditions obtaining before the crime was committed, necessitated the victim's death. It is equally clear, I think, that the police are not asking for information that would indicate the degree of expectation that a rational person would have had regarding the crime's occurrence, before the fact.[9]

Indeed, I think that there are cases in which the Hempelians' paradigm request for explanation, "Why did E, the event to be explained, occur?" is clearly intended to request a process explanation. Consider Dray's engine example, above. I think that it would be natural for the owner of the car to ask the mechanic, "Why did the engine seize up?" In this case, there is no ambiguity about what the car owner is asking for—the series of events leading to the engine's failure, not reasons that it had to occur, or reasons for having a certain degree of rational expectation that it would do so. This request for explanation

---

[9]The crime scene example was suggested to me by Peter Achinstein (pers. comm.).

could also be equally well phrased "What happened to the engine?" or "How did the engine seize up?"

I want to conclude this exposition of process explanation by commenting on the claim that such explanations are causal explanations. As I have already stated, the explanatory narratives provided in response to the kind of requests for explanation that I have in mind cite causally relevant events leading up to the event E to be explained. In order to see how this works, recall the crime scene example. Suppose that the witness answered the police by stating that Orion is visible in the night sky over the crime scene, or by declaring that he or she (the witness) had eaten a particularly spicy meal just prior to seeing the crime. I think that the police would judge the witness either to be deranged, or to have an odd and indeed rather malicious sense of humor; or perhaps they would believe him or her to be evading the question, and to possibly be involved in the crime. This is because these events most likely have no bearing on the cause of the victim's death, and can bear no explanatory weight regarding it.

I will have more to say about causation and causal explanation in the next chapter;[10] for now, I just want to reiterate my claim that process explanations are intended to be a species of causal explanation, indicating a series of causally linked events that led to the event of interest.

Before concluding this description of process explanation, I would like to provide some further examples of process explanations used in the sciences; this constitutes the third and final aim of this section. I want to call attention to three particularly striking examples: the use of the concept of the natural history of a disease, in medicine; George

---

[10]See page 75.

Gaylord Simpson's macroevolutionary studies; and developmental theories in psychology. In addition to these three examples, a multitude of other sciences use process explanations, some of which I will describe briefly, after presenting the three examples just mentioned.

First, there are many examples of process explanations of the natural history of a disease.[11] For instance, consider the natural history of AIDS.

1. Contact with a person or object that transmits HIV to the infected person;

2. A period during which the infected person carries HIV, but does not have any diseases due to it;

3. The onset of full-blown AIDS, that is, the contraction of infections brought on by the infected person's weakened immune system; and

4. Death due to AIDS.

The explanatory question that is answered by describing the natural history of a disease D is "How do the symptoms of an organism with D progress?" or, supposing D to be fatal, "How do organisms with D die?" These questions clearly differ from the Hempelian "Why do people with D proceed through the stages that they do?" or "Why do people with D eventually die?"

It is important to note that, in many cases, a Hempelian answer to the latter questions cannot be produced, because, in many cases, the laws governing the transition

---

[11]Instances of the explanatory use of the concept of the natural history of a disease can be found in numerous sources in recent literature concerning a wide range of medical conditions, including irritable bowel syndrome [41], HIV [74], asthma [149], depression resulting from brain injury [36], plague [52], chronic subdural hematoma [87], hepatitis C [4], epilepsy [85], and diverticulitis [49]. A *PubMed* search performed in March of 2005 for publications with the term "natural history" in the title that have been published between 1985 and 2005 returned 4,328 results.

from one stage of an illness to another are not known. Furthermore, if Hempel's own version of Hempelianism were correct, many diseases' natural histories would be inexplicable. This is because many diseases only proceed from one stage to the next infrequently, and Hempel requires high probability for explanation.[12]

Second, Simpson's famous *Tempo and Mode in Evolution* [134] contains many excellent examples of process explanations. One of the most famous such examples is Simpson's [134, 83 - 93] classification of the different modes of selection and his use of them to explain the evolution of horses. Simpson classifies the ways in which selection might change the direction of evolution in a population by introducing the notions of centripetal and centrifugal selection. As he indicates in a compelling set of diagrams [134, 84, 90], centripetal selection occurs when natural selection removes new variants from a population, maintaining existing variants; centrifugal selection occurs when natural selection removes existing variants from a population, promoting new variants.

These modes of selection describe the sequence of states that an evolving population would pass through under the influence of a given mode of selection, and they are used in process explanation to answer the question "How did population P come to exhibit the state of variation that it now does?" Simpson did not intend for the modes of selection to be used to show that certain changes are to be expected, or should be judged to have a certain probability.

Third, developmental psychologists make use of process explanation. To focus on just one of the many developmental theories used by psychologists, consider theories of moral development. These theories answer the question "How does an individual come to

---

[12]Scriven's famous paresis case exemplifies this phenomenon. See below, pages 70 – 71.

63

have a full sense of the relationship between him- or herself and others?" These explanations proceed in narrative fashion, indicating the view of the self-other relationship that someone in a given stage possesses.

For instance, Gilligan [57] famously differs with Kohlberg [84] and other theorists by claiming that men and women differ in the stages of moral development that they pass through as they mature. Kohlberg claims that all people pass through various stages, including a stage in which they are motivated to be good to others for fear of punishment, and that those who reach the highest stages of moral development are able to follow abstract principles of right and wrong. Gilligan's contrary view is that there is an alternative sequence of stages, in which the terminal stage is reached when one is able to balance caring for others with caring for one's self, as opposed to acting from duty.

Theories of moral development are not the only psychological theories that can be used in process explanation. Freud's [50] famous stages of infantile sexuality were used by him in a similar manner, to answer the question, "How does a child mature?" As well, other theories, such as the theories of cognitive development proposed by Piaget, have been used in process explanations.

Finally, there are many other disciplines in which explanation-seeking questions demand narratives. Consider the following:

Cosmology In theories of cosmology, such as those concerning the Big Bang, scientists explain what has happened since the origin of the universe. As Peebles [114] and his coauthors indicate, they do so by describing its progression from a small, dense, primordial object containing all the matter in the universe, through an initial period in

which the fundamental properties of matter as we know them today were established, and continuing to the present, in which the universe is filled with stars and galaxies.

Geology Geologists [117] explain how rock strata have come to be in their relative present-day vertical positions by the so-called "principle of superposition," according to which processes of subsidence cause older layers to appear below younger layers.

Paleoecology Paleoecologists [108] explain how the remains of ancient environments, exposed to forces such as subsidence and weather, come to be distributed vertically in the present-day rock record.

Taphonomy Taphonomy ([96] and [6]) is the study of how living organisms become buried and, eventually, fossilized. A taphonomist might explain how a deer becomes fossilized by describing its death in the forest; its burial and compression under falling leaves; and the chemical processes that cause the organic matter of the deer's skeleton to be replaced with rock.

Evolutionary developmental biology Many biologists today are working toward accounting for the role of morphogenesis in evolutionary change. For instance, Oster and Alberch [113] ask, "How do animals grow feathers, as opposed to hair?" The answer to this question traces the genetics and physiological processes that give rise to both, which overlap considerably, suggesting that feathered and hair-covered organisms share a common ancestor.

Social work Social workers ([53], [109, chs. 12 - 15], and [129]) explain how therapy groups become social units directed toward the mutual aid and healing of their members by

identifying stages through which they proceed. The interest of these explanations is that they describe stages of group development leading to important events that may occur, for instance, the scapegoating of a group member, or the end of the group. They do not show whether, or how much, these events should be expected to occur.

2.3   Explaining Processes

Though my aims in this chapter have been modest, they are important. I have described Hempelianism and the views of some of its adherents. This is crucial because Hempelianism and the universalism it entails form one central target for my critical arguments in this dissertation. I have also described the nature of process explanation and provided some examples of it from the sciences. My central claim is that process explanations are called for in contexts in which someone asking an explanation-seeking question wants to know how the state of affairs of interest came about by learning the sequence of events that caused that state of affairs. This view forms the core of my claims that evolution occurring due to random drift can be explained: I believe that explaining such evolution requires process explanation. This entails the falsity of universalism.

I have not endeavored to defend contextualism or process explanation against criticisms that Hempelians are sure to want to level at those doctrines. As I suggest above, my aims in this chapter have been modest. I am not guilty of avoiding controversy, however: the task of the next chapter is to consider objections that Hempelians will want to make against my views, and to respond to them.

Chapter 3

A Defense of Process Explanation and Contextualism

Any Hempelians that have read this far will be thoroughly disappointed but perhaps also impressed that I should have the audacity to affirm the views I describe in the previous chapter—views that, on the Hempelians' account, are certainly mistaken. My aim in this chapter is to give the Hempelians their due, considering powerful objections to the warrant for both process explanation and contextualism. My view, of course, is that each objection is mistaken, a claim I argue for at length in this chapter, which consists of four subsections, as follows.

Section 3.1 concerns an objection that I term "the incompleteness objection:" the Hempelian claim that process explanations are best regarded as sketches of Hempelian explanations. Section 3.2 represents a further stage in the dialectic between myself and the Hempelian over this issue. In section 3.3, I consider fundamental issues concerning the

justification for universalism. Section 3.4 provides a brief set of concluding remarks.

3.1 The Incompleteness Objection

As I suggest in the previous chapter (page 57), the challenge that faces the contextualist in the argument against the Hempelian is as follows: identify and describe contexts in which requests for explanation do not demand laws. I have identified a few such contexts in my examples. The Hempelian should be expected to respond in kind. He or she would want to argue that cases in which laws do not appear to be required only in fact *appear* that way; laws are required if process-explanation narratives are to have any explanatory force. I think that the Hempelian would want to frame his or her attack on process explanations by claiming that they are *incomplete.* The Hempelians' view is that whatever explanatory force process explanations possess stems from the fact that they can be reconstructed to form Hempelian explanations. Let me explain the Hempelian view about how this works.

Consider a process explanation of an event E that consists of a sequence of events S; suppose that S is composed of $i = 1 \ldots N$ events S_i, where $S_N = E$. The Hempelians' proposal is that each of the S_i's should be viewed as events to be explained in their own right, in the Hempelian manner. For each stage S_i, causal laws should be provided, as should conditions obtaining in stage S_{i-1}. These conditions and causal laws together should satisfy the conditions for explaining S_i of whichever Hempelian model of explanation is preferred, i.e., that of Hempel, Salmon, Railton, or some other Hempelian. The Hempelian believes that this procedure of reconstruction suffices to bring what are supposedly process explanations back into the fold of Hempelianism, and moreover, into that of universalism

about scientific explanation more generally.

What is my response to this claim on the part of the Hempelian? Of course, I believe that the Hempelian is mistaken: process explanations are complete as stated, and need no further reference to laws, or any other conditions. My argument against the Hempelian proceeds by way of arguing against two statements to the effect that laws are essential to causal explanation. As might be expected, I draw heavily on the central tenets of contextualism about scientific explanation: the criteria for a good explanation are fixed by context, and differ across contexts. The first of the two statements that I want to argue against is as follows.

Statement 3.1 *Anyone who knows that a particular event c (of kind C) is the cause of a particular event e (of kind E) knows a causal law relating the occurrence of events of kind C to the subsequent occurrence of events of kind E.*

I do not know what most Hempelians would say regarding statement 3.1. Nonetheless, it is clear that, if this statement is true, then process explanations are, as the Hempelian suggests, incomplete. As I have emphasized, process explanations do not contain laws; if statement 3.1 is true, then it would follow that laws are required for discriminating which events cause which others. To speak metaphorically, this would break apart the causal chains that I believe are reflected in process explanations: the sequences of events that I claim are explanatory are, in fact, nothing more than a disconnected set of events of uncertain relation to one another, and so, clearly, cannot be explanatory. To join these events into an explanatory whole, the Hempelian would urge, the process explanation narrative must be supplemented with causal laws.

69

I think that it is clear that statement 3.1 is not correct, that is, I think that there is a wide range of cases, many of them appropriate to process explanation, in which no law is known, but in which the cause of an event of interest is clearly identifiable as such. My argument for this is that there are a number of widely-accepted strategies for establishing causal relationships among particular events that do not require that laws be invoked. I list these strategies below, providing a brief description of each.

Experiments In experiments, a situation is engineered in which the range of possible causes of some event e is restricted, so that if e occurs, it is known to be caused by another particular event c. Even Hume, who is famous for advocating a regularity account of causality, admits that a single critical experiment can reveal a causal relationship between two particular events.[1] Mill [3, 190 - 207] incorporates the idea that experiments can be used to identify causal relationships in his four methods. As Mill's methods are intended to provide the means of discovering causal laws, it is clear that the methods cannot presuppose knowledge of such laws.

Retrospective causal analysis Scriven [130, 479] articulates the canonical statement of this form of reasoning in terms of a well-known example, which is as follows. Paresis only appears in syphilitic patients, although it is rare even among them. When a syphilitic patient contracts paresis, it is correct to say that it is due to syphilis, because it is known that there are no causes of paresis except syphilis. No information about the frequency with which syphilitics contract paresis is required. More generally, suppose that an event e of a kind E is observed to occur, and it is known that only

---

[1]Ducasse [39, 144] discusses Hume's claim that a single instance can be evidence of a causal relationship between two particular events, providing references to the relevant passages of Hume.

events of kind C_i (where $i = 1 \ldots N$) cause events of type E.[2] Suppose that it is also known that an event c of type C_1 was present at some time just before e occurred, and that none of the other C_i's were. Even if nothing is known about the frequency with which any of the C_i's cause events of type E, it is still reasonable to claim that, in the case in question, e was due to c.[3]

Common causes Suppose that ten people attend a picnic, and all of them are afflicted with food poisoning the next day. It is much more reasonable to suppose that the tuna salad that all ten people ate caused their sickness, as opposed to thinking that they each contracted their illnesses separately. This is so, even with no information about how frequently tuna salad causes food poisoning.[4]

Everyday experience To take another example from Scriven [131, 68], consider the following. As I reach across my desk, my elbow hits an ink bottle; it falls to the floor, staining the carpet. There can be no question in my mind, or in that of others who learn of the spill and its circumstances, that it was my elbow that caused the bottle to fall and stain the carpet. Nonetheless, neither myself nor most other people have any knowledge whatever about causal laws governing the fall of ink bottles—or, more generally, about any of the medium-sized physical objects generally encountered

---

[2]Someone might point out that "only events of kind C_i (where $i = 1 \ldots N$) cause events of type E" is a law. If so, it does not help the Hempelian. It is disjunctive; it permits only the inference that some C_i caused e. Additionally, in cases such as the syphillis-paresis case, it cannot serve as the basis for a prediction. Suppose that $N = 100$, and that each C_i indicates a different form of syphilis, each of which is equally probable, none of which can be distinguished from any other by any means possessed by physicians; and suppose that only C_1 and C_2 cause paresis. Given the information that someone has syphilis and that "only events of kind C_i (where $i = 1 \ldots N$) cause events of type E," all that can be claimed is that there is a 2% chance that the person in question will contract paresis—a value falling far short of Hempel's requirement of practical certainty.

[3]Collins [24, 136] also provides some examples of this kind of reasoning.

[4]Sober [138] provides a description and defense of common-cause reasoning, also providing references to works by both Reichenbach and Salmon, who made common-cause inferences famous.

in everyday life. Nonetheless, many people know many correct claims about causal relationships among such objects.

This concludes my argument against statement 3.1, which I consider to be conclusive: establishing an event as the cause of another does not require subsuming the events under any laws. Now I want to consider another statement that must be shown to be false, if process explanation is to survive the Hempelians' claim that process explanations can be reconstructed as Hempelian explanations.

Statement 3.2 *Anyone who knows that a particular event c (of kind C) causally explains the occurrence of a particular event e (of kind E) knows a causal law relating the occurrence of events of kind C to the subsequent occurrence of events of kind E.*

The idea is that the Hempelian has capitulated to my attack on statement 3.1, admitting that it is possible to identify one event as the cause of another without recourse to laws. The Hempelian is not prepared to give up the fight, however, and falls back to statement 3.2. Indeed, statement 3.2 represents the very heart of Hempelianism. Let me explain.

As I argued above in my account of the Hempelians' theories of explanation, Hempelians believe that the models of explanation that they propose provide explanatory force in one of two ways.

1. Showing that the event to be explained was to be expected (Hempel).

2. Showing that the event to be explained was to be regarded with some certain degree of rational expectation (Salmon and Railton).

In causal explanation, causal laws provide the source of this explanatory force. In the case of Hempel, the causal laws form the basis for a deductive or inductive argument to the effect that the event to be explained will occur. In the case of Salmon and Railton, the causal laws provide a probability that is used to determine the degree of rational expectation that one ought to have in the event's occurrence. My suggestion that statement 3.2 is the very heart of Hempelianism amounts to the following: the causal laws required by statement 3.2 provide the link to rational expectation that Hempelians identify with explanatory power.

The Hempelians' suggested remedy for process explanation, therefore, is to supplement the explanatory narratives of process explanation with causal laws, so that the degree of rational expectation that one ought to have concerning the occurrence of each stage S_i of the process explanation can be determined. This reflects an idea I suggested above (page 68): the Hempelian would view process explanation as fundamentally incomplete, and would adopt the strategy of reconstructing process explanation in Hempelian terms.

So, if process explanation is not to be assimilated to Hempelianism, the claim that there can be causal explanations in the absence of laws—statement 3.2 above—must be conclusively defeated. I would now like to propose an argument that I believe does just that. The crucial issue concerns the *level of organization* at which explanations ought to be formulated. The Hempelian wants to insist that explanations ought only to be formulated in terms of levels of organization at which there are laws. In contrast, my view is that the level of organization at which explanations should be formulated is fixed by context, and

that there are many contexts in which a level of organization at which there are no known laws is called for. To begin my discussion of this issue, let me provide a rough sense of what I mean by "level of organization."

By this phrase, I intend to indicate the degree of complexity of a group of entities, in the following sense. One set of entities H exists at a higher level of organization than another set, L, if and only if entities in H are composed of entities in L, and the behavior of the entities in H depends in some way on the behavior of the entities in L. Atoms and molecules versus medium sized physical objects and neurons versus beliefs and desires exemplify pairs of sets of entities that have what appears to be a higher level-lower level relationship to one another.

Now what I would like to do is to state how the issue about explanatory completeness arises in connection with levels of organization. I begin with some observations that the Hempelian would make about process explanations. First, the Hempelian would point out, the narratives characteristic of process explanation are often formulated at a level of organization at which there are no laws governing the phenomena in question. Call this level of organization L_1. Second, in many cases, the behavior of entities at L_1 depends upon the behavior of entities at a lower level, which will be designated level L_0. Third, in many cases, there are laws that govern the behavior of entities at level L_0.

The Hempelian believes that these three observations are important because together, he or she believes, they entail that the behavior of entities at level L_1 (at which there are no laws) can be described, at least in principle, in terms of the behavior of entities at level L_0 (at which there are laws). From this, the Hempelian concludes that explanations

formulated in terms of entities at level L_1 are incomplete, compared to explanations formulated at level L_0. Process explanations formulated at higher levels of organization therefore constitute explanation sketches awaiting completion, in the Hempelian manner, in terms of a lower level of organization.

My argument against this incompleteness charge goes as follows. First, one of the conditions for a good explanation that is fixed by context is the level of organization at which the explanation should be formulated. Second, there are many contexts in which the level of organization is fixed at a higher level rather than at a lower one. Thus, sometimes—indeed, quite often—an explanation formulated at a higher level is complete as stated, and does not require any further elaboration. This is so, regardless of whether there are laws at the higher level, the lower level, neither, or both. I believe that Scriven encapsulates the contextualist view about causal explanation nicely as follows:

> But the giving of causes, and of scientific explanations and descriptions in general, is not the giving of "complete" accounts; it is the giving of useful and enlightening partial accounts.... The search for a *really* complete account is never-ending, but the search for causes is often *entirely* successful, and someone who saw a man killed by an automobile but refused to accept the coroner's statement that this was the cause of death on the grounds that some people survive being hit by car, does not understand the term *cause*. The coroner is perfectly correct, even though other factors are involved. [130, 479]

What I would like to do now is to illustrate both the Hempelian charge of incompleteness and my contextualist response in terms of a coin tossing example derived from a scene in Tom Stoppard's play *Rosencrantz and Guildenstern are Dead* [144]. As the play opens, Rosencrantz and Guildenstern are on stage "*passing the time in a place without any visible character.... Each of them has a large leather money bag*" [144, 11].[5] The scene is

---

[5]I have maintained the script-writer's convention of italicizing stage instructions.

further described by Stoppard as follows.

> *Guildenstern's bag is nearly empty. Rosencrantz's bag is nearly full. The reason being: they are betting on the toss of a coin, in the following manner: Guildenstern (hereafter "GUIL") takes a coin out of his bag, spins it, letting it fall. Rosencrantz (hereafter "ROS") studies it, announces it as "heads" (as it happens) and puts it into his own bag. Then they repeat the process.* [144, 11]

As the play develops, the audience learns that it is not just tossing the coin that is repeated, but the entire sequence of events, from the time a coin is selected to the time it lands: a coin is selected from Guildenstern's bag; the coin is tossed; the coin lands heads; the coin is placed in Rosencrantz's bag.

> ROS: Heads.
> *He picks it up and puts it in his bag. The process is repeated.*
> Heads.
> *Again.*
> Heads.
> *Again.*
> Heads.
> *Again.*
> GUIL: *(flipping a coin)*: There is an art to the building up of suspense.
> ROS: Heads.
> GUIL: *(flipping another)*: Though it can be done by luck alone.
> ROS: Heads.

The process is repeated a total of 92 times [144, 16], each with the same result: the coin selected from Guildenstern's bag lands heads. This set of events may be described as follows.

Statement 3.3 *All the coins selected from Guildenstern's bag—ninety-two in all—landed heads.*

The crucial passages of the play, which I would now like to consider, concern Rosencrantz and Guildenstern's discussion, after 85 coins have landed heads, of how the strange events described by statement 3.3 should be explained. The key point is that they discuss the issue of *cheating.* Guildenstern is surprised that Rosencrantz is not more disturbed by the result; Rosencrantz's first thought, as it turns out, is for his winnings, remarking, "Well, I won—didn't I?" after stating that "You [Guildenstern] spun them yourself" [144, 14]. The idea seems to be that Guildenstern would not have stocked his bag with two-headed coins if he were going to bet on tails.

> GUIL: *(approaches him—quieter)*: And if you'd lost? If they'd come down against you, eighty-five times, one after another, just like that?
>
> ROS: Eighty-five in a row? *Tails*?
>
> GUIL: Yes! what would you think?
>
> ROS: (*doubtfully*): Well . . . (*Jocularly.*) Well, I'd have a good look at your coins for a start!
>
> GUIL: (*retiring*): I'm relieved. At least we can still count on self-interest as a predictable factor . . . I suppose it's the last to go. [144, 14]

There are two claims I want to make about Rosencrantz and Guildenstern's discussion of cheating.

1. *Upon determining that the coins are not biased, they give up trying to explain why the coins from Guildenstern's bag landed heads. No such explanation why is possible.* The context of the discussion—a game of chance—fixes a requirement on explaining the outcome of the game. This requirement is that results disproportionately favoring one player must be explained by indicating whether the coins and tossing apparatus were tampered with or fixed so as to favor the supposedly lucky player. As a consequence of this requirement, an explanation of why the coins from Guildenstern's bag landed

heads is acceptable only if formulated in terms of "macro" sized objects, that is, the coins and the tossing mechanism.

2. *The absence of cheating does not preclude explaining* how *the coins landed heads.* The appropriateness of explaining how the coins landed heads accounts for the role of process explanation in cases such as the one at issue here. The process explanation explains how the coins landed heads by describing a mechanism, viz., a set of chance events that occurred in sequence to generate the run of coins that landed heads. This is compatible with the requirement that the explanation be formulated at the "macro" level.

Let me illustrate these claims in greater detail by considering the kind of conversational exchanges about the explanation of the coins' behavior that would be generated by the following question.

Question 3.1 *What happened? (Asked in an agitated, alarmed tone, after witnessing the coin-tossing game played by Rosencrantz and Guildenstern, and indicating Guildenstern's depleted bag of coins, in contrast to Rosencrantz's enlarged bag.)*

Rosencrantz and Guildenstern would most likely want to answer this question with something such as "Guildenstern has been terribly unlucky," which is a claim to the effect that the events described in statement 3.3 are chance events. This suggests the question, "Does chance explain anything about the coins' behavior?" My response to this question, as my discussion in the passages just prior to question 3.1 is intended to suggest, may be encapsulated in the following three points. On the one hand, (a) chance does not explain

why the result described in statement 3.3 obtained. On the other hand, (b) chance provides a limited but significant explanation *how* the result described in statement 3.3 obtained. Furthermore, (c) it is possible to improve upon the limited explanation provided by reference to chance by formulating a process explanation.

To begin my account of these three points, consider that "Guildenstern has been terribly unlucky" does not explain *why* the coins behaved as they did. My view is that some events, such as highly improbable chance events, occur for no reason at all. This claim must be understood in a contextualist manner; it is compatible with determinism, which I understand, following Earman [40], to be the view that for any state of of the universe U_T at a time T, there is one and only one state of the universe $U_{T+\triangle}$ that necessarily obtains at time $T + \triangle$. Let me explain this in more detail.

Perhaps it is possible to determine the state of Guildenstern's coins before each is tossed at a scale small enough that this information, together with laws of kinematics, could be used to deduce that each would land heads. Nonetheless, I do not believe that this information would be of any use to anyone who wants a causal explanation of the results of the coin-tossing game. As I suggest above (page 75), citing Scriven, the factors that one cites in causal explanations depend upon one's beliefs and intentions. The description of the coin's "micro-states" before each toss and the kinematic theory required would most likely be too complicated for anyone understand. Additionally, this information would not satisfy any kind of curiosity that anyone is likely to have, because it would not bring to light any systematic connections among the coins; it is highly unlikely that any two coins share the same state before they are tossed. Thus, I do not deny that each result of heads

has some physical antecedent that necessitates it; what I deny is that, in the context of a game of chance, knowledge of that antecedent would provide a substantial answer to the question, "Why did 92 of 92 coins land heads?"[6]

In taking the view that highly improbable chance events do not have causal explanations, that is, that reasons cannot be given why they occur, I follow Aurthur Collins.

> Suppose someone tosses a coin five times and gets five heads. Suppose he is astonished. We might mitigate his astonishment by telling him that over 3 per cent of such suequences are all heads. Not having reflected he had thought of such an outcome as much less likely. Whether by extension, analogy, or in its own right, I think all this might be presented as "an explanation of what happened".... [Nevertheless,] the explanation here is not condensable as, "You got five heads because it was a sequence of five tosses of a normal coin," paralleling, "x got the rash because he was injected with penicillin." [24, 133 - 134]

Collins's view is that there is a disanalogy between the explanation of the rash and the (putative) explanation of the result of the coin tosses. The disanalogy is that the former is a legitimate causal explanation, while the latter is not. To indicate why Collins claims this, let me provide some additional information. Previous to the passage I cite above, Collins states that "A small number of people who are injected with penicillin develop an allergic skin reaction." How many is "a small number?" "I invent a figure: 2 per cent" [24, 130]. He then describes and evaluates a brief conversational exchange of the kind that might occur concerning a case in which someone injected with penicillin breaks out in a rash.[7]

> Why did he [the person injected with penicillin, "x," break out in a rash]? Well, by hypothesis, because he received an injection of penicillin. Does this

---

[6]I want to note that Joseph Keller [77] and Persi Diaconis have developed models of the kinematics of coin tossing and other chance phenomena; however, I think they bring us only a trivial distance toward the "laws of coin tossing," if they do so at all.

[7]All references to Collins in this paragraph are from page 130 of his "Use of Statistics in Explanation" [24].

make sense? Is the explanatory claim illogical? No, the explanation would be accepted and rightly accepted in most situations in which an explanation for x's getting a rash is sought. [24, 130]

Looking back to the passage of Collins's that I cite before this one, it is clear that his view is as follows. Notwithstanding the low probability of the rash's occurrence, it is explained by the injection. The explanation is causal: the injection caused the rash. In contrast, citing a normal coin-tossing process does not explain why its improbable result obtained. This is because a sequence of coin tosses does not cause its result, and so cannot be cited in a causal explanation of it.[8]

This concludes my argument that reference to chance does not explain why all 92 coins landed heads; I now turn to the claim that reference to chance provides a limited explanation of *how* all 92 coins landed heads. "Guildenstern has been terribly unlucky" explains how the coins landed as they did by describing the *kind of mechanism* that produced the unlucky result. As I suggest above, the central issue in a game of chance in which one person loses repeatedly is whether that person has been cheated. "Guildenstern has been terribly unlucky" eliminates precisely this possibility: it is not the case that the coins have two heads; that Rosencrantz surreptitiously turned over all coins that landed tails when Guildenstern had his back turned; that the coins are painted on one side and so favor heads; and so on. To be clear about it, I want to emphasize that, while I believe that reference to chance does indeed exert explanatory force, I see it as limited in degree. Reference to chance provides no information concerning the nature of the mechanism that is actually responsible for the outcome of the coin-tossing game, and may be regarded as

---

[8]Achinstein [2, 179], who also cites Collins for the reason that I do here, makes a similar argument, using a similar example. I am also indebted to Achinstein [2] for some of the ideas I express concerning causality, above.

what Achinstein (pers. comm.) terms the "zero point" of process explanation.[9]

This brings me to the third and final point that I want to make about the coin-tossing example. The question may be asked, Is there anything more that can be said in answer to the question "What happened?" that would provide further explanation of the state of affairs at issue? As chance provides a limited explanation, what further information, if any, can be provided that would explain how the coins landed as they did? My answer to these questions is that there is indeed something further that can be said, as follows: a substantive process explanation provides further insight into the mechanism by which the coins landed heads. The process explanation fills in the details unspecified by an empty reference to chance, describing a sequence of coin tosses that, together, resulted in 92 heads.

Such a process explanation would go something like the following in the case at hand.

Explanation 3.1 *A balanced coin C_i ($i = 1 \ldots 91$), with heads on one face and tails on the other, was selected by Guildenstern from his bag of coins. C_i was tossed in the usual manner and allowed to land on a flat, hard surface. C_i landed heads. The entire sequence of events, from Guildenstern's selection of C_i to its landing heads, was observed by both Rosencrantz and Guildenstern, who both noted the result of the toss. C_i was placed in Rosencrantz's bag. This procedure—selecting a coin, tossing it, its landing heads, noting the result of heads, and placing the coin in Rosencrantz's bag—was repeated with coin C_{i+1}.*

---

[9]Perhaps mentioning the probabilities would add to the explanatory force of "Guildenstern has been terribly unlucky" by indicating just how unlucky he has been. Guildenstern's streak of bad luck is indeed precipitous. The probability of getting 92 heads in a row by tossing fair coins is approximately $2.02 \times 10^{-26}\%$. Although I do not find the analogy entirely suitable, I think that a comparison to the age of the universe helps to make this probability more tractable to the imagination. Astronomers believe that the universe has been in existence for approximately 15 billion years. An elapsed time of one second constitutes $2.11 \times 10^{-16}\%$ of this time span—a percentage that is still 10 orders of magnitude (1 billion times) larger than the chance of 92 of 92 fair coins landing heads.

To conclude my assessment of the coin-tossing case, the key point that I want to make here is as the follows. *Pace* Hempel, explanation 3.1 is not incomplete because it lacks laws or detailed information about the coin. There is no need for more information about the cause of the coins' odd behavior, which, as Guildenstern remarks, is "a spectacular vindication of the principle that each coin spun individually ... is as likely to come down heads as tails and therefore should cause no surprise each individual time it does" [144, 16]. There is no need for a theory of the kinematics of coin tossing, together with information that could be used to deduce each coin's fate, showing that all 92 should have been expected to land heads. It might be the case that such a low-level theory of coin tossing could be developed, but no one losing a game of chance such as that played by Rosencrantz and Guildenstern would care about such a theory.

In contrast with the Hempelian view, my view is that requirements about the level of organization are implicit in the context of the discussion in which a request for explanation is made. In this case, the discussion is about a game of chance, which requires a higher level explanation that accounts for the fairness of each toss, rather than why each coin landed as it did. This requires a process explanation answering "What happened?"—explaining *how* the coins behaved—which should be formulated at the "macro-level." Like all process explanations, the explanation of how the coins landed heads describes a sequence of events, in this case, a sequence of repeated events of the kind described in explanation 3.1.[10]

---

[10]I want to to acknowledge that Rosencrantz and Guildenstern do, in fact, consider further reasons why the coins behaved as they did, for instance, the claim that "time has stopped dead, and the single experience of one coin being spun once has been repeated ninety times" [144, 16]. I think that this does not show that they would not be not satisfied with a process explanation such as explanation 3.1. Rather, this absurd explanation and the others they consider highlight the notion that there are events that occur without reason, and indicate Rosencrantz and Guildenstern's sense of foreboding. Careful reading of the opening passages of the play confirms that they believe that the coins provide a sign of something exceptional in their future—something that will occur apparently without reason. Of course, this event is the order that

The explanation of certain events resulting from certain processes of sexual reproduction also exemplifies the point that the level of organization at which an explanation should be formulated is specified by context. Suppose that Mr. and Mrs. Smith each have the genotype A_1A_2, and that they produce three children, each of which has the genotype A_2A_2. Assuming that these children are conceived by normal processes of Mendelian reproduction, the probability of this event is $1/4 \times 1/4 \times 1/4$, which is between 1% and 2%—a diminishingly small value.[11] I think that answering the question "What happened?" by affirming that the children were conceived in the usual manner and that Mr. and Mrs. Smith have normal reproductive physiology would satisfy the curiosity of most friends and family. In particular, answering "What happened?" only requires reference to higher level entities such as the general character of the reproductive physiology of each parent. It does not require reference to lower level entities such as the particular sperm and eggs that produced each child.

Considerations parallel to those I raised in the Rosencrantz and Guildenstern case apply in this case, as well. Suppose someone sees the Smith family together, noting that the children look like one another but that they differ markedly from their parents. Suppose, furthermore, that this person asks "What happened?" To answer in an explanatory manner, it would be sufficient to say, "It happened by chance." This eliminates "biasing mechanisms" such as adoption, genetic engineering, a disease that causes all of Mr. Smith's sex cells to carry the A_2 allele, and the like. One might elaborate, providing a process explanation

___

they be killed, which is inexplicable to them. "GUIL: But why? Was it all for this?. . . No—it is not enough. To be told so little—to such an end—and still, finally, to be denied an explanation—" [144, 122]. This is in accord with my view, which is that they do not want a lower-level explanation of why the coins behaved as they did, and that since there is no explanation *at the level of interest to them*, no answer to the relevant why-question will satisfy them, even one at a lower level.

[11]The precise value is 1/64, which is 0.015625.

describing Mendelian reproduction. This is as much of an explanation as can be given, because there is no explaining *why* the Smith children have the traits they do: they have them by chance. This is so, determinism notwithstanding. Except perhaps for physicians studying sex, no one would want to know or be able to know the small-scale physical facts and laws that make it necessary that each child has the genotype in question.[12]

To bring my discussion of levels of organization and explanatory completeness to a close, I want to assert that the coin tossing case and the case of sexual reproduction I discuss above exemplify a general phenomenon. Causal explanation is called for at one level of organization at which there are no known laws; but such explanations are nonetheless complete. Biology and medicine provide some conspicuous examples of this. Consider a complex biological mechanism, such as photosynthesis. There are differing levels of organization at which this mechanism can be described. Some provide more detail than others; some, perhaps, are even framed in terms of the fundamental properties of matter, and provide a lawlike basis for predicting the result of any given step of the process. However, explanations of photosynthesis that do not provide such detail are not, for that reason, incomplete. Disease processes also exemplify this phenomenon. Consider Scriven's paresis example, cited above. Perhaps there is some characteristic, yet to be discovered, that distinguishes syphilitics who will contract paresis from those who will not. Nonetheless, the claim that someone had syphilis at time T explains that person's having paresis at time $T + 1$.

Looking to everyday events, a similar point can be made concerning the explana-

---

[12]I discuss the issue of the appropriate level of organization for explaining the outcome of sexual reproduction in greater depth in section 5.2.2.

tion of a flat tire. It is in principle possible to provide a statistical-mechanical description of the behavior of the air molecules in the tire when it is punctured. Presumably, this explanation would invoke laws, and show that the tire should have been expected to lose air pressure. However, this would be a very poor explanation; "We ran over a nail" would be better.[13] Explaining why a square peg does not fit into a round hole raises similar issues.[14] On the one hand, someone could try to answer this question by citing molecular-level interactions among fundamental constituents of the peg and the material into which the hole is carved; on the other hand, someone could simply describe the shapes of the peg and the hole.

This concludes my argument against statement 3.2 above, the claim that a complete causal explanation requires knowledge of laws. To recap, the argument is that there are many contexts in which a causal explanation is complete, even if it invokes no laws; and it is clear that it is complete, because the contexts in which it is advanced do not demand laws: one requirement for a good answer to an explanation-seeking question fixed by context is the level of organization at which the explanation should be formulated. As I have illustrated with various examples, the fact that it is possible to formulate an explanation in terms of a level of organization at which there are laws does not show that explanations formulated at higher levels are incomplete, even if no laws can be formulated for objects at those higher levels.

What would a committed Hempelian believe that the arguments I have just been making really show? While a Hempelian might agree that statement 3.1 is false, he or she

---

[13]The tire example is due to Peter Achinstein (pers. comm.). Salmon [126, sec. 11] believes that explanations of such cases (i.e., the distribution of gas molecules in space) are paradigm statistical explanations.

[14]I believe that I first encountered this example in a paper by Hilary Putnam, but I have not been able to locate that paper.

certainly would not agree that statement 3.2 is false, claiming rather that it is certainly true. In the next section, I consider the Hempelian defense of that statement, a defense which, as will be seen, invokes the famous "demon" described by Laplace; and I argue that the Hempelian defense of statement 3.2 fails.

3.2   Exorcising the Laplacian Demon

To be sure that the focus my current disagreement with the Hempelian is clear, let me reproduce the statement about which we disagree, statement 3.2, introduced above (page 72):

> Anyone who knows that a particular event c (of kind C) causally explains the occurrence of a particular event e (of kind E) knows a causal law relating the occurrence of events of kind C to the subsequent occurrence of events of kind E.

The Hempelians would begin their defense of statement 3.2 by pointing out that there is a criterion of adequacy, fixed by context, for all of the process explanations that I have considered so far. Moreover, the Hempelian would point out that this criterion of adequacy is fixed by the context in which *any* explanation-seeking how-question is asked, as long as that question is directed to a human being. The criterion of adequacy that the Hempelian has in mind is as follows.

Statement 3.4 *A process explanation E is adequate, even though the information contained in E is limited in extent and quality in the following sense: E includes only information that could have been obtained by someone subject to the conditions that usually limit a human being—perhaps even one augmented with artificial devices for measuring*

the environment and for calculating—in his or her capacity for gathering and analyzing

information.

I will term any context in which the criterion of adequacy for an explanation described in statement 3.4 applies a "human-centered context."

I agree that whether a context is human-centered might seem to be irrelevant to whether a process explanation is adequate, because it might seem that all contexts in which explanation-seeking questions are asked are human-centered. It might seem to some that while statement 3.4 accurately describes a requirement for good explanation, it is trivial. Surely, it cannot be a subject of serious philosophical debate: would it be reasonable to require that an explanation express information that could only be gathered by someone whose capacities for information gathering and analysis exceed those usually possessed by human beings, even human beings aided by artificial sensors and calculating devices?

Indeed, as will be seen, this is precisely what the Hempelian requires for an adequate explanation. In fact, the Hempelian makes the following claim about explanations of phenomena explained by process explanation: such explanations require information that can only be obtained by someone with capacities for information gathering and analysis that far exceed those imaginable for any human being, judged by today's standards. This is the case, the Hempelian believes, even if that human being's abilities are supplemented by detection and calculation tools. Before enlarging on this point, let me describe some of the human limitations at issue, to be clear about how human-centered contexts should be understood.

Some of the limitations that the Hempelian has in mind are imposed upon us by the construction of our sense organs and our cognitive capacities. Things far away are difficult to see, as are very small things; the human mind, unaided, can only hold within it a few steps of a complex line of thought at a given time. A further set of limitations on our capacity to gather and analyze information are imposed by the environment itself. For instance, we cannot know what is happening in a distant galaxy until some signal from that galaxy reaches us, and fundamental physical laws constrain the rate at which a signal can be transmitted—so there is a limit on when we can know what is happening in the galaxy in question. Similar constraints limit our ability to know what happens among very small objects.

Of course, one of the most notable aspects of human beings is that, to a large extent, we have managed to overcome many of the limitations that our constitution places on us. We have constructed instruments for detecting the properties of the environment, and we have invented formal techniques and calculating machines that extend our powers of analysis. Despite this, of course, we are painfully aware of the limitations that our instrumentation and calculating power have yet to overcome. For instance, the complexity of many objects of study in the biological and social sciences puts them well out of reach of even our most powerful observational and computational tools. The limitations placed on us by the environment itself cannot be overcome at all.

This is all I want to say about the limitations on human cognitive abilities; now what I would like to do is to indicate the relevance that the Hempelian believes that those limitations have for the adequacy of process explanation. In order to do so, recall the coin

tossing case described above (pages 75 - 83): 92 of 92 fair coins taken from Guildenstern's bag and tossed in a fair manner each land heads; Rosencrantz and Guildenstern puzzle over the explanation of this strange result.

The suggestion that I articulate above is that, in the context of a game of chance such as that played by Rosencrantz and Guildenstern, a process explanation that highlights the fair manner in which the coins were tossed would be satisfactory. Upon verifying that the coins and their manner of tossing are not biased, Rosencrantz and Guildenstern have no further interest in *why* the coins behaved so strangely. In particular, they do not want to know the state of each coin prior to being tossed, from which, together with laws of nature, it follows that each coin should have been expected to land heads. Similarly, they do not want information from which they could deduce how rational it would have been to expect each coin to land heads.

The Hempelians are not satisfied with this account of the coin tossing case. The reasons for their dissatisfaction can be described in two stages, the first of which is as follows. The Hempelians would point out that Rosencrantz and Guildenstern have replaced an explanation-seeking why-question about the coin with an explanation-seeking how-question about it because they know that it is not possible for human beings to acquire the information needed for answering the why-question. This is so even for human beings augmented with the best tools for measurement and mathematical analysis presently available. Someone asking explanation-seeking questions within a human-centered context restricts those questions, accepting that some explanation-seeking questions simply cannot be answered within such contexts.

The second stage of the Hempelians' response to my contextualist account of the coin tossing case builds on the first. The Hempelians agree that replacing an explanation-seeking why-question with an explanation-seeking how-question is reasonable in the coin tossing case. No one would expect someone S to ask a question Q, if it were known that S believed that Q could not be answered by the person to whom it was to be addressed. Nevertheless, Hempelians would point out that process explanations, which are incomplete by Hempelian standards, do not become any more complete for being posed within human-centered contexts: regardless of the context in which they are posed, process explanations are at best sketches of more complete explanations.

The Hempelians' view is that the most important explanation-seeking question about the coins is, "*Why* did 92 fair coins, tossed in a fair manner, each land heads?" The Hempelian agrees that there are indeed contexts in which some people are satisfied with answers to explanation-seeking how-questions about this very event. What accounts for the willingness of these people to accept the answers to such how-questions, according to the Hempelians, is that those people recognize the limits on human powers of cognition. They know that the corresponding why-questions cannot be answered. What their willingness does not show, according to the Hempelians, is that process explanations are complete. The correct attitude toward process explanations is that they fail to live up to the correct standard of adequacy for explanations—not that those standards should be weakened to account for human limitations.

The Hempelians see the conclusions that they draw concerning the coin tossing case to extend to all similar cases. Hempelians recognize a strict distinction between pragmatic

and non-pragmatic criteria of adequacy for explanation, strongly valuing the non-pragmatic. The set of requirements that make a context a human-centered one[15] exemplify pragmatic requirements for an explanation's adequacy. These pragmatic requirements are weaker than the non-pragmatic criteria that Hempelians believe apply across contexts, that is, the requirements that explanations contain laws, and that they conform to the various other requirements stated by Hempel, Salmon, or Railton that I describe in section 2.1 above.

To be clear, there are cases in which restricting one's explanation-seeking questions to human-centered contexts does not require that one lower one's standards for what one is willing to accept as an adequate explanation. Many explanation-seeking why-questions can be answered according to the highest standards of the most rigorous Hempelians, even given the limited cognitive resources available to human beings. For instance, many explanation-seeking why-questions about the interactions of medium- and large-sized physical objects moving at slow speeds can be completely answered using the framework of Newtonian mechanics. In such cases, the Hempelians would assert, explanation-seeking how-questions are superfluous.

The powerful implications of the Hempelians' universalism about explanation may be illustrated by a famous passage of Laplace's *Essai Philosophique sur les Probabilités*. In this passage, cited by Sober [139, 120], Laplace describes what has come to be known as "Laplace's demon," a being of extraordinary powers of perception and computation.

> Given for one instant an intelligence which could comprehend all the forces by which nature is animated and the respective situation of the beings who compose it—an intelligence sufficiently vast to submit these data to analysis—it would embrace in the same formula the movements of the greatest bodies of the universe and those of the lightest atom; for it, nothing would be uncertain and the future, and the past, would be present to its eyes.

---

[15]See statement 3.4 (pages 87 – 88).

The Laplacian demon is also what might be termed a "Hempelian demon." Because of its powers of cognition, the demon can deduce the state S of any object O at any time T, past, present, or future; Laplace says this much in the passage I cite above. The important point is that, on Hempel's own Hempelianism, this means that the Laplacian demon can also explain why any object O is in any state S at any time T. The demon's abilities far outstrip those of human beings; nonetheless, on the Hempelians' universalist account of explanation, only a being such as the demon can hope to possess a complete explanation of a given event.[16]

I do not want to obscure the subtleties of the Hempelians' understanding of the relationships between pragmatic and non-pragmatic standards for evaluating explanations. The Hempelians would agree that—for an audience consisting of human beings—a process explanation of the coin tossing case is clearly superior, from a pragmatic point of view, to the kind of explanation that the Laplacian demon might provide. Nevertheless, the process explanation fails fundamental Hempelian non-pragmatic conditions of adequacy: a committed Hempelian would not admit that it is an explanation at all, but would claim that it is an explanation *sketch*. Unfortunately, human beings are not able to understand much more, in the coin tossing case. Additionally, there is no reason for a Hempelian to ignore pragmatic concerns when assessing explanations that meet essential Hempelian non-pragmatic criteria for adequacy. For instance, a Hempelian would allow that an explanation E_1 is superior to an explanation E_2, if both are equally good from a non-pragmatic point

---

[16]Because Salmon and Railton's views about explaining low-probability events differ from Hempel's, their formulations of Hempelianism inform a different understanding of the Laplacian demon's explanatory capacities. Salmon and Railton believe that highly improbable events can be explained, even though they cannot be predicted (see pages 50 and 54, respectively). On Railton or Salmon's account, to have a complete explanation of why an object O is in a state S at any time T, the demon would only have to draw on its knowledge about the *probability* that O is in S at T.

of view, but E_1 is superior from a pragmatic point of view.

I have now concluded my account of the Hempelians' defense of their view in light of my arguments against statements 3.1 and 3.2 (pages 69 and 72, respectively), and what I would like to do now is to respond to the Hempelians. My strategy is as follows. I take it that the Hempelian believes that he or she has shown that my attack on statements 3.1 and 3.2 fails to show that the Hempelian view is incorrect because, according to the Hempelian, I have failed to provide examples of phenomena that cannot be accounted for on the Hempelian view. The Hempelians' response to my arguments is to re-assert their position more strongly, claiming that there are some phenomena such as coin tosses that can only be explained by a Laplacian demon, and that, for everyone else, the process explanations that I claim to be adequate are indeed incomplete. Having made the scope of their view clear, the Hempelians believe themselves to have successfully defended their view by showing that it is coherent, my arguments notwithstanding.

I do not have any conclusive arguments against the coherence of the Hempelians' view. Nevertheless, I can offer what I believe to be compelling reasons for abandoning their view. Further elaboration of the conceptual foundations of contextualism, themselves highly plausible, shows that, from the contextualist point of view, Hempelianism is positively outlandish. What I want to show is that, on some eminently reasonable claims about the nature of explanation, contextualism provides a better account than Hempelianism of the relevant "data" for the philosophical analysis of explanation, viz., practices of scientists and others that are correctly described by "explaining" and other related terms. The problem that I want to point out is that the Hempelians posit incompleteness where there does not

appear to be any. Let me begin my argument by indicating two important elements of conceptual background against which contextualists form their position, the first of which is as follows.

The central observation that motivates contextualism about explanation is the view that, as Achinstein [1, 135] suggests, "explanations are human inventions, serving human purposes. Their most important ... use is in acts of explaining to achieve a state of understanding in an audience." The notion that the primary aim of explanation and explaining is to place an audience in a state of understanding is also strongly reflected in Bromberger's [13] theory of explanation, according to which the goal of explanation is to remove someone from what is termed a "p-predicament," which is, roughly speaking, a state of not understanding some phenomenon. Bromberger suggests the following, reflecting on an historical account of the meaning of "explain."

> [The historical account of the meaning of "explain"] reminds us that sentences of the form '[person] A explained [proposition] W to [person] B' are aptly chosen to report episodes in which a tutor [person doing the explaining] turns someone who could truly have said 'I don't understand W' into someone in a position to assert 'I know W'. [13, 34]

To be clear, I want to note that from this point forward, when I indicate that the aim of explanation is to generate understanding, I mean the following. The aim of producing understanding may be directed at an actual audience, or a potential audience. The case of the former is clear enough: an audience poses an explanation-seeking question of someone or some group of people, who attempts to answer that question. The case of the latter, although not quite as obvious as the former, is nonetheless easily grasped. The idea is that, when someone produces an explanation that is not prompted by someone else's explanation-seeking question, he or she addresses a potential audience. For instance,

suppose a professor is practicing his or her lectures. Although no students have posed any questions of the professor on the topic of the lecture—they have not heard it yet—the lecture contains explanations directed at them; they, or someone with the same beliefs and interests as them, are the potential audience.

The second point of background against which the contextualists form their position is as follows. Just as there are many ways in which someone might want to understand some event E, there are many ways in which someone can have that event explained to him or her. This encompasses a range of explanation-seeking questions. As Bromberger suggests, "*Explain* and its cognate *explanation of* admit of interrogative sentences ... as their complements.... They admit most why-questions, how-questions, [and] what-is questions as their object" [14, 3].

Moreover, the variety of ways in which someone might fail to understand an event is not limited to the range of "wh-" and how-questions that one can ask about it: explanation-seeking questions are uttered against a background of requirements for their answers that Achinstein terms "instructions." Consider the following question, which Achinstein [1, 53] uses to explain the notion of instructions.

Question 3.2 *What caused Smith's death?*

Suppose, as Achinstein [1, 53] suggests, that the following two statements are proposed as explanations of Smith's death, in answer to question 3.2.

Explanation 3.2 *The cause of Smith's death was his contracting a disease.*

Explanation 3.3 *The cause of Smith's death was his contracting a disease involving a bacterial infection.*

Explanations 3.2 and 3.3 both meet basic criteria for answering question 3.2, because both indicate a cause of Smith's death. Nevertheless, they differ in the sense that they are formulated to satisfy alternative sets of instructions, adapted from Achinstein [1, 54], as follows.

Statement 3.5 (Instructions for explanation 3.2) *Say in a general way what caused Smith's death, e.g., whether it was caused by contracting a disease, or by some accident that befell him, or by an act of suicide.*

Statement 3.6 (Instructions for explanation 3.3) *As well as following instructions in statement 3.5, if a disease is cited, indicate something about what it is involves, e.g., whether it is bacterial or viral.*

The idea is that the first set of instructions above, which explanation 3.2 satisfies, calls for less detail than the second set, which explanation 3.3 satisfies.

Having provided these examples, I would like to offer some further clarification of the notion of instructions. The question of which set of instructions applies on a given occasion is settled by the intentions and cognitive states of the person asking the explanation-seeking question at issue. For instance, whether explanation 3.2 or 3.3 is the best answer to give in response to question 3.2 depends upon whether the person asking that question wants a more detailed explanation or a less detailed one, and whether the person understands the difference between bacterial and viral infection. Additionally, the relevant intentions and cognitive states need not be explicitly articulated by the person asking the explanation-seeking question. This does not mean, of course, that these intentions and cognitive states

are beyond detection; indeed, they very often are correctly detected, as indicated by the high frequency with which requests for explanation are fully met.

At this point, I am in a position to indicate how the central claim of contextualism can be generated by the various elements of background that I have just introduced. This represents the start of the final stage of my argument against the Hempelians, which, as I suggest above (pages 94 - 95), is as follows: I want to show that, given the conceptual foundations of contextualism (which I have just sketched), Hempelianism is a highly unusual view that fails to account for the phenomena that a philosophical account of explanation ought to be expected to account for. The central point, as I state above, is that Hempelians posit incompleteness where there does not appear to be any. In order to argue this, let me gather together the highlights of the discussion of Bromberger and Achinstein above into a brief line of thought that provides the rationale for contextualism about explanation.

First, consider the point that the aim of explaining is to generate understanding. From this, it follows that the worth of an explanation of an event E can be measured by how well it accomplishes this goal, that is, by whether it places its intended audience in the appropriate state of understanding concerning E. Second, consider the point that there are many ways in which someone might fail to be in a state of understanding about a given event E. This is because, as Bromberger suggests, explanation-seeking questions can be "why-questions, how-questions, [and] what-is questions" [14, 3]. Also, this is because the conditions for answering any such question depend upon which instructions are in place at the time the question is asked.

Third, consider the claim that which set of instructions applies on a given occasion depends upon the interests and cognitive states of the person asking the explanation-seeking question, on the occasion in question. From this, together with the previous two points, it follows that there is a broad range of criteria that are used to evaluate explanations, and that which criteria are used on a given occasion depends upon the interests and cognitive states of the person asking the explanation-seeking question at issue—the main claim of contextualism.

Now what I would like to do is to evaluate the Hempelians' universalism from the contextualist point of view that I have just elaborated. Reconstructed in contextualist terms, Hempelianism may be described as follows.

Statement 3.7 *If and only if someone fails to understand a particular event E in the context of a scientific inquiry, that person fails to understand why E occurs, that is, the way in which the person fails to understand E is that he or she fails to understand why E occurs.*

To be clear, I believe that statement 3.7 reflects the commitments of Hempelianism for the following reasons. As I have just indicated, on contextualism, explanation-seeking questions express an audience's lack of understanding, and accordingly, explanations aim at generating understanding. As I indicate in my account of Hempelianism above, Hempelians believe that why-questions are the only kind of explanation-seeking questions that scientists ask about particular events. From this, it follows that Hempelians see scientific explanations of particular events as always and only having the goal of remedying the audience's lack of understanding about why those events occurred.

Having indicated the sense in which statement 3.7 reflects Hempelianism, I would now like to evaluate that statement. My view is that it is highly implausible: it is strongly disconfirmed by what might be termed the "data" for the philosophical analysis of explanation, viz., the explanatory practices of scientists and others. This "data set" may be characterized by the following three claims.

1. There is a broad diversity of practices widely believed to be explanatory, including, but not limited to, explaining why.

2. These practices aim at generating understanding, and they often succeed at doing so.

3. When these practices do succeed at generating understanding, many of their audiences do not hesitate to enthusiastically affirm this point, showing no sign that the putative explanation in question is the least bit incomplete.

The idea is that, on contextualism, these three points obtain for the following reasons. As a matter of linguistic fact—as I indicate above, citing Bromberger—an explanation need not be prompted by a why-question. Additionally, there is a broad diversity of instructions, fixed by context, whether implicitly or explicitly, that audiences specify. This represents a broad range of ways in which someone might fail to understand an event. Audiences' positive assessments of responses to questions across this broad range of putatively explanation-seeking questions strongly suggest that their states of not understanding are adequately remedied by those putative explanations. Because generating understanding is the *sine qua non* of explaining, I take it that the diversity of putatively explanation-seeking questions I have been considering really are explanation-seeking questions, and that the responses given to them really are explanations.

In light of the diversity of understanding-generating practices that are recognized as falling into the category of explanation, I think that the following question must be put to the Hempelian. On what grounds, except that it would be in accord with his or her philosophical analysis of explanation, does the Hempelian posit incompleteness in cases of process explanation, or other explanations not conforming to Hempelianism? Perhaps the Hempelians have a clear intuition that explanations deviating from the Hempelian model are incomplete, while those conforming to it are complete. Contextualists do not share this intuition, but this is not simply a clash of intuitions; the contextualist view is motivated by the diversity of explanatory practices that I have just been discussing. Indeed, the contextualist view is formulated to account for this very diversity, describing explaining and explanation in terms of their fundamental aim, generating understanding.

This argument is a naturalistic. I believe that the philosopher's task is primarily descriptive. As Wittgenstein suggests, "The work of the philosopher consists in assembling reminders for a particular purpose" [150, no. 127], and "What we [philosophers] are supplying are really remarks on the natural history of human beings" [150, no. 415]. The idea is that the explanatory practices of scientists and ordinary people are for the most part in good working order, and the problem for the philosopher is to discern what they are, characterize their logic, and provide for their justification. From this point of view, it seems to me that the Hempelians' insistence that complete explanation requires a Laplacian demon is a badly mistaken assessment of the explanatory practices of scientists. The Hempelians make the mistake of taking one of the many uses of "explain" in science to be its canonical use.

This last point, my criticism of the Hempelians' assessment of the explanatory practices of scientists, is in accord with the following observation of Bromberger's about universalists generally, including the Hempelians and others.

> Most writers on the subject implicitly limit their attention to some subfamily of the questions admissible as objects of *explain*, and limit themselves to different ones. So Duhem, for instance, limits himself to the *what is the physical structure underlying such and such* subfamily; Mill, Hempel, and other adherents of the covering law view limit themselves to the *why* subfamily.... They are like the blind men who each reported (perhaps correctly) on a different part of the elephant. But unlike the blind men, they follow a reasonable strategy, if one assumes, as most of them seem to do, that their object should be to display truth-conditions distinctive of answers admitted by *explain* and its cognates. That cannot be done in one fell swoop for such a heterogenous family. It is therefore reasonable to concentrate on some particularly challenging subfamily. Of course that does not justify the widespread attitude that only one of these subfamilies is legitimate. [14, 3 - 4]

It may be the case that the Hempelians do not share my naturalistic methodology. Perhaps they want to introduce a new ideal for explanatory practices in science. If this is the case, the Hempelians need to provide arguments for accepting this ideal. As far as I can tell, the Hempelians do not offer any such arguments, and it is unclear what such arguments might look like. They would have to support the conclusion that the use of "explain" and related terms ought to be restricted to accord with Hempelianism. This would require scientists to give up explanations such as process explanations that, to all outward appearances, seem adequate. In light of the widespread use of process explanations, which I document in the previous chapter, I think that whatever arguments Hempelians provide in favor of the proposed ideal would have to be enormously compelling.

To conclude this line of argument in defense of contextualism, let me state what I believe I have accomplished by following the line of thought I take above: I believe that I have shifted the burden of argument back to the Hempelian. Consider the case of

process explanation, the species of explanation falling under the contextualist rubric that is most important for my arguments in this dissertation. The question is not, "Are process explanations complete?" Rather, the question is, "Why should process explanations be deemed to be incomplete, despite appearances to the contrary?" Let me explain, at the same time recapitulating key points of the discussion.

I agree that the Hempelian need not abandon his or her position, if the main worry is incoherence or inconsistency. Furthermore, I agree that, according to Hempelianism, process explanations are indeed incomplete: on Hempelianism, only a Laplacian demon can explain any given event in a complete manner. Nevertheless, coherence and consistency—virtues which contextualism also possesses—are not the only virtues of a philosophical theory. There is the issue of how many of the phenomena at issue are accounted for, and how well the theories at issue account for them. On this additional criterion, contextualism is superior to Hempelianism, a point which I believe obtains for the following reasons.

Contextualism accounts for the observation that there is a diversity of linguistic practices and associated concepts that are typically classified as explanatory, and that frequently accomplish the aim of causing audiences to understand. Causing audiences to understand, as I believe Hempelians would readily agree, is the *sine qua non* of explaining and explanations. This places the Hempelians in the position of having to argue that their position should be adopted as a description of the canonical use of "explain" and related terms and concepts (applied to particular events). That is, it is unclear why one should be a Hempelian instead of a contextualist: according to the latter, there is no single canonical use for "explain" and related terms and concepts, which is the view most strongly suggested

by a careful assessment of those terms and concepts.

3.3   Against Universalism

In the previous two sections, I responded to the Hempelians' incompleteness objection. I believe that I have met this objection: the examples I present above, interpreted in light of the idea that the aim of explaining is to create understanding, create enough doubt about the Hempelian position to weaken it beyond a point at which it is minimally credible. A Hempelian would not want to surrender his or her position nonetheless, because he or she would contend that a significant element of it remains untouched by my arguments: universalism. In this section, I consider how a Hempelian would want to respond to my naturalistic claim that there is a diversity of explanatory practices, and I defend contextualism against the Hempelian.

The variety of explanatory practices, the Hempelian would argue, reflect what might be termed "understanding relative to a context." The Hempelian readily admits that different contexts call for different kinds of explanations, and would even go so far as to express appreciation for the work that contextualists do to describe explanatory practices appropriate to different contexts. Nevertheless, the Hempelian claims, there is a kind of understanding that is context-free. The Hempelian would invoke the analogy, mentioned at several points in this dissertation so far, between explanation and mathematical proof.[17]

The idea is that whether the conclusion of a proof follows from its premises does not depend upon the proof's audience. Analogously, there is a sense of "understanding" that

---

[17]See, for instance, page 40.

refs to context-free understanding. The mark of a good explanation is that it generates understanding of this sort, which is the kind of understanding that scientists aim at creating when they are not working toward satisfying the needs of a particular audience. Moreover, the Hempelian would assert, this is the most important aim of scientific explanation.

The Hempelian line of counter-attack that I am considering here may also be described as follows. In addressing the concerns raised by the notion of the Laplacian demon (page 102 above), I attacked the Hempelian for introducing into science an ideal for explanatory practice. My argument was that, given the diversity of explanatory practices in science, there is no reason to think that scientists are motivated by the kind of ideal the Hempelian affirms.

The Hempelian would articulate the following response to my claim that the "Hempel-Laplace ideal" is illegitimate. He or she would not deny that there are a diversity of contextually-fixed requirements for evaluating explanations. As I have suggested above, the Hempelian would admit that utterances made in accord with such requirements produce understanding relative to a context. Indeed, the Hempelian would also admit that some scientists some of the time—or perhaps even all scientists some of the time—produce and aim at producing explanations in accord with such contextually-fixed requirements. Nevertheless, the Hempelian would claim, these scientists recognize that such explanations fall short of the Hempelian ideal for explanation, which they seek to attain whenever possible. The Hempelian does not see their inability to attain this ideal in some cases as a reason to abandon the ideal; rather, it is a reason to affirm the ideal all the more strongly, recognizing the place of contextually-fixed requirements for explanation in the structure of

aims that guide scientists in their work.

I do not believe that there is a context-free kind of understanding, but I do not know how to argue for this. Instead, I would like to argue for a weaker claim that, if true, is still quite strong—strong enough to defeat the Hempelian. This claim is that scientists do not and should not aim at producing context-free understanding: the diversity of explanatory practices indicates that scientists do not aim at the Laplacian ideal for explanation proposed by Hempelians. On the one hand, this should create significant doubts about whether there in fact exists any such thing as context-free understanding. On the other hand, more importantly for my position against the Hempelians, if such an ideal plays no role in guiding explanatory practices in science, then there is considerably less motivation for the Hempelians' universalism as an account of scientific explanation. Contextualism is more reasonable.

To begin my argument, let me introduce some new ideas, adapted from Achinstein [1, ch. 4, pt. 1]: illocutionary and non-illocutionary standards for evaluating explanations. The former may be understood as follows.

Analysis 3.1 (Illocutionary evaluation) *A standard S for evaluating an explanation E is an illocutionary standard if and only if the following obtains: E is adequate according to S only if E satisfies appropriate audience instructions.*

To be clear about what analysis 3.1 states, let me recall how Achinstein's notion of instructions should be understood; as well, I would like to explain what "appropriate instructions" means. Instructions, as I indicate above (pages 96 - 98), are contextually-fixed requirements for the adequacy of an explanation that relativize the success of the

106

explanation to the audience's beliefs and intentions. For example, as I suggest above in my explanation of instructions, different people might have different interests and cognitive states that influence the requirements they place on the explanation of someone's death. Some might want to know whether the person's death is explained by disease or something else such as an accident or suicide. Others, upon learning that a disease is responsible, might want to know whether the disease was bacterial or viral.

The notion of appropriate instructions protects illocutionary standards from trivialization. Suppose someone requires that all explanations be submitted to the Federal Reserve Board and reviewed and signed by Alan Greenspan himself. This is irrelevant to generating knowledge, except perhaps concerning monetary policy, and in any case, the existence of such cases means that audience instructions alone cannot be necessary or sufficient for illocutionary evaluation. A stronger condition is needed: appropriate instructions. Roughly speaking, instructions are appropriate if and only if they reflect an audience's interest in obtaining knowledge that it lacks or that would be of some interest or value to it. Because audience instructions will frequently be appropriate as stated, the requirement that instructions be appropriate does not represent a particularly onerous condition.[18]

Now, let me describe the second important notion I want to introduce here, that of non-illocutionary evaluation.

Analysis 3.2 (Non-illocutionary evaluation) *A standard S for evaluating an explanation E is a non-illocutionary standard if and only if the following obtains: whether E is adequate according to S does not depend in any way upon the beliefs or intentions of E's*

---

[18]I am deeply indebted to Achinstein [1, 112 - 116] for this account of appropriate instructions. To be clear, the term "appropriate instructions" is his.

audience.

Hempelians believe that their models of explanation constitute non-illocutionary criteria, and furthermore, that those models describe context-free understanding. Take Hempel's D-N model, for example. Suppose that argument A satisfies the D-N model. If Hempel is correct to claim that the D-N model provides necessary and sufficient conditions for explanation, A is adequate as an explanation from a non-illocutionary point of view. This is because, according to the D-N conditions, whether an explanation is a good one does not depend upon whether anyone in the audience has a particular belief or intention. All that is required is that A contain at least one law, and that its conclusion, which must follow validly from its premises, be a statement that the phenomenon to be explained obtains. Neither of these conditions require that the audience of A have any particular beliefs or desires; indeed, these conditions can obtain even if A were not told to anyone in particular, or existed in some abstract sense in a possible world in which there were no beliefs and desires.

Hopefully it is clear enough that the distinction between illocutionary and non-illocutionary standards matches up with that between contextualism and universalism. As I have indicated previously,[19] contextualists and universalists disagree about the role of the audience in determining whether an explanation is a good one. Contextualists believe that the audience's intentions and cognitive states play an essential role in formulating criteria for evaluating explanations; universalists believe that they do not play any such role. Given this, it should be clear that contextualists believe that illocutionary criteria play an essential role in evaluating explanations: audience intentions and cognitive states

---

[19]See the opening passages of the previous chapter (page 36).

are embodied in audience instructions, which form the essence of illocutionary criteria. In contrast, universalists believe that the correct criteria for evaluating explanations are non-illocutionary; audience intentions and cognitive states play no role in such criteria.

Having introduced the distinction between non-illocutionary and illocutionary standards for evaluating explanations, I would now like to describe some of the theory of population genetics central to an example that I will use in my argument against the Hempelian. What I want to consider are two descriptions of the same phenomenon, the change in allele frequencies due to natural selection across a single generation, represented by $\triangle_s p$.

The first description of $\triangle_s p$, which may be found in Gillespie's population genetics text [56, 52, eqn. 3.1], is as follows.

$$\triangle_s p = \frac{pqs[ph + q(1 - h)]}{\bar{w}} \tag{3.1}$$

(Let p indicate the frequency of the A_1 allele; q, the frequency of the A_2 allele; s, the selection coefficient, a ratio of the fitness values of the A_1 and A_2 alleles; h, the heterozygous effect, a measure of how much a heterozygote's fitness differs from either homozygote; and let $\bar{w}$ represent the mean fitness of the population.)

The second formula I want to consider, which also appears in Gillespie's text [56, 59, eqn. 3.5], describes $\triangle_s p$ as follows.

$$\triangle_s p = \frac{pq}{2\bar{w}} \frac{d\bar{w}}{dp} \tag{3.2}$$

To be clear, I want to emphasize that these two equations describe the same

phenomena: both can be derived using relatively simple mathematics from the same set of claims about natural selection, fitness, and Mendelian genetics. Neither is controversial, and both are frequently presented in introductory population genetics texts as fundamental models of natural selection.

In my argument against the Hempelian, it will be useful to take into account remarks made by Gillespie concerning the two formulas I have just cited. Gillespie claims that, on the one hand, equation 3.1 "is probably the single most important equation in all of population genetics and evolution," but that, on the other hand, "it isn't pretty, being a ratio of two polynomials with three parameters each" [56, 52]. In contrast, equation 3.2 has the virtue of describing otherwise disparate phenomena in common terms. He elaborates as follows.

> There is something unsatisfying about the description of the three forms of natural selection. They come off as a series of disconnected cases. One might have hoped for some unifying principle that would make all three cases appear as instances of some more general dynamic. In fact, Sewall Wright found unity when he wrote ... [equation 3.1] in the more provocative form [of equation 3.2].
> [56, 59]

The three forms of natural selection he mentions in this passage are known as directional, balancing, and disruptive selection. Gillespie describes each of these in terms of relationships among the values of the parameters of equation 3.1 by indicating what values each parameter must take on if one or the other forms of selection is to occur. The point he is making in the passage above is that while equation 3.2 describes the factors affecting changes in allele frequency in a manner that makes relationships between these three forms of selection easier to see, equation 3.1 describes those factors in a manner that makes those relationships more difficult to see. This is because equation 3.2 describes changes in

110

allele frequencies as a function of the mean fitness of the population, showing that allele frequencies will always change in a way that increases the mean fitness of the population.

I am now in a position to argue against the Hempelians' claim that scientists aim at producing context-free understanding. The central claim that I want to make in favor of this conclusion is as follows.

Statement 3.8 *As well as requesting explanations for the purpose of understanding the event to be explained, virtually all audiences request them in order to advance some further purpose or other that is integral to science.*

Let me explain how I understand this claim before indicating its significance for the argument against the Hempelian. I think that statement 3.8 is best explained in terms of an audience's *reasons for requiring a set of instructions I*. Let me elaborate. As I have indicated at several points so far, instructions are contextually-fixed requirements for an explanation that an audience puts in place in order to specify the particular aspect of the phenomenon to be explained. The point I want to make is that there are a number of reasons why an audience might require that an explanation meet a set of instructions I, many of which reflect the audience's pursuit of goals of central importance to the conduct and aims of science. Two goals that stand out as particularly important and that serve, in many cases, as reasons for specifying a set of instructions I include (a) promoting further research, and (b) providing what I will call—for lack of a better term—"metaphysical insights." Let me describe each of these goals and elaborate on the sense in which explanations meeting illocutionary standards advance them in terms of Wright's description of $\triangle_s p$.

First, consider the goal of promoting further research. Explanations formulated

using important parameters can promote further research by suggesting the explanation of other phenomena. Suppose that Dr. Smith, an evolutionary biologist, is studying two populations: one is evolving by directional selection, the other, by disruptive selection. Suppose, furthermore, that Dr. Smith does not know Wright's equation 3.2, only knowing the less perspicacious equation 3.1. Dr. Smith might puzzle over the case of disruptive selection, unable to understand the evolutionary dynamics that produce it. The key insight may come for Dr. Smith when he learns of equation 3.2, which enables him to see the relationship between directional and disruptive selection. This reflects instructions such as "explain the directional selection case in a compact but highly general manner."

Promoting further research can be an overriding concern, particularly when the subject under study is not well known. William Whewell, claiming that "The character of the true philosopher is, not that he never conjectures hazardously, but that his conjectures are clearly conceived, and brought into rigid contact with facts" [3, 155], suggests that, at the start of an investigation, aiming to generate good explanations can be more useful than aiming to generate well-confirmed theories.

> Hence he who has to discover the laws of nature may have to invent many suppositions before he hits upon the right one; and among the endowments which lead to his success, we must reckon that fertility of invention which minister to him such imaginary schemes, till at least he finds the one which conforms to the true order of nature. [3, 154]

Second, consider the goal of generating what I will call "metaphysical insights." Scientific explanation has important consequences for what were once considered key questions of metaphysics. Sewall Wright's description of $\triangle_s p$ provided in equation 3.2 contributes to answering one such question about progress in nature. The insight that (under certain conditions) natural selection will increase a population's mean fitness has impor-

tant consequences for the issue of whether there is some sense in which later generations represent progress over earlier generations, an issue that was of great concern to Victorians— and that concerns many in the present day. Wright's equation 3.2 makes this clearer than equation 3.1: explanations of $\triangle_s p$ formulated using the former contribute more to our understanding of this broader issue than do those formulated using the latter. This reflects instructions such as "explain $\triangle_s p$ in a way that shows its relation to mean fitness."

This completes my explication and defense of the claim that there are some important reasons that scientists have for selecting the instructions that they do for a given explanation. What is the consequence of this for Hempelianism? I would now like to answer this question. First, I want to suggest that the Hempelian affirms what I will call "the Hempelian conditional."

Statement 3.9 (The Hempelian conditional) *For all audiences A and all explanations E of particular events, If A's member or members evaluate E using an illocutionary standard, then they do so because they know that, due to limitations on human cognitive abilities, the appropriate non-illocutionary standard S cannot be met.*

Hempelians, as I indicate above,[20] believe that scientists aim at the ideal of scientific explanation represented by Laplace's demon. As I also indicate above,[21] universalism may be understood as a claim that there exists (or exist) some non-illocutionary criterion (or criteria) for evaluating explanations, and that no non-illocutionary criteria are correct, except for evaluating whether an explanation is good in a pragmatic or contextual sense. Deviations from the ideal represented by non-illocutionary standards, according to

---

[20]See pages 104 - 106.
[21]See pages 108 – 109.

Hempelians, are due to the inability of human beings to gather the information that would be needed for a complete explanation. That is, the use of illocutionary criteria for evaluating explanations marks an audience's recognition that, as human beings, we are bound to request explanations in what I have termed "human-centered contexts" (see pages 87 – 88). This accounts for the Hempelians' affirmation of the Hempelian conditional: either someone uses a non-illocutionary standard; or else does not, and believes that it is impossible to do so, because of the limitations of the human-centered context.[22]

Now let me indicate the relevance for Hempelianism of my claim about the reasons that scientists have for choosing the instructions that they do. My position is that audiences using illocutionary standards do not do so because they believe that it is impossible to meet an appropriate non-illocutionary standard. Rather, they use illocutionary standards because they want to attain some end that motivates their requests for explanation. This end is what shapes their instructions. Recall the example I introduce above. The idea is that someone would choose to use an explanation incorporating equation 3.2 (Wright's formula) rather than an explanation incorporating equation 3.1 not because the former provides a more complete picture of the truth, but because it is better suited for the ends he or she has in mind. As I suggest above, someone might want to promote further research, or shed light on a question of general significance (i.e., what I have termed a "metaphysical" issue).

How would a Hempelian respond to this line of thought? He or she would argue that, contrary to what I claim, I invoke audience-independent goals of science in order to

---

[22]As I have stated it, the Hempelian conditional may be vacuously true. It allows that E has no referent, that is, that the audience not provide any explanation at all. In this case, presumably, the antecedent is not satisfied, as the audience uses no illocutionary standard—no standard is employed at all. To foreclose this possibility, the conditional may be qualified by requiring that E not refer in an empty manner, i.e., that some explanation actually be given.

justify contextualism. Promoting further research and addressing metaphysical concerns are aims intrinsic to science. They are not, as I claim, ends that individuals can choose among in shaping their investigations. If a scientist takes these ends as reasons for requiring a certain set of instructions I, he or she cannot be said to generate understanding suited for ends that depend essentially on his or her own beliefs and desires. Rather, he or she aims to generate understanding of a context-free kind—a kind of understanding described by one or another of the models proposed by the Hempelian.

I view this as a weak response. Suppose that promoting research and generating what I have termed "metaphysical insights" are indeed universal aims. Even so, it does not follow that audience instructions play no role in determining how explanations advanced to promote these ends are evaluated. The problem is that these aims are too broad to directly inform the explanatory practices of scientists: subsidiary aims that are derived from audience instructions determine how these broader aims are to be realized in a given case. Let me show how this works in terms of the example from population genetics I presented above.

A biologist wants to promote further research; which explanation does he or she choose: an explanation incorporating equation 3.1, or an explanation incorporating Wright's equation 3.2? In my discussion of these explanations above, I construe the former as more illuminating and the latter as less illuminating, claiming that the former explanation better advances the goal of promoting further research. What I would like to emphasize now is that these assessments of relative adequacy depend on aims of the investigator beyond whether he or she wants to promote further research. A consequence of this is that investigators sharing

the goal of promoting further research may nonetheless differ on whether an explanation incorporating equation 3.2 is better than one incorporating equation 3.1. In the present case, this plays out as follows.

In the example above, the explanation incorporating Wright's formula is deemed to be better because it is assessed from the point of view of someone who wants, in addition to promoting further research, to attain the following aim: to see the relationships among the various forms of selection. What I would like to point out is that this is not the only aim that someone wanting to promote further research might have for wanting to explain $\triangle_s p$, and that explanations incorporating equation 3.1 would be better for some of these other aims. Note that the parameter h does not appear in Wright's formula; it is a component of $\bar{w}$. Suppose that someone is interested in promoting research, but is not as interested in understanding relationships among the forms of selection as he or she is in understanding the role of the heterozygous effect h in natural selection. From this person's point of view, an explanation incorporating Wright's formula would be less illuminating than an explanation incorporating equation 3.1.

My point here is that the claim that someone wants to promote further research does not imply that there exists an unique standard for evaluating a given explanation: further more particular aims can affect which standard applies in a given case. More generally, all that the contextualist requires is that there are cases in which there exist *some* contextually fixed criteria for evaluating explanations, a requirement that is compatible with the existence of criteria that do not depend on anyone's beliefs or intentions. A scientist may be motivated by aims intrinsic to science such as promoting further research or generating

metaphysical insights; nonetheless, other concerns particular to the context dictate how these aims are to be attained in a given case. The universalists needs to argue for a stronger claim, that is, that *all* aims used by scientists to determine how explanations are to be evaluated are independent of context.

This concludes my argument against the Hempelians' universalism. To recap, that argument has been as follows. I believe that there is no such thing as context-free understanding. Nonetheless, I have not endeavored to argue for this claim here. Rather, I have argued for the claim that scientists are not motivated to pursue explanations for the purpose of generating such understanding. I formulate my argument in terms of a distinction between illocutionary and non-illocutionary standards, a distinction that can be used to formulate contextualism and universalism. Contextualism is the view that an audience's instructions play a central role in establishing standards for evaluating explanations, while universalism is the view that they play no role at all in doing so. With this distinction in hand, the argument against the universalist is as follows.

The universalist position is that the use of illocutionary standards reflects scientists' inability to generate explanations that meet the higher standard embodied in Hempelian non-illocutionary theories of explanation. According to Hempelians, these standards represent what is required to attain context-free understanding. I argue that the use of illocutionary standards is not a recognition of human limitations. Rather, the use of such standards reflects the range of goals advanced by explanations, goals that cannot be served by explanations that meet Hempelian requirements such as including laws or showing why an event occurred. I illustrate this claim by comparing the explanatory value of two theories

of the change in allele frequency $\triangle_s p$. The difference between the two theories is not that one is closer to an ideal of completeness, but that it serves other important aims that scientists have. I conclude that even if there does exist some form of context-free understanding, scientists do not seek this understanding—they seek understanding fit for human ends.

3.4 Process Explanation Vindicated

The Hempelian posits incompleteness where there does not seem to be any. Moreover, he or she claims that the warrant for doing so is that there exists a kind of understanding that obtains free of any context, a view that, if true, has been overstated in its importance by the Hempelian, because scientists do not often aim at attaining understanding free of context. To revive the Hempelian position, the Hempelians must provide reasons for ignoring contextually-imposed criteria—criteria that do not require laws—for the evaluation of explanations. This is because ignoring these contextually-imposed criteria in favor of some Hempelian set of criteria, which are supposed to represent the conditions under which context-free understanding obtains, is what creates the appearance of incompleteness. I believe that the arguments of this chapter and the examples and explication of process explanation in the previous chapter signify the defeat of Hempelianism and the vindication of contextualism and process explanation, and I do not think that reviving the Hempelian position will be an easy task.

Chapter 4

The Probability Account of

Indiscriminate Sampling

In this chapter, my aim is to argue that a process that John Beatty [9] terms "indiscriminate sampling" is a mechanism of drift, and that what I term "the probability account" of indiscriminate sampling is correct. Indiscriminate sampling shares its essential properties with a canonical model for chance processes, which "indiscriminate sampling" is intended to suggest: a blindfolded person drawing beads from an urn containing beads of a variety of colors. This kind of sampling is indiscriminate in the sense that the blindfolded person cannot see the beads, and so cannot discriminate beads of one color from beads of another color. The bead-drawing process can also be described using another sense of "discriminate:" because he or she cannot see them, the person cannot intentionally select beads of one color rather than another—for instance, favoring red beads over green for his or her bead collection. As will be seen, the blindfolded person, the beads, the urn, and the

sample of beads created by the activities of the blindfolded person have clear biological and ecological parallels in processes of random drift.

The account of indiscriminate sampling that I develop in this chapter provides essential background to chapter 5, the next chapter, which represents the culmination of my work in this dissertation. In chapter 5, I describe various research programs in evolutionary biology that aim at explaining, by process explanation, phenomena that occur by random drift. This represents the culmination of my work in this dissertation because it shows the application of narrative-style explanations to chance phenomena in evolutionary biology, defying the Hempelian view that explanations always include laws of nature. The phenomena explained by scientists pursuing the research programs that I describe in chapter 5 occur not just by drift, but by a particular mechanism of drift, that is, indiscriminate sampling. Without a clear sense of what indiscriminate sampling is, I would not be able to accurately describe the phenomena at issue, nor the process explanations that scientists provide for them.

The account of indiscriminate sampling that I develop in this chapter is also intended to deflect controversy away from issues having to do with the nature of drift, in order to focus it on issues having to do with explanation. Philosophers of biology do not agree on how to best understand drift, as a recent dispute in print between Brandon and Millstein exhibits ([105], [12], and [106]). Nonetheless, philosophers of biology do generally agree that indiscriminate sampling is a form of drift. This is why I have decided to formulate the examples of the next chapter in terms of that process. I hope that the theory of indiscriminate sampling that I develop in this chapter cements the limited consensus about

the nature of drift, and that, by framing my arguments about explanation in terms of that theory, I am able to direct attention toward explaining drift.

In this chapter, I proceed as follows. In section 4.1, I provide further background to indiscriminate sampling, drawing the conclusion (among others that I also draw) that John Beatty is correct that there are two types of indiscriminate sampling, "indiscriminate parent sampling" and "indiscriminate gamete sampling." In sections 4.2 and 4.3, I elaborate the probability account of indiscriminate sampling; the former concerns indiscriminate parent sampling, and the latter concerns indiscriminate gamete sampling. Section 4.4 concerns the relationship between indiscriminate sampling and evolution. I conclude, in section 4.5, with a brief summary and overview of my work in the chapter.

4.1   Background to Indiscriminate Sampling

The notion of indiscriminate sampling takes its central meaning from a metaphor: blindly drawing beads from an urn to create a sample. This metaphor also serves as a model for practices that are widespread in applications of probability and statistics such as opinion polling, census-taking, quality control, and the clinical testing of medicines. Beatty explains the analogy as follows.

> What does it mean to attribute ... [gene- and genotype-frequency changes to random drift]? Since as early as 1932 ... a popular approach to the exposition of random drift has been via a classic means of modelling [sic] chance processes—namely, the blind drawing of beads from an urn. The beads in this case are alleles—the different alleles are the different colors, but they are otherwise indistinguishable by the blindfolded sampling agent. One urn of beads represents one generation of alleles—a finite number, characterized by particular allele frequencies. The next generation of alleles is determined by a blind drawing of beads from an urn. This second generation of alleles fills a new urn, blind drawings from which determine the frequencies in the third generation.

121

And so on. The frequencies of alleles may differ from urn to urn—generation to generation—as a result of the fact that frequencies of otherwise indistinguishable beads sampled by blind drawings may not be representative of the frequencies in the urns from which the samples were drawn. [9, 188]

Beatty understands biological sampling processes to occur in different stages of the life cycle, juvenile and adult. He terms sampling in the juvenile stage "parent sampling," which he describes as "*the process of determining which organisms of one generation will be parents of the next, and how many offspring each parent will have*" [9, 188; Beatty's emphasis]. Sampling in the latter stage of the life cycle, which Beatty terms "gamete sampling," occurs as a part of sexual reproduction, and is described by Beatty 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*" [9, 189; Beatty's emphasis].

Beatty claims that "[parent sampling] might be 'indiscriminate' in the sense that any physical differences between the organisms of one generation might be irrelevant to differences in their offspring contributions. A forest fire, for instance, might so sample parents—killing some, sparing some—without regard to physical differences between them" [9, 188 - 189]. The analogy with blindly drawing beads from an urn is that in the case of the urn, as in the case of organisms, "any physical differences . . . between the entities in question are irrelevant to whether or not they are sampled" [9, 188 - 189]. He makes a similar claim concerning gamete sampling, which he sees in a precisely analogous manner [9, 189].

Two simple examples illustrate the central idea behind indiscriminate parent sampling. The first example is the "colorblind predator" example [105, 35, 38]. Suppose that,

122

in population P, organisms that differ from one another in color are preyed upon by a colorblind predator. This is analogous to a blindfolded person drawing beads from an urn in the following sense. Differences in color among the beads cannot make a difference to whether they are selected, just as differences in color among organisms in P cannot make a difference to whether they are killed by a colorblind predator.

The second example is a variation on Beatty's [9, 192] "twins" example. Suppose that organisms in population P^* are phenotypically identical, but differ in genotype. One generation, random lightning strikes kill several organisms in P^*. Just as the differences in color among beads cannot make a difference to whether the blindfolded person selects them from the urn, differences among the genotypes of the organisms in P^* cannot make a difference to whether they escape the lightning. This is because the genotypes are "invisible" to the lightning, "cloaked" by their bearers' identical phenotypes.

Beatty provides three interpretations of indiscriminate sampling, each of which differs from the others. Although I do not think it is intuitively clear that these interpretations of indiscriminate sampling differ from one another, I do not think it is intuitively clear that they are identical, either. In any case, concern for the length and complexity of this chapter prevents me from arguing that they do in fact differ; the argument is both lengthy and technical. Rather, I will simply state the three interpretations for parent sampling, asserting that parallel interpretations exist for indiscriminate gamete sampling.

1. Indiscriminate parent sampling occurs only if organisms fail to differ in *fitness*. Beatty states that "*selection* is ... a sampling process that discriminates, in particular, on the basis of *fitness* differences" [9, 190], making similar statements elsewhere [9, 191].

2. Indiscriminate parent sampling occurs only if organisms fail to differ in *physical properties.* This interpretation is suggested by the passages I cite in my initial discussion of indiscriminate sampling, above.

3. Indiscriminate parent sampling occurs only if organisms fail to differ in their *probability* of survival [9, 190].

I conclude that Beatty does not have a settled view about which of these three interpretations he believes to be correct; moreover, I think it is not clear that he recognizes that there are ambiguities in his account: he does not discuss any sense in which they differ from one another, and he uses them as though they are interchangeable. In a striking passage that exhibits his deep ambivalence, he combines the first and second interpretations above, referring to a lightning strike as sampling "with regard to *physical fitness* differences" [9, 192; emphasis mine]. As will be seen, my account differs substantially from each of these three alternatives, although it contains elements of each.

Beatty presupposes (but does not explicitly advance) a further claim about the nature of indiscriminate sampling. This claim is that, in each case of indiscriminate sampling, there is some condition, structure, or event in the environment or mating system that does the sampling, and that it is this condition, structure, or event to which the predicate "indiscriminate" is properly applied. The idea is that, when describing a case of parent or gamete sampling, the environment and mating system are to be described piecemeal, decomposed into various conditions, structures, or events that interact with organisms and gametes, and that each may be discriminate or indiscriminate independently of others.

For instance, although a forest fire might cause more organisms to die than any other source of mortality, it need not be the only such source. Suppose that before the forest fire organisms had to contend with the usual causes of death such as predation, starvation, sickness, weather, and so on. These represent further agents of parent sampling; each may be discriminate or indiscriminate.

Likewise, different factors can be isolated in the mating system, and some may be indiscriminate, while others are not. I think that the mating system is not best understood on analogy with an ecological setting, as a kind of arena in which events occur. Rather, I see the mating system as a continuous physical structure extending between generations. It encompasses physiological traits, behavioral traits, and spatial relationships among organisms; and it culminates with the union of the gametes in sex, which forms the connection of chromosomal inheritance between generations.

For this reason, I think that it is more intuitive to speak of a *mating system structure*, rather than of a sampling agent of the mating system. In his discussion of indiscriminate gamete sampling [9, 189], Beatty mentions Mendelian reproduction, which is a good example of a mating system structure that samples gametes. Other such structures include behaviors that determine mate choice and physiological conditions promoting fertility.

The evidence that I believe supports my interpretation of Beatty on this point about what might be called the "decomposability" of the environment and mating system is the set of examples he provides. He always considers a particular agent, for instance, a forest fire [9, 189, 193], a lightning strike [9, 192], or predation [9, 193]. His discussion of

the physical basis for fitness [9, 193] also suggests that he has this view. I take Millstein's discussion of the interaction of drift and selection among some hypothetical snails [105, 44] as evidence that she also holds this view.

I view Beatty's account of indiscriminate sampling as a qualified success of a kind that I believe to be typical of efforts to introduce a new idea into the literature: it provides a novel and useful framework for thinking about a significant issue, but lacks critical details. On the positive side, developing the bead-drawing model more explicitly than it had been before in philosophy is a notable contribution to the problem of understanding the nature of selection and drift. Additionally, I view Beatty's division of the life cycle into parent and gamete sampling as useful and important, and I integrate this division into my own account of indiscriminate sampling, below. Similarly, I integrate his "decomposability" claims into my own account.

On the negative side, Beatty's account does not advance beyond the metaphorical level. Clearly, there is some analogy between a blindfolded person drawing colored beads from an urn and a colorblind predator selecting its prey, or lightning killing organisms with identical phenotypes. The question remains, however, What are the causal relationships that justify these analogies? Furthermore, while provocative, Beatty does not develop his suggestion that, when indiscriminate parent sampling occurs, important relationships obtain between the roles of probability, fitness, and physical properties. Indeed, as I point out above, he seems unaware that there is any need to clarify these relationships.

4.2 Indiscriminate Parent Sampling

The central claim of the probability account of indiscriminate parent sampling is as follows. For a parent sampling agent to sample organisms indiscriminately in virtue of their bearing a given trait, it is both necessary and sufficient that, *ceteris paribus*, organisms with alternative variants of the trait have the same probability of being killed by that parent sampling agent. This reflects the central property of the canonical model of indiscriminate sampling, blindly drawing beads of different colors from an urn: the color of a bead does not affect its probability of being selected by the person drawing them. The aim of this section is to develop a formal account of the circumstances under which a variant of a trait fails to make a difference to its bearers' probability of being killed by a parent sampling agent, relative to organisms bearing other variants of the trait.

In section 4.2.1, I introduce important background to the account. Section 4.2.2 describes the central provision of the account, which I term "the core probabilistic equality." In section 4.2.3, I describe an additional provision of the account that, together with the core probabilistic equality, constitutes a complete set of necessary and sufficient conditions for indiscriminate parent sampling.

4.2.1 Essential background

In this section, I describe the scope of the theory of indiscriminate parent sampling that I want to advance, and I introduce a set of definitions that are essential for formulating that theory.

The scope of the theory

I intend for the probability account of indiscriminate parent sampling to apply only to a limited set of populations and parent sampling agents, as follows.

Statement 4.1 (Scope of the theory) *As I formulate it in this dissertation, the probability account of indiscriminate parent sampling applies only to synchronously mating organisms that have non-overlapping generations, and to indiscriminate parent sampling agents that act by killing organisms during the juvenile stage of the life cycle.*

By "juvenile stage," I mean to indicate the period between the time at which organisms in a population are born and the time at which they become sexually mature. Synchronous mating occurs if and only if organisms in a population mate at the same time, that is, during a mating season. Such generations are non-overlapping if and only if each organism in the population only mates once in its lifetime.

Limiting my account to parent sampling agents that kill organisms means that it does not apply to indiscriminate parent sampling that occurs in connection with mating and reproductive success. While I recognize that a complete analysis of indiscriminate parent sampling would require an account of the latter, I do not believe that leaving these topics out imposes a serious limitation on the usefulness and importance of my theory. Indiscriminate sampling associated with mortality is a significant phenomenon, of which there are many important examples.

Furthermore, I think that the probability account can be extended to apply to cases in which indiscriminate parent sampling occurs in connection with mating and reproductive success; and I also think that it applies to cases in which mating is not synchronous, and

128

in which generations overlap. In any case, the reason for limiting the scope of my account in each of the ways I mention above is that, without such limitations, my account would be considerably longer and more complex.

Essential definitions

The following definitions, essential to the theory, apply implicitly to any generation G in any population P meeting the conditions described in statement 4.1.

Definition 4.1 *"O_i" refers to the ith organism in the population at the start of the juvenile stage of the generation in question, where there are N organisms in the population at that time, and where $i = 1 \ldots N$; and "O_j" refers to the jth organism in the population at the start of the juvenile stage of the generation in question, where $j = 1 \ldots N$.*

Definition 4.2 *"V_v" refers to the vth variant of trait T, where V_v is heritable, and where there are a total of V variants in the population, and where $v = 1 \ldots V$; "V_w" refers to the wth variant of trait T, where V_w is heritable, and where $w = 1 \ldots V$.*

I note that definition 4.2 requires that the variants in question be heritable. I address the issue of why I require heritability in this way in section 4.2.3 below (pages 147 - 148).

Definition 4.3 *"$O_i[V_v]$" means "Organism O_i has variant V_v;" "$O_j[V_w]$" means "Organism O_j has variant V_w."*

Definition 4.4 *A parent sampling agent is referred to by "S_P."*

Definition 4.5 *"$K[S_P, O_i]$" means "Parent sampling agent S_P kills organism O_i." Accordingly, "$K[S_P, O_j]$" means "Parent sampling agent S_P kills organism O_j."*

Definition 4.6 *"$\forall(O)$" means "$\forall(O_i)\,\forall(O_j)$".*[1]

Definition 4.7 *"$\forall(V)$" means "$\forall(V_v)\,\forall(V_w)$".*

4.2.2 The core probabilistic equality

In this section, I develop the central claim of the probability account of indiscriminate parent sampling. Because of its central role in my theory, I term this claim the *"core probabilistic equality;"* because of its content, I term it the "core *probabilistic equality;"* it is a statement of identity between two probability statements.[2] I proceed by first developing a "naïve" statement of it, which I then modify to generate a necessary condition for indiscriminate parent sampling. The reason I proceed in this manner is as follows. On the one hand, the naïve statement of the equality most clearly expresses the key idea of my theory; on the other hand, it is neither necessary nor sufficient for indiscriminate parent sampling, although a descendant of it is.

I proceed as follows, in further subdivisions of this section. First, I formulate the naïve statement of the core probabilistic equality. Second, I formulate what I term a "sophisticated" statement of the core probabilistic equality, which is generated by modifying the naïve statement. Third, I qualify the sophisticated statement of the equality in order to formulate a necessary condition for indiscriminate parent sampling.

---

[1]In definitions 4.6 and 4.7, I use "$\forall(\phi)\,(\ldots P\phi\ldots)$" to mean "for all ϕ, $\ldots P\phi\ldots$". I will continue to use this notation in the remainder of the dissertation.

[2]For economy and ease of expression, I will sometimes refer to the core probabilistic equality as "the core equality" or just "the equality."

A naïve statement of the core probabilistic equality

The naïve statement of the core probabilistic equality reflects the central notion of the probability account, as I describe it in the introduction to this section: if and only if parent sampling by a given parent sampling agent is indiscriminate, bearing a given variant of the trait at issue does not make any difference to an organism's probability of being killed by that parent sampling agent, relative to organisms with other variants, *ceteris paribus*. The naïve statement of the equality is as follows.

Statement 4.2 (Core equality, naïve) *In any population P in any generation G in which the conditions described in statement 4.1 obtain, for any parent sampling agent S_P, any pair of organisms O_i and O_j, and any pair of variants V_v and V_w, the probability that O_i is killed by S_P given that O_i has variant V_v is equal to the probability that O_j is killed by S_P given that O_j has variant V_w.*

That is, the following is true: in any population P in any generation G in which the conditions described in statement 4.1 obtain,

$$\forall(S_P)\,\forall(O)\,\forall(V)\,(p(K[S_P,O_i]|O_i[V_v]) = p(K[S_P,O_j]|O_j[V_w])).$$

The colorblind predator case, introduced above, can be used to illustrate how statement 4.2 constrains across-variant probabilities of an organism's being killed by a given parent sampling agent. Consider the following additional definitions.

Definition 4.8 *"$K[Pr,O_i]$" means "Predators kill organism O_i," and "$K[Pr,O_j]$" means "Predators kill organism O_j."*

Definition 4.9 *"$O_i[D]$" means "Prey organism O_i has dark coloration."*

131

Definition 4.10 *"$O_j[L]$" means "Prey organism O_j has light coloration."*

Having set out these definitions, I am now in a position to describe the circumstances under which the core probabilistic equality would be satisfied in the case of colored organisms and colorblind predators in the instance in question.

Statement 4.3 *In population P_{CP} in generation G_{CP} in which the conditions described in statement 4.1 obtain, for any pair of prey organisms O_i and O_j, the probability that O_i is killed by a predator given that O_i has dark coloration is equal to the probability that O_j is killed by a predator given that O_j has dark coloration.*

That is, the following is true, for population P_{CP}, in generation G_{CP}:

$$\forall(O)\,(p(K[Pr, O_i]|O_i[D]) = p(K[Pr, O_j]|O_j[L])).$$

Statement 4.3 means that, regardless of whether an organism in population P_{CP} has dark or light coloration, it has the same probability of being killed by predation as any other organism in the population. The reason for this is that predators cannot determine the color of their prey, and so the color of the latter does not affect their probability of being killed by the former.

A "sophisticated" statement of the core equality

As I intend for its name to suggest, the *naïve* statement of the core probabilistic equality (statement 4.2, page 131) is not adequate for the analysis of indiscriminate parent sampling. On the one hand, as I suggest above, I believe that the naïve statement of the equality captures the central idea of indiscriminate parent sampling: indiscriminate parent

132

sampling occurs if and only if the variant borne by an organism does not make a difference to its probability of being killed by the parent sampling agent in question, *ceteris paribus*. On the other hand, there is an important class of counterexamples to statement 4.2 that show that it cannot be a necessary condition for indiscriminate parent sampling. My aim in this section is to explain this class, and to produce a "sophisticated" statement of the core equality that is not affected by cases in this class.

Note that I am not claiming that the sophisticated statement of the equality is a necessary condition for indiscriminate parent sampling; this requires a further qualification, which I describe in the next subsection ("Relativizing to a time scale," page 138).

I begin with an account of what I term "correlated variants." Different traits, even those whose purposes are unrelated, can modify one another's fitness. Suppose that organisms that have a variant V_1 of a trait T also tend to have a variant V_1^* of another trait T, while organisms that have variants other than V_1 tend not to have V_1^*. I will say that the variants V_1 and V_1^* are "correlated," a term that I define as follows.

Definition 4.11 (Correlated variants) *In a population of interest, in a generation of interest, variant V_1 of a trait T is correlated with variant V_1^* of a trait T^* if and only if, in the population of interest, in the generation of interest, for any pair of organisms O_i and O_j, the probability that O_i has variant V_1^* given that O_i has variant V_1 is greater than the probability that O_j has variant V_1^* given that O_j does not have variant V_1.*

That is, the correlation obtains if and only if, in the population of interest, in the generation of interest, the following obtains:

$$\forall(O)\,(p(O_i[V_1^*]|O_i[V_1]) > p(O_j[V_1^*]|\neg O_j[V_1])).$$

133

Note that, if an organism in the relevant population does not have variant V_1, it has some other variant of the trait, e.g., V_2 or V_3.

An example derived from the colorblind predator case illustrates the notion of correlated traits. I will take a case which might be called "extreme correlation:" all and only organisms bearing V_1 bear the correlated trait, V_1^*. Suppose that, among the prey of the population with the colorblind predators, all and only organisms with dark coloration have bad eyes, that is, cannot see very well. Furthermore, suppose that prey that have bad eyes have a high probability of being eaten by a predator, because they are unable to notice a predator until it is about to attack, leaving no time to escape. In contrast, suppose that all and only organisms with light coloration have good eyes: they can see quite well, and as a consequence, can notice predators before they attack, leaving plenty of time to escape. I indicate these relationships on table 4.1.

Table 4.1: Example of correlated variants

All and only prey with ...	Also have ...	And so have a ...
Dark coloration	bad eyes	high probability of being eaten
Light coloration	good eyes	low probability of being eaten

This might occur for any one of several reasons, for instance, gene linkage: the variants for color might be on the same chromosome as genes controlling visual acuity, so they are inherited together. Given enough time, meiosis would break the linkage; but the time required for this to occur might be longer, even, than the lifetime of the species, depending upon how strong the linkage is and how fast natural selection for visual acuity operates. Another reason for a correlation between the inheritance of variants of a trait might be that they share a developmental pathway. Assortative mating might also cause

134

this kind of relationship. Perhaps organisms with bad eyesight prefer to mate with their dark-colored conspecifics, while those with good eyesight prefer to mate with their light-colored conspecifics.

Having introduced the idea of correlation between variants, I am now in a position to describe a counterexample to the naïve statement of the core equality. I continue to frame my argument in terms of the modified colorblind predator example. To begin with, I want to point out that, if color variants and visual acuity are correlated with one another in the manner I indicate above, then color variants do not satisfy the naïve statement of the equality. The probability that an organism is killed by predation will differ depending upon which color variant it has: color variants are correlated with visual acuity variants, which affect an organism's chance of being killed by predation.

The next step in my argument is as follows. Even though the naïve statement is not satisfied by the color variants, parent sampling by predation for color variants should still be regarded as indiscriminate. Although I have added to the example by stating that color variants are correlated with visual acuity variants, I have not changed another central property of the example: the predators are colorblind. Just as in the original presentation of the example, this means that an organism's color plays no role in causing it to be captured and eaten, or to escape from predators. Any differences among organisms' probabilities of being eaten are due to the correlation, not the causal role of color in attracting or distracting predators. I conclude that, in the modified colorblind predators case that I have presented here, parent sampling of color variants by predators is indiscriminate, even though the naïve statement of the core equality is not satisfied.

In order to replace the naïve statement of the core equality with the sophisticated one, the probabilities represented in the equality must be conditionalized on a statement to the effect that variants of interest are not correlated with any other variants of any other traits, that is, they must meet what I call a "no-correlation condition," which I indicate by "NC_P". In order to state the no-correlation condition NC_P, I need to introduce an additional definition, which provides notation for an additional set of variants.

Definition 4.12 "$V_{v^*}^*$" refers to the v^*th variant of trait T^*, where there are a total of V^* variants in the population, and where $v^* = 1 \ldots V^*$.

Now, let me introduce the no-correlation condition that I propose to integrate into statement 4.2.

Statement 4.4 (Condition NC_P) For any pair of organisms O_i and O_j, and for all variants $V_{v^*}^*$ of any trait T^*, variant V_1 of trait T is not correlated with $V_{v^*}^*$. That is, the following obtains:

$$\forall(O)\,\forall(V_{v^*}^*)\,(p(O_i[V_{v^*}^*]|O_i[V_1]) = p(O_j[V_{v^*}^*]|\neg O_j[V_1])).$$

Statement 4.4 is a generalization, stating that there are no variants of any trait with which the variant V_1 is correlated. In the example above, I consider whether coloration is correlated with a variant of another trait, visual acuity; for condition NC_P to be satisfied for coloration, it would have to be the case that coloration is not correlated with visual acuity or any other variant of any other trait. Of course, V_1 can refer to an arbitrary variant of the trait of interest.

Additionally, note that I will consider condition NC_P to be satisfied in what might be termed "virtual" cases, that is, I allow what might be termed "virtual satisfaction" of NC_P. Such cases occur if there is a correlation among the variants in question, but scientists are able to create a probability model that "corrects" for the correlation by computing what the relevant probabilities would be, if no such correlation existed.

Having now described the no-correlation condition, I must introduce one more definition before stating the sophisticated formulation of the core equality.

Definition 4.13 *"$\forall(Q_P)$" means "$\forall(S_P)\,\forall(O)\,\forall(V)\,\forall(V_v^*)$".*

Statement 4.5 (Core equality, sophisticated) *In any population P, in any generation G, in which the conditions described in statement 4.1 obtain, for any parent sampling agent S_P, any pair of organisms O_i and O_j, any pair of variants V_v and V_w of any trait T, and any variant $V_{v^*}^*$ of any trait T^*, the probability that O_i is killed by S_P given that O_i has variant V_v, and given that the appropriate instance of NC_P is true is equal to the probability that O_j is killed by S_P given that O_j has variant V_w, and given that the appropriate instance of NC_P is true.*

That is, the following is true: in any population P in any generation G in which the conditions described in statement 4.1 obtain,

$$\forall(Q_P)\,(p(K[S_P, O_i]|O_i[V_v]\ \&\ NC_P) = p(K[S_P, O_j]|O_j[V_w]\ \&\ NC_P)).$$

(A quantifier for V_v^* is necessary because it is implicitly referred to in NC_P.)

This states that, *in the absence of confounding correlations*, the probabilities that organisms with variants of the trait in question are killed by the sampling agent in question

are equal. Note that I am not suggesting that indiscriminate parent sampling can only occur in environments and among organisms in which there are no such confounding correlations. Rather, what I am suggesting is that any probabilities used in determining whether a particular parent sampling agent is indiscriminate must be conditional on the absence of such confounding correlations. Even if such correlations exist, they can be accounted for using statistical models that "subtract out" the effects of such correlations on the probability of survival of organisms with a given set of variants. This is what I mean to accomplish by allowing that the no-correlation condition is satisfied if there is a correlation between traits whose effects on the relevant probabilities can be accounted for, statistically; I have called this "virtual satisfaction" of NC_P, above.

Relativizing to a time scale

At this point, I have shown that the core equality can be protected from counterexamples having to do with correlated traits. Nevertheless, the sophisticated statement of the equality is not a necessary condition for indiscriminate parent sampling, because it requires further qualification. In this section, I would like to describe the further qualification I have in mind.

The need for qualification arises in response to a further objection to the claim that the core equality, even in its sophisticated form, is not necessary for indiscriminate parent sampling. The objection, raised by Peter Achinstein (pers. comm.), is constructed around cases in which the relevant probabilities are not precisely equal to one another, but are very nearly so. The objection can be formulated in terms of the following statement,

which refers to the colorblind predators example.[3]

Statement 4.6 *Among the colorblind predators and their prey, in some generation,*

$$\forall(O)\,\forall(V^*)\,(p(K[Pr,O_i]|O_i[D]\ \&\ NC_P) = \frac{999}{1000}\,p(K[Pr,O_j]|O_j[L]\ \&\ NC_P)).$$

The only difference between the instance of the core equality that would describe the colorblind predators example (statement 4.3 supplemented with condition N_{CP}) and statement 4.6 is that, in the latter, the term to the right of the equality symbol is multiplied by 999/1000.

Now, let me elaborate the objection by reference to statement 4.6. Surely parent sampling by predation in the population and generation in question is indeed indiscriminate, the factor of 999/1000 notwithstanding. Such a small difference in probability is insignificant. If the core equality is to be taken as a necessary condition for indiscriminate parent sampling, it is trivializing. The claim is that if someone insisted that the sophisticated statement of the equality is necessary for indiscriminate parent sampling, he or she would have to accept that there are hardly any cases of it. This is because there are probably very few cases in which the across-variant probabilities of death due to a given parent sampling agent are precisely equal to one another—that is, there are probably very few cases in which the sophisticated statement of the core equality is satisfied, if such cases even exist at all.

To state this objection more precisely, and to provide some useful background to my answer to it, consider the following new term.

Definition 4.14 *"$\triangle P^P$" should be understood in accord with its role in the following state-*

---

[3]See definitions 4.8 - 4.10 and statement 4.3 (pages 131 − 132).

ment:

$$\forall(Q_P)\,(\triangle P^P = p(K[S_P, O_i]|O_i[V_v] \,\&\, NC_P) - p(K[S_P, O_j]|O_j[V_w] \,\&\, NC_P)),$$

such that, if they differ, the larger of the two probabilities is placed to the left of the "−"

sign, so that $\triangle P^P \geq 0$.

To be clear, $\triangle P^P$ is just the difference between the two probabilities that appear on either side of the equality symbol in the core equality, providing a measure of the distance between those probabilities.

This can be used to state the objection as follows. To claim that the core equality is necessary for indiscriminate parent sampling without further modification is to claim that, in all cases, $\triangle P^P = 0$. However, it is absurd to claim this, because there are very few cases, if there are any at all, in which $\triangle P^P = 0$; that is, there are very few cases in which the relevant probabilities are precisely equal to one another. Claiming that $\triangle P^P = 0$ across all cases commits one to the view that there are hardly any cases of indiscriminate parent sampling.

Now, I would like to elaborate my response to this objection, framed in the terms I have just introduced. I believe that there are in fact some deviations from $\triangle P^P = 0$ that are trivial, in the sense that they are not significant enough for parent sampling to be discriminate. I would put the matter as follows. For every case of parent sampling, there is some maximum value $\triangle P^P_{max}$ of $\triangle P^P$ such that, if $\triangle P^P \leq \triangle P^P_{max}$, that case of parent sampling is indiscriminate, even though the core equality is not satisfied. Note that my suggestion is not that $\triangle P^P_{max}$ is a constant value; rather, my suggestion is that it is a function of a range of parameters describing the population at issue.

140

The next question I want to consider—a question which follows quite naturally from the discussion of $\triangle P_{max}^P$ in the previous paragraph—is, What function determines what $\triangle P_{max}^P$ is, for any given case? Answering this question in a general way is beyond the scope of this chapter. However, I would like to suggest, in broad terms, how I think $\triangle P_{max}^P$ is best understood. The central consideration is that $\triangle P_{max}^P$ must be small enough so that it is highly improbable that the parent sampling agent in question causes any trends to arise in the evolution of the population in question. Although there may be other parameters on which $\triangle P_{max}^P$ depends, such as the overall rate of mortality in the population, I see $\triangle P_{max}^P$ as primarily a function of the *length of time* at issue. To see this, consider the following example.

Suppose that $\triangle P^P$ is minute, as in statement 4.6 above (page 139), which concerns the organisms preyed upon by colorblind predators, and in which $\triangle P^P = 1/1000$. Over a short period of time, no trend would be expected to arise in the population, i.e., there is a low probability that, after a short period of time, either kind of organism would dominate the population. However, over a very long period of time—tens or hundreds of thousands of generations—dark organisms would be expected to increase in frequency over light organisms, *ceteris paribus*. This is because they have an advantage of 1/1000 over their lighter-colored conspecifics which would begin to manifest itself, given enough time, in the absence of other confounding influences. Over a long period of time, even very small values of $\triangle P^P$ will exceed $\triangle P_{max}^P$; over shorter periods of time, small values are unimportant, and will fail to exceed $\triangle P_{max}^P$.

To be clear, the account of indiscriminate parent sampling that I am developing

is intended to apply to a single generation, that is, my account is intended to describe necessary and sufficient conditions under which parent sampling occurring in the time span of a given generation is indiscriminate. My proposal is that each generation should be viewed in the context of a longer time span, so that the issue of whether parent sampling in a given generation is indiscriminate depends in part on the length of the larger time span at issue.

I integrate the idea that parent sampling is indiscriminate relative to a time scale into my account by making use of a new notion: satisfying the core equality *relative to a time scale* Θ. The idea is that the core equality is satisfied relative to a time scale Θ if and only if, in the case at issue, $\triangle P^P \leq \triangle P^P_{max}$. This means that the across-variant probabilities of being killed by the parent sampling agent in question in the case in question do not differ from one another in an amount greater than $\triangle P^P_{max}$. More precisely, I understand this notion as follows.

Statement 4.7 *The sophisticated statement of the core probabilistic equality describing any generation G, any population P, any parent sampling agent S_P, any pair of organisms O_i and O_j, any pair of variants V_v and V_w of any trait T, and any variant $V^*_{v^*}$ of any trait T^* is relativized to time scale Θ if and only if, in the episode of parent sampling in question, the following obtains:*

$$\triangle P^P \leq \triangle P^P_{max}.$$

Having indicated how I believe the core probabilistic equality should be qualified in order to relativize it to a time scale, I am now in a position to state a necessary condition for indiscriminate parent sampling.

Statement 4.8 (Necessary condition) *In any population P in any generation G in which the conditions described in statement 4.1 obtain, for any time frame Θ, and for any parent sampling agent S_P, any pair of organisms O_i and O_j, any pair of variants V_v and V_w of any trait T, and any variant $V_{v^*}^*$ of any trait T^*, parent sampling is indiscriminate only if the sophisticated statement of the core probabilistic equality, applied to parent sampling, relativized to the time frame Θ, is satisfied.*

4.2.3   A further condition

A further condition, if added to statement 4.8, forms a complete set of necessary and sufficient conditions for indiscriminate parent sampling. This further condition is needed to rule out spurious parent sampling agents. In successive subsections of this section, I elaborate and argue for this condition; state the completed probability theory of indiscriminate parent sampling; and make an additional remark concerning heritability.

A further condition

In order to explain what spurious parent sampling agents are, and why they present a problem for the claim that statement 4.8 is sufficient for indiscriminate parent sampling, let me present an example, which is constructed around the following definitions.[4]

Definition 4.15 *"U_i" refers to the ith member of an undiscovered species of beetle, whose home is in the Amazon; "U_j" refers to the jth member of that same species.*

Definition 4.16 *"$K[Y, U_i]$" means "By winning the 2006 World Series, the New York*

---

[4]Note that, as well as the terms I introduce here, I will make use of terms already defined; see definitions 4.1 - 4.7 (pages $129 - 130$).

Yankees kill organism U_i," where there are U beetles in the species, and where $i = 1 \ldots U$;
and *"$K[Y, U_j]$" means "By winning the 2006 World Series, the New York Yankees kill
organism U_j,"* where $j = 1 \ldots U$.

Definition 4.17 *"$U_i[L]$" means "Undiscovered beetle U_i is large."*

Definition 4.18 *"$U_j[S]$" means "Undiscovered beetle U_j is small."*

Definition 4.19 *"$U_j[V^*]$" means "Undiscovered beetle has variant V^* of trait T^*."*

Definition 4.20 *"$\forall(U)$" means "$\forall(U_i)\,\forall(U_j)$".*

According to statement 4.8, the necessary condition for indiscriminate parent sam-
pling I state above (page 143), the Yankees' victory in the 2006 World Series indiscriminately
samples the undiscovered beetles on the basis of their size only if the following instance of
the core probabilistic equality is satisfied, relative to the appropriate time scale.

Statement 4.9 *In the undiscovered beetles population in the generation of interest,*

$$\forall(U)\,\forall(V^*)\,(p(K[Y, U_i]|U_i[S]\,\&\,NC_P) = p(K[Y, U_j]|U_j[L]\,\&\,NC_P)).$$

I would now like to argue that statement 4.9 is, in fact true. The first step in my
argument is to point out that, as it happens, the following is the case:

Statement 4.10 *In the undiscovered beetles population in the generation of interest,*

$$\forall(U)\,\forall(V^*)\,(p(K[Y, U_i]|U_i[S]\,\&\,NC_P) = p(K[Y, U_j]|U_j[L]\,\&\,NC_P) = 0).$$

The idea is that the beetles are completely causally isolated from the Yankees'
victory, and so it has no probability whatever of causing any of the undiscovered beetles to

144

die, i.e., as indicated in statement 4.10, the relevant probabilities are equal to zero. It follows from that that the core probabilistic equality is satisfied in their case, and that, furthermore, the necessary condition for indiscriminate parent sampling expressed in statement 4.8 is satisfied.

To conclude my argument, what I would like to point out is that, the truth of statement 4.9 and its consequences notwithstanding, there is no reason to regard the Yankees' victory as an indiscriminate parent sampling agent. For instance, suppose that the blindfolded person who is supposed to select beads from the urn refuses to do so. This does not make the person an indiscriminate sampler; for this, it is required that the person actually attempt to sample beads. In this case, as in the beetles' case, the putative sampling agent satisfies the core equality trivially, because it is a vacuous case.

I propose to remedy this situation by adding the following condition to statement 4.8 as a further necessary condition for indiscriminate parent sampling.

Statement 4.11 *In any population in any generation, for any parent sampling agent S_P, and for any organism O_i, S_P indiscriminately samples O_i on the basis of any trait T only if*

$$p(K[O_i, S_P]) \neq 0.$$

Statement 4.11 rules out the Yankees case and other similar cases by limiting indiscriminate parent sampling agents to events that have some probability of causing the death of an organism in the population.

In conclusion to this discussion of the further condition that I believe is required for indiscriminate parent sampling, I want to mention that the initial impetus for my thinking

on this issue was provided by Timothy Shanahan's [132, 145 - 146] remark that "location with regard to Greenland, Margaret Thatcher, and the Hollywood Bowl" might be claimed by someone to be indiscriminate parent sampling agents. (It is unclear whether Shanahan himself wants to make such a claim.) By making this remark, he seems to be intending to point out that many objects irrelevant to the fate of any of the organisms in a given population can serve, or can be claimed to serve, as indiscriminate parent sampling agents of the organisms in the population in question. If this is the case, it seems to me that the notion of indiscriminate parent sampling is hopelessly flawed—a state of affairs I intend for my work in this section to show does not obtain.

The completed theory

Finally, then, taking into consideration the various modifications I have introduced to the naïve statement of the core probabilistic equality, the probability account of indiscriminate parent sampling may be stated as follows.

Analysis 4.1 *In any population P in any generation G in which the conditions described in statement 4.1 obtain, for any time frame Θ, and for any parent sampling agent S_P, any pair of organisms O_i and O_j, any pair of variants V_v and V_w of any trait T, and any variant $V_{v^*}^*$ of any trait T^*, parent sampling is indiscriminate if and only if the following conditions obtain.*

1. *The relevant instance of the sophisticated formulation of the core probabilistic equality, applied to parent sampling, relativized to the time frame Θ, is satisfied; and*

2. *$p(K[O_i, S_P]) \neq 0$.*

A note about heritability

To conclude my account of indiscriminate parent sampling in this section, I want to make a few remarks about heritability in connection with indiscriminate parent sampling. I require heritability for indiscriminate parent sampling, although I do not consider it essential. Whether indiscriminate sampling requires heritability depends upon whether indiscriminate sampling is sufficient for evolution. This is because evolution requires heritability; for within-generation changes to be transmitted across generations, they must have some hereditary basis. I regard the difference between evolutionary and non-evolutionary conceptions of indiscriminate parent sampling to be stylistic, and I see the choice about which conception to use as pragmatic. Drawing on Endler's [43, ch. 1] discussion of similar issues, let me explain why I hold this view.

If supplemented by a separate theory of heritability, it is possible to employ a non-evolutionary conception of indiscriminate sampling to predict and explain evolution. When observing within-generation change, heritable change is not distinguished from non-heritable change; the evolutionary effect of within-generation change is determined by a heritability factor. This is the approach often taken by quantitative geneticists, who often do not have access to the underlying genetics of the traits they study. For them, a non-evolutionary conception of indiscriminate sampling is warranted, as it suits the aims and limitations of their studies.

Generally speaking, however, there is no reason for a population geneticist to separate indiscriminate sampling from heritability; alleles—the heritable entities generally studied in population genetics—are intrinsically heritable. My decision to require heritabil-

ity for indiscriminate parent sampling reflects my focus on theories of population genetics.[5]

4.3 Indiscriminate Gamete Sampling

This section forms a companion to the previous one: here, I articulate the probability account of indiscriminate gamete sampling. I proceed in this section in a series of further subdivisions that parallel those of the previous section: I provide the definitions needed for stating the theory (section 4.3.1); the core probabilistic equality, applied in the context of gamete sampling (section 4.3.2); and a response to an objection analogous to the "Yankees counterexample" presented above (section 4.3.3). Because many of the arguments and explanations are precisely analogous to those I presented above in the context of indiscriminate parent sampling, I proceed more directly and in a less circumspect manner in the present section than I do in the previous one. I illustrate the account by applying it to Mendelian reproduction.

4.3.1 Essential background

As in the case of the probability account of indiscriminate parent sampling, the present theory requires the formulation of a series of definitions, which are as follows. I intend for the following definitions to apply implicitly to any generation G in any population P meeting the conditions described in statement 4.1.

Definition 4.21 *"The gamete pool of population P in generation G" designates the collection of gametes borne by all organisms alive at the start of the mating season in population*

---

[5]My position on this issue represents a middle ground between Shanahan's [132, 142] and Millstein's [105, 37].

P in generation G.

Definition 4.22 *"G_i" refers to the ith gamete in the gamete pool, where there are N_G gametes in the gamete pool, and where $i = 1 \ldots N_G$; and "G_j" refers to the jth gamete in the gamete pool, where $j = 1 \ldots N_G$.*

Definition 4.23 *"A_a" refers to the ath allele at a gene locus L, where there are a total of A variants in the population, and where $a = 1 \ldots A$; "A_b" refers to the bth allele at L, where $b = 1 \ldots A$.*

Definition 4.24 *"$G_i[A_a]$" means "Gamete G_i has allele A_a;" and "$G_i[A_b]$" means "Gamete G_i has allele A_b."*

Definition 4.25 *"S_G" refers to a mating system structure.*

Definition 4.26 *"$T[S_G, G_i]$" means "Gamete G_i is transmitted through mating system structure S_G," that is, that gamete G_i is not stopped from being passed on to the next generation by mating system structure S_G; and "$T[S_G, G_j]$" means "Gamete G_j is transmitted through mating system structure S_G," that is, that gamete G_j is not stopped from being passed on to the next generation by mating system structure S_G.*

Definition 4.27 *"$\forall(G)$" means "$\forall(G_i)\,\forall(G_j)$".*

Definition 4.28 *"$\forall(A)$" means "$\forall(A_a)\,\forall(A_b)$".*

4.3.2 The core probabilistic equality

As in my discussion of indiscriminate parent sampling, I develop the core probabilistic equality in a series of further subsections, in which I aim at achieving the following

goals, respectively: formulating a naïve statement of the equality; formulating a sophisticated statement of the equality; and indicating how it should be relativized to a time scale.

A "naïve" statement of the core equality

The naïve statement of the core equality, as it applies to indiscriminate gamete sampling, is as follows.

Statement 4.12 (Core equality, naïve) *In any population P in any generation G in which the conditions described in statement 4.1 obtain, for any structure of the mating system S_G, any pair of gametes G_i and G_j, and any pair of alleles A_a and A_b, at any gene locus L, the probability that G_i is transmitted through S_G given that G_i has A_a is equal to the probability that G_j is transmitted through S_G given that G_j has A_b.*

That is, the following is true: in any population P in any generation G in which the conditions described in statement 4.1 obtain,

$$\forall(S_G)\,\forall(G)\,\forall(A)\,(p(T[S_G, G_i]|G_i[A_a]) = p(T[S_G, G_j]|G_j[A_b])).$$

This reflects the central property of the broader phenomenon of indiscriminate sampling: gamete sampling in the gamete pool of a given population in a given generation is indiscriminate, if and only if the allele borne by a gamete at a given gene locus does not make any difference to whether that gamete is transmitted to the next generation, in comparison with other gametes in the gamete pool, *ceteris paribus*. If and only if gametes are sampled indiscriminately, an allele is analogous to the color of a bead in the urn, in the sense that the allele borne by a gamete does not have any causal influence on whether it is

passed on to the next generation; this is just as the blindfolded person cannot use the color of a bead in his or her deliberations about which bead to select from the urn on a given draw. Statement 4.12 describes conditions under which this occurs by setting up an equality between gametes' probabilities of being transmitted through to the next generation, each probability conditionalized on each gamete's bearing an alternative allele at the locus in question.

To see how this works, consider Mendelian reproduction, a paradigm case of indiscriminate gamete sampling. Mendelian reproduction requires that a population be in a variety of particular states, all at the same time. These states include mating randomly (organisms do not exhibit mate preferences), the absence of mutation among sex cells, the absence of gene linkage, and the "fair" production of gametes by meiosis in all organisms. Meiosis is "fair" in a given heterozygotic organism if and only if the relative frequency of each allele among the gametes contributed to the gamete pool by that organism is 50%; for Mendelian reproduction to occur, meiosis must be fair for all heterozygotes in it. Homozygotes only carry one allele, so the relative frequency of this allele (barring mutation) among its gametes is always 100%.

Note that I do not claim that the Hardy-Weinberg equilibrium is required for Mendelian reproduction. The Hardy-Weinberg equilibrium obtains in a sexually reproducing, diploid population if and only if, in addition to Mendelian reproduction, a variety of other conditions obtain, for instance, the absence of natural selection, mutation, or migration during the adult stage of the life cycle. Mendelian reproduction is a mating process, and does not place requirements on other events of the life cycle. I believe that this use of

"Mendelian reproduction" accords with other uses of this phrase and other similar phrases that appear in the literature, for instance, Sewall Wright's use of "Mendelian population" in his famous "Evolution in Mendelian Populations" [156]. Additionally, note that the predicate phrase "Mendelian reproducer" applies to a gene locus: a population is Mendelian *at a given gene locus*. Accordingly, a population may be Mendelian at some loci, but not Mendelian at others.

As I have suggested at various points above, indiscriminate parent sampling requires that whether a gamete bears one allele rather than another does not influence whether it is passed to the next generation, and I would now like to explain why Mendelian reproduction exemplifies this property. Broadly speaking, it does so because, during Mendelian reproduction, the gametes are physically symmetrical to one another—the alleles they bear do not alter any of their characteristics that are causally relevant to their (the alleles) being passed on. For example, if mating is random, organisms do not choose mates on the basis of the allele they bear. As a consequence, they do not influence any of the the causal processes of mate choice that determine whether they are passed on. The fairness of meiosis has a similar effect. Alleles do not play any causal role in determining how many gametes they are distributed into, or which gametes they are distributed into.

Having provided this background about Mendelian reproduction, let me state an instance of the naïve statement of the core equality that applies to it. The following definitions are used in the statement.

Definition 4.29 *"$T[M, G_i]$" means "Gamete G_i is transmitted to the next generation by Mendelian reproduction;" "$T[M, G_j]$" means "Gamete G_j is transmitted to the next gener-*

ation by Mendelian reproduction."

Definition 4.30 *"$G_i[A_1]$" means "Gamete G_i bears allele A_1;" "$G_i[A_1]$" means "Gamete G_j bears allele A_2."*

The appropriate instance of the naïve statement of the core equality, applied to Mendelian reproduction, is as follows.

Statement 4.13 *In population P_M, in generation G_M, for any pair of gametes G_i and G_j, the probability that G_i is transmitted to the next generation by Mendelian reproduction given that G_i bears allele A_1 is equal to the probability that G_j is transmitted to the next generation by Mendelian reproduction given that G_j bears allele A_2.*

That is, the following is true, for population P_M, in generation G_M:

$$\forall(G)\,(p(T[M,G_i]|G_i[A_1]) = p(T[M,G_j]|G_j[A_2])).$$

A "sophisticated" statement of the core equality

As with indiscriminate parent sampling, a sophisticated treatment of the core equality (as opposed to the naïve treatment I have just developed) is required in the case of indiscriminate gamete sampling. The problem is that—again, in parallel with indiscriminate parent sampling—examples of what I term "correlated alleles" show that the naïve statement of the core equality does not hold for indiscriminate gamete sampling. I begin my exposition of the sophisticated statement of the core equality for indiscriminate gamete sampling with a discussion of correlated alleles; this parallels my exposition of the sophisticated statement of the core equality for indiscriminate parent sampling, which begins with a discussion of correlated variants (see page 133).

Suppose that gametes bearing an allele A_1 at a gene locus L also tend to bear an allele A_1^* at a different locus L^*, while gametes that have alleles other than A_1 at L tend not to have A_1^* at L^*. If this obtains, I will say that the alleles A_1 and A_1^* are "correlated." This may be defined for the general case as follows.

Definition 4.31 (Correlated alleles) *In a population of interest, in a generation of interest, allele A_1 at locus L is correlated with allele A_1^* at locus L^* if and only if, in the population of interest, in the generation of interest, for any pair of gametes G_i and G_j, the probability that G_i bears A_1^* given that G_i has A_1 is greater than the probability that G_j has A_1^* given that G_j does not have allele A_1.*

That is, the correlation obtains if and only if, in the population of interest, in the generation of interest, the following obtains:

$$\forall (G) \, (p(G_i[A_1^*]|G_i[A_1]) > p(G_j[A_1^*]|\neg G_i[A_1])).$$

To see how this works, consider a case in which gene linkage together with so-called "meiotic drive" operate to determine the fate of the alleles at a locus of interest. When meiotic drive occurs, gametes with what is termed a "driver" allele disable gametes with the alternative gene, so that the latter have no chance of being passed on: as Crow [27] puts it, these driver genes "violate Mendel's rules," destroying the "fairness" of Mendelian reproduction. Any genes that are linked to the driver genes will have their pattern of inheritance determined in part by the drivers. So even though a linked gene might be entirely passive as it moves through meiosis, causing no changes in its bearer's probability of being passed on, it might exhibit striking trends in its pattern of inheritance. It would not be correct to say that it is either a driver gene or the disabled "victim" of a driver gene;

its main mode of inheritance, rather, is reduction-division (meiosis) and random mating, in addition to linkage. These relationships are further illustrated on table 4.2.

Table 4.2: Example of correlated alleles

All and only gametes with ...	Also have a ...	And so have a ...
Allele A_1	"driver" allele	high prob. of being passed on
Allele A_2	"victim" allele	low prob. of being passed on

The argument that this case is a counterexample to the claim that the naïve statement of the core equality (statement 4.12, page 150) is a necessary condition for indiscriminate gamete sampling is as follows. The A_1 and A_2 alleles in the example above do not satisfy statement 4.12, because their fate is tied to that of the drivers and victims, respectively. The probability that an A_1 allele is passed on depends on the probability that the driver to which it is linked is passed on. This differs from the probability that an A_2 allele is passed on, which is linked to a victim. All this notwithstanding, I think it is clear that the A_1 and A_2 alleles are sampled indiscriminately. Neither has any causal responsibility for the difference in probability between them; each is passive, so to speak, in its connection to the driver and victim alleles.

Having pointed out the kinds of cases that the naïve statement of the core equality cannot handle when applied to indiscriminate gamete sampling, I would now like to formulate a no-correlation condition, analogous to the no-correlation claim that I introduced in statement 4.4 (page 136). I need to introduce a further definition, to describe alleles in correlated pairs.

Definition 4.32 "$A^*_{a^*}$" refers to the a^*th allele at locus L, where there are a total of A^*

alleles at that locus, and where $a^* = 1 \ldots A^*$.

Now, let me introduce the no-correlation condition.

Statement 4.14 (Condition NC_G) *For any pair of gametes G_i and G_j, and for all alleles $A^*_{a^*}$ at any gene locus L^*, allele A_1 of locus L is not correlated with $A^*_{a^*}$. That is, the following obtains:*

$$\forall(G)\,\forall(A^*_{a^*})\,(p(G_i[A^*_{a^*}]|G_i[A_1]) = p(G_j[A^*_{a^*}]|\neg G_j[A_1])).$$

Incorporating this into statement 4.12, the following obtains, which is the sophisticated formulation of the core equality, as it applies to indiscriminate gamete sampling. A further definition is required, for economy of expression.

Definition 4.33 *"$\forall(Q_G)$" means "$\forall(S_G)\,\forall(G)\,\forall(A)\,\forall(A^*_a)$".*

Statement 4.15 (Core equality, sophisticated) *For any structure of the mating system S_G, any pair of gametes G_i and G_j, any pair of alleles A_a and A_b at any gene locus L, and any allele $A^*_{a^*}$ at any gene locus L^*, the probability that G_i is transmitted through S_G given that G_i has A_a, and given that the appropriate instance of NC_G is true is equal to the probability that O_j is transmitted through S_G given that G_j has A_b, and given that the appropriate instance of NC_G is true. That is, the following is true:*

$$\forall(Q_G)\,(p(T[S_G, G_i]|G_i[A_a]\ \&\ NC_G) = p(T[S_G, G_j]|G_j[A_b]\ \&\ NC_G)).$$

Relativizing to the time scale

Just as with my account of indiscriminate parent sampling, the objection can be raised to my account of indiscriminate gamete sampling that it is trivializing. If the claim

is that the core probabilistic equality is a necessary condition for indiscriminate gamete sampling, the objection goes, this amounts to a claim that there are hardly any cases of indiscriminate gamete sampling. This is because the core probabilistic equality is hardly ever satisfied, as there are very few cases, if any, in which the relevant probabilities are precisely equal to one another. Just as in the case of indiscriminate parent sampling, I believe that the proper response to this objection is to relativize the sophisticated formulation of the equality to an appropriate time scale. The central principle behind this response is that gamete sampling is indiscriminate as long as the difference between the relevant probabilities is low enough so that there is a very low probability that any trend can arise during the time interval in question. In order to articulate this response in a precise manner, I need to define some terminology analogous to that articulated in definition 4.14 (pages 139 – 140).

Definition 4.34 *"$\triangle P^G$" should be understood in accord with its role in the following statement:*

$$\forall (Q_G)\,(\triangle P^G = p(T[S_G, G_i]|G_i[A_a]\,\&\,NC_G) - p(T[S_G, G_j]|G_j[A_b]\,\&\,NC_G)),$$

such that, if they differ, the larger of the two probabilities is placed to the left of the "$-$" sign, so that $\triangle P^G \geq 0$.

If and only if the core probabilistic equality is satisfied, $\triangle P^G = 0$. My view is that there are some deviations from $\triangle P^G = 0$ that are not significant, that is, that are too small for gamete sampling to be discriminate. As in the case of indiscriminate parent sampling, some additional terminology is useful for formulating this claim. For every case of gamete sampling, there is some maximum value $\triangle P^G_{max}$ of $\triangle P^G$ such that, if $\triangle P^G \leq \triangle P^G_{max}$, that

157

case of gamete sampling is indiscriminate, even though the core equality is not satisfied. As with the similar claim concerning parent sampling, my view is that $\triangle P_{max}^G$ is not a constant, but rather, is a function of a range of parameters describing the population at issue.

As I suggested in the case of the similar claim concerning parent sampling, providing a general account of how $\triangle P_{max}^G$ is determined is beyond the scope of this chapter; and as I have already suggested, I believe that $\triangle P_{max}^G$ depends in a critical way on the time scale. In order to integrate this into my theory, I want to introduce the idea that the core probabilistic equality as it applies to gamete sampling can be *satisfied relative to a time scale* Θ. In a given case, the core equality is satisfied relative to a time scale Θ if and only if, in the case at issue, $\triangle P^G \leq \triangle P_{max}^G$. This means that the across-allele probabilities of being transmitted by the relevant mating system structure in the case at issue do not differ from one another in an amount greater than $\triangle P_{max}^G$. This may be stated more precisely as follows.

Statement 4.16 *The sophisticated statement of the core probabilistic equality describing any generation G, any population P, any mating system structure S_G, any pair of gametes G_i and G_j, any pair of alleles A_a and A_b of any gene locus L, and any allele A_{a*}^* of any gene locus L^* is* relativized to time scale Θ *if and only if, in the episode of gamete sampling in question, the following obtains:*

$$\triangle P^G \leq \triangle P_{max}^G.$$

Having appropriately modified and qualified the core probabilistic equality as it applies to indiscriminate gamete sampling, I would now like to state a necessary condition

158

for indiscriminate gamete sampling constructed around the core equality.

Statement 4.17 (Necessary condition) *In any population P in any generation G in which the conditions described in statement 4.1 obtain, and for any time frame Θ, any mating system structure S_G, any pair of gametes G_i and G_j, any pair of alleles A_a and A_b of any gene locus L, and any allele A_{a*}^* of any gene locus L^*, gamete sampling is indiscriminate only if the core probabilistic equality, applied to gamete sampling, relativized to the relevant time frame Θ, is satisfied.*

4.3.3 A further condition

In this section, I introduce a further condition, required to rule out spurious gamete sampling agents, that completes the set of necessary and sufficient conditions for indiscriminate gamete sampling; I then state the completed theory. I carry out these tasks in further subdivisions of this section, respectively. Note that, in contrast to the case of indiscriminate parent sampling, no discussion of heritability is required, because alleles are intrinsically heritable—indiscriminate gamete sampling necessarily occurs among heritable entities.[6]

A further condition

I present an example that illustrates the problem that spurious, trivializing gamete sampling agents pose for the claim that the necessary condition for indiscriminate gamete sampling that I propose above (statement 4.17) is both necessary and sufficient. I begin with a definition that is required to formulate the example; the definition describes a term

---

[6]See pages 147 - 148 for my discussion of requiring heritability for indiscriminate parent sampling.

that indicates a putative structure of the mating system of the Amazon beetles' population discussed above in connection with indiscriminate parent sampling, in section 4.2.3 (page 143).[7]

Definition 4.35 *"$T[Y, G_i]$" means "By winning the 2006 World Series, the New York Yankees transmit gamete G_i to the next generation;" and "$T[Y, G_j]$" means "By winning the 2006 World Series, the New York Yankees transmit gamete G_j to the next generation."*

According to statement 4.17, the Yankees' victory in the 2006 World Series indiscriminately samples the gametes of the population in question on the basis of which allele they bear at the L locus only if the following instance of the core probabilistic equality is satisfied, relative to the appropriate time scale.

Statement 4.18 *In the Amazon beetles' population, in the generation of interest,*

$$\forall(G)\,\forall(V^*)\,(p(T[Y, G_i]|G_i[A_1]\;\&\;NC_G) = p(T[Y, G_j]|G_j[A_2]\;\&\;NC_G)).$$

Note that I frame the example in terms of the arbitrary alleles A_1 and A_2, which correspond to the large and small size variants I used to frame the corresponding example for indiscriminate parent sampling in terms of (page 144).

To add a further detail to the example, I want to suppose that the beetles are a Mendelian population at locus L (see pages 151 - 153). This means that the following instance of the sophisticated statement of the core probabilistic equality is satisfied, relative to the appropriate time scale:

---

[7]Note that, as well as the additional term I define here, I will make use of some terms already defined; see definitions 4.21 - 4.28 (pages 148 – 149).

Statement 4.19 *In the Amazon beetles' population, in the generation of interest,*

$$\forall(G)\,\forall(V^*)\,(p(T[Y,G_i]|G_i[A_a]\ \&\ NC_G) = p(T[Y,G_j]|G_j[A_b]\ \&\ NC_G)).$$

This just states that the core probabilistic equality, relativized to the appropriate time scale, is satisfied in this case by the Yankees' victory, which satisfies the necessary condition for indiscriminate gamete sampling expressed in statement 4.17. Nevertheless, what I would now like to argue is that this does not mean that the Yankees' victory should be regarded as an indiscriminate gamete sampling agent—that is, that statement 4.17 is not a sufficient condition for indiscriminate gamete sampling. As in the case of indiscriminate parent sampling, the Yankees' victory trivially satisfies the relevant instance of the core probabilistic equality, because it represents a vacuous case. The argument for this is as follows.

The satisfaction of the instance of the core probabilistic equality I formulate above in statement 4.19 depends entirely upon the fact that the beetle population is Mendelian at the locus L. Because the beetles are causally isolated from the Yankees, the victory of the latter in the 2006 World series has no influence on them; this extends to the success of gametes in mating, just as it extends to the fate of organisms in the population. This means that the Yankees' victory does not alter the relevant probabilities, which remain equal (relative to the time frame in question) to one another regardless of the outcome of the game. However, it seems clear that, for a structure of the mating system to be an indiscriminate gamete sampling agent, it must have some causal influence on the population in question: the structure must actually sample gametes.

The analogy with the urn case is as follows. Suppose that, during the interval of time that it takes the blindfolded person to draw a series of beads from the urn in an indiscriminate manner, the beetles are mating. This means that an instance of the core probabilistic equality (relativized to an appropriate time frame) can be constructed in which the beetles are construed as an indiscriminate sampler of the beads. This is absurd, of course, because the beetles are causally isolated from the urn and the person drawing the beads from it.

I propose to remedy this situation by adding the following condition to statement 4.17 as a further necessary condition for indiscriminate gamete sampling.

Statement 4.20 *In any population in any generation, for any structure of the mating system S_G, and for any gametes G_i, S_M indiscriminately samples O_i on the basis of any allele A_a only if*

$$p(T[S_G, G_i]) < 1.$$

The idea behind this condition is that, for a structure of the mating system S_G to be an indiscriminate gamete sampling agent, there must be some chance that S_G *prevents* a gamete from being passed on, i.e., the probability that S_G transmits gametes to the next generation must be less than unity. This condition is clearly not satisfied in the case of the Yankees and the beetles: there is no chance that the Yankees' victory in the 2006 World Series will prevent any gamete from being passed on, because the outcome of the Series has no causal connection with the fate of any gamete in the beetles' population. In contrast, the condition expressed in statement 4.20 is satisfied by Mendelian reproduction; some of

162

the beetles' gametes will fail to be "picked" in the "lottery of sex."[8]

The completed theory

The probability account of indiscriminate gamete sampling is as follows.

Analysis 4.2 *In any population P in any generation G in which the conditions described in statement 4.1 obtain, for any time frame Θ, for any structure of the mating system S_G, any pair of gametes G_i and G_j, any pair of alleles A_a and A_b at any gene locus L, and any allele $A^*_{a^*}$ of any gene locus L^*, gamete sampling is indiscriminate if and only if the following conditions obtain.*

1. *The relevant instance of the sophisticated formulation of the core probabilistic equality, applied to gamete sampling, relativized to the time frame Θ, is satisfied; and*

2. *$p(T[S_G, G_i]) < 1$.*

4.4   Indiscriminate Sampling and Evolution

My efforts in this chapter so far have concentrated on clarifying the notion of indiscriminate sampling, and this has taken me away from the discussion of drift itself. I would like to remedy that state of affairs in the present section, in which my main aim is to describe how I see the relationship between random drift and indiscriminate sampling. As well, I consider the relationship of random drift to evolution. The highlight of my view is that indiscriminate sampling is sufficient but not necessary for drift, which I understand

---

[8]As in the case of indiscriminate parent sampling, the thinking that led me to believe that statement 4.20 is necessary for indiscriminate gamete sampling was given its initial impetus by some comments of Shanahan's (see page 146 for further details on this point).

pluralistically. I articulate this view in section 4.4.1. In section 4.4.2, I assess the impact of the claim that indiscriminate sampling is a mechanism of evolution on some discussions of evolution, drift, and natural selection in recent literature in philosophy of biology.

4.4.1 A pluralistic view of drift

To begin my exposition of my pluralistic view of drift, I want to state my view on a fundamental point.

Statement 4.21 (Ind. sampling and drift) *Evolution occurring by indiscriminate sampling is evolution by accident in precisely the sense required for random drift.*

I think that statement 4.21 should be easy to accept because indiscriminate parent and gamete sampling are structurally so similar to the paradigm case of drawing beads from an urn, as well as other canonical chance processes such as tossing fair coins and drawing cards from a well-shuffled deck. I believe that the core probabilistic equality can be applied to these canonical chance processes, if interpreted for the appropriate entities and events, perhaps requiring sophisticated (as opposed to naïve) formulations, analogous to those I describe above, that are appropriate for beads, coins, or cards. I will not argue for statement 4.21 here, because I believe that it has such strong intuitive plausibility, and because doing so would add considerable length and complexity to the chapter.

Now I would like to describe my views about random drift, understood broadly, beginning with a statement of my understanding of how indiscriminate sampling figures into a complete description of the concept of random drift. I see indiscriminate sampling as sufficient for random drift, in the sense that if either indiscriminate parent sampling

or indiscriminate gamete sampling occurs in a population of the appropriate kind, then random drift of the relevant entities (variants or alleles) occurs in that population, in that generation. The following represents a more precise statement of this claim.

Statement 4.22 (A sufficient condition for drift) *Random drift occurs in population P in generation G in which the conditions described in statement 4.1 obtain, if either*

1. *In P in G, for any parent sampling agent S_P, any pair of organisms O_i and O_j, any pair of variants V_v and V_w of any trait T, any variant V_{v*}^* of any trait T^* and any time frame Θ, indiscriminate parent sampling of V_v and V_w by S_P occurs in P in G, relative to the time frame Θ; or*

2. *In P in G, for any structure of the mating system S_G, any pair of gametes G_i and G_j, any pair of alleles A_a and A_b at any gene locus L, any allele A_{a*}^* at any gene locus L^* and any time frame Θ, indiscriminate gamete sampling of A_a and A_b by S_G occurs in P in G, relative to the time frame Θ.*

I allow that random drift can occur even if no evolution occurs. For instance, suppose that indiscriminate parent sampling occurs, resulting in the deaths of a significant proportion of organisms. That is, suppose that a parent sampling agent that is indiscriminate kills a large number of organisms; for instance, suppose that the parent sampling agent in question is an earthquake, that is, catastrophic drift occurs. Suppose that during the same generation, assortative mating alters the gene pool so that allele frequencies at the start of the next generation are precisely the same as they were before the earthquake. In this case, drift occurred but was "canceled out" by assortative mating. In such cases, I would say that random drift contributes to the maintenance of an equilibrium state.

165

While I believe that indiscriminate sampling is sufficient for drift, I do not believe that it is necessary. As I suggest in the opening paragraph of this section, I have a pluralistic view of drift. The central claim of this pluralistic view is that there are three mechanisms of drift, including indiscriminate sampling. The point of commonality among these three mechanisms is that each is a kind of evolution by accident, in both the causal and statistical senses I describe in chapter 1 (pages 16 - 18). Although I prefer to say that there is one concept of drift—evolution by accident—and three mechanisms of it, it would be acceptable to say that there are three concepts of drift, although I think that this is not as accurate. While is beyond the scope of this chapter to describe each of the further mechanisms of drift in full, I would like to describe them here in broad outline as follows; I provide a more detailed account of these two mechanisms of drift in appendix B.

Random fluctuations in evolutionary parameters Changes in fitness, rates of mutation, and rates of migration across generations can cause allele frequencies to change from one generation to the next, but leave the mean allele frequency unchanged. If this occurs, the generation-by-generation changes in allele frequencies are drift.

Idiopathic events "Idiopathic events" are those events that are so rare that they occur only once in the lifetime of a species or population. This can result in a significant change in the direction of evolution that is uncorrelated with any traits' abilities to carry out their purposes; such changes are drift. Some cases of catastrophic drift fall into this category, although idiopathic events need not be dramatic.

There are significant precedents in the philosophy of science for this kind of pluralism. Carnap [15] is well known for proposing that there are two concepts of probability,

statistical and evidential; many philosophers and scientists take a pluralistic view of species [148] and concepts of function [5]; and Kitcher and Sterelny [83] advocate pluralism about the units of selection.

4.4.2 Matthen and Ariew's hierarchical realization model

Matthen and Ariew's "Two Ways of Thinking About Fitness" [97] is one of several papers that have appeared recently as a part of the published record of what I will term the "theory of forces" debate. This name reflects the central issue of the debate: is Elliot Sober correct to claim that evolutionary theory is what he terms a "theory of forces?"[9] Among the papers by protagonists in the debate, Matthen and Ariew's stands out as particularly important to this chapter: the central claim of Matthen and Ariew's paper, an analysis of the concepts of natural selection and evolution they term the "hierarchical realization" model, is incompatible with the view that evolution occurs by indiscriminate sampling.

The aim of this section is to indicate the nature of the conflict between the view that evolution occurs by indiscriminate sampling and the hierarchical realization model, and to argue that, because of this conflict, the hierarchical realization model should be abandoned. I carry out the work to be done in the remainder of this section in four further subsections. The first of these has the aim of describing the hierarchical realization model; the second aims at elaborating a criticism of the model. In the third section, I account for the relationship of my argument against Matthen and Ariew to a similar argument advanced by Bouchard and Rosenberg. Finally, in the fourth section, I elaborate an additional argument against Matthen and Ariew.

---

[9]Sober makes this claim in *The Nature of Selection* [139, ch. 1].

Before taking on the main subject matter of this section, I want to make a remark concerning the extent of literature on this issue. Walsh, Lewens, and Ariew [146] also contribute a paper to the theory of forces debate, taking a position similar to that of Matthen and Ariew. Nonetheless, I do not address their position directly, because I believe that their position is similar enough to Matthen and Ariew's that my arguments against the latter also tell against Walsh, Lewens, and Ariew.

Matthen and Ariew's hierarchical realization model

I begin my discussion of the hierarchical-realization model with an important element of background. The following claim by Matthen and Ariew forms a central part of the model [97, 72].

Statement 4.23 *In a subdivided population, the rate of change in the overall growth rate is proportional to the variance in growth rates.*

Although statement 4.23 is of great interest to evolutionary biology because of its connection with Fisher's famous "fundamental theorem of natural selection," I do not want to consider its meaning beyond pointing out that it applies only to what Matthen and Ariew term "subdivided populations," by which I take it that they mean the following.

Analysis 4.3 *A population P is subdivided if and only if it contains at least two types of individuals which differ in their rates of growth.*

Although Matthen and Ariew do not provide an explicit definition of "rate of growth," it is clear that something like the following will suffice: the rate of growth G_T of a

type of individual T in a population P at time τ over a time span δ units of time in length

is the mathematical expectation of the number of descendants that an individual of type

T has in the population at time $\tau + \delta$. For instance, if individuals of type A produce, on

average, 7 offspring per generation, then their rate of growth is 7 individuals per generation.

If two types of individuals differ in their rates of growth, the ratio of those rates

deviates from unity, the degree to which it does so providing a useful index of how much

faster one type of individual grows than the other. Suppose that individuals of type B

produce, on average, 5 offspring per generation. The ratio of growth rates of A individuals

to B individuals is 7:5 in this case, a ratio favoring A individuals. If a population contains

both of these types, that population is subdivided in the sense of "subdivided" intended by

Matthen and Ariew.

I have now introduced enough background to introduce the fundamental notion of

the hierarchical-realization model of evolution, the natural selection formula:[10]

> A natural selection formula is one of the form ($L\&C$), where L is the antecedent
> of ... statement 4.23—that is, L posits a population subdivided by growth
> rates—and C is a *substrate specification* which states properties of [that] popu-
> lation (including properties of its members or of their parts), and/or the causes
> of differential growth rates in these populations and their parts, and/or the
> conditions of inheritance, development, and environmental interaction.

Matthen and Ariew [97, 76] explain that "corresponding to each of these natural

selection formulae is the set of possible histories that satisfy the formula," adding that they

will "call such a set of possible histories a *natural selection type*." They state that "each

history in a natural selection type is a *concrete realization* of [natural selection], subject

to the substrate-specification C.... Functional types have subtypes.... The subtypes of

---

[10]Matthen and Ariew hyphenate this term, i.e., they write "natural-selection formula." I do not follow
them in this practice.

natural selection are sets of histories that satisfy a particular substrate specification. These are *kinds* of [natural selection]—[natural selection] with Mendelian inheritance, with sexual reproduction, and so on. These types have sub-types, too. This is why we call our model a *hierarchical-realization scheme* " [97, 76].[11]

To illustrate the model, I want to present an example, as follows:

Statement 4.24 (Example of the H-R model) *In population P, at locus G, which is a Mendelian locus, there are two alleles, A_1 and A_2, the former of which codes for an enzyme that is twice as effective at metabolizing a key nutrient; and, over a period of time $\triangle t$ units of time in length, the ratio of A_1's growth rate to A_2's is 2:1.*

Because A_1 and A_2 differ in growth rate, a statement describing their growth rates is an instance of the antecedent of statement 4.23. This is an instance of L in a natural selection formula that applies to population P; it (the instance of L) would be appropriately formulated in a manner something like the following: "Population P is subdivided in growth rates at gene locus G, the ratio of the growth rate of A_1 to A_2 at G being 2:1." The substrate specification C in this case describes the biological and ecological conditions that form the physical basis for the differences in growth rates described in the instance of L described immediately above. These conditions include Mendelian reproduction and the physiology of each allele.

---

[11]In this passage, I have replaced Matthen and Ariew's use of "Li selection" with the more modest "natural selection," as indicated by the presence of brackets in the text. Matthen and Ariew term statement 4.23 "Li's theorem," which I find confusing. It does not seem that usage among biologists has become entrenched around the name "Li's theorem" for statement 4.23; and "Li selection" seems to be Matthen and Ariew's coinage. Thus, I do not think that any confusion will result from the terminological conventions I propose to adopt. I hesitate to use Li's name in such close connection with the hierarchical realization model because I am not sure that using it in this way is the best way to honor the great geneticist.

Matthen and Ariew believe that the hierarchical realization model is central to understanding evolutionary theory, claiming that it provides a conceptual framework into which any explanation of evolution whatever can be fitted. The following passage should dispel any doubt about whether I am overstating this point about how strongly they feel about their model.

> In this way of looking at things, the distinction between evolution (the total change of gene frequencies due to all causes), and natural selection (the portion of evolution due to differences in competitive advantage) is unmotivated. Natural selection is, as ... [statement 4.23] tells us, the aggregative result over time of differential growth rates in a population. These growth rates are explained by considering *all* of the factors posited by the most specific relevant natural selection formula, competitive advantage acting in concert with all the others. In histories that conform to this formula, certain trends get established at the ... [population level] as accumulations of multiple concrete events—births, deaths, mate choices, as well as events at the cellular and molecular level. There is no difference between these trends and evolution. [97, 78]

The views expressed in this passage are quite strong, but given Matthen and Ariew's theory, it is clear enough why they hold them. I think it is appropriate to describe their view as a kind of holism, a description they themselves apply to their position [98, 362]. To see what I mean by attributing a form of holism to them, consider the following. Matthen and Ariew do not believe that the kinds of statistical changes in a biological population usually attributed to natural selection can be explained by reference to the kinds of physical causes usually cited. This is because they believe that what are usually understood as distinct kinds of physical processes are in fact dependent upon one another for their existence. For instance, they believe that there is no fact of the matter about whether the death of a given organism contributes to drift, or to selection. This is the idea behind their position against Sober's claim that evolutionary theory is a "theory of

171

forces."[12]

The correct approach, as they see it, is as follows. Statistical changes in biological populations typical of natural selection are to be explained in terms of the physical system on which they supervene, taking that system as a whole. This requires unifying a diversity of sub-systems that are usually distinguished from one another, combining those sub-systems into a complex whole. As a consequence, statistical changes usually attributed to drift, mutation, and migration are explained by the same process that explains the kind of progressive evolution usually believed to be explained by fitness differences. This means that there is no sense in distinguishing among physical processes responsible for drift, mutation, migration, and natural selection, because all are different forms of the latter.

It is important to note that Matthen and Ariew do not deny that there is a difference between drift, mutation, migration, and natural selection. To deny this would be absurd; it would be fair to suspect someone interpreting Matthen and Ariew in this way of constructing a straw man. The important point, according to Matthen and Ariew, is that there is no difference *at a fundamental physical level* among these processes. For instance, consider the following.

> [W]e want clearly to acknowledge that it *is* legitimate to ask, *in a statistical sense*, how much of the causation of B is due to competitive advantage.... This question is similar to the questions asked when determining insurance premiums. For example: How much does being a male youth contribute to road accidents in which male youths are involved? There is some sort of answer to this question in statistical correlations. [97, 78][13]

---

[12]They argue against Sober in the early sections of their paper [97, secs. II - IV].

[13]They also make the following claims. "Let us define a *stochastic property* as one that belongs to ensembles as a *mathematical* (note: not *nomic*) consequence of the ... properties of individuals in that ensemble. Further, define a *trend* as a change of an ensemble over a period of time with respect to one or more of its stochastic properties. The claim we want to make is that while predictive fitness values are predictors of trends in populations, and may thus be considered *probabilistic causes*, they are not causes in the sense appropriate to fundamental processes" [97, 81]. "[S]ome may think that we are asserting that, if

The hierarchical realization model and natural selection

At this point, I have said all that I want to say in order to explain the model. In particular, I want to postpone further discussion of the manner in which Matthen and Ariew believe that an instance of *L&C* explains an episode of evolution.[14] What I would like to do now is to argue for two fundamental criticisms of the model, each of which follows from the claim that evolution occurs by indiscriminate sampling. These criticisms are as follows.

1. The hierarchical realization model is not a correct analysis of the concept of natural selection, that is, it does not describe what is commonly referred to as "natural selection" by evolutionary biologists.

2. The hierarchical realization model is not a correct analysis of the concept of evolution, that is, it does not describe what is commonly referred to as "evolution" by evolutionary biologists.

My view is that these claims, if true, are fatal to Matthen and Ariew's ambitions for their theory. As I discuss above, I believe that Matthen and Ariew intend for the hierarchical realization model to occupy a central place in our understanding of both of the concepts at issue, natural selection and evolution. I begin by arguing for the first claim listed above, an effort which forms the subject matter of the remainder of this section; my

a class of properties S supervenes on base properties B, then since all changes in properties S are wholly determined by properties B, there are no genuine causal relations at level S. In fact, we have not relied on the supervenience relation between . . . [supervenient] and . . . [base-level] properties in making our point. We have distinguished two kinds of causal relations, fundamental and stochastic. We concede that stochastic causation occurs at the S level, but deny that process causation occurs at this level" [97, 82].

[14]See section 5.1.4, in which I argue that Matthen and Ariew are what I term "Hempelian evolutionists," a phrase that I define explicitly in chapter 1 (page 27).

argument for the second claim listed above forms the subject matter of the section after the next.

The argument for the first claim listed above requires a new example, which I formulate by modifying the example described by statement 4.24.[15] The new example is described by the following statement.

Statement 4.25 (Modified example of the H-R model) *In the population described in statement 4.24 above, the A_2 allele increases in frequency over a time period T (which is $\triangle t$ units of time in length), a census of the population showing that the A_2 allele exhibited a rate of increase twice that of the A_1 allele. During T, a 100-year storm occurred.*

I now want to argue the following, against Matthen and Ariew: although biologists would suspect that the evolutionary process at work in these cases is drift, in the form of indiscriminate sampling, there is no way that the hierarchical realization model can account for this possibility, that is, the possibility that drift is the process by which the evolution described in these cases occurs. Before following out my main line of argument, I want to make an important point of clarification: the hypothetical case described by statement 4.25 maintains the supposition, expressed in statement 4.24, that the ratio of A_1's rate of growth to A_2's is 2:1. That is, in both cases, I suppose that the mathematical expectation of A_1's representation in the population after $\triangle t$ units of time is twice that of the A_2 allele's.

To be clear, the idea is that, in the example described by statement 4.25, contrary to what was expected, it happened that, in the amount of time in question, the A_2 allele

---

[15]Statement 4.24, which I introduce on page 170 above, is as follows: "In population P, at gene locus G, which is a Mendelian locus, there are two alleles, A_1 and A_2, the former of which codes for an enzyme that is twice as effective at metabolizing a key nutrient; and, over a period of time $\triangle t$ units of time in length, the ratio of A_1's growth rate to A_2's is 2:1."

increased at a rate twice as fast as the A_1 allele. Though this is highly improbable, it is not impossible. Rates of growth, as I understand them—and as I believe that Matthen and Ariew understand them—are not *actual* census counts, but are *expected* census counts, that is, they are probabilistic. The point of this example is that, in the population described in it, what obtains is evolution that proceeds in a manner like Rosencrantz and Guildenstern's coin-tossing game, discussed in chapter 3 (see pages 75 - 83).

Now I want to pick up my main line of thought against Matthen and Ariew. To begin that line of thought, I would now like to argue that the case described by statement 4.25 does indeed fall into the extension of the concept of natural selection, according to the hierarchical realization model. As I indicate above, Matthen and Ariew's view is that, for an episode of evolution to be an instance of natural selection, appropriate instances of L and C must be satisfied. In the example I describe in statement 4.25, this works in the following manner.

For an appropriate instance of L to apply to a population, that population must be subdivided by growth rates, a condition satisfied in the case at issue. There are two points I want to make regarding why this is so. First, there are differences in the expected growth rates of the A_1 and A_2 alleles; as I have indicated at several points above, the former is expected to grow twice as fast as the latter. This supervenes on differences in the alleles that are responsible for differences in the enzymes they produce, which differ in the rate at which they metabolize a key nutrient.

The second point I want to make regarding the relative rates of growth in this case is of central importance to my argument. This point is as follows: *the storm does not*

make any difference to the relative rates of growth of the variants at issue. This reflects a general property of indiscriminate sampling. Mechanisms of indiscriminate sampling do not introduce any across-variant differences in probability distributions of survivorship or offspring contribution, because they affect an individual's probability of survival or of bearing offspring in the same way, regardless of what allele that individual bears.

This accounts for the instance of L; now, consider the instance of C. An appropriate instance of C is generated by describing the biological and ecological setting for the 2:1 ratio of growth rates of individuals with the A_1 and A_2 alleles, respectively. Paralleling the discussion of L in the previous paragraphs, there are two points that I want to make here. First, C must make reference to differences in the ability of the protein products of the alleles to metabolize the key nutrient, as I indicate in statement 4.24. As I have just noted, this difference is in turn a consequence of physical differences between these alleles.

Second, if C is to be a complete description of the biological and ecological setting for the 2:1 ratio of growth rates, then it must incorporate the possibility of rare storms of extraordinary intensity, which can be responsible for indiscriminate sampling. I take it from Matthen and Ariew's own examples, in which they refer to Mendelian reproduction, that indiscriminate sampling agents are to be described in instances of C. This is confirmed by consulting the passages from Matthen and Ariew that I cite above (pages 169 - 170). In these passages, Matthen and Ariew refer to systems of inheritance generally and to Mendelian inheritance by name. This indicates that they believe that one indiscriminate sampling agent, Mendelian reproduction, ought to be included in the substrate specification; I take it that they would agree that others such as storms ought to be included as well.

I have now shown that the case described by statement 4.25 meets the criteria for natural selection set out in the hierarchical realization model. Together with the instance of L described above, the instance of C outlined in the previous paragraph describes a natural selection type. On Matthen and Ariew's view, this means that any evolution occurring in this case is natural selection, and is explained by reference to the appropriate instance of $L\&C$. At present, I will not discuss how Matthen and Ariew think that this explanation works, although I do so in chapter 5.[16] The important point is just that Matthen and Ariew would agree that the case I describe is a clear case of natural selection. All that is left now for me to do in my argument against Matthen and Ariew is to show that the conclusion I want to obtain follows from this: the hierarchical realization model badly mischaracterizes the case described by statement 4.25, because it cannot account for the action of indiscriminate sampling agents such as storms.

My argument for this claim proceeds in two steps, the first of which concerns Matthen and Ariew's analytic ambitions. I believe that Matthen and Ariew see themselves as providing an analysis of a concept common to Darwin, population geneticists since the 1930's and through to today, and game-theory-oriented behavioral ecologists such as Richard Dawkins and John Maynard Smith. Although some of the theoretical framework may have changed since the 1860's, the idea of the survival of the fittest remains; Matthen and Ariew see themselves as providing a correct explication of that idea, not as replacing it with another.[17] More precisely, the property of natural selection that they see themselves as

---

[16]See section 5.1.4.

[17]More precisely, they would claim that they are providing a correct analysis of fitness, as it appears in formal theories of population genetics. They might want to draw a clearer boundary between fitness as Darwin understood it (a non-quantitative concept) and the formal notion of fitness than I have attributed to them here. Nonetheless, I think it is clear that they see all of the various fitness concepts as bearing a strong family resemblance to one another.

clarifying is as follows. Suppose that a population P is subdivided into $s = 1 \ldots S$ sub-populations P_s, and that each P_s differs in growth rate from every other P_s; natural selection occurs if and only if there is some change in the relative sizes of the P_s's, and *this change represents differences in the intrinsic rates of growth across the S sub-populations P_s of P.*

Given the analytic ambitions that I describe in the previous two paragraphs, I think it follows quite naturally that Matthen and Ariew would assent to the following: the hierarchical realization model classifies as natural selection the vast majority (if not all) central cases usually classified as such by scientists in their intuitive assessments, i.e., in the course of their normal practice. Although Matthen and Ariew believe that they have shown something new and important about causation in the context of natural selection, they see themselves as describing a concept with roughly the same extension as the more intuitive notion of selection described by "survival of the fittest." As I will now show, in light of the case described by statement 4.25, this is certainly not the case. This brings me to the second and final step of my argument, which is as follows.

The storm in the example I present above provides a strong reason for thinking that changes in the frequencies of the alleles observed during the time span at issue are not due to intrinsic rates of increase across sub-populations. When such a storm occurs, it renders differences in intrinsic rates of growth irrelevant; this is just what it means for a condition in the environment to be an indiscriminate sampling agent. Nonetheless, as I argued above, the hierarchical realization model classifies this case squarely in the domain of natural selection. This means that the hierarchical realization model cross-classifies cases of random drift with cases of natural selection.

To put it another way, it follows from Matthen and Ariew's view that whatever changes in allele frequencies that occur in a population subdivided by growth rates is natural selection. While Matthen and Ariew require that there be some instance of L that applies to a population if natural selection is to occur within it, they place no further requirements on the actual outcome of evolution; any at all is compatible with natural selection. My view is that it is false that whatever changes in allele frequency occur in a population are natural selection. Some changes in allele frequency—those occurring by means of indiscriminate sampling agents—are drift. Because the hierarchical realization model describes natural selection entirely in terms of the relative growth rates of variant alleles, it is not powerful enough to distinguish between evolution resulting from conditions that induce differences in growth rates from those that do not. Rather, it lumps them all together under the heading of evolution resulting from conditions that do induce such differences.

This is a serious problem for the hierarchical realization model. The cases that the model cross-classifies with natural selection are unlike the latter in essential respects. The descriptive and explanatory rationale for identifying the process of natural selection is that it is the means by which progressive trends in a population arise. Darwin's central insight is that this occurs as a result of fitness differences, which Matthen and Ariew understand in terms of intrinsic differences in growth rates across sub-populations of a larger, subdivided population. In cases of indiscriminate sampling such as are caused by a 100-year storm, it is highly improbable that any progressive trends arise, and whatever evolution occurs does not do so because of differences in growth rates across sub-populations. Rather, evolution occurs in spite of such differences, in a random manner.

This kind of case cannot be dismissed as a borderline case or as a hypothetical case not likely to be encountered in the actual world. Indiscriminate sampling is a central and important mechanism of evolution. As I report in chapter 1 (page 12), the paleontologist Steven Stanley believes that catastrophic drift plays an important role in the history of life. Additionally, I believe that indiscriminate gamete sampling that occurs in Mendelian processes of sexual reproduction is of enormous importance in evolution: it plays a central role in genetic recombination, and occurs widely across major taxonomic groups. In the next chapter, I describe a variety of phenomena that occur by indiscriminate gamete sampling that are of great importance in evolution. Population size is particularly important: in a small population, indiscriminate sampling can radically alter the direction of evolution. When this happens, the evolution that occurs is nothing like natural selection. Such changes are random in direction, and are uncorrelated with any differences in propensity for survival and reproduction.

It is especially important to point out that indiscriminate sampling, particularly indiscriminate gamete sampling, can be particularly strong in small populations. Perhaps Matthen and Ariew would want to claim that cross-classifying some cases of indiscriminate sampling with natural selection should be allowed by their theory because chance fluctuations in allele frequencies are, by their very nature, rare. Moreover, Matthen and Ariew might want to claim that if chance fluctuations in allele frequencies do occur, they are most likely unimportant to evolution. A philosopher can always come up with a counterexample; but how many are relevant to scientific practice? Matthen and Ariew might want to insist that indiscriminate sampling can safely be ignored because it is irrelevant to the practicing

population geneticist.

This view is seriously mistaken, although it is clear how someone whose perspective on evolution that is strongly informed R. A. Fisher's views would come to have it. Indeed, Matthen and Ariew do seem to have such a perspective on evolution, a claim that is confirmed by a glance at the biologists cited by Matthen and Ariew, which include Li, Price, and Edwards, all of whom are in the Fisherian school. Only someone unaware of lines of thought developed by Sewall Wright and the role of drift in speciation would feel at ease dismissing the importance of indiscriminate sampling. Let me elaborate.

Fisher emphasized the importance of natural selection in large populations in which individuals mate randomly. His view was that, due to the large size of most natural populations, chance fluctuations in allele frequencies are rare, and that those that do occur are short-term aberrations that disappear quickly. The allele frequencies in a population could be counted on to reach their expected values given adequate time, which would usually be short. This contrasts with Sewall Wright's views: Wright emphasized the importance of population structure, pointing to the influence of random drift in larger populations structured into local isolates. Wright believed that, rather than "balance out" over the long run, chance fluctuations in allele frequencies often radically alter the direction of evolution, and are essential to the evolution of adaptation. Wright consistently claimed that these kinds of chance fluctuations were often caused by indiscriminate gamete sampling occurring in sex.

Drift in small populations is also of central importance in the origin of species. Rapid changes in the direction of evolution due to drift can prevent the extinction of an

incipient species, which begins as a small sub-population at the geographic periphery of a wide-ranging species. The members of this small sub-population may require new adaptations to deal with the peripheral environment, which may be quite different from the environments generally found in the species' range. Favorable traits can be created by drift by means of indiscriminate sampling during sex. Once such traits appear, they boost the mean fitness of the sub-population to a level sufficient to forestall its extinction. At the same time, genetic incompatibility between the organisms of the sub-population and its parent cause the biological separation of the two into distinct species. In this case, as in the case of the shifting balance process, chance fluctuations do not "balance out." Rather, chance fluctuations play a key role in establishing a new direction for evolution, a direction contrary to what would be established under the influence of differences in intrinsic rates of growth.[18]

I see the inability of the hierarchical realization model to distinguish between cases of indiscriminate sampling and natural selection as fatal to Matthen and Ariew's project of reconstructing evolutionary theory entirely in statistical terms. As I have indicated at several points above, indiscriminate sampling agents make no difference to the relative growth rates of variant alleles: there is no purely statistical means of distinguishing evolution due to indiscriminate sampling agents from that due to discriminate sampling agents. I take it that this means that the only way to distinguish the action of such agents is to look to a lower level of organization, viz., the physical construction of organisms and their causal interaction with the environment.

---

[18]I consider the shifting balance theory and the origin of species in greater detail in the next chapter (see sections 5.3.2 and 5.3.3, respectively).

Suppose a biologist asks, What makes evolution by a storm drift, as opposed to natural selection? The answer to this question is that no one organism is constructed any better than the other for surviving the storm, however well constructed some may be for surviving threats to mortality that are not as rare. Suppose a biologist asks, What makes evolution due to the Mendelian mechanism of reproduction drift, as opposed to natural selection? The answer to this question is that the Mendelian mechanism isolates each allele from the causal processes responsible for inheritance. There are important statistical consequences of these physical interactions among organisms and alleles; nevertheless, these statistical consequences cannot be identified with drift. Correlatively, *pace* Matthen and Ariew, natural selection cannot be identified as whatever evolution occurs in a subdivided population.

Bouchard and Rosenberg's argument

I would now like to consider the relationship of the argument I have just made to a similar argument directed against Matthen and Ariew, elaborated by Bouchard and Rosenberg in two recent papers.[19] My position is that, while some of the ideas Bouchard and Rosenberg introduce in the course of their argument provide a useful framework for stating certain claims essential to my argument, the two arguments have fundamentally different strategies. Moreover, I believe that my argument is more powerful than theirs. Let me elaborate on these claims, starting with a brief account of Bouchard and Rosenberg's position.

---

[19]One of these papers appeared in *BJPS* [10], the other in *Biology and Philosophy* [123]. Bouchard is first author of the former; Rosenberg, of the latter. I give them equal credit in the text below by listing their names in alphabetical order. Additionally, note that, because my discussion of their work draws solely on relatively short passages of their work ([10, 704 - 710] and [123, 350 - 353]), I do not cite them below.

Bouchard and Rosenberg's construal of drift is built around their understanding of the causal relationships characteristic of evolution, which are as follows. Suppose that, in some population at issue, the expected value of the frequency of the A_1 allele is p. Let there be G_g generations of this population, where $g = 1 \ldots n$, and suppose that in generation G_1, the frequency of the A_1 allele is $p_1 \neq p$. What explains why this is so, that is, what explains why $p_1 \neq p$? To describe Bouchard and Rosenberg's answer to this question, let me introduce a few new terms and ideas.

Bouchard and Rosenberg believe that the frequency of the A_1 allele at the end of any given generation G_g is causally dependent upon the state of all the physical objects, including the organisms in the population and all aspects of the environment, at a point in time just before the start of G_g. This is an expression of their belief that biological populations behave deterministically. For economy of expression, I will refer to all of the relevant physical systems simply as "the population," understanding this to include biological and non-biological aspects of the environment as well as states internal to the organisms themselves, their spatial relationships to one another, their behavior, and whatever else Bouchard and Rosenberg would believe to be relevant.

As the population is quite complex, there is a large number of states that it can be in. To represent this formally, I would like to use some new notation. Suppose that a complete physical description of the state of the population at a point in time just before a given generation G_g is represented as $D_d[g]$, which may be thought of as an enormous conjunction, each conjunct of which describes some part of the population. Let g be understood as indicated in the previous paragraph, and let there be $d = 1 \ldots m$ possible

states of the population. Using this notation, the set of all the possible physical states of the population just before a generation G_g, taking into account all the physical variables, is

$$\{D_1[g], D_2[g], D_3[g], \ldots, D_{m-2}[g], D_{m-1}[g], D_m[g]\}.$$

As it turns out, in general, the $D_d[g]$'s are more "fine-grained" than allele frequencies, so that there is a mapping that connects some set of $D_d[g]$'s to a given allele frequency. Take the extreme case of a set of $D_d[g]$'s that have the following property: if any of the $D_d[g]$'s in the set in question obtain, the allele frequency $p_g = 0$. For instance, suppose that one of the physical variables mentioned in $D_d[g]$ is whether humanity decides to embark on nuclear war, destroying all life on Earth. If this instance of $D_d[g]$ is realized, then the frequency of the A_1 allele in the relevant generation is zero. Similarly, one of the variables mentioned in $D_d[g]$ is whether an asteroid destroys all life in the neighborhood of the population. In this case, as well, the allele frequency of the A_1 allele will be zero in generation G_g.

These unfortunate examples could be multiplied; the important case, which is less dramatic, is the set of descriptions $D_d[g]$ that *over the long run, appear most frequently—* that is, the set that corresponds to the allele frequency p. I will call this set of states $D_p[g]$, to indicate that it is the set of states that causes the population to have the allele frequency p in generation g. This idea holds the key to Bouchard and Rosenberg's answer to the question posed above: Why is it the case that $p_1 \neq p$? That is, why, in generation G_1, did the allele frequency of the population deviate from the mean frequency p?

The answer to this question, according to Bouchard and Rosenberg, is simply that

the state $D_d[1]$ of the population that obtained before generation G_1 is not in the set of states $D_p[g]$. Perhaps it includes a 100-year storm; perhaps it includes a constriction in population size that provides an opportunity for sex to reshape the population by indiscriminate gamete sampling. Bouchard and Rosenberg's view is (roughly speaking) that any evolution resulting from any state of affairs other than an element of the appropriate instance of $D_p[g]$ may be classified as drift, supposing that there is no mutation or migration. The idea is just that states of affairs other than those in $D_p[g]$ are unrepresentative of what usually obtains, and also unrepresentative of what can be expected to obtain, in the long run. Accordingly, they lead to allele frequencies that differ from the expected value p—this is why Bouchard and Rosenberg categorize such cases as drift.[20]

These ideas form the basis for Bouchard and Rosenberg's criticism of Matthen and Ariew, which I will now detail. Suppose, the argument begins, that biologists have census information about the population, and that this information includes allele frequencies. This information dates back many generations into the past, and includes the present generation; it may be judged to be of excellent quality. Suppose, furthermore, that biologists infer, on the basis of this census information, that the population is subdivided by growth rates into two sub-populations, and that they determine the ratio of growth rates to be 7:3. Additionally, suppose that the population in question has, for several generations in the present, exhibited allele frequencies deviating from this 7:3 ratio, exhibiting a different ratio, say, 1:3.[21]

---

[20] As an historical note, I want to point out that the framework employed by Bouchard and Rosenberg for describing the causes of evolution bears a strong similarity to suggestions made by Rosenberg ([121] and [122]) in his earlier work about the supervenience fitness. It seems reasonable to believe that Bouchard and Rosenberg's work in the papers I cite here represents the further development of this earlier work.

[21] Bouchard and Rosenberg use a different ratio. I believe that a 1:3 ratio is more instructive, so I use it rather than the one provided by Bouchard and Rosenberg. It does not change the strategy of the argument

The problem, Bouchard and Rosenberg argue, is that this statistical information is not adequate for determining whether the change in allele frequencies represents the effects of drift, or whether it represents a change in the fitness values of the population, that is, a change in the relative intrinsic growth rates of the alleles. Bouchard and Rosenberg explain this in terms of the following example.

> Suppose we measure the fitness differences between . . . [variants in the population] to be in the ratio of 7:3, and suppose further that in some generation, the actual offspring ratio is . . . [1:3]. There are three alternatives: (a) the fitness measure of 7:3 is right but there was drift—i.e., the initial conditions in this generation are unrepresentative. . . ; (b) the fitness measure of 7:3 was incorrect and there was no drift; (c) both drift and wrong fitness measure. [123, 352]

Bouchard and Rosenberg's argument is that none of the alternatives (a) - (c) can be eliminated without understanding something about the interactions of the organisms in question with their environments: census information is not enough. The problem, they suggest, is that there is a regress. Where does the 7:3 figure come from? For all anyone knows, all the census data used to arrive at this value was collected during a time period during which unusual ecological conditions obtained. What is needed is some way of showing that 7:3 ratio represents the rates of survivorship that represent the realization of differences in intrinsic rates of growth across variants. The only way to learn this, according to Bouchard and Rosenberg, is to determine the ecological and biological conditions under which the 7:3 ratio arose, that is, to eliminate 100-year storms, sex, and other aberrations as causes of it.

Using the terminology I develop above, Bouchard and Rosenberg's suggestion is that scientists need to know whether the census information used to generate fitness values

---

in any way, or its degree of credibility.

fails to result from states of affairs in $D_p[g]$. I take it that this is what they mean when they say that "the initial conditions in this generation are unrepresentative." If the conditions prior to the generation in question are described by some instance of $D_d[g]$ that is not in $D_p[g]$, the allele frequencies will not behave as expected, because the conditions responsible for them are unrepresentative.

This concludes my account of Bouchard and Rosenberg's argument, and what I would like to do now is to indicate the similarities and differences between my argument and theirs. I begin with the similarities, which I think are rather conspicuous. My argument may be looked at as a special case of Bouchard and Rosenberg's. To show this, let me begin by describing indiscriminate sampling such as would occur due to a 100-year storm in Bouchard and Rosenberg's framework for thinking about causation in evolution. I will do this by reconstructing the claim that the storm is an indiscriminate sampling agent, which obtains because each organism in the population, regardless of which allele it bears, has an equal chance of being killed by the storm.

Let S_{100} represent the set containing all of the descriptions $D_d[g^*]$ in which a 100-year storm occurs during the relevant generation, which I will call g^*. Each $D_d[g^*]$ differs in its account of what happens during the storm. For instance, if $D_1[g^*]$ obtains, the storm knocks down a large tree; if $D_2[g^*]$ obtains, it leaves the tree standing, but floods a shallow stream; if $D_3[g^*]$ obtains, it does both; and so on. In some cases, a large number of organisms of one kind are killed; in others, a small number of organisms of that same kind are killed. The important point is that, however probability is distributed over the members of S_{100}, it is also equally distributed over organism survival.

Given this reconstruction of indiscriminate sampling in Bouchard and Rosenberg's terms, I think that they would want to say something like the following. Recall that, for a given population in a generation g, $D_p[g]$ represents a set of states of affairs described by a range of $D_d[g]$'s. This range of descriptions $D_d[g]$ is such that if any one of the states of affairs described by any one of them obtains, then the allele frequency in the generation of interest is p. Accordingly, as I indicate above, Bouchard and Rosenberg's view is that evolution occurring as a result of circumstances in which one of these states of affairs fails to obtain is drift. Just as it applies to other cases of drift, this analysis applies to indiscriminate sampling such as might occur by a storm, which is one instance of drift among many: the states of affairs in S_{100} do not fall within the set $D_p[g]$.

What does this brief discussion of the affinity between my argument and Bouchard and Rosenberg's show? In particular, does it show that my argument against Matthen and Ariew is superfluous? Indeed, someone might want to argue for that conclusion. It is irrelevant, the argument might go, whether drift occurs in the context of indiscriminate sampling due to a storm, or for any other reason. The essential point is that some state of the population not in $D_p[g]$ obtains at the relevant point in time; indiscriminate sampling is a special case of such a state of affairs. Thus, it follows that my argument is superfluous. I claim that, contrary to Matthen and Ariew, statistical information is not adequate to determine whether drift or selection has occurred, because one must know whether a discriminate or indiscriminate sampling process has occurred. According to Bouchard and Rosenberg, the right question to ask is whether the state $D_d[g]$ of the population prior to the generation at issue was in $D_p[g]$ or not.

I think that this is not a correct assessment of the significance of my argument. I agree that some of the central points of my argument can be reconstructed using Bouchard and Rosenberg's framework for thinking about causation in evolution; nevertheless, I do not agree that my argument is superfluous. I believe that my argument is more powerful than Bouchard and Rosenberg's, for two reasons. The first of these reasons is as follows: while Bouchard and Rosenberg's argument depends upon biologists not having complete certainty about the relative rates in a subdivided population, my argument applies even for a Laplacian demon. Let me explain.

Matthen and Ariew could respond to Bouchard and Rosenberg by calling attention to the following. A premise of Bouchard and Rosenberg's argument is that biologists cannot be sure that the census data they possess represents a true subdivision in the population. The biologists observe a 1:3 survivorship ratio in the present generation; this differs from the 7:3 ratio observed in previous generations. How can they know that the 7:3 value represents a subdivision in the population, rather than itself being an aberration? The argument, according to Bouchard and Rosenberg, is that data from any generation can be called into question in this manner. A regress follows. There is no generation that can be used to reliably determine whether a population is subdivided by growth rates.

I think that the line of attack that Matthen and Ariew would want to take is of the "whole is greater than the sum of its parts" variety. They would want to point out that the regress that Bouchard and Rosenberg believe threatens only does so if census data are consulted one generation at a time. Surely it would be a strange aberration that persisted for many generations in a row; to take the numbers from the example above, the more

consistently a 7:3 ratio is observed, the greater one's confidence ought to be that it really does represent a difference in intrinsic growth rates in the population. More generally, the more census data used to generate the estimate of differences in growth rates, the more reliable that estimate is. Moreover, in many cases, the precise degree of an estimate's reliability can determined by hypothesis testing.

Matthen and Ariew cannot raise this objection to my argument, which does not posit any uncertainty in estimates of fitness. My argument depends only on the fact that indiscriminate sampling agents do not induce any relative differences in probability distributions of survival or offspring contribution across variants. What my argument shows is that even a Laplacian demon would have to consult the mechanism of evolution in a given generation in order to determine whether drift or selection occurred in that generation: the statistical information that Matthen and Ariew would provide the demon simply is not sufficient for making such a determination. The idea is that an apparently aberrant result of evolution (Bouchard and Rosenberg's 1:3 result, for instance, in contrast with the expected 7:3 result) cannot be classified as not being due to intrinsic differences in rate of growth, just given the kind of information that Matthen and Ariew allow, that is, just given the fact that the population is subdivided by growth rates and the substrate specification for that population. Something further is required, viz., the knowledge that the mechanism that produced the 1:3 ratio is an indiscriminate one rather than a discriminate one.

This concludes my account of the first reason why I think that my argument is more powerful than Bouchard and Rosenberg's; now, let me indicate the second. I think that my argument demonstrates just how poorly the hierarchical realization model describes

191

evolution, and how poorly it describes the structure of evolutionary biology, in a way that Bouchard and Rosenberg's does not. I do not want to be pedantic about this. Nonetheless, I cannot conclude my assessment of Matthen and Ariew's theory of natural selection without mentioning that failing to account for indiscriminate sampling in such a theory is about as serious a mistake as someone can make. As I indicate above, it seems that Matthen and Ariew have formulated the hierarchical realization model with the intention that it be applied only to large randomly mating populations in which indiscriminate sampling can be disregarded. As I also indicate above, this seriously limits the applicability of the theory. Indiscriminate sampling cannot be disregarded in the shifting balance process and in populations that give rise to new species: in both cases, there are small populations in which indiscriminate sampling can have important consequences.

In conclusion, I think that Bouchard and Rosenberg's argument is important because it shows that reference to the physical structures on which evolution supervenes can play a role in determining fitness values. I see my own argument as making a more fundamental criticism of Matthen and Ariew's view, however: on their model, theories and issues of central importance to evolutionary biology cannot be reconstructed.

The hierarchical realization model and evolution

I now turn to the second of the two arguments against the hierarchical realization model that I listed at the start of the previous section (page 173): the argument that the model cannot be an analysis of the concept of evolution. The problem with the hierarchical realization model is that, although it is clear that there is evolution among selectively neutral alleles, the hierarchical realization model entails that there is none, and so cannot

be a necessary condition for evolution. The problem is that the model only applies to populations subdivided by growth rates. To see how this objection works in detail, consider the following.

As I explain in the section before last,[22] Matthen and Ariew claim that the hierarchical realization model applies to subdivided populations, that is, populations that contain sub-populations that differ in their growth rates. The degree of difference in growth rates across alleles, together with the substrate specification—instances of L and C, respectively—determine which natural selection type a population falls under. This means that, if a population is not subdivided by growth rates, there is no natural selection type that it fits into, because there is no instance of L that is satisfied by the population. Matthen and Ariew identify evolution with natural selection, which forces them into the position of affirming the following: if there is no natural selection type into which a population fits, then no evolution can occur in that population.

This position is untenable, a point that accrues to the disadvantage of the claim that the hierarchical realization model can serve as an analysis of the concept of evolution. The evidence for this is overwhelming: arguments about evolution among alleles that do not differ in their intrinsic rates of increase form some of the most important debates in evolutionary biology. I will briefly consider two such debates here: the neutralism debate, and the Fisher-Wright debate.

Neutralism is the view that most alleles at the molecular level do not differ in their intrinsic rate of increase. Supposing neutralism to be false, it is nonetheless not false *a priori*; but this is just what it would be, if Matthen and Ariew were correct to claim

---

[22]"Matthen and Ariew's hierarchical realization model," starting on page 168.

that the hierarchical realization model describes the concept of evolution. As it happens, it appears that even if it is false that *most* alleles at the molecular level do not differ in their intrinsic rate of increase, there are some do not so differ, an important point that explains the so-called molecular clock, widely used in establishing taxonomies. Finally, the so-called "nearly neutral" theory is affirmed by some who have rejected the neutral theory; the nearly-neutral theory posits evolution among alleles that do not differ in intrinsic rates of increase.[23]

Sewall Wright and R. A. Fisher differed with one another over the conditions under which a population is most likely to reach its maximum mean fitness. As I outline above, Fisher believed that this would most likely occur in a large panmictic population. In contrast, as I also outline above, Wright disagreed. As embodied in his so-called "shifting balance theory," Wright believed that a population would advance to its fittest state if it were large, but subdivided into small local isolates connected to one another by a consistent but low level of migration. Wright held this view because he believed that, in such populations, random drift would create favorable gene combinations that could not arise under the control of selection. The key point for my argument here is that Wright consistently emphasized that a process that operates on alleles that do not differ in their intrinsic rate of increase—drift associated with Mendelian reproduction—is an important mechanism by which this could occur.[24]

In parallel with my argument about the neutral theory, the issue here is not whether evolution in fact occurs by drift of alleles that do not differ in their intrinsic rates

---

[23]See section 5.2.3 for a more complete account of the neutral theory.

[24]As I indicate above, section 5.3.2 provides a detailed account of Wright's shifting balance theory.

of increase. The important point is that the Fisher-Wright debate is not meaningless. This is just what would follow, if Matthen and Ariew were correct that evolution always and only occurred in populations subdivided by growth rate. Thus, I conclude that the hierarchical realization model cannot be an analysis of the concept of evolution.

4.5   Concluding Remarks

I would like to conclude this chapter by reviewing some of the main claims that I have made in it. The first of these concern the nature of indiscriminate sampling. Indiscriminate sampling, a mechanism of drift that John Beatty first brought to the attention of philosophers of science, operates during two stages of the life cycle. In the period between birth of the organisms in the population in question and their sexual maturity, indiscriminate sampling is known as "indiscriminate parent sampling;" in the period between sexual maturity and the birth of the organisms in the next generation, it is known as "indiscriminate gamete sampling."

Second, I believe the probability account of indiscriminate sampling to be correct. According to this account, indiscriminate sampling requires a relationship of near-equality between certain probabilities. This relationship is most clearly and intuitively described by what I have called the "naïve formulation of the core probabilistic equality." To avoid a class of fatal counterexamples to the naïve formulation of the core probabilistic equality, it must be modified, resulting in a "sophisticated" statement of it. This forms the central claim in a set of necessary and sufficient conditions for indiscriminate sampling, which is a follows. If and only if the sophisticated statement of the core probabilistic equality is satisfied relative

to a time scale, and a further principle meant to eliminate spurious sampling agents is satisfied, indiscriminate sampling occurs.

Finally, I claim that indiscriminate sampling is sufficient but not necessary for random drift. The concept of drift is pluralistic, in the sense that there are three mechanisms of it, including indiscriminate sampling. These are unified under the rubric of drift because each is a mechanism of evolution by accident. I argue that Matthen and Ariew are incorrect to claim that their hierarchical realization model describes the essential elements of evolutionary theory. Their claim that evolutionary theory is essentially statistical is incorrect. My argument against Matthen and Ariew, which I also believe applies to similar claims by Walsh, Lewens, and Ariew, depends on the claim that indiscriminate sampling is an important mechanism of evolution, and that it cannot be accounted for by the hierarchical realization model.

This chapter is a cornerstone of this dissertation. Philosophers of biology generally agree that indiscriminate sampling is a form of drift, and I intend for my account of indiscriminate sampling in this chapter to further solidify that general agreement. Although contributing to the understanding of the nature of drift is a beneficial result of doing this, I regard this as a by-product of my efforts in this chapter. My aim is rather to deflect controversy away from questions about the nature of drift in order to focus attention on questions about explaining it. My complaints are against the Hempelian evolutionists, who claim that drift cannot be explained because it is a chance process. The present chapter secures the claim that drift occurs by indiscriminate sampling, and clarifies the nature of that process, so that issues about applying process explanation to instances of it may be brought

to light and resolved. It is to this task that I now turn, taking up the most pressing and most central issues I mean to confront in this dissertation in its next and most important chapter.

Chapter 5

Process Explanations of Drift

This chapter is the culmination of my work in this dissertation: I argue that evolutionary biologists explain, by process explanation, evolutionary events occurring by random drift. I articulate my view in opposition to those I term "Hempelian evolutionists." The Hempelian evolutionists, as I indicate in chapter 1 (pages 23 - 27), take Hempelianism as a motivation for what I term the "exclusivity thesis." The exclusivity thesis is the view that, if and only if an evolutionary event can be explained, natural selection explains it by providing the reason why the event occurred. This entails that evolution by drift cannot be explained, which is incompatible with my view.

The Hempelian evolutionists' opposition to my claim on behalf of drift is unwarranted, however. My chapter 3 conclusion that Hempelianism is not correct deprives the Hempelian evolutionists of their motivation for the exclusivity thesis. This amounts to depriving them of their reasons for opposing the claim that process explanations can be used to explain evolution occurring by drift—the very claim I argue for in this chapter. My

argument is casewise: there are five kinds of evolutionary events, each occurring by random drift, each of which is of fundamental importance to active research programs central to evolutionary biology, and each of which is explained using process explanation.

These phenomena, which I sketch in chapter 1 (pages 30 - 32), include the following: the chance elimination of a rare but favorable allele; molecular evolution; the shifting balance process; the origin of species; and punctuated equilibrium's effects on the shape of phylogeny. I consider instances of these events that occur by indiscriminate sampling, and to describe each, I draw on my chapter 4 description of that mechanism of drift.

As I suggest in chapter 1 (pages 30 - 31), I divide these events into two classes. In one class, I include the chance elimination of a rare but favorable allele and molecular evolution, whose process explanation does not refer to population size N. In the other class, I include the remaining three kinds of events, whose process explanation does refer to population size N. I draw this distinction because drift significantly alters the evolutionary dynamics of small populations, creating an important role for population size in process explanations of phenomena in the latter class.

The major subdivisions of this chapter are as follows. In section 5.1, which concerns the Hempelian evolutionists, I recall my work in chapter 1 on the exclusivity thesis. I review the thesis and my account of its connection with Hempelianism, extending my work on the thesis by arguing that the ranks of the Hempelian evolutionists include Daniel Dennett, Richard Dawkins, and Matthen and Ariew.

Sections 5.2 and 5.3 form the crux of the chapter, and indeed, of the dissertation. In these sections, against the Hempelian evolutionists, I argue that evolutionary biologists

explain, in the narrative manner of process explanation, the five kinds of evolutionary events I describe above. Section 5.2 concerns phenomena whose explanation need not refer to population size N, while section 5.3 concerns events whose process explanation must refer to population size N. The chapter concludes in section 5.4 with a brief overview of some of the main conclusions I reach in the chapter.

5.1   The Hempelian Evolutionists

In this section, I argue that Daniel Dennett, Richard Dawkins, and Matthen and Ariew are Hempelian evolutionists, a notion I introduced in chapter 1 (analysis 1.5, page 27).

A person is a Hempelian evolutionist if and only if he or she affirms both the exclusivity thesis and Hempelianism, and affirms the former because he or she affirms the latter.

In section 5.1.1, I recall my work in chapter 1 concerning the exclusivity thesis; in sections 5.1.2 and 5.1.3, I argue that Dennett and Dawkins, respectively, are Hempelian evolutionists. In section 5.1.4, I argue that Matthen and Ariew have the view. Throughout this section, for economy of expression, I will understand "trait" to indicate genes and gene sequences as well as physiological, morphological, and behavioral traits.

5.1.1   The exclusivity thesis

I begin by recalling the the exclusivity thesis (analysis 1.4, page 26).

Statement S expresses a proposition that explains an evolutionary event if and only if (1) S describes the process of natural selection responsible for adaptation A's having spread in P to the extent that it did between the time of its initial appearance and time T, and (2) S answers the question, "Why, between the time of its initial appearance in population P and time T, did adaptation A spread to the extent that it did, in P?"

The central idea expressed by the thesis, as I suggest above, is as follows: if and only if an evolutionary event can be explained, natural selection explains it by indicating why it occurred. Clearly, someone who affirms the exclusivity thesis also affirms that no evolutionary events occurring due to random drift can be explained. As I stated in chapter 1 (page 26), the thesis need not be stated in terms of propositions, and I urge those who have alternative views about the metaphysics of explanation to reformulate the thesis as they see fit.

To be clear, let me recall the meanings of two central terms used in the thesis, discussed in chapter 1. By "evolutionary event," I mean to indicate any event whose occurrence constitutes evolution. As I indicate in chapter 1 (page 24), this includes, say, cross-generational changes in allele frequencies, and it does not include, say, sex among organisms. As I also indicate in chapter 1 (page 24), by "adaptation" I mean to indicate any variant that makes the greatest contribution to its bearers' fitness, relative to other variants of the trait in question. Contrary to Sober [139, ch. 6] and most other philosophers, and in accord with Reeve and Sherman [120], I do not require that a trait have a history of selection to be an adaptation.

The relationship between the exclusivity thesis and Hempelianism, as I explain in chapter 1 (pages 23 - 27), is as follows. The Hempelian evolutionists believe that the principle of natural selection is the only lawlike statement that it is possible to make about evolution. The idea is that the principle of natural selection describes a lawlike relationship between fitness and evolutionary change: *ceteris paribus*, traits with greater fitness proliferate in greater degree than others. Furthermore, no other theoretical principles describe

lawlike relationships among any biological properties and evolution; for instance, there are no "laws of migration," according to which organisms of a certain kind will always, or always with some certain probability P_M, migrate from one population to another.

If one is a Hempelian—which the Hempelian evolutionists, naturally, are—and one believes that natural selection is the only law of nature describing evolution, then the exclusivity thesis follows readily. Hempelianism incorporates the view that laws of nature are required for explaining particular events; and if natural selection is the only law of nature describing evolution, it follows readily that it alone can explain evolution. This is what the first provision of the thesis amounts to. The second provision, according to which explanation-seeking questions about evolution are "why-questions" of a particular form, also follows readily from Hempelianism, according to which all explanation-seeking questions about particular events in science have just the form in question.

5.1.2 Dennett

The main source of evidence that Dennett's Hempelianism informs his adherence to the exclusivity thesis is his general view of evolutionary explanation, which is that evolutionary explanations fit a teleological model that he terms "reverse engineering." My aim in this section is to describe the central principles of the reverse engineering strategy, articulated by Dennett in *Darwin's Dangerous Idea* [34, chs. 8 and 9]; to indicate how the strategy fits into the Hempelian framework; and to indicate how it represents Dennett's commitment to the exclusivity thesis. I begin by describing the aims of the strategy.

Using the reverse engineering strategy, a biologist answers explanation-seeking questions of the following kind.

Question 5.1 *Why do organisms of type O have adaptation A?*

A biologist using the reverse engineering strategy answers such questions by describing why an intelligent designer would have constructed the adaptation such as it is in fact constructed, were that designer intending to promote its bearers' survival, reproduction, or both. Dennett believes that this represents why the trait would have succeeded in natural selection over alternatives; so, he believes that this answers questions with the form of question 5.1. Dennett believes this, in turn, because he believes that adaptation is due to natural selection, and that natural selection produces the same kind of results that an intelligent designer would.[1]

Question 5.1 reflects a commitment to Hempelianism for the following reason. As I have pointed out previously (pages 19 - 20 and 39), Hempelians believe that the aim of scientific explanations of particular events is to answer explanation-seeking why-questions by citing laws of nature. An organism of a certain type possessing an adaptation of a certain type is the kind of event that Dennett believes that evolutionary biologists seek to explain by reference to laws. As question 5.1 indicates—it is a why-question—Dennett believes that evolutionary biologists ask why-questions concerning such events.

The reverse engineering strategy also embodies the Hempelian claim that laws are essential to explanation. The laws in question personify natural selection as an intelligent agent that acts to maximize fitness against a background of constraints. These constraints include interactions with other traits, developmental constraints, and design constraints such as those imposed by the structural materials and basic design geometry of the traits

---

[1]I would like to note that Lewens [88] provides an excellent account of how Dennett understands reverse engineering.

at issue.

The laws have the form of an imperative: "To attain state of affairs x under constraints y, do z." The state of affairs x is the maximum attainable fitness under the constraints described by y; z is a certain value of the design parameter of the trait of interest. The idea is that a biologist using the reverse engineering strategy assumes that organisms are fitness-maximizing systems operating within a set of constraints, and uses the appropriate teleological law to predict or explain the trait's taking on state z. Because of its teleological nature, many models of natural selection exemplifying the reverse engineering strategy are borrowed from economics.[2]

Reverse engineering explanations do not entail that there is a conscious designer. Rather, according to the reverse engineering strategy, the teleological account of the trait's behavior corresponds to a mechanical, non-teleological account of the trait's evolution by natural selection. This mechanical, non-teleological account substitutes advantage in survival and reproductive success for benefit or usefulness in the eyes of a designer: rather than maximize value according to an artificer, the trait maximizes the organism's fitness. This strategy reflects the idea that natural selection is a kind of evolution by design, which I outlined in chapter 1.

To see that Dennett's commitment to the reverse engineering strategy reflects his commitment to the exclusivity thesis, the important point to take note of is that Dennett sees reverse engineering as a completely general strategy for explaining evolution. That is, he believes that it is the only legitimate strategy for explaining evolutionary events,

---

[2]Maynard Smith [99] and Dawkins [29, ch. 3] provide touchstone statements of the view that organisms may be viewed as fitness-maximizing agents, against a backdrop of constraints; Beatty [8] discusses philosophical issues connected with this notion.

affirming the following statement.

Statement 5.1 *Any evolutionary event is either explained by reverse engineering, or it cannot be explained at all.*

Because only traits that evolve by natural selection can be explained by reverse engineering, this means that, according to Dennett, only traits that evolve by natural selection can be explained. Together with the claim that reverse engineering answers questions of the form "Why did the event to be explained occur?" it follows that someone with Dennett's level of commitment to reverse engineering is also committed to the exclusivity thesis.

The following passage from *Darwin's Dangerous Idea*[3] exemplifies Dennett's view that reverse engineering is a general strategy; it also indicates the importance of Hempelianism for motivating his position.

> Perhaps nobody cares [whether hinges are on the left or the right], so a coin is flipped, and hinges on the left get installed.... [Other] builders copy the result unthinkingly, establishing a local tradition.... "Why are all the doors in this village hinged on the left?" would be a classic ... [evolutionary biologist's] question, to which the answer would be: "No reason. Just historical accident." [34, 276]

In accord with Hempelianism, Dennett believes that biologists ask why-questions about the existence of certain traits; recall that the first main tenet of Hempelianism is that explanations of particular events answer why-questions. This is represented in the passage above by an explanation-seeking why-question about doors with hinges on the left rather than the right. According to Dennett, these kinds of questions are teleological: what design goal did the engineers have in mind that they believed would be achieved by putting the hinges on the left?

---

[3]Many passages of Dennett's can be interpreted in a manner similar to the manner in which I interpret the passage I consider here ([33, 386] and [34, 198, 238, 247, 250 – 251, 403 – 404, 421]).

For the biologist using the reverse engineering strategy, such teleological questions are to be answered by assuming that the organisms in question are fitness-maximizing agents operating against a background of design constraints; this assumption warrants the further claim that the organisms can be described by laws of the form "To attain state of affairs x under constraints y, do z." As I explain above (page 204), x is the maximum fitness possible under constraints y, x being attained by taking on a certain state z of the trait in question. In this case, the state z is left-sided hinge-placement rather than right-sided. By invoking these laws, a scientist conforms to the second main tenet of Hempelianism, that is, the claim that explanations of particular events in science require laws.

In the case described in the passage above, the hinges do not result from a process of design analogous to natural selection, but from a chance process—a coin toss. As a consequence, the assumption fundamental to the reverse engineering strategy is incorrect: it is wrong to regard organisms as fitness-maximizing systems. This means that, according to Dennett, who believes that reverse engineering provides the only form of explanation acceptable for explaining evolution, no explanation of the hinges' left-sided placement is possible. This is reflected in his claim that there is "no reason" for their location. This reflects the dichotomy indicated in statement 5.1 (page 205), which someone who believes the exclusivity thesis affirms.

5.1.3 Dawkins

I will begin my account of Dawkins by arguing that he holds the provision of the exclusivity thesis according to which explanation-seeking questions about evolution are why-questions about the spread of adaptation, and that he holds this view in part

because of his Hempelianism.[4] Note that I consider Dawkins's position on the two tenets of Hempelianism separately; I postpone consideration of his belief that explanation requires laws until after I consider his stance on the role of why-questions (page 208). Similarly, I divide my treatment of his affirmation of the exclusivity thesis in two, first considering the "why-questions" provision along with my account of Dawkins's Hempelianism.

The first premise of my argument about Dawkins's stance on the role of why-questions in evolutionary explanation is that he believes that all explanation-seeking questions about evolution have the following canonical form, which is a question about why an adaptation has spread throughout a biological population by natural selection.

Question 5.2 *Why do organisms of type O characteristically have adaptation A, which is intricate and highly complex?*

Note that, although similar in form to question 5.1 (page 203), which Dennett takes to be canonical, question 5.2 differs slightly from the former. Dawkins believes that evolutionary biologists ask explanation-seeking questions about adaptations that are so complex and intricate that it seems impossible that they were not designed by someone of great intelligence.[5] Dennett does not believe that evolutionary explanations must be restricted to intricate and complex adaptations; he sees the reverse engineering strategy as applying to all adaptations, simple and complex alike.

The Blind Watchmaker [31] provides a broad base of evidence for the claim that Dawkins takes questions with the form of question 5.2 to be canonical. In *The Blind*

---

[4]As will be seen, my claims about Dawkins's Hempelianism in this section draw primarily on *The Blind Watchmaker* [31]. However, though I will not argue for this point in this dissertation, I believe that my claims apply to *The Selfish Gene* [30] and *The Extended Phenotype* [29], as well.

[5]I owe this point to Reeve and Sherman [120, 6 - 7].

Watchmaker, Dawkins [31, 4 - 6] responds to natural theologians, who claim that organic purposes provide evidence for a divine artificer. At the heart of this issue are explanation-seeking why-questions such as the following.

Question 5.3 *Why do human beings have organs so well-suited for seeing, i.e., eyes?*

Dawkins agrees that this question is puzzling and important; however, unlike a natural theologian, he believes that natural selection provides its answer. The point I want to make here is that, by seeing explanation-seeking questions about evolution in the same way that the natural theologian does, Dawkins accepts the provision of the exclusivity thesis according to which all explanation-seeking questions about evolution are why-questions about the spread of an adaptation, i.e., questions having the form of question 5.2.

Does Dawkins hold the provision of the exclusivity thesis about why-questions because he is a Hempelian, which is required of someone who is a Hempelian evolutionist? The answer to this question is an unqualified "Yes." The evidence for this claim is implicit in his discussion of explanation in *The Blind Watchmaker* [31, 11 - 18]. His discussion of explaining adaptation, which requires answers to why-questions, is seamless with his discussion of how a steam engine works, and also, with a more general discussion of explanation in physics. He makes no distinction between the kinds of explanation-seeking questions asked across disciplines: all are why-questions.

Now I want to argue that Dawkins affirms the provision of the exclusivity thesis according to which only natural selection explains evolution, and also, that this belief stems from the further belief that scientific explanation of particular events requires laws. Before doing so, there is some background I need to provide: I need to sketch Dawkins's general

208

view of evolutionary explanation.

Dawkins [31, 9 - 18] claims that explanation-seeking why-questions about the spread of adaptation are answered by a strategy that combines the principle of natural selection with a strategy of explanation that resembles what Robert Cummins [28] terms "functional analysis." According to Dawkins, functional analysis of an adaptation A of an organism O describes the laws governing a causal capacity of A that contributes to some state S of O. Dawkins believes that answering explanation-seeking why-questions about the presence of an adaptation A requires functional analysis of A's contribution to either of the following two states S of an organism O: (a) O's ability to survive, reproduce, or do both, or (b) any other capacity of O that contributes to its ability to survive, reproduce, or do both.

To explain why an adaptation A is present, its functional analysis must be supplemented by the principle of natural selection. The functional analysis explains why the adaptation does whatever it does that happens to contribute to its bearers' survival, reproduction, or both; but it does not explain why A is characteristic of organisms in the population. This is explained by the additional claim that A provides a fitness advantage over alternatives, and that this fitness advantage caused it to succeed in natural selection.[6]

This concludes my general account of Dawkins's view of evolutionary explanation. I am now in a position to argue the following: Dawkins's view entails that evolutionary explanations always and only cite selection as the reason why an adaptation spreads. This

---

[6]Although they were apparently formulated independently of one another, Dawkins's account of evolutionary explanation bears a strong resemblance to Paul Griffiths's [63] account of biological functions, according to which (very roughly) the function of a trait is described by the functional analysis of whatever that trait does that contributes to a certain proportion of its bearers' fitness; this idea is originally due to Larry Wright [151]. Griffiths's contribution is to combine functional analysis with natural selection in the analysis of the concept of function.

is the provision of the exclusivity thesis that I have not yet considered in connection with Dawkins. My argument is as follows.

As I argue in section 1.1.3, only natural selection results in organic purposes. Furthermore, as I have just been arguing, Dawkins believes that evolutionary explanations proceed by functional analysis of what a variant does that makes a fitness difference to its bearers, that allowed it to succeed in natural selection. Like Dennett, Dawkins believes that natural selection can be represented by teleological laws of the form "To attain state of affairs x under constraints y, do z," where x is the maximum fitness under the constraints described by y, and where z is a certain value of the design parameter of the trait in question. However—and this is the key point—traits not evolving by natural selection do not have purposes, and do not contribute to fitness differences in the required manner. So if Dawkins is correct, then no appropriate questions can be asked about them, and they cannot be explained.

Finally, I would like to argue that Dawkins's function-analytic strategy is informed by his belief that laws are necessary for explanation—that is, he is Hempelian about laws in explanation. Dawkins, an avowed reductionist, rejects emergentist strategies of explanation, seeing his view as the only viable alternative [31, 9 - 18]. He takes emergentism to be the view that explaining an event E, e.g., a trait's presence, requires showing that E is a part of some other entity, e.g., an organism. The problem with emergentist explanations, as Dawkins sees it, is that they appeal to the wrong level of organization. Consider the following.

First, note that fundamental laws of physics describe the lowest possible level of the

organization of matter. Second, note that emergentist explanations of an event E require reference to objects at a *higher* level of organization than E itself; this is because emergentist explanations require reference to an entity of which E is a part. These two points suggest that emergentism fails to contribute to showing how E relates to fundamental laws of physics. This is crucial for Dawkins, because he believes that this is what an explanation ought to show.

In contrast to emergentism, Dawkins sees the function-analytic strategy as reductionist in the appropriate sense. As I argue above, Dawkins's view is that the reason why an organism has an adaptation is provided by laws describing the behavior of its main components, together with natural selection, which Dawkins believes is a law; indeed, he claims that any adaptation anywhere in the universe is due to natural selection ([31, ch. 11] and [32]). In turn, the main components of an adaptation are explained by laws that apply to their components, which are explained by laws that apply to theirs, and so on, the explanation of an adaptation's very smallest components appealing to fundamental physical laws.

Thus, Dawkins's Hempelianism about laws is embodied in a view like the following. Whether a given explanation is correct depends on whether there exists an appropriate set of laws reaching across all levels of biological, chemical and physical organization, laws at each level depending on the laws of the next level down, all depending on the lowest level, the laws of microphysics.

5.1.4 The hierarchical-realization model

I would now like to argue that Matthen and Ariew, who advocate the hierarchical-realization model of evolution that I considered in chapter 4 (pages 167 - 195), are Hempelian evolutionists. Matthen and Ariew's Hempelianism derives from their acceptance of a strategy that is something like Salmon's statistical relevance account of explanation (see section 2.1.2). Consider the following passage, taken from the paper in which Matthen and Ariew introduce their hierarchical-realization model.

> We arrive at an adequate explanation of the evolution of a biological phenomenon by subsuming it under the *most specific* formula that applies to it, that is, the formula that posits *all* the substrate factors relevant to it. The probability of the target phenomenon is estimated relative to histories that constitute the corresponding *least inclusive* natural selection type. We understand why the phenomenon came about by comparing this probability with those yielded by natural selection formulae which impose relevantly different substrate specifications. For instance, we understand why a deleterious hereditary condition like sickle-cell anemia was not eliminated, by comparing the probabilities in its natural selection formula with those with relevantly different ones—ones in which malaria was not a factor, ones in which reproduction is not sexual, and so on. [97, 77]

To understand their view, recall how they understand natural selection types and natural selection formulas.[7] A natural selection type is a set of possible histories that a biological population might have. Whether a given population's actual history falls into a given natural selection type depends upon its natural selection formula. A population's natural selection formula depends upon the following. First, what are the relative growth rates of the variants in the population? Second, what is the physical basis for these growth rates?

---

[7]I discuss natural selection types and natural selection formulas in chapter 4. See the pages referred to immediately above, in the opening paragraph of this section.

The latter characteristic of a population is known by Matthen and Ariew as its "substrate specification." Substrate specifications generate the hierarchy of types, a phenomenon that I will explain by way of example. All possible histories of all possible populations that reproduce sexually fall into what might be called the "sexual reproduction natural selection type." Within this natural selection type, there are histories of populations whose sexual reproduction is Mendelian; those whose sexual reproduction is not Mendelian because of assortative mating; those whose sexual reproduction is not Mendelian because of meiotic drive; and so on. These three conditions represent kinds of sexual reproduction, and so populations possessing one or another of these properties fall into natural selection types subordinate to the broader type that consists of all populations that reproduce sexually.

Now, let me describe Matthen and Ariew's account of the explanatory value of the hierarchy of natural selection types. Suppose that a given population P is characterized by organisms with a certain variant V, a condition that I will term "population P's having V." By "characterized by organisms with a certain variant," I mean to say that nearly all the organisms in the population have this variant, and that it came to be widespread by way of natural selection, that is, variant V is the main variant in the population.[8] Suppose, as well, that population P is in natural selection type T_1, and that there are T_n other natural selection types, where $n = 2 \ldots m$. Let the probability that P has V given that it is in the natural selection type T_1 be represented by p. Now, suppose that someone asks the explanation-seeking question, "Why does population P have variant V?"

According to Matthen and Ariew, to answer this question, someone must consider all natural selection types T_n; and that person must describe natural selection type T_n if

---

[8]See page 7 for discussion of how I understand "main variant."

and only if the probability p^* that a population has V given that that population is in type T_n differs from P's actual probability p of having V; and the person must also state the probability p^*. This amounts to pointing out all of the natural selection types that are statistically relevant partitions of the reference class for the probability that a population have V.

For instance, suppose that variant V is a camouflage pattern of coloration possessed by organisms in population P. Additionally, let A mean "There is heavy predation by sharp-sighted predators in the environment," and suppose that A is true of population P. Statement A is a part of the substrate specification for the natural selection type into which P falls, together with a long conjunction B of statements describing other conditions in P's environment. Suppose that the probability that P has V, given $A\&B$, is p.

Now, let A' mean "There is no predation by sharp-sighted predators in the environment," and suppose that the probability that P has variant V, given $A'\&B$, is p^*. Statement A' is false of P, so $A'\&B$ describes a natural selection type that population P is not a member of. Moreover, the probability of a population in this alternative natural selection type having V differs from the probability of a population in P's natural selection type having V. Pointing out these differences in probabilities and substrate specifications contributes to explaining why population P has V. Doing so for all correspondingly similar natural selection types constitutes a complete answer to the explanation-seeking question, "Why does population P have V?"

Having outlined Matthen and Ariew's views on evolutionary explanation, I would like to argue that they are Hempelian evolutionists. This is not a difficult task, because of

214

their position's affinities with Salmon's statistical relevance view, and because of the significance of natural selection formulas in their view of evolutionary explanation. Regarding the former, they clearly hold that to explain a particular event is to show why it occurred by indicating the degree of rational expectation that someone ought to have had that it occur. This is the motivation, I take it, for their view that statistical relevance is central to explaining evolution.

Regarding the latter, their picture of evolutionary theory forces upon them the view that only evolution occurring by natural selection can be explained. They believe that any episode of evolution can be explained by subsuming it under an appropriate natural selection formula. As I argue in the previous chapter (pages 167 - 195), they believe that even evolution occurring by indiscriminate sampling may be subsumed under a natural selection formula. This, together with the view I cited in the previous paragraph about explanation, is the basis on which I claim that they are Hempelian evolutionists.

5.2   Explaining Drift Independently of N

My chapter 3 conclusion that Hempelianism is incorrect robs the Hempelian evolutionists of their reason for affirming the exclusivity thesis. In turn, this robs them of their reason for opposing my claim in favor of random drift—that is, my claim that process explanation is used to explain events in evolution that occur by indiscriminate sampling. The description of these events and their process explanation by drift forms my casewise argument for this claim, which I elaborate in this section and the next.

My aim in this section is to describe two phenomena that occur by drift and their

process explanation: the chance elimination of rare but favorable alleles and molecular evolution. Because the process explanation of those events does not require mention of small population size, I consider them apart from other events in evolution occurring by drift that are explained by process explanation, which I treat in the next section. I divide the two kinds of cases from one another because small populations evolving due to drift exhibit distinctive evolutionary dynamics, creating a distinctive role for population size in the process explanation of their evolution.

I argue that the chance elimination of rare but favorable alleles is explained by process explanation in section 5.2.1. In section 5.2.2, I respond to a Hempelian objection that the process explanations of the chance elimination of rare but favorable alleles that I describe are better understood on the Hempelian model. In my response to the Hempelian, I formulate what I term "the irreducibility thesis." The irreducibility thesis is important because I refer to it throughout the chapter in response to similar objections by the Hempelian concerning other phenomena that I claim are explained by process explanation. In section 5.2.3, I argue that the evolution of neutral alleles is explained by process explanation.

5.2.1 The chance elimination of rare but favorable alleles

In the two subdivisions of this section, respectively, I give an account of the following: (a) some theoretical findings of population geneticists concerning the influence of drift on rare but favorable alleles, and (b) the explanation, by process explanation, of some events described by these findings.

The cost of rarity

One of the most intriguing conclusions that population geneticists have reached concerning evolution is that an advantageous allele, if it is rare, has a high probability of disappearing due to chance. For instance, suppose that, before generation G, the A locus in some population P has no alleles; but during G, a mutation occurs, and a single organism comes to possess an allele, A_2. Suppose that the A_2 allele confers an advantage on its bearer for surviving disease, but that, in processes of indiscriminate sampling, any one of the following events occurs: its bearer dies in a forest fire; fails to find a mate; or, finds a mate, but is not passed on in sex. Despite the advantage it confers on its bearer, A_2 will disappear from the population.

This phenomenon occurs because the strength of selection for an allele depends in part on its frequency in the population. This means that, for rare alleles, drift can overwhelm selection. More precisely, the rate of change of an allele depends upon the frequency of the heterozygote. This is represented in some of the fundamental models of population genetics, such as the following.

Let p indicate the frequency of the A_1 allele; q, the frequency of the A_2 allele; s, the selection coefficient; $\bar{w}$, the mean fitness of the population; and h, the heterozygous effect. The following model represents change in the frequency of the A_1 allele due to natural selection, across generations.

$$\triangle_s p = \frac{pqs[ph + q(1-h)]}{\bar{w}} \tag{5.1}$$

What I want to call attention to is the *pqs* term in the numerator of equation 5.1.

217

This term reaches its maximum when both p and q have values of 0.5, the value at which the heterozygote is at its most frequent.[9] These relationships are represented pictorially in figure 5.1, which shows that the rate of change in frequency of an allele increases until its frequency reaches 0.5, at which point it begins to decrease.
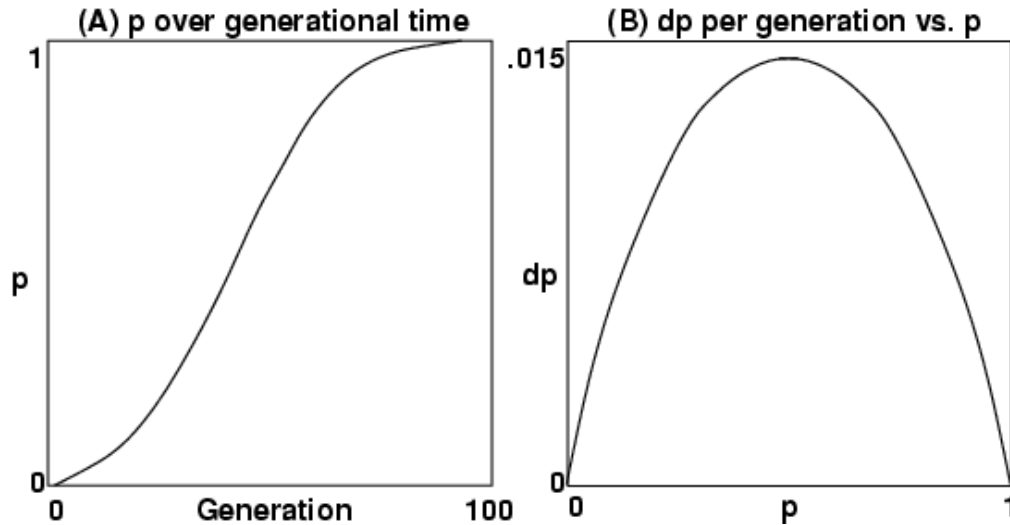

Figure 5.1: Selection, $h = 0.5$, $s = 0.1$. (A): frequency p of A_1 over generations; (B): change in p ("dp") for one generation, as a function of p. In (A), slope is steepest at $p = 0.5$; in (B), dp peaks at $p = 0.5$, showing p's influence on selection intensity. Acknowledgment: Gillespie, John H. *Population Genetics: A Concise Guide.* p. 54, Fig 3.3. © 1998 The Johns Hopkins University Press. Reprinted with permission of The Johns Hopkins University Press.

A simple model representing the chance elimination of an A_1 allele that has only one copy in the population ($p = \frac{1}{2N}$) is as follows [56, 77, eqn. 3.25]. The fixation probability[10] of the A_1 allele is represented by π; h represents the heterozygous effect; and s

_____

[9]This is because, at Hardy-Weinberg equilibrium, there are $2pq$ heterozygotes, a term that reaches its maximum when both p and q have values of 0.5.

[10]An allele is *fixed* or *reaches fixation* in a population if and only if its frequency is 100%, that is, if and only if it is the only allele at its locus in the population. Its fixation probability is just the probability that it becomes fixed.

represents the selection coefficient.

$$\pi \approx 2hs \qquad\qquad (5.2)$$

To illustrate the consequences of equation 5.2, Gillespie suggests the following.

[A] new mutation with a 1 percent advantage when heterozygous, $hs = 0.01$, has only a 2 percent chance of ultimately fixing in the population. A 1 percent advantage represents rather strong selection. In a very large population, say $N = 10^6$, 1 percent selection will overwhelm drift once the allele is at all common. Yet, 98 percent of such strongly selected mutations are lost. Think of all the great mutations that failed to get by the quagmire of rareness! [56, 78]

Explaining the chance elimination of a rare but favorable allele

The following is an explanation-seeking how-question that a biologist might ask whose answer takes the form of a process explanation invoking random drift, in the form of indiscriminate sampling.

Question 5.4 *How did the A_1 allele, the fittest of those at its gene locus in population P, disappear from P by generation G_1?*

This question presupposes the following, which I will assume to be true.

Statement 5.2 *The A_1 allele, the fittest of those at its gene locus in population P, disappeared from P by generation G_1.*

Additionally, as further background to the explanation, suppose that the following is the case.

Statement 5.3 *In population P, $hs = 0.01$, and the frequency of the A_1 allele $p = \frac{1}{2N}$, i.e., only a single copy of the A_1 allele exists in the population.*

If indiscriminate gamete sampling were the mechanism by which A_1 were eliminated from the population, then the following process explanation would correctly and appropriately answer question 5.4.

Explanation 5.1 *In generation G_1, organism O inherited a copy of the A_1 allele, and was the only organism in population P to do so, so that $p = \frac{1}{2N}$. O's gametes were created by a normal process of meiosis, so that half of O's gametes carried the A_1 allele and the other half carried another allele. O survived to adulthood, and found a mate. Sexual reproduction in population P proceeded normally, by a process of indiscriminate gamete sampling. This preserved the symmetry established by meiosis, i.e., no "bias" was introduced into the process of inheritance: no differences in the probability that different gametes are passed on were caused by the alleles carried by the gametes in question. During sex (indiscriminate gamete sampling), a gamete carrying an allele other than A_1 was passed on by O to the next generation. As no other organisms carried copies of the A_1 allele, it did not appear in generation G_2, i.e., it was eliminated from the population by that time.*

Question 5.4 is highly idealized: except under rare conditions, evolutionary biologists are not likely to have the information necessary for asking it. During the A_1 allele's brief appearance in the population, only one organism carried it, and it was unlikely to be discovered; after it disappeared, of course, it could not be discovered. Clearly, no explanation-seeking questions can be asked of events that no one is aware of. However, biologists do know that there are mutations that appear at a single locus of a single organism and nowhere else. As well, they know that it follows from equation 5.2 that there is a small probability that such alleles become fixed. Therefore, biologists are in a position to know

the following.

Statement 5.4 *For any allele A_N that has been present in population P for the interval between T and $T - \triangle$, and that is the fittest allele at locus L at time T, there is a probability P_L that during the interval between T and $T - \triangle$, there existed an allele A_L, lost due to chance, that was more fit than A_N.*

I do not know the value of the probability P_L, and I doubt that anyone else does, either. Its theoretical maximum has the same order of magnitude as the rate of mutation, and its actual value is probably several orders of magnitude smaller. Nonetheless, the truth of statement 5.4 means that there are some cases—indeed, over the history of life, a good many cases—in which rare but favorable alleles have been lost due to chance. From this, it follows that there are a good many explanations such as explanation 5.1 that are correct and would be appropriate answers to questions such as question 5.4. For this reason, I consider explanation 5.1 to be an important model for processes explanations in evolutionary biology.

5.2.2 The irreducibility thesis

A Hempelian looking at my work in the previous section might take the following view of it. The Hempelian might agree that the exclusivity thesis is false, and that the Hempelian evolutionists are wrong to believe that the warrant for the exclusivity thesis can be found in Hempelianism. That is, the Hempelian might agree that chance phenomena such as the chance elimination of a rare but favorable allele can indeed be explained, even though they cannot be explained by natural selection. Nevertheless, the Hempelian would not accept that such explanations can be provided by the strategy of process explanation:

221

true to the name, the Hempelian believes that Hempelianism is true. As I have been emphasizing throughout this dissertation, Hempelianism is incompatible with the claim that process explanations carry explanatory force. The Hempelian would claim that there are Hempelian explanations of the chance elimination of rare alleles.

There is a particular objection that the Hempelian would want to raise to my claim that explanation 5.1 adequately answers the question "How did the A_1 allele, the fittest of those at its gene locus in population P, disappear from P by generation G_1?" (question 5.4, page 219). The objection is that explanation 5.1 is *incomplete*. I consider the general strategy of this objection in chapter 3, and the Hempelian would want to press the issue in the case of the elimination of rare alleles. I have already argued in chapter 3 that the general strategy of this response is mistaken. To eliminate any doubts about the adequacy of my chapter 3 response to the Hempelian, however, I want to respond to the objection as it might be raised in the case of the elimination of rare alleles.

I proceed now to articulate the Hempelian objection, and then, in response, to articulate what I term the "irreducibility thesis," a claim that embodies my response to the Hempelians' incompleteness objection. The reason I cast my response to the incompleteness objection in the form of an explicitly stated thesis is that I will refer to it in subsequent sections when I consider objections similar to the one raised here.

As I indicate in chapter 3 (page 68), the Hempelian would explain the claim that process explanations are in general incomplete in the following manner. Suppose that a process explanation of an event E describes a sequence of events S, and that S consists of $i = 1 \ldots N$ events S_i, where $S_N = E$. The Hempelian believes that each S_i should be

viewed as an event to be explained in its own right, in the Hempelian manner—that is, by citing the appropriate causal laws, together with a description of S_{i-1}, to show why S_i occurred. Putting the explanation of each of these events together explains the process by which E occurred, the Hempelian believes, by showing why each stage in the process resulting in E occurred. The idea is that a process explanation P is incomplete in the sense that it requires supplementation with the further information contained in the Hempelian explanations used to reconstruct it.

Applied to the chance elimination of rare alleles, this strategy plays out in the following manner. Explanation 5.1 describes a sequence of events leading up to the chance elimination of an advantageous allele from a population: the distribution of the A_1 allele into approximately half of O's gametes by meiosis; O's survival to adulthood; O's finding a mate; and so on. According to the Hempelian, placing these events in the appropriate sequence does not provide enough information to be explanatory. What is required is to describe the causal laws that account for why the A_1 allele was inherited by O, given conditions antecedent to O's birth; why the A_1 allele was distributed into approximately half A's gametes, given the conditions antecedent to meiosis; why O survived to adulthood, given the conditions in its environment and its phenotype; and so on.

The Hempelians' incompleteness charge is powerful in this case because it suggests that, unless human knowledge and capability for observation and computation expand to levels now unthinkable, it is not possible for humans to explain the elimination of an advantageous allele from the population: except with extraordinary knowledge, powers of detection, mathematical analysis, and simulation far beyond what what we currently possess, the

information needed to render explanation 5.1 as a series of linked Hempelian explanations, each complete in its own right, cannot be obtained by anyone. In order to see this, consider the event described by the following statement, which appears in explanation 5.1.

Statement 5.5 *During sex (indiscriminate gamete sampling), a gamete carrying an allele other than A_1 was passed on by O to the next generation.*

As I have just been arguing, the Hempelians' view is that explaining the chance elimination of a favorable allele requires a Hempelian-style explanation of the event described by statement 5.5, which is one of the stages described in explanation 5.1 above. This is just what the Hempelian believes to be impossible, given the state of knowledge today. This can be seen by considering what information would be required to formulate a good candidate for a Hempelian D-N explanation of statement 5.5. I think that a Hempelian would agree that the following further statements are essential to a D-N explanation of the event in question.[11]

Statement 5.6 (Organism O's reproductive physiology) *At time T_i, as a result of the processes of reproductive physiology that conduct alleles through meiosis and into the gamete pool, an allele other than an A_1 allele is conducted into one of O's gametes G_g, which is located in space at location L_{ig}, and had other physical, chemical, and biological properties P_{ig}; and the mating system of O's species S_s is characterized by physical, chemical, and biological state P_{is}.*

Statement 5.7 (Causal law of fertilization) *If at time T_I, gamete G_G occupies location L_{ig} and has properties P_{ig}, and the mating system at issue has properties P_{is}, then*

---

[11] I discuss D-N explanation in chapters 1 (pages 20 - 21) and 2 (page 42).

there exists an unique gamete G_H that G_G fuses with at time T_F (later than T_I) to form a

zygote.[12]

As the annotation preceding statement 5.6 indicates, that statement concerns the conditions local to O just prior to the time that mating begins, describing the position in space of one of O's gametes, as well its properties at several levels of organization; additionally, statement 5.6 describes the properties of the mating system in O's species. Statement 5.7 indicates that it is nomically necessary that if the conditions described in statement 5.6 obtain, the gamete described in the latter will fuse with a gamete of O's mate. Crucially, the gamete in question does not carry the A_1 allele. This is just the event whose explanation is of interest in this case, that is, the event described by statement 5.5.

Together, statements 5.6 and 5.7 form the premises of a valid deductive argument whose conclusion is statement 5.5. According to Hempel, this explains the event described in statement 5.5 by showing why that event occurred, that is, by showing why that event ought to have been expected. Similar probabilistic statements could be constructed to satisfy requirements of the models of other Hempelians such as Salmon or Railton, or to satisfy those of Hempel's own I-S model.[13]

I would now like to take up the issue of how probable it is that human beings will ever be in a position to gather the information that would be required to confirm statements 5.6 and 5.7. I think it is clear that there is at best a diminishingly small

---

[12]Statements 5.6 and 5.7 need not be understood to refer to *particular* locations; rather, they may be construed as referring to locations relative to other organisms in O's population, or relative to some other property of the mating system that is characteristic of the kind of mating system in O's species. In the same way, Kepler's laws need not be construed as referring to the Earth and Sun, but about any bodies that stand in appropriate relationships to one another. This is important because a statement S cannot describe a law of nature if S refers to a particular object or location.

[13]I discuss these models in chapter 2.

probability that anyone will ever be in a position to know statements such as statement 5.6 and 5.7. It is unclear how anyone could place him or herself in a position to learn the spatial location of a gamete before its bearer mates, or any causal laws about how a gamete behaves during sex.

This is precisely why models of sex are probabilistic: there are too many parameters, and too many are inaccessible, so that at best, all that can be said is that there is a certain probability that allele frequencies will take on a certain value after one generation of mating, except in limiting cases.[14] The Mendelian mechanism of inheritance maintains a physical symmetry among gametes, so that no gamete differs from any other in its probability of being passed on.[15] It seems reasonable to believe that there are some physical differences between gametes that account for one's being passed on and others not being passed on. Nevertheless, there is no reason to think that these differences are systematic or biologically interesting in any way.

To see what I mean by this, consider a thought experiment: suppose that, in some Mendelian species, there were some way of determining the various mechanical, chemical, and biological properties of all of the gametes that were successfully transmitted to the next generation. This thought experiment would show that the gametes in question do not exhibit any one property or set of properties in common that caused them to be passed on. Rather, a variety of differences in a broad range of otherwise unrelated parameters account for their being passed on. This range of parameters includes spatial location in the

---

[14]One limiting case obtains if an allele has a frequency of 100%; assuming that no mutation occurs, there is a probability of 100% that allele frequencies do not change as a result of mating. The other obtains if the frequency of one allele is 0%.

[15]Recall that I argued in chapter 4 (section 4.3) that the absence of differences in the probability of being passed on across gametes is the distinguishing feature of indiscriminate gamete sampling.

mating system; very small scale irregularities in their environment, including, in some cases, environments internal to the reproductive physiology of organisms during sex; differences in motility, for sperm; and slight differences in their mechanical and biochemical properties.

Indiscriminate sampling is analogous to coin tossing, which I consider in chapter 3 (pages 75 - 83) in connection with the game played by Stoppard's Rosencrantz and Guildenstern. It is not clear what the "laws of coin tossing" would look like, but there is no doubt that they would be be enormously complex, and that they are beyond the current capabilities of human observation and calculation as well as any that we are likely to obtain. The laws of coin tossing would have to be framed in terms of the parameters causally relevant to the outcome of a coin toss, which include "micro-states," of the coin, such as the distribution of the mass of the coin and its shape; the nature of the tossing device; and the environment of the toss, including properties of the air into which the coin is tossed.[16]

I have been considering the case of Mendelian reproduction here; however, I believe that the same considerations apply to indiscriminate parent sampling, and that it is impossible for human beings to know statements analogous to statements 5.6 and 5.7 in the case of the latter. Take catastrophic drift. Suppose that a storm is an agent of indiscriminate parent sampling, causing catastrophic drift; consider a thought experiment analogous to the one that I described above in the case of indiscriminate gamete sampling: all the organisms that survive the storm are assessed for the properties causally relevant to their survival. I do not see that this assessment would reveal any systematic or biologically interesting differences among the organisms. The differences among them would have to do

---

[16]I believe that Keller's [77] model of the kinematics of coin tossing do not appreciably increase the chances that the "laws of coin tossing" will ever be formulated. I comment on this in chapter 3 (note, page 80).

with their location relative to various hazards created by the storm, which do not depend on their phenotype in any way.

This concludes my exposition of the incompleteness objection, as the Hempelian would make it in the context of the chance elimination of a favorable allele. To be clear, let me summarize my account of the objection and its significance, before answering it. The Hempelian believes that a process explanation does not explain an event by describing a sequence of events leading up to that event. Rather, each of the events in the sequence of events referred to in a process explanation requires explanation in and of itself. In order to provide this explanation, further information is required, viz., laws of nature and particular facts required to explain each event in the sequence S in a Hempelian manner.

The Hempelian sees the incompleteness objection as particularly damning for the claim that indiscriminate sampling can be explained by process explanation. The argument is that limitations on human knowledge-gathering abilities place indiscriminate sampling beyond the reach of explanation entirely. There is no way that human beings will ever be able to formulate the appropriate laws of nature, or to learn the appropriate facts about particular events to be explained. This is the point of my discussion immediately above of whether there are causal laws of Mendelian reproduction.

My response to the incompleteness objection as it is raised in the context of indiscriminate gamete sampling parallels the response I articulated in chapter 3 to the objection as it is raised in the general context. The problem is the Hempelians' universalism. In chapter 3 (pages 73 - 86), I suggest that the *level of organization* at which an answer to an explanation-seeking question is to be answered is fixed by context; this suggestion applies

in the context of the chance elimination of a favorable allele, as well. Let me explain.

No evolutionist asking question 5.4, "How did the A_1 allele, the fittest of those at its gene locus in population P, disappear from P by generation G_1?" wants to know the particular circumstances under which a particular gamete met the conditions required by the "laws of sex" for being transmitted to the next generation, and that that gamete lacked the A_1 allele; nor does he or she want to know what those laws are. This information is at too low a level of organization to be of concern. Rather, the information requested is, Which evolutionary process eliminated the allele from the population? Mutation? Migration? Drift? Natural selection clearly cannot be responsible, as it culminates with the continued survival and proliferation of fitter alleles, not with their elimination from the population. The appropriate level of organization in this case is the "macro-level:" what condition of the mating system obtained, causing the probabilistic relationships in virtue of which indiscriminate gamete sampling occurs? This is just the question that explanation 5.1 answers; it does not need to be supplemented with laws or with any facts about a particular episode of mating.

Although I have formulated my argument here in terms of indiscriminate gamete sampling such as would result in the chance elimination of a rare but favorable allele during sex, I believe that it applies to indiscriminate parent sampling, as well. What would be required to formulate the argument for indiscriminate parent sampling would be to state it in terms of organisms and variant traits rather than gametes and alleles, and in terms of parent sampling agents, rather than structures of the mating system.

A Hempelian would want to extend his or her argument against my contextualist

response in a manner that I consider in detail in chapter 3. My contextualist response, the Hempelian would claim, does not show that the Hempelian view is incoherent. At most what my response shows is that I have failed to appreciate the full scope of Hempelianism: Hempelianism entails that only a Laplacian demon can have a complete explanation for any given event. Human beings, at least in their present state of knowledge and knowledge-gathering abilities, must recognize that there are limits to what can be explained.

As a consequence of these limits, the Hempelian would claim, many of the phenomena that I claim are explained by process explanation, including the chance elimination of a favorable allele, cannot be explained by human beings at all. It does not matter that the appropriate process explanations cannot be reconstructed as Hempelian explanations, given the constraints of the contexts in which they are asked: as Hempelians make no allowance for context, they would point out, it is a mistake for the contextualist to claim that context be invoked in an argument against their position.

Let me sketch my chapter 3 response to the Hempelian on this point. Hempelianism fails to account for an important phenomenon: a wide range of linguistic practices (and their accompanying concepts) are typically considered explanatory; these practices and concepts are directed toward generating understanding, and often succeed at doing so. This is important because generating understanding is the *sine qua non* of explanation. Hempelianism, according to which explanations have a single, canonical form, cannot encompass this diversity of practices, in contrast with contextualism, which is formulated precisely with this diversity of practices in mind.

This shifts the burden of argument in the direction of the Hempelian. The chal-

lenge for the Hempelian is to argue that this diversity of explanatory practices ought to be ignored: of all the putatively explanatory utterances, only those conforming to Hempelianism really are in fact explanatory. Alternatively, perhaps the Hempelian is proposing an ideal for explanatory practices. As I indicate in chapter 3, I find it highly unlikely that the Hempelian will be able to bear this burden of argument. In the context of the present section, this means that my account of the process explanation of the chance elimination of a favorable allele has been vindicated.

I consider my response to the incompleteness objection concerning the chance elimination of a favorable allele to be a particularly striking defeat for the Hempelian. By showing that the Hempelian is mistaken about the incompleteness of process explanation, I show that events that the Hempelian believes cannot be explained at all by human beings can in fact be explained by them. Admittedly, it may not be possible to explain *why* certain events occur—why was the favorable allele eliminated?—but it is possible to explain *how* those events occur by process explanation.

A persistent Hempelian will want to raise this kind of incompleteness objection to each of the claims that I make below to the effect that some evolutionary event occurring by indiscriminate parent or gamete sampling can be explained by process explanation. My response in each case will be the same. As I have just argued, process explanations of events occurring by indiscriminate sampling are complete, requiring no further information in the form of laws or facts obtaining prior to the event to be explained; the "macro-level" statistical relationships among the appropriate entities (i.e., gametes or organisms) that obtain due to a mechanism of indiscriminate sampling are sufficient.

In order to simplify my response to the Hempelian in subsequent sections of this chapter, I would like to formulate a thesis that embodies this claim about the completeness of process explanations of indiscriminate sampling, as well as the conclusion of the general argument that I make in chapter 3. I term this thesis "the irreducibility thesis."

Statement 5.8 (The irreducibility thesis) *If an explanation-seeking question is appropriately answered by a process explanation, then it cannot be appropriately answered by a Hempelian explanation or by any set of Hempelian explanations.*

The idea is that a process explanation cannot be adequately reformulated as a Hempelian explanation or set of Hempelian explanations, because doing so would violate contextually imposed requirements on the level of organization at which the explanation should be formulated. Hempelian explanations are pitched at too low a level.

5.2.3   Molecular evolution

In two subdivisions of this section, respectively, I provide an account of (a) theories about evolution at the molecular level, and (b) how one such theory is used in process explanations of such evolution.

The DNA revolution in evolutionary biology

As Dietrich [35] and Gayon [54, sec. 10.3] indicate in their historical accounts, the molecular revolution in biology took approximately fifteen years to reach evolutionists. Until the mid-1960's, evolutionists studied the phenotypic characters of organisms, using techniques of breeding and genetic dissection to infer the genetics of individuals and popula-

tions. New molecular techniques permitted the direct assessment of the genetic structure of populations. Using these techniques, evolutionists were able to determine which nucleotide base pairs many organisms possess at various gene loci, and which amino acids are encoded by the genes at those loci. Key early work in this area was carried out by Lewontin [89], Lewontin and Hubby ([90] and [69]), and Zuckerandl and Pauling [163].

Discoveries about molecular evolution motivated J. L. King and T. H. Jukes [81] and Motoo Kimura ([78], [79], and [80]) to formulate what has come to be known as the "neutral theory of molecular evolution." This theory, also termed the "theory of non-Darwinian evolution," differs sharply from the Darwinian explanation offered for phenotypic evolution. Proponents of the theory claim that the majority of differences between genes at the molecular level are nonadaptive, and that random drift and mutation are the main mechanism of molecular evolution. The neutral theory has been extended by Ohta ([110], [111], and [112]) into what is called "the nearly neutral theory," which posits a balance of neutral, slightly deleterious, and selectively advantageous alleles. Nei [107] and Gillespie [55] provide comprehensive accounts of many of the key issues concerning the neutral theory and subsequent developments, and are widely regarded as touchstones in the discussion.

It is important to note that the neutral theory is not concerned only with "junk DNA," that is, DNA that is apparently functionless, playing no role in the life of the organism. Kimura [79, 100] states that "the neutral theory, I should make clear, does not assume that neutral alleles are functionless but only that various alleles may be equally effective in promoting the survival and reproduction of the individual." Nei [107, 411] concurs, stating that "neutral alleles are not functionless genes but are generally of vital

233

importance to the organism." Indeed, examples of neutral alleles include those responsible for elements of the hemoglobin molecule and cytochrome c, which have important functions.

Indiscriminate gamete sampling is an important mechanism of drift occurring among neutral alleles. Sex occurs each generation, introducing an element of chance into the evolution of a population regardless of ecological conditions. Thus, the claims of the neutral theory do not depend upon the regular occurrence of catastrophes or short-term fluctuations in the environment.

In my account of process explanation of molecular evolution, I will focus on a simple formulation of the neutral theory. The key claim of this simple formulation of the neutral theory is as follows. Let k represent the probability of fixation of a particular allele of interest in a given population of interest, and μ represent the rate at which alleles of that kind appear in the population by mutation, per generation [56, 30, eqn. 2.11].

$$k = \mu \qquad (5.3)$$

Equation 5.3 asserts that the fixation probability of a neutral allele is the rate at which new copies of it enter the population by mutation each generation.

Process explanation of the evolution of neutral alleles

An explanation-seeking how-question that calls for a process explanation of the evolution of neutral alleles at the molecular level, and that invokes drift, in the form of indiscriminate parent and gamete sampling, is as follows.

Question 5.5 *How, by generation G in population P, at gene locus L, did the A_1 allele become fixed in P?*

This question presupposes the following statement, which I will assume to be true.

Statement 5.9 *In generation G, the A_1 allele at the L locus is fixed.*

Additionally, consider the following statement.

Statement 5.10 *A_1 confers neither an advantage nor a disadvantage for survival or reproduction on its bearers, relative to other alleles at the L locus—A_1 and other alleles at the L locus are neutral with respect to one another.*

Supposing that statement 5.10 is true, the process explanation that answers question 5.5 is as follows.

Explanation 5.2 *In generation $G - \triangle$, A_1 alleles entered the population by mutation at a rate of μ alleles per generation. There were some changes in A_1's frequency during the juvenile stage of the life cycle in generation $G - \triangle$ due to indiscriminate parent sampling. Sex in generation $G - \triangle$ proceeded normally: meiosis distributed alleles in equal proportions among the gametes; mating was random; and there was no meiotic drive. This resulted in changes in the frequency of the A_1 allele due to indiscriminate gamete sampling. Each of the events in generation $G - \triangle$ occurred again in every generation between generation $G - \triangle$ and generation G, change in A_1's frequency during that interval being sufficient to drive A_1 to fixation.*

A Hempelian would argue that explanation 5.2 is incomplete, a conclusion he or she arrives at by extending, to explanation 5.2, the argument that process explanations are in general incomplete.[17] The Hempelians' view is that each event in the sequence of events

---

[17]See my discussion of the incompleteness objection in section 5.2.2.

indicated in explanation 5.2 requires explanation in its own right: each must be explained by reference to conditions that obtained prior to its occurrence and to laws of nature, none of which are mentioned in explanation 5.2. The Hempelian takes this supposed incompleteness to be fatal to explanation 5.2, because adding this information to explanation 5.2 would result in replacing it with a set of linked Hempelian explanations. Indeed, the Hempelian believes that the incompleteness objection is fatal to the endeavor of explaining the evolution of neutral alleles, because he or she does not believe that such a replacement can be carried out in this case. As in all cases of indiscriminate sampling, the required facts and laws are beyond the reach of human comprehension.

I do not believe that the incompleteness objection poses a serious threat to explanation 5.2, or to the project of explaining the evolution of neutral alleles. On the strength of my section 5.2.2 argument for the irreducibility thesis, I conclude that explanation 5.2 requires no further information. It meets contextually-supplied criteria for completeness because it cites the correct level of organization, viz., the "macro-level" of the Mendelian mechanism of reproduction, as opposed to the "micro-level" of the fate of particular organisms and gametes.

5.3 Explaining Drift by Reference to N

My aim in this section is continuous with my aim in the previous section, because both form a part of my argument that, contrary to the Hempelian evolutionists, there are events that occur by drift that are explained by process explanation. Unlike the previous section, however, this section concerns events whose process explanation makes reference to

236

population size N. I group these events together because the distinctive patterns of evolution that drift induces in small populations warrant special treatment in and of themselves, and occupy a prominent role in process explanation of the evolution of such populations.

Before describing each of these phenomena and their process explanation by drift, I provide a general account of the theme that unifies them, drift in small populations (section 5.3.1). In the three subsequent sections, I describe three kinds of events and their process explanation: the shifting balance process (section 5.3.2), the origin of species (section 5.3.3), and punctuated equilibrium and its effects on the shape of phylogeny (section 5.3.4).

5.3.1 Drift in small populations

The evolutionary dynamics of small populations present one of the most striking phenomena in evolution: small populations tend to evolve rapidly and unpredictably, often in unchanging environments. It turns out that drift is in large part responsible for these dynamics, the nature of which I explain in this section. I begin by considering a theoretical model of a limiting case of small population size, elaborated by Sewall Wright in 1931.

> The extreme case is that of a line propagating by self fertilization which may be looked upon as a self contained population of one. In this case, 50 percent of the factors with equal representation of two allelomorphs (that is, in which the individual is heterozygous) shift to exclusive representation of one of the allelomorphs in the following generation merely as a result of random sampling among the gametes. [156, 107]

The case described by Wright is a hypothetical "population" consisting of a single organism that is self-fertilizing, hermaphroditic, heterozygotic, and that is not subject to mutation, migration, or natural selection. Wright's suggestion is that, in a single generation

of self-fertilization, there is a 50% chance that this individual will produce a homozygote, which will cause one of the following two outcomes.

1. An A_2 homozygote is created, so the A_2 allele reaches a frequency of 100%; or

2. An A_1 homozygote is created, so the A_1 allele reaches a frequency of 100%.

In either case, drift irreversibly destroys all genetic diversity in the population—the most drastic evolutionary change that a population can undergo, short of its extinction or transformation into a new species. Moreover, this radical change occurs in the time span of a single generation, and because each allele has a 50% chance of being eliminated, its direction is random. In contrast, in a population in which a large number of offspring are produced, there is a diminishingly small chance that such radical changes occur in so short a time. It is a matter of combinatorics. Many permutations of homo- and heterozygotes are possible in such populations, so they can easily maintain a diversity of alleles.

The following model, described by Gillespie [56, 44, eqn. 2.18], describes the relationship between drift and population size in a general manner.

$$Var\{\triangle p\} = \frac{pq}{2N} \tag{5.4}$$

Equation 5.4 states that the variance of the change in the frequency of the A_1 allele per generation under the influence of drift alone is a function of the frequency of both alleles, pq, and population size, $2N$. Attending more closely to the latter, equation 5.4 states that the variance of $\triangle p$ each generation is inversely proportional to population size: greater deviations from the mean are to be expected due to drift in small populations than in large ones. To see the consequences of this, consider the following.

Consider two ensembles **I** and **II**, each consisting of populations of heterozygotes subject only to drift; however, suppose that **I** consists only of large populations, while **II** consists only of small populations. According to equation 5.4, more populations in **II** will exhibit evolution in any given amount $\triangle p$ each generation than will populations in **I**. Alternatively, equation 5.4 means that some allele or other will become fixed in a given number of populations of **II** sooner than in the same number of populations in **I**. Both of these ways of understanding equation 5.4 suggest that drift acts more quickly in small populations.

5.3.2 The shifting balance process

In Sewall Wright's view, what is termed the "shifting balance process"—a combination of drift, selection, and migration—is most conducive to rapid adaptation; natural selection and mutation alone are not as effective. In keeping with the general theme of this section, the dynamics of evolution in small populations due to drift play a central role in causing adaptive evolution by the shifting balance process.

In each of the three subdivisions of this section, respectively, I describe (a) the shifting balance process; (b) a visual model of the process, the "adaptive landscape;" and (c) how the shifting balance process fits into process explanations of adaptive evolution.

The shifting balance process

Throughout his career, Sewall Wright advanced the shifting balance theory of evolution. He stated the basic tenets of the theory in the early 1930's in his first publications on evolutionary theory; these ideas are presaged in his work on domestic breeding, much of

239

which he completed before 1930. He continued to advance the theory, with relatively little modification, until his last publications, some fifty years later. Provine [118, ch. 12] argues convincingly that changes that Wright made to the theory in the exchange with Fisher in the late 1940's amount to refinements on the major elements of the theory, which remained constant. The theory continues to be actively discussed today.[18]

Wright claims that the theory describes conditions under which a population will reach its maximum mean fitness in the shortest possible time. For the shifting balance process to occur in a population P, P must be large, but structured into small sub-populations ("local isolates" or just "isolates"). Though the shifting balance process cannot occur if there is extensive interbreeding among the isolates, it requires that they be linked by a low level of migration ("gene flow").

The shifting balance process has three stages. The first stage occurs when, in one of the isolates, random drift creates a gene combination that is more fit than any already existing in it. For any isolate I, there is a high probability that a large change in the frequency of alleles at many loci in I occur in the time span of one or a few generations. This is because the isolates are small in size: they will exhibit the rapid, extreme changes in the magnitude and direction of evolution characteristic of drift in such populations that I describe in section 5.3.1.

---

[18]Provine [118, chs. 2 - 5] reviews Wright's ideas on domestic breeding, which Wright published during his tenure at the USDA ([154], [158], [159], and [160]). His early formulations of the theory appear in the early 1930's ([156, 158], [162, 168]), a later formulation appearing in 1980 [153, 630]. The relevant part of the controversy with Fisher spans a range of publications ([47], [48], [157], and [161]); Skipper ([136], [137]) reviews philosophical issues raised by the Fisher-Wright debate. Some of the more prominent scientists working on the theory today include Wade and Goodnight [145], who claim to have confirmed the theory empirically, and who exchange views with Coyne and his coauthors ([25], [60]). Mallet and his collaborators ([92], [93], [94], and [95]) suggest that mimicry and warning colors in butterflies evolve by the shifting balance process.

According to Wright, change in many loci at once is particularly important for adaptive evolution. Wright views traits that affect fitness to be under multi-locus control: novel traits require a novel gene *combination*; the increase in frequency of a single rare allele is not sufficient. As I suggest above, drift in the isolate (the first stage of the shifting balance process) produces just such a novel gene combination; the small population size causes many gene loci to fluctuate randomly all at once.

Indiscriminate gamete sampling (sex) is a particularly important mechanism of drift in the shifting balance process. This is because, as I pointed out in connection with the neutral theory, sex occurs regularly, introducing an element of chance into the life cycle of a population regardless of its ecological conditions. Additionally, indiscriminate gamete sampling is non-lethal, unlike indiscriminate parent sampling that might occur in connection with catastrophic drift; events causing mortality required for catastrophic drift in a large population would most likely destroy the small isolates required for the shifting balance process.

The second stage of the shifting balance process occurs if organisms with the novel gene combinations created in the first stage are able to spread their genes throughout the isolate by natural selection. Wright terms this stage "intrademe selection." The stage of intrademe selection represents a point of difference between the shifting balance theory and the neutral theory of molecular evolution. According to the latter, drift is a mechanism by which alleles appearing by mutation are driven to fixation. In contrast, according to the shifting balance theory, drift is a mechanism for generating new gene combinations; natural selection is the mechanism of their increase.

Intrademe selection sets up the conditions for the third stage of the shifting balance process, "interdeme selection" (as opposed to the *intra*deme selection of the second stage). Interdeme selection works as follows. As organisms with the favorable gene combination become more frequent in the isolate in which they originated, some will migrate to nearby isolates. Just as they succeeded by selection in the isolate in which they originated, they will do so in isolates to which they spread by migration. As each isolate is colonized by organisms bearing the favorable combination of genes, the mean fitness of the population rises rapidly.

A visual model: The adaptive landscape

The shifting balance theory provides the conceptual and theoretical backdrop to an especially influential[19] and vivid visual metaphor for evolution, "rugged fitness landscapes" or "adaptive landscapes" (figure 5.2, page 243). The "landscape" is an N-dimensional space in which $N - 1$ dimensions represent allele frequencies at each of $N - 1$ gene loci. The Nth dimension is the mean fitness of the population, $\bar{W}$. A point in the space represents a possible genetic structure of a population, graded according to its mean fitness. Such a multidimensional space is most easily envisioned as a "landscape" in the highly idealized case of $N = 3$ dimensions, that is, for $N - 1 = 2$ gene loci. Each point along the X-axis represents the frequency of genes at one locus; each point along the Z-axis represents

---

[19]As Gayon [54, 347] suggests, the concept of the adaptive landscape has been enormously influential in the development of evolutionary biology since the 1930's. Provine [118, 307 - 8] points out that they appear in each of the many editions of Dobzhansky's *Genetics and the Origin of Species* [37], a text read by virtually all evolutionists in the period between 1930 and 1970. Additionally, they figure prominently in Simpson's landmark *Tempo and Mode in Evolution* [134, 89 - 93], which incorporates a phenotypic interpretation of the landscape, a formal theory of which is developed by Lande [86]. The landscapes are also central to Kauffman's provocative and influential *Origins of Order* [76], which describes the dynamics of self-organizing complexity in terms of the landscapes.

the frequency of genes at the other locus. The Y-axis—the "altitude" of the landscape—represents mean fitness $\bar{W}$.
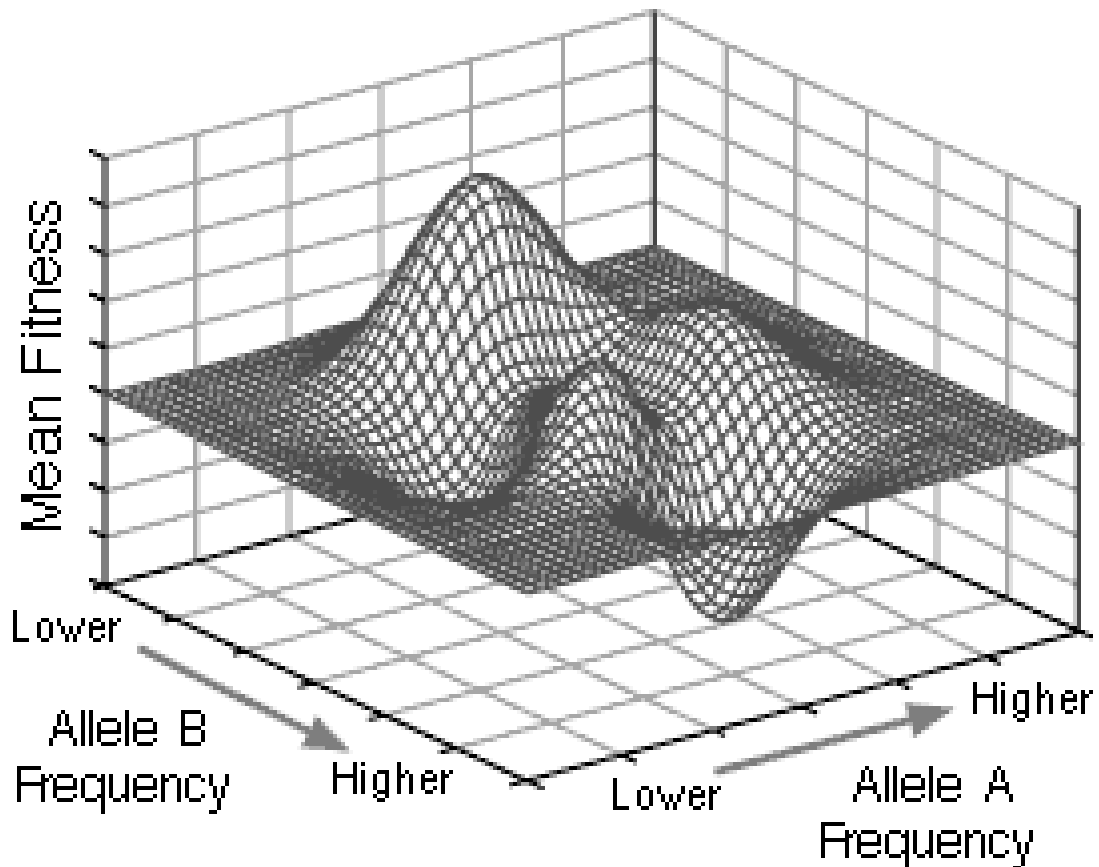

Figure 5.2: Wright's adaptive landscape and the shifting balance process. The X- and Z-axes represent the frequencies of alleles at different gene loci α (alleles A and a) and β (Alleles B and b). The Y-axis represents the mean fitness of the population as a function of the frequencies of alleles at these loci. Natural selection causes a population to ascend the nearest fitness peak, but cannot carry the population across a valley to a higher peak, even if such a peak exists. For this, according to Wright, the combination of random drift, natural selection, and migration characteristic of the shifting balance process is required. See text for further explanation. Image provided by the UC Museum of Paleontology, from a graph by Rodney Dyer.

Each point in such a three-dimensional space may be described by an ordered triple $\langle x, y, z \rangle$. Each triple describes a possible state of the population, grading each set of

allele frequencies according to its fitness. The first and third elements of this triple are the frequencies of alleles at each of the two loci, and the second element of the triple is the mean fitness of a population with those allele frequencies. Projected into three dimensions, a point described by a fitness-element that is greater than that of neighboring points is a fitness "peak," while a point described by a fitness-element that is less than that of neighboring points is a fitness "valley."[20]

An adaptive landscape with three dimensions might have only one or two peaks. However, Wright [162, 163] suggests that adaptive landscapes with the number of dimensions required to represent an animal species in nature are quite "rugged" in the sense that they have many peaks and valleys. Selection, which, as a general rule,[21] increases the mean fitness of a population, will cause the population to ascend to a local maximum: "in a rugged field of this character, selection will easily carry the species to the nearest peak" [162, 163]. However, the local maximum may not be a global maximum. "[T]here may be innumerable other peaks which are higher but which are separated by 'valleys' " [162, 163]. This informs Wright's provocative suggestion that "the problem of evolution ... is that of a mechanism by which the species may continually find its way from lower to higher peaks in such a field" [162, 163 - 164].

The surfaces of selective value provide a powerful visual tool for understanding the shifting balance process (see figure 5.2, page 243). Representing a natural population on an adaptive landscape would require a large number of dimensions $N - 1$, corresponding to the

[20]Provine [118, 307 - 317] points out that Wright sometimes claims that the X- and Z- axes represent alleles, so that a point on the landscape represents the genes possessed by an organism.

[21]I say that selection increases the mean fitness of a population "as a general rule" because, in cases of frequency-dependent selection, it does not do so. The landscapes are of no use for representing selection of this kind.

large number of gene loci in the population. As I suggest above, Wright claims that such an adaptive landscape would have many peaks and valleys. Represented on such a surface, the shifting balance process would look as follows.

At the start of the process, the population might occupy a peak other than the tallest on the landscape, i.e., a local maximum of mean fitness that is not a global maximum; or, it might occupy a saddle between local maxima, or a valley beneath such local maxima. Because of the population's proximity to local maxima and its separation by valleys from the global maximum, the latter is inaccessible by selection, which will take the population to the top of the local peak, causing it to remain there. Drift, in contrast, can carry a population across a valley to the slopes of a global peak, as drift is not constrained to increase the mean fitness of the population, as natural selection is.

This is what occurs in the first stage of the shifting balance process: organisms with a novel favorable gene combination created by drift in one of the isolates will break away from the rest of the population, moving to the slopes of a global maximum. If these organisms spread by natural selection in their isolate, they will cause that isolate to climb the global peak; this is intrademe selection. Next, suppose that migrants spread the favorable gene combination to other isolates. Then the other isolates will ascend the global peak as the favorable gene combination spreads throughout each. This is interdeme selection, the third stage of the process. Thus, the entire population can be made to ascend the global peak. In this manner, as Wright suggests, the shifting balance process provides "a trial and error mechanism by which in time the species may work its way to the highest peaks in the general field" [162, 167]—trial and error, because random changes in the population due to

drift provide the means for escaping local maxima and attaining the global.

Explaining adaptive evolution with drift

An explanation-seeking how-question a biologist might ask whose answer requires reference to the shifting balance process—and therefore, random drift—is as follows.

Question 5.6 *How did species S ascend to the global peak in its adaptive landscape? (Asked at time T.)*

I will suppose that the following is true.

Statement 5.11 *Species S, which "inhabits" a rugged adaptive landscape, is at the global peak in that landscape at time T, and at time $T - \triangle$, it was at a low-lying local peak, surrounded by deep valleys.*

Statement 5.12 *Species S ascended to the global peak by the shifting balance process, which occurred between time T and $T - \triangle$.*

Given statements 5.11 and 5.12, an answer to question 5.6 is as follows.

Explanation 5.3 *At $T - \triangle$, species S was large, and was structured into small isolates connected by occasional migrants; as indicated in statement 5.11, it occupied a low-lying local peak in the adaptive landscape, surrounded by deep valleys. As indicated by statement 5.12, the shifting balance process occurred in S, proceeding as follows. Within an isolate I, which exhibited the special evolutionary dynamics associated with drift in small populations, drift occurring by indiscriminate gamete sampling created a new and highly favorable gene combination (stage 1, random drift). Natural selection of the favorable gene combination lifted*

246

I to the global peak (stage 2, intrademe selection). Migrants from I carried the favorable gene combination to other isolates, lifting all of them to the global peak by time T (stage 3, interdeme selection).

On account of the role of indiscriminate gamete sampling in explanation 5.3, a Hempelian might raise the incompleteness objection that I have described above in section 5.2.2. I would respond to this objection in the case of explanation 5.3 in the same way that I responded to it previously: on the strength of my argument for the irreducibility thesis (statement 5.8, page 232), such an incompleteness objection would be mistaken, and poses no threat to the claim that explanation 5.3 provides an adequate answer to question 5.6.

5.3.3 The origin of species

The origin of species (or just "speciation")—a phrase now heavily laden with historical significance—occurs when a new biological individual, a species, comes into existence by way of separation from another, its parent species. In this section, I will focus on Hampton Carson's "flush-crash" theory of the origin of species, which posits a nonadaptive phase during speciation.

In the subdivisions of this section, respectively, I describe (a) the conception of species that Carson believes to be correct; (b) Carson's "flush-crash" theory of speciation; (c) my position in a controversy about a phenomenon known as "the founder effect;" and (d) the role of drift in process explanations of speciation invoking Carson's theory.

The biological species concept

Carson's theory applies most clearly to what is termed the "biological species concept" [20, 83 - 84]. The central claim of the biological species concept may be stated as follows.

Analysis 5.1 *Population P is a species if and only if the following obtains: all and only organisms of the appropriate sex in P can mate with one another to produce viable offspring.*

Analysis 5.1 rules out populations whose members can mate with organisms not in the population to produce viable offspring; and it rules out populations that have members that cannot do so with others of the appropriate sex that are in the population. The condition described by analysis 5.1 is usually called "reproductive isolation." The biological species concept describes what might be termed a "unit of evolution," a role for which it is well-suited because of the hereditary nature of evolutionary change. Because evolutionary change is hereditary, a group of organisms is a part of an evolving biological entity if and only if they can exchange genes among one another.

As it turns out, analysis 5.1 does not precisely describe many of the entities that scientists and non-scientists alike consider to be species. Pandas in the zoo are notorious for not breeding with one another; but surely they remain members of the same species. Likewise, many clearly recognized species hybridize with other species regularly. The biological species concept does not apply in a clear way to many plants and microorganisms.

In order to remedy these problems, many have proposed alternative species concepts, a phenomenon documented by Wilson [148] and Ereshefsky [44]. However, despite these problems, there is a large range of important cases to which the biological species

248

concept applies, and it is widely used among animal taxonomists because it captures the notion of a unit of evolution so well. Additionally, many philosophers and scientists take a pluralistic view, allowing for a range of species concepts, each of which applies to an alternative range of biological phenomena. This allows for the use of the biological species concept in a limited range of cases, such as sexually reproducing animals.

Carson's "flush-crash" theory of speciation

In a series of publications, Carson ([16], [17], [18], [19],[20], and [21]) has elaborated the flush-crash theory of speciation, sometimes referring to a "flush-crash-founder cycle" rather than just a flush and crash. As the name of the theory suggests, Carson believes that drastic changes in population size figure prominently in speciation. What is important for my purposes is that there are episodes of drift corresponding to both flush and crash alike.

Carson's theory is geographical [17, 1414, 1418]: he believes that speciation begins when two populations of a single species become separated from one another in space, one population occupying territory at the edge of the species' range. Because the habitat of the latter population is at the edge of the species' range, it is termed "peripheral," and it differs ecologically from the species' usual habitat. The peripheral population—which I will term "the isolate" in the remainder of this section—is the incipient species.

In the next stage of speciation, the isolate experiences a rapid increase in population size, known as the "flush stage" of speciation. Carson understands the flush to be a normal part of cycles of population growth, possibly due to a sudden increase in food sources or seasonal fluctuations [20, 88]. The population flush is facilitated by a relaxation

249

of natural selection in the peripheral environment, where the isolate takes advantage of what Carson ([17, 88], [20, 1418]) calls "permissive" conditions. The idea is that unfilled ecological niches exist in the peripheral environment: no competitors exist, food resources are plentiful, and there are no predators. The flush can be dramatic. Carson [17, 1417] claims that some fruit fly species on the island of Hawaii were founded by single individuals, gravid females, from Maui. This would have resulted in a massive initial expansion of these populations in the generations immediately following the arrival of the founder individual.

Carson believes that random drift plays an important role in the isolate's evolution during the flush ([17, 88], [20, 1418]). During the flush, genetic recombination due to sex creates new gene combinations; this is drift in the form of indiscriminate gamete sampling. Because natural selection is weak or nonexistent during the flush, it may be thought of as a period of genetic experimentation. Note that, because a population flush is an *increase* in population size, the population-genetic dynamics typical of small populations do not apply.

After a small number of generations, the crash stage occurs: the isolate's size decreases drastically, perhaps leaving only a single gravid female ([17, 1418], [20, 89]). This individual begins to repopulate the isolate. Carson sees drift as playing an important role in the isolate's evolution during the crash [20, 88 - 89]; intense natural selection is not responsible for it. Catastrophic drift by indiscriminate parent sampling determines the identity of the surviving individual or individuals. As the isolate begins to recover from the crash, indiscriminate gamete sampling introduces drastic shifts in the direction and magnitude of evolution characteristic of drift in small populations (see section 5.3.1 above). Additionally, the drift that occurs during the flush stage exerts a delayed effect in the crash

stage. This is because the surviving individual may have the gene combinations that it in fact has because of drift during the "experimentation" of the flush stage.

Carson believes that the flush-crash cycle can occur in a small number of generations, ten or less ([17, 1418], [20, 89]). The end of the flush-crash cycle is signaled by the initiation of stronger natural selection. By the end of the crash stage, the isolate consists of a small number of individuals in a lineage consisting only of descendants of the sole survivor or survivors of the crash. If the isolate is to avoid extinction, the organisms in this lineage must adapt to their environment ([17, 1418], [20, 89]). This is possible only if drift during the flush-crash cycle supplied the lineage with gene combinations that will enable it to increase its mean fitness at a rate fast enough to maintain the isolate's size as the environment becomes harsher.

Indeed, if the isolate is to persist, the radical changes that drift creates in it during the flush-crash cycle must create the genetic basis for what Mayr [100, 170] calls a "genetic revolution:" a complete reorganization of the genome. The "genetic revolution" results in the creation of distinctive new adaptive traits as well as reproductive isolation from organisms in the population from which the isolate originated. That is, the two populations do not or cannot interbreed. On account of both of these kinds of changes, the process of speciation is complete, the isolate having obtained the biological separation from its parent required to be a new species [17, 1418].

In conclusion to my discussion of Carson, I want to note that Mayr's ([100], [101], and [102]) theories of speciation can also be explained in terms of the occurrence of random drift. I do not consider Mayr's views here because they present special problems of interpre-

tation: his opinions about the importance of drift changed significantly over his long career. Like many other evolutionists, he seems to have emphasized drift more heavily before the late 1940's than he did after that time. Provine [118, 406 - 408, 429 - 431, 453] suggests that the shift in Mayr's views reflects general trends among evolutionists, who were reacting partly to the Wright-Fisher dispute.

Nevertheless, my view is that Mayr always held that drift plays a role in speciation, even if he placed less emphasis on it in his presentation of his views. Moreover, many evolutionists today seem to be unaware of shifts in his view, and they approve of a view, attributed to Mayr, which accords a substantial role for drift in speciation. Thus, I believe that a review of Mayr's work and its impact would show that drift plays an important role in how speciation is conceptualized today.

Is the founder effect a mechanism of drift?

Carson's theory, as well as others that posit a small number of founders such as Mayr's ([100], [101], and [102]), entails that some of the marked differences between parent and daughter species are due to the "founder effect." Philosophical commentators on evolution sometimes identify the founder effect with drift, and I would now like to argue that this is a mistake. I understand the founder effect as follows.

Analysis 5.2 *Daughter species D diverges from its parent P by the founder effect if and only if some of the differences between D and P are caused by D's having been founded by a small number of organisms from P whose allele frequencies differed considerably from those in P considered as a whole.*

On Mayr and Carson's theories, a species receives its initial store of genetic diversity from the isolate from which it is derived. Thus, characteristics particularly frequent in the isolate will be particularly frequent in the early generations of the daughter species, before extensive mutation and recombination produce new traits. Some characteristics of the isolate may never disappear from the daughter species, an event which is useful for creating taxonomies from DNA samples. The situation Carson [17] describes in which some species of fruit flies on Hawaii are derived from a lineage founded by a single gravid female from Maui represents an extreme case of the founder effect.

Because the founding population of the isolate is small in size, it is likely to differ markedly in its allele frequencies from the parent population. Extending the metaphor of parent and gamete sampling to the formation of the isolate, suppose that the isolate is thought of as a sample of the organisms in the parent population. Theorems of probability suggest there is a low probability that a small sample match its source in statistical properties such as allele frequency.

Millstein [105, 34] and Sober [140, 112] believe that the founder effect is a kind of random drift, viewing it as a kind of indiscriminate parent sampling. I agree that there is some commonality between drift and the founder effect, particularly if the latter is conceived of as a kind of sampling process, as I suggest above. However, many evolutionists—including Mayr ([101, 529ff.], [102, 237]), Carson [17, 1418], Stanley [142, 165 - 168], and Brandon [11, 44, note 19]—do not consider the founder effect to be a kind of drift. I agree with the latter for the following reason, which they do not articulate. The founder effect brings about rapid divergence because the "sample" (i.e., the isolate) differs from the parent species. This is a

mechanism that operates *across* populations. In contrast, drift occurs *within* the isolate.

I do not want to claim that the two phenomena are completely unrelated. The reduced variability of the isolate due to the founder effect makes radical changes in gene frequencies due to drift more likely. With a smaller number of alleles entering the Mendelian "lottery" of recombination, and with fewer mates, there is a good chance that homozygotes will be formed. However, I think that it is important to recognize that evolutionists draw a clear distinction between the initial separation in space of the isolate from the original population of its parent (the founder effect), and the evolution of the isolate after separation (drift).

Process explanation of the origin of species

An explanation-seeking how-question that a biologist would pose concerning the origin of species, and whose answer invokes drift, is as follows.

Question 5.7 *How did daughter species D separate from its parent species P, becoming reproductively isolated and phenotypically distinct in key respects from the latter?*

This question presupposes the following, which I will suppose to be the case.

Statement 5.13 *Species D is the daughter of P, and is phenotypically distinct in key respects from the latter.*

Given statement 5.13, Carson's flush-crash theory informs the following answer to question 5.7.

Explanation 5.4 *Isolate I became spatially isolated from the population P_P containing the other organisms in species P, moving to a new environment at the periphery of P_P's range. Relaxation of selection pressures resulted in a period of increasing population number, during which random drift due to indiscriminate gamete sampling created novel gene combinations in I (flush stage). Then, due to indiscriminate parent sampling of a catastrophic nature, the population size N decreased rapidly, perhaps shrinking to unity, I's sole member being a single gravid female. As I built itself up again, indiscriminate gamete sampling created still more new gene combinations, exhibiting the distinctive dynamic of rapid evolution in changing directions characteristic of drift in small populations (crash stage). These gene combinations differentiated I from P_P phenotypically, and also, caused its reproductive isolation. When natural selection intensified, some of these gene combinations proved advantageous, facilitating its continued survival. Retrospectively, I may be identified as the population from which D descends.*

I would respond to the charge that explanation 5.4 is incomplete in the same way that I have responded to such charges previously. As I argue above (section 5.2.2), the irreducibility thesis (statement 5.8, page 232) completely blocks such incompleteness objections.

5.3.4   The shape of phylogeny and punctuated equilibrium

In three subdivisions of this section, respectively, I describe (a) what I mean by "shape of phylogeny;" (b) the phenomenon of punctuated equilibrium and the role of drift in it; and (c) the role of drift in the process explanation of the shape of a phylogeny shaped

by punctuated equilibrium.

The shape of phylogeny

Phylogeny—ancestor-descendant relationships among species and other taxa—can exhibit patterns. One kind of phylogenetic pattern concerns the timing of morphological change: when, in relation to the divergence of a species from its parent, do the morphological changes that distinguish parent and daughter evolve? Because of the way phylogeny is represented visually, a phylogeny exhibiting a given timing pattern has a characteristic "shape." Let me explain.
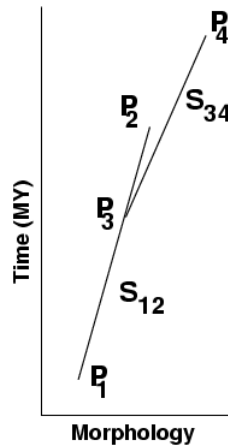

Figure 5.3: Representation of phylogenetic patterns in two dimensions. See text for explanation.

Suppose that a two-dimensional Cartesian space is given the following interpretation, as in figure 5.3.

Vertical axis Time, in millions of years. This suggests the geological record, in which fossils from earlier periods are found in deeper strata than those from later periods, as long as the strata have not been overturned.

256

Horizontal axis Units measuring some morphological trait, for instance, the number or

radius of the whorls of a snail's shell; the shape of a bivalve; or the height of an Irish

Elk's antlers.

As represented in figure 5.3, a species S_{12} is represented in this space by a line L_{12}

with endpoints P_1 and P_2. The vertical distance between P_1 and P_2 represents the length

of time between S_{12}'s origin and its extinction; the horizontal distance between P_1 and P_2

represents the amount of morphological change that S_{12} undergoes during its lifetime. A

daughter of S_{12} is represented by a line with an endpoint on L_{12}. The daughter species, S_{34},

is represented by endpoints P_3 and P_4; P_3 falls on L_{12}. The horizontal distance between

P_3 and P_4 represents the total amount that species S_{34} diverges morphologically from S_{12}

during the former's lifetime.

The theory of punctuated equilibrium

According to proponents of the theory of punctuated equilibrium, random drift

plays an important role in generating patterns of phylogeny, because of its role in speciation.

In order to indicate why this is so, let me first describe two competing claims about the

shape of the phylogeny of all species in the history of life, first elaborated in an explicit

manner by Eldredge and Gould [42].

Punctuated equilibrium Morphological differences between parent and daughter species

arise almost entirely at the time of speciation. Thus, the morphological differences

between a parent and its daughter species evolve by a characteristic two-step pattern:

rapid change at speciation ("punctuation"), followed by a long period of little or
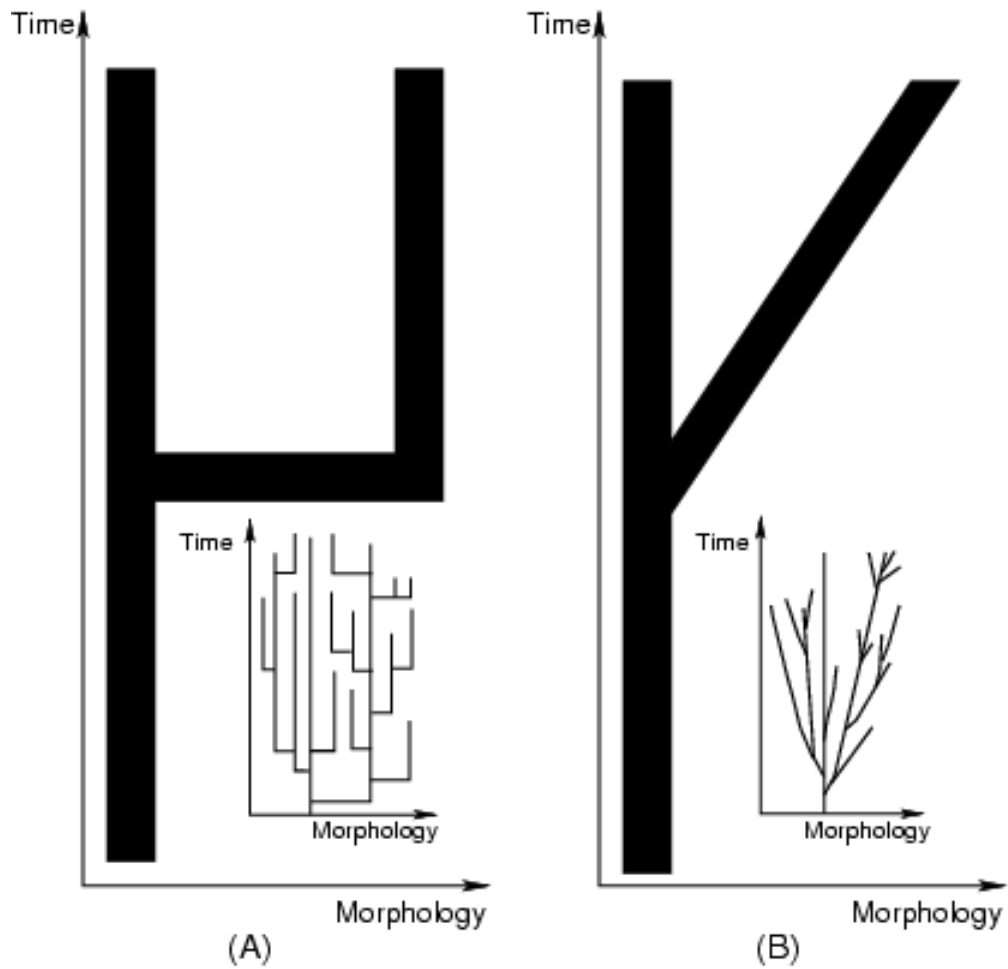
Figure 5.4: Contrasting images of the shape of phylogeny. (A) represents the punctuational pattern. The daughter species branches off rapidly; changes in morphology that distinguish it from its parent occur in an intense episode concentrated at speciation. (B) represents the gradualist pattern. Morphological differences between daughter and parent species that evolve over the span of any given interval of time occur at a uniform rate during that interval of time. Models of the shape of phylogeny of the history of life as a whole entailed by each view are depicted in the inset figures, respectively. (A) and (B) and the inset figures are idealized. Punctuationists do not claim that morphological change has no phyletic component, and gradualists do not claim that rates of phyletic evolution always remain uniform. See text for further discussion.

no significant change, lasting for the remaining life of the species ("equilibrium" or "stasis"). (See (A), figure 5.4, page 258).

Phyletic gradualism The majority of morphological differences that exist between parent and daughter species at time $T + \triangle$ evolve at a uniform rate, between the time of divergence T and $T + \triangle$. This kind of morphological change—change occurring in a species that is already established, and not associated with speciation—is termed "phyletic evolution." "Gradualism" is intended to suggest the constant rate of such evolution. (See (B), figure 5.4, page 258).

Punctuationism, or the theory of punctuated equilibrium, emphasizes the pattern of punctuated equilibrium over that of phyletic gradualism. The central claim of the theory may be stated as follows.

Statement 5.14 (Punctuated equilibrium) *Differences in the morphology characteristic of organisms in a species and the morphology characteristic of organisms in its daughter species arise at the time of speciation, as a general rule.*

Before its presentation in 1972 by Eldredge and Gould [42], statements like statement 5.14 may have appeared in the literature, or been implied by claims made in the literature, particularly by Mayr. Nevertheless, there is no dispute that Eldredge and Gould's formulation of the theory is its most explicit, forceful, and provocative; they are generally given credit for its invention. The theory, now central to paleontology, is well supported by a rich literature extending to a variety of taxa, and it constitutes an active frontier of research.[22]

---

[22]Gould and Eldredge [61] and Stanley [141] provided early articulation and defense of the theory; they review the evidence for the theory after 20 and 26 years, respectively ([62] and [142]).

Note that the theory does *not* assert or entail that all species exhibit the punctuational pattern: the claim is that *if* there are significant morphological differences that distinguish an ancestral species from one of its descendants, *then* those changes arise at speciation. Also, punctuationists recognize that there are exceptions to the pattern they describe. This is why I have appended "as a general rule" to statement 5.14. As well, note that the theory concerns morphological change. This is because the theory concerns the fossil record, which does not preserve genetic or behavioral changes.

Numerous well-documented examples of stasis provide the most striking evidence for punctuated equilibrium. Many ancestor-descendant pairs exhibiting significant morphological differences coexist with one another, unchanged, for significant proportions of their lifetimes, periods lasting millions of years. This restricts the timing of their divergence to a short interval near the time of the descendant's origin, just as the theory of punctuated equilibrium suggests. Notable examples of this include bryozoans and bivalves ([22], [23], [73], and [143]).

Drift's role in generating punctuated equilibria occurs in speciation. As Stanley [142, 164 - 168] suggests, speciation occurring in conformity with Carson's flush-crash model generates the kind of rapid speciation characteristic of the punctuational stage of punctuated equilibrium. As I indicate above (pages 249 - 251), according to Carson, a new species diverges in morphology from its parent because its founding isolate diverges from that parent. As I also indicate above, the role of drift in this divergence is to promote the "genetic revolution" producing variants that enable the isolate's survival, distinguish it from its parent, and cause reproductive isolation. Because the time required for this initial

divergence is quite small relative to the species' life span, drift "punctuates" otherwise stable lineages with an abrupt shift in morphology, timed at a daughter species' initial separation from its parent.

Note that Carson's theory concerns speciation only, and so is only relevant to the punctuational component posited by the theory of punctuated equilibrium. Claims about subsequent stasis require additional theories about the stability of species after they form. As well, note that the flush-crash cycle is not the only mechanism of rapid speciation, although it is one that requires random drift.

Explaining patterns of phylogeny with drift

Suppose that a hypothetical genus, *Hypotheticus*, exhibits the pattern of diversification typical of punctuated equilibrium: species branch away from one another in an amount of time that represents a small fraction of their total life span, as represented in the appropriate diagrams in figure 5.4 above (page 258). The following is an explanation-seeking how-question that an evolutionist might ask about such a lineage.

Question 5.8 *How do daughter species of the genus* Hypotheticus *diverge rapidly from their parent species in a short time (relative to their total life span), generating the punctuationist pattern?*

In addition to the claim that *Hypotheticus* does in fact exhibit the punctuational pattern, which I have already stated that I will suppose to be true, there is another statement that I will suppose to be true.

Statement 5.15 *Daughter species of* Hypotheticus *diverge from their parents by the mechanisms described by the flush-crash theory of speciation.*

Given this, the following process explanation answers question 5.8.

Explanation 5.5 *For each daughter species D with a parent species P of the genus* Hypotheticus, *there exists a time T at which an isolate I formed at the periphery of P's range. By time $T + \triangle$, the following had obtained by way of the flush-crash cycle, a process explained in detail above (explanation 5.4, page 254). Drift, in part by the distinctive dynamics it creates in small populations operating during the crash stage, caused a "genetic revolution" to occur in I. The "revolution" resulted in advantageous variants that differed markedly from any that existed in P at time T; the "revolution" also resulted in I's reproductive isolation from the other organisms in P. In retrospect, it is possible to identify I as the isolate from which daughter D descended. After its origin, D changed little, relative to the changes it underwent at the time of its origin. Because the time span between T and $T + \triangle$ was a small fraction of D's total lifespan, a "punctuated" shape of* Hypotheticus's *phylogeny resulted.*

Note that explanation 5.5 does not explain stasis in *Hypotheticus*; it explains only the rapid divergence of species. The stability of the species is presupposed, but left unexplained. This is appropriate, as question 5.8 requests an explanation for the rate of divergence, but not an explanation of the species' stability.

Finally, I want to indicate that I would respond to any incompleteness objections concerning explanation 5.5 such as I describe above (section 5.2.2) by citing my argument in favor of the irreducibility thesis (statement 5.8, page 232), which I believe shows that

262

incompleteness objections lack any credibility whatever as claims that process explanations lack explanatory force.

5.4   Explaining Drift by Process Explanation

This chapter represents the culmination of my efforts in this dissertation. In chapter 3, I argue for the justification of process explanation; in chapter 4, I describe a mechanism of drift, indiscriminate sampling. In this chapter, I argue that indiscriminate sampling causes drift that is explained by process explanation. I identify five such kinds of events, which I divide into two classes: one for cases whose process explanation does not require mention of small N, and one for cases whose process explanation does require it. Each of these kinds of events are central objects of study in ongoing research programs deeply connected with historically important questions about evolution.

My arguments in this chapter show that random drift, a chance process, has considerable explanatory power. This is at odds with the Hempelian evolutionists, who, as I argue in this chapter, take Hempelianism as a reason to affirm the exclusivity thesis. They believe that only natural selection explains evolution. I believe that the arguments against this position in this chapter and in previous chapters are conclusive, and that anyone wishing to argue for either Hempelianism or the exclusivity thesis is best advised to abandon the effort before investing too much in it.

Chapter 6

Chance and Explanation in Evolutionary Biology

As I suggested in the opening paragraphs of chapter 1, my aim in this dissertation is to argue that a range of important phenomena that occur by random drift are explained by evolutionary biologists using the strategy of process explanation. I pursue this aim so that I can achieve the further aim of showing that those who hold what I term "the exclusivity thesis," according to which natural selection alone explains evolution, are mistaken. I have now achieved these aims. I would like, in this concluding chapter, to review the central points of my argument.

To begin my review of the argument, let me recall what is at stake. The issue of whether there are evolutionary events that occur by random drift and that can be explained by process explanation depends, in part, on the resolution of a conceptual issue. This conceptual issue is, What is the warrant for process explanation? This issue requires confronting

the views of an important group of philosophers of science that I term "Hempelians." The Hempelians affirm a form of what I term "universalism" about explanation, while process explanation requires what I term "contextualism" about explanation.

As I indicate at various points in the dissertation (for instance, sections 1.2.1 and 1.2.2), universalism and contextualism are views about the nature of the criteria used to assess whether an explanation is a good one. Universalism is the view that all good explanations meet a single set of criteria, criteria independent of the audience's interests or cognitive states. In contrast, contextualism is the view that, because audience interests and cognitive states play a role in determining whether an explanation is a good one, there is no single set of criteria that all good explanations meet.

The Hempelians' preferred variety of universalism places two requirements on the scientific explanation of particular events. First, the explanation must answer an explanation-seeking why-question. Second, the explanation must do so by showing that, according to laws of nature, either (i) the event was to be expected (Hempel), or (ii) there was some degree of expectation that someone ought to have had that the event would occur (Salmon and Railton). I describe these requirements in detail in chapter 2 (section 2.1).

Process explanations, in contrast, need not make reference to laws of nature, and do not answer why-questions. Although there is no canonical form for process explanations, many of the central cases respond to explanation-seeking how-questions, that is, questions with the form, "How did E, the event to be explained, occur?" I elaborate on the nature of process explanation at length in chapter 2 (section 2.2).

Process explanations answer such how-questions by describing a sequence of events

that are causally relevant to the occurrence of the event to be explained. For example, citing Dray, I indicated in chapter 2 (pages 58 - 59) that an auto-repair-shop owner might explain the seizure of a car's engine to its owner by describing the events leading up to the seizure: a hole formed in the oil reservoir; oil ran out of the reservoir; the pipes through which oil is brought to the engine dried up; oil failed to reach the pistons; and, lacking lubrication, the pistons froze in place in the engine block. This explanation is formulated for a particular audience, the car's owner, who lacks extensive knowledge of auto engines; a different explanation would be offered to an auto mechanic, who would have a set of interests and beliefs different from the car's owner.

Like all process explanations, the explanation of the engine's seizure does not show why the event to be explained occurred. Moreover, the explanation of the engine's seizure does not cite laws of nature, which process explanations need not do. For these reasons, process explanations are particularly well-suited for answering explanation-seeking questions about chance events. The contextualist framework is essential to justifying the use of process explanation for this purpose. Let me illustrate this point with an example, taken from Stoppard's *Rosencrantz and Guildenstern are Dead* [144], that I first presented in chapter 3 (pages 75 - 83).

Suppose that 92 fair coins, each taken from a bag of coins and tossed in a fair manner, each land heads; this represents a serious loss for one of two players in a game of chance, who bets each coin as it is tossed that it will land tails. Someone discovering this state of affairs would no doubt be amazed, and he or she might ask, "How did this happen?" or simply "What happened?" A process explanation describing repeated instances

in which fair coins subject to the normal mechanical forces are tossed and lands heads would answer these questions. In general, in the contexts in which this kind of question would be asked, someone would want to know whether there was any bias or cheating that was responsible for the coins' unusual run of heads. Few people would want to know the particular circumstances of each toss that, together with the relevant laws of nature, necessitated (or necessitated in some degree) a sequence of 92 heads. These circumstances would doubtlessly be so diverse, and the laws so complicated, that knowing them would not illuminate the coins' behavior. In chapter 3, I formulated this response in terms of "levels of organization" (page 73). The context calls for an explanation at a higher level—the coins and the tossing mechanism—rather than an explanation at a lower level—the fundamental physical laws governing the coins' behavior.

This contextualist account of the coin tossing case provides the grounds for responding to a Hempelian who, considering the claims I have just made, might formulate an incompleteness objection of the kind that I have described at various points in the dissertation (pages 68 and 222). The Hempelian believes that a process explanation derives whatever explanatory force it has from a set of Hempelian explanations. These Hempelian explanations take each stage in the sequence of events mentioned in a process explanation as events to be explained in their own right; the idea is that each stage can be explained in a Hempelian manner by laws of nature and particular circumstances prior to each. In this way, the Hempelian believes that any process explanation can be reconstructed in terms of Hempelian explanations.

This is incompatible with contextualism. As I suggest in the paragraph before last,

such Hempelian "explanations" would consist of a myriad of unconnected details about the circumstances of the coin prior to being tossed each time as well as laws of extraordinary complexity. Nonetheless, neither these details nor these laws are called for in the context in which the question "What happened?" is asked. This question is directed toward discovering whether there were any irregularities, such as a two-headed coin, that might lead to such an improbable event. If there are no such irregularities, there is no *reason why* the coin landed heads so many times in a row; however, it can still be explained *how it happened* that it did so.

A Hempelian would protest loudly against this contextualist line of argument, taking an approach that I detail in chapter 3. My contextualist argument, the Hempelian would claim, does not threaten the coherence of the Hempelian view, which entails that a Laplacian demon—but not human beings, as long as they have their current epistemic capabilities—can have a complete explanation of any given event. The Hempelians accept the pessimistic view that a large number of events, including many of those that I claim are explained by process explanation, lie outside the limit of what human beings can hope to explain.

The central claim of my chapter 3 response to the Hempelian is that Hempelianism fails to account for the following fundamental fact about explanation: there are many linguistic practices and associated concepts widely recognized to be explanatory, and that are used for the purpose of generating understanding, and that often succeed at doing so. This is important because generating understanding is the *sine qua non* of explanation. Hempelianism posits a single canonical set of criteria for the adequacy of an explanation,

and so cannot account for this diversity of practices. Contextualism can, because it was designed precisely in order to do so.

This puts the burden of argument onto the Hempelian. To recover their position, the Hempelians must argue that the apparent diversity of explanatory practices ought to be discounted, and that, of all the putatively explanatory utterances, only those conforming to Hempelianism are in fact explanatory. I do not think Hempelianism can bear this burden of argument, and I view contextualism as the stronger position.

Hempelians would want to respond to this line of argument by pointing out that they need not sacrifice their universalism to acknowledge the work that contextualists have done to describe the many different standards scientists use for evaluating explanations. The universalist draws a distinction between context-dependent and context-free understanding. Contextualist studies of the various standards for evaluating explanations represent surveys of the forms of context-dependent explanation in use among scientists. In contrast, what the Laplacian demon is able to obtain—and what humans are not—is context-free understanding. The nature of such context-free understanding is described by the various models of explanation elaborated by the various Hempelians. The important point that the Hempelian would want to make is that whether an explanation meets the standards reflected in these models does not depend in any way on anyone's intentions or beliefs.

I do not believe that these points strengthen the Hempelian position. As I argue in chapter 3 (section 3.3), scientists do not aim at the kind of explanations that the Hempelians describe. Scientists do not seek understanding for its own sake; rather, they pursue explanations to advance further purposes, such as promoting research or answering

269

deeper questions about the nature of the universe. These further purposes, in turn, shape the standards by which an explanation is evaluated. This weakens the Hempelian position. In chapter 3 (section 3.3), I illustrate this claim with an example from population genetics: explanations incorporating different descriptions of the same phenomena would be evaluated differently, depending upon the particular aims of the audiences for the explanations in question.

Finally, the Hempelian response that promoting research and answering deeper questions about the nature of the universe are intrinsic to science, and so are independent of the audience, is a weak one. I say this because there are always particular aims of investigators on a given occasion that shape *how* such context-independent aims are to be attained on that occasion. Promoting further research is intrinsic to science, and so, not dependent upon anyone's particular desire to do so. Nonetheless, having the aim of promoting further research does not completely inform the kinds of explanatory pursuits that someone might carry out on a given occasion. This requires the input of the audience for the explanations in question, and must answer to their particular needs, as they exist on that occasion.

Having sketched my defense of process explanation, let me turn now to the utility of process explanation for explaining drift. The reasons that process explanation is so useful for explaining events such as coin tossing also make it useful for explaining evolution resulting from random drift. In chapter 1 (section 1.1), I argue that random drift is a kind of evolution by accident, in contrast with natural selection, which is a purposive process. As long as the environment remains stable, organisms bearing variants of a trait that have succeeded in

natural selection will exhibit greater evolutionary success compared with organisms bearing other variants. The activity that a variant carries out that causes this evolutionary success is the purpose of the trait; paradigm cases of such advantageous traits such as sense organs have the degree of complexity associated with products of human design: such traits appear goal-directed, because their complexity facilitates the activity that causes bearers of the favored variant to exhibit greater evolutionary success than bearers of other traits.

In contrast, when drift occurs, differences in rates of reproductive success among organisms bearing alternative variants are not due to differences in their ability to carry out their purposes. As well, drift does not introduce any trend toward greater adaptedness in a population. Drift breaks down complex traits, producing what appear to be mishaps and technical failures among organisms characterized by adaptation to their environments: the sophisticated machinery evolved by natural selection fails to promote their evolutionary success.

Throughout the dissertation, I illustrate the process of drift by reference to two paradigm cases of it. The first is what I term, following Stanley (pers. comm.), "catastrophic drift." Such drift occurs if an event such as a storm or earthquake kills a large number of organisms in a population. The magnitude of the event prevents differences in variants' abilities to carry out their purposes from making any difference to their bearers' survival, which depends upon their location relative to hazards created by the event. Second, evolution due to sex in Mendelian populations is a paradigm of drift. This is because the Mendelian process isolates alleles from the causal processes responsible for inheritance. Variant alleles' differing abilities to create variants differing in ability to carry out their

purposes do not make a difference to whether those alleles are passed on.

In chapter 4, I elaborate a theory of indiscriminate sampling. As I discuss in chapter 4 (section 4.4), indiscriminate sampling is one of three processes of drift that I believe exist. Indiscriminate sampling is modeled after a canonical chance process, blindly drawing beads from an urn. My conception of indiscriminate sampling is centered around what I term in chapter 4 "the core probabilistic equality" (sections 4.2.2 and 4.3.2). The core probabilistic equality has different formulations for organisms and gametes, each corresponding to processes of indiscriminate sampling that occur at different times in the life cycle.

To take a simple case, the core probabilistic equality is satisfied when the probability that an organism with one variant V_1 survives is equal to the probability that an organism with an alternative variant V_2 survives. The core probabilistic equality must be qualified substantially to generate a complete description of indiscriminate sampling. Nevertheless, these qualifications do not attenuate the connection between the canonical chance process of drawing beads from an urn and indiscriminate sampling in biological contexts.

As I detail in chapter 5, indiscriminate sampling plays a key role in a set of five evolutionary processes that are well-studied by contemporary evolutionary biologists, and that are deeply rooted in historically important questions in evolutionary biology. This set of five processes may be divided into two classes: those whose evolutionary dynamics do not depend upon population size N and those whose evolutionary dynamics do (sections 5.2 and 5.3, respectively). The former class includes (a) the chance elimination of a rare but favorable allele and (b) the evolution of neutral alleles. The latter class includes (c) the shifting balance process, (d) the origin of species, and (e) the shape of phylogeny in lineages

evolving by punctuated equilibrium.

Much in the same way that process explanation can be used to explain how 92 of 92 fair coins, each tossed in a fair manner land heads, process explanation—as I indicate in chapter 5—is used to explain each of the five phenomena that I name above. For instance, the chance elimination of a favorable allele (section 5.2.1) is explained by describing the sequence of events that occur between the inheritance of a new favorable allele by a single organism and the loss of that allele during the indiscriminate sampling that is a normal part of Mendelian reproduction. The loss of the allele occurs as the result of a chance process, and it may well be the case that there is no reason why the allele disappeared from the population. Nonetheless, it is explained by process explanation.

Like the process explanation in the coin example above (page 266), process explanations of drift are formulated in terms of "higher level" entities such as organisms, the mating system, and gametes, rather than being formulated in terms of "lower level" entities such as fundamental constituents of matter. Also as in the case of the coin example, the requirement that process explanations of drift be formulated at a particular level of organization is fixed by context. This constitutes an effective response to a Hempelian who wants to argue that, because process explanations of drift can be replaced with a set of Hempelian explanations that are formulated at a lower level of organization, those process explanations are incomplete. This response to the Hempelian is embodied in what I term the "irreducibility thesis" (page 232).

This brings the argument to its conclusion. Phenomena occurring by random drift are indeed explained using process explanation. From this, it follows that the exclusivity

thesis is incorrect. The exclusivity thesis, as I indicate at various points previously in the dissertation (pages 26 and 200), is the claim that the only explanation-seeking questions that are asked about evolution are why-questions about the spread of an adaptation, and that natural selection is the only evolutionary process that can be invoked to answer those questions. Those affirming this thesis are in error. Explaining nonadaptive evolution such as occurs by random drift is a central aim of evolutionary biologists, and the explanatory strategy they employ in pursuit of this aim, process explanation, is methodologically sound.

Appendix A

Images of Random Drift

In chapter 1 (section 1.1), I articulated the idea that drift is a kind of evolution by accident, in contrast with natural selection, which is purposive. My aim in this appendix is to further illustrate this point by reference to two images, each of which models evolution due to drift.

The first of these images (figure A.1, page 276) represents a computer simulation of the change in allele frequencies (evolution) of five populations of twenty individuals each. Generational time is represented along the X-axis; the Y-axis represents the frequency p of an allele in each of these populations, each of which starts with a frequency $p = 0.2$. Any allele will do; I will suppose that it is an arbitrary A_1 allele.

The simulated evolution of these populations is highly idealized. It is supposed that the populations are genetically identical at all gene loci, and that they each live in identical environments. Because of these idealizing assumptions, figure A.1 may be interpreted as representing a set of five possible histories that a single population might take across possible

worlds. Because it is easier to speak about a comparison of the histories of five populations than to speak about the comparison of five possible histories of a single population across possible worlds, I will use the former manner of speaking in the discussion that follows.
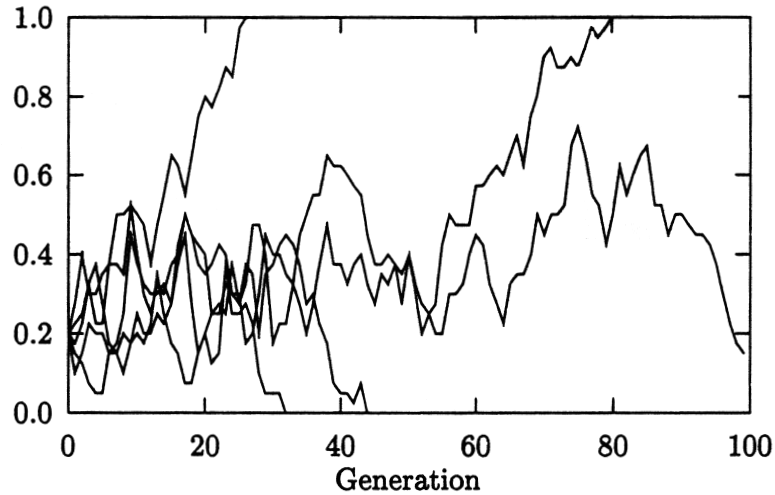

Figure A.1: Allele frequencies of identical populations evolving under drift alone. Frequency of the A_1 allele is represented on the Y-axis. See text for discussion. Acknowledgment: Gillespie, John H. *Population Genetics: A Concise Guide.* p. 21, Fig 2.1. © 1998 The Johns Hopkins University Press. Reprinted with permission of The Johns Hopkins University Press.

The divergence of the populations is quite striking, exemplifying the trendless, directionless, random "drifting" from which random genetic drift takes its name. By about the 40th generation, the A_1 allele has become extinct in two of the populations, reaching a frequency of $p = 0.0$; in another, the A_1 allele has become the only allele at the locus in question, reaching a frequency of unity. In the populations in which the allele has become extinct, it does not become so immediately. Rather, its frequency shifts drastically before eventually plummeting to $p = 0.0$. In the population in which the allele reaches a frequency of unity, the path to fixation is similarly indirect.

The A_1 allele remains present in the other two populations well past the 40th generation. In one, the allele reaches a frequency of unity at 80 generations; in the other, the allele seems to be headed for extinction at 100 generations, returning to its initial frequency of $p = 0.2$ and apparently beginning a downward trend. Nonetheless, this trend begins only after the population has continued for some time on a path nearly identical to the population just mentioned, that is, the one in which the allele becomes fixed at 80 generations. Perhaps the A_1 allele will persist for some time in this population, "drifting" about in frequency before disappearing—or perhaps it will become fixed in the population after another 100 generations.

This contrasts with what would be expected if natural selection were at work. Under the influence of natural selection, identical populations in identical environments would be expected to respond similarly, diverging only slightly, if at all. If the A_1 allele were advantageous, it would be expected to increase in frequency; if it were disadvantageous, it would be expected to disappear. Interpreting figure A.1 to represent a set of possible histories that a single population might take across possible worlds, the idea is that, if natural selection were at work, the population would be expected to take a roughly similar path in each possible world.

Now I would like to indicate the sense in which I think that figure A.1 accords with the view that natural selection is a kind of purposive evolution and that drift is a kind of evolution by accident. On the one hand, the histories of the five populations that would most probably obtain if natural selection were at work is characteristic of design. In retrospect, it would make sense to claim that the A_1 allele exhibited a trend, for instance,

becoming fixed in all five populations at around the same point in time, as though the populations were each aiming at that result.

On the other hand, figure A.1 represents evolution that is characteristic of accident, no trend being exemplified: the A_1 allele becomes extinct in two populations and moves to fixation in another two, each at different times; in the remaining population, it persists at its initial frequency. I think that this is what would be reasonable to expect, if there were no pattern or direction to the evolution of these populations. Similarly, there is no cluster of similar histories that a single population would be expected to take on, across possible worlds; rather, there is a nearly one-to-one ratio of possible worlds to histories.

Now, I would like to argue that figure A.2 below (page 279) illustrates the idea that drift is a kind of evolution by accident, in contrast with natural selection, which is a kind of purposive evolution. First, let me explain the figure, which shows a set of four probability distributions. The X-axis represents three states, each corresponding to the genotype of a single individual at a single gene locus of interest. The height of each bar represents the probability that this individual has the corresponding genotype, i.e., the genotype indicated beneath it. The topmost distribution represents the state of the population at generation G_0, and subsequent distributions represent the state of the population in successive generations. The idea is that the figure represents changes in the probability distribution over genotypes through generational time in a lineage created by a single individual that produces a single offspring each generation.

It is important to be clear about the biological situation modeled by figure A.2. Each distribution represents the chance that a single hermaphroditic individual producing
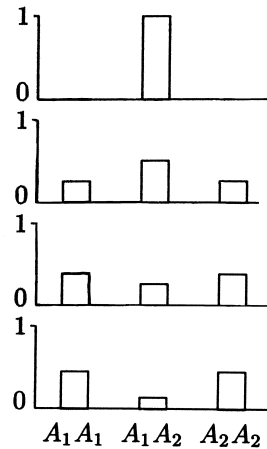
Figure A.2: Probability distribution of genotypes in a self-propagating lineage, $N = 1$, evolving due to drift alone. The probability that an individual has a given genotype is represented on the Y-axis. See text for discussion. Acknowledgement: Gillespie, John H. *Population Genetics: A Concise Guide.* p. 22, Fig 2.2. © 1998 The Johns Hopkins University Press. Reprinted with permission of The Johns Hopkins University Press.

a single offspring has any one of the genotypes listed. The first (topmost) distribution indicates that the ancestor of subsequent individuals in the lineage is a heterozygote, because all probability is concentrated above the "A_1A_2" label. The second distribution accords some chance to the possibility that the single individual produced to carry on the lineage past its progenitor is a homozygote; each subsequent distribution represents the chance that the sole member of the population, produced by self-fertilization of its parent, has a given genotype. Clearly, this is a highly idealized state of affairs, primarily useful for understanding the dynamics of drift by representing the behavior of an extreme case, although not necessarily for representing any real biological populations.

Consider how the probability that the individual in the lineage has a given genotype changes over time. In the topmost distribution, which represents the initial generation,

there is a probability of unity that the individual is a heterozygote—that is to say, the individual starts in the heterozygotic state. The second distribution shows that the chance that the second-generation individual will have the $A_1 A_1$ genotype is equal to the chance that it will have the $A_2 A_2$ genotype, while the chance remains considerably higher that the individual will be a heterozygote.

The third and fourth distributions indicate that there is an increasing chance that the individual will be a homozygote, and a decreasing chance that the individual will be a heterozygote. It is crucial to note that the probability that the individual takes on a given homozygotic genotype is the same for both such genotypes: no homozygote is favored. If subsequent distributions were shown, they would indicate that the chance of homozygosity continues to increase in this way, with no "favored homozygote," while the chance of heterozygosity decreases. In a "long-run" of n generations, there is a probability near unity that the nth individual in the lineage will be a homozygote.

Because of its characteristic shape, this distribution is known as "the U-shaped distribution." The reason that this pattern obtains is that, once an organism becomes homozygotic, it cannot, under the influence of drift alone, produce heterozygotic offspring. Mutation would be required; one of its genes would have to mutate into the other allele. This falls outside of the scope of the model used to generate figure A.2, one of the idealizing assumptions of which is that only drift occurs. The U-shaped distribution obtains for lineages with more than one individual; the difference is that, in larger populations, the decrease in the chance of heterozygosity is slower. This is important to note because it means that figure A.2 represents a general pattern of populations evolving due to drift, not

just a pattern exemplified by small populations.

Figure A.2 represents a model described by Sewall Wright in 1931; I have already cited the following passage (page 237), in which he refers to a "self contained population of one," and I reproduce it here.

> The extreme case is that of a line propagating by self fertilization which may be looked upon as a self contained population of one. In this case, 50 percent of the factors with equal representation of two allelomorphs (that is, in which the individual is heterozygous) shift to exclusive representation of one of the allelomorphs in the following generation merely as a result of random sampling among the gametes. [156, 107]

The situation Wright is referring to concerns the transition between the topmost distribution and the second-from-topmost. In the topmost distribution, the individual is a heterozygote, i.e., has a 100% chance of having the A_1A_2 genotype. The second distribution shows that there is a 50% chance that the individual will have either the A_1A_1 genotype or the A_2A_2 genotype, which represent "exclusive representation of one of the allelomorphs."

Now consider what figure A.2 indicates about drift. The important point is that the outermost states grow at the same rate: as generations pass, equal proportions of individuals pass from the inmost state to the outermost states, with none passing from either outermost state. This exemplifies the trendless, directionless "drifting" of random genetic drift. If differences in adaptedness were responsible for the shape of the probability distributions generation after generation, it would be expected that one of the states would grow faster than others, as less-well-adapted organisms fail to pass on their genes. This would be characteristic of design, for instance, if someone designed the organisms in the population for a certain function best carried out by the "favored" genotype. However, this is not what occurs.

281

Appendix B

Pluralism About Drift

In chapter 4 (page 166), I claim that there are two mechanisms of drift other than indiscriminate sampling. While each such mechanism is a kind of evolution by accident, each brings about such evolution in a manner different from the others. In this appendix, my aim is to describe these alternatives to indiscriminate sampling in greater detail than I was able to in chapter 4.

The first mechanism forms a broad family; in order to describe it, I need to introduce a bit of vocabulary, as follows. The evolution of allele frequencies in a population is determined by various parameters, such as the fitness values of the various alleles, the rate of mutation, and the rate at which organisms migrate to and from the population. I will term these the "evolutionary parameters," and I call this family of mechanisms "random fluctuations in the evolutionary parameters."

The idea is as follows. Each of the evolutionary parameters can fluctuate in time. Some one of the family of processes that I have in mind occurs if and only if some one of the

evolutionary parameters fluctuates *randomly* across the span of two or more generations. Roughly speaking, an evolutionary parameter fluctuates randomly across the span of some number of generations in a population if and only if, over the course of that number of generations, it affects the variance but not the mean of the probability distribution of allele frequencies in that population.

To see how this mechanism works, consider an example of random fluctuations in fitness values of two alleles. There are two requirements that must be met, one on environments, the other on fitness values. First, consider the former. There must be two states of the environment, and it must spend the same amount of time, on average, in each state. For instance, suppose that an environment of a hypothetical population shifts between wet (rainy) and dry (drought). It need not be the case that the environment fluctuate in a regular manner; all that is required is that on average, the population spend as much time in a wet environment as it does in a dry one. So, for instance, there might be three wet generations followed by five dry ones, which are followed by another two wet generations, for a total of five generations in each state.

Second, consider the requirement on fitness values. It must be the case that each allele in the population is advantageous in one type of environment, but disadvantageous in the other, so that fluctuations in the environment lead to reversals in the fitness values of each allele. I represent these relationships on table B.1.

Table B.1: Example of fitness relationships in fluctuating environments

	Wet Environment	Dry Environment
Fitness of the A_1 allele	w	$.95w$
Fitness of the A_2 allele	$.95w$	w

Under these conditions, evolution proceeds as follows. In wet environments, *ceteris paribus*, natural selection will favor A_1. However, in dry environments, the pendulum will swing the other direction, so to speak: A_2 will be favored. Across a span of many generations, no net advantage accrues to either allele, and allele frequencies will move back and forth across a stable mean value as gains by one allele in one environment are wiped out by gains by the other allele in the other environment. Any net change that occurs, *ceteris paribus*, is due to differences in the amount of time that the population spends in each environment. Most likely, the population will not spend *precisely* the same amount of time in each environment, resulting in a small amount of net change. Both the generation-by-generation fluctuations and the net change that I have just described may be attributed to drift.

Note that there is an important point of contrast between indiscriminate sampling and the random fluctuations of evolutionary parameters: while indiscriminate sampling can operate across a single generation G_N and its immediate successor G_{N+1}, random fluctuations in the evolutionary parameters require several generations in which to cause drift.

Millstein [104, 16 - 17] does not agree that random fluctuations in evolutionary parameters constitute drift, because she believes that indiscriminate sampling is both necessary and sufficient for drift. This disagreement between us can be put aside, however, for the purposes of this dissertation. My argument in chapter 5 only cites indiscriminate sampling; if Millstein or someone who shares her view disagrees with the central conclusions that I reach here, it should not be because we disagree about the nature of drift.

The second mechanism of random drift other than indiscriminate sampling that I want to introduce here—not a family—consists of what I term "idiopathic events." I intend for this term to indicate events that recur rarely or are unique in the history of a species. This accords with the use of "idiopathic" in medicine, in which the term is used to indicate disease states that arise for unknown reasons, and connotes both uniqueness and spontaneity of the disease states to which it is applied.

In the evolutionary context in which I intend to apply "idiopathic events," severe storms, epidemics, and tidal waves fall into the category; often, such events are termed "bottlenecks" because they represent a rapid, significant constriction in population numbers, and catastrophic drift falls into this category. However, idiopathic events need not be dramatic. Ecological circumstances that are rare but not particularly marked in their effects fall into the category of idiopathic events. These events cause short-term fluctuations in allele frequencies that often occur, causing evolution in the absence of any net pressure from mutation, selection, or migration.

I count John Beatty's [9] thought experiment concerning some moths as an example of evolution due to idiopathic events. The background to this example is as follows. The organisms in question come in two forms, light-colored and dark-colored; and they live in an environment in which there are trees with light and dark bark. There is a greater proportion of trees with dark bark, and so dark-colored moths are fitter, because they can camouflage themselves against the dark background. The population has reached equilibrium due to selection, and it is assumed that the environment is stable. This means that, *ceteris paribus*, no evolution should be expected to occur in the population. Nevertheless, there is

one generation in which evolution occurs, as follows.

> [Suppose] that in one particular generation, we find, among remains of the moths killed by predators, a greater proportion of dark moths than was characteristic of the population as a whole, and a smaller proportion of light moths than was characteristic of the population.... [F]inally, let us say that we find the remains so distributed in the areas of dark and light trees that we have reason to believe that the dark moths were perched on light trees when attacked, and the light moths on dark trees.... [T]he dark moths chanced to land on light trees more frequently than on dark trees, even though the frequency of dark trees was greater. [9, 195]

I interpret the passage above to be about an event that occurs rarely, and so is an idiopathic event—and so, on my view, is an example of random drift.

Millstein [105] does not believe that drift occurs among the moths, because she does not believe that any indiscriminate sampling has occurred among them; as I mention above (page 284), above, she takes indiscriminate sampling to be both necessary and sufficient for random drift. As I also mention above, my disagreement with Millstein on this point is not a reason for her or someone who shares her view to disagree with me over the central claims about explaining drift that I make in this dissertation.

These two mechanisms of drift were first identified as drift by Sewall Wright in the appendix of a famous 1948 paper, "On the Roles of Directed and Random Changes in Gene Frequency in the Genetics of Populations" [152], written in response to Fisher and Ford. Wright describes these mechanisms in his typical style: highly compressed mathematical models with minimal prose explanation. Provine [118, 426] points out that, prior to 1948, Wright believed that indiscriminate sampling, which Wright terms "sampling drift," was the only mechanism of drift, and that he did not recognize the other forms of drift before 1948; and that after 1948, Wright always considered drift to be tripartite.

Bibliography

[1] Achinstein, Peter. *The Nature of Explanation*. New York: Oxford University Press (1983).

[2] ———. "The Symmetry Thesis." In James H. Fetzer, ed., *Science, Explanation and Rationality*. New York: Oxford University Press (2000).

[3] Achinstein, Peter, ed. *Science Rules*. Baltimore: Johns Hopkins University Press (2004).

[4] Afdhal, Nezam H. "The Natural History of Hepatitis C." *Seminars in Liver Disease* 24, supplement 2 (2004): 2 – 7.

[5] Allen, Colin, Marc Bekoff, and George Lauder, eds. *Nature's Purposes: Analyses of Function and Design in Biology*. Cambridge: MIT Press (1998).

[6] Allison, Peter A. and Derek E. G. Briggs. *Taphonomy: Releasing the Data Locked in the Fossil Record*. New York: Plenum Press (1991).

[7] Ayala, Francisco. "Teleological Explanations in Evolutionary Biology." In Allen *et al.* [5] (1998).

[8] Beatty, John. "Optimal-Design Models and the Strategy of Model Building in Evolutionary Biology." *Philosophy of Science* 47 (1980): 532 – 561.

[9] ———. "Chance and Natural Selection." *Philosophy of Science* 51 (1984): 183 – 211.

[10] Bouchard, Frédéric and Alexander Rosenberg. "Fitness, Probability, and the Principles of Natural Selection." *British Journal for the Philosophy of Science* 55 (2004): 693 – 712.

[11] Brandon, Robert. "Biological Teleology: Questions and Explanations." In *Concepts and Methods in Evolutionary Biology.* New York: Cambridge University Press (1996).

[12] ———. "The Difference Between Selection and Drift: A Reply to Millstein." *Biology and Philosophy* 20 (2005): 153 – 70.

[13] Bromberger, Sylvain. "An Approach to Explanation." In *On What We Know We Don't Know* [14] (1992).

[14] ———. *On What We Know We Don't Know.* Chicago: University of Chicago Press (1992).

[15] Carnap, Rudolf. *Logical Foundations of Probability.* Chicago: University of Chicago Press (1950).

[16] Carson, Hampton L. "The Population Flush and its Genetic Consequences." In Richard C. Lewontin, ed., *Population Biology and Evolution.* Syracuse: Syracuse University Press (1968).

[17] ———. "Chromosome Tracers of the Origin of Species." *Science* 168 (1970): 1414 – 1418.

[18] ———. "Speciation and the Founder Principle." *Stadler Genetics Symposia* 3 (1971): 51 – 70.

[19] ———. "Reorganization of the Gene Pool During Speciation." In Newton E. Morton, ed., *Genetic Structure of Populations.* Honolulu: University Press of Hawaii (1973).

[20] ———. "The Genetics of Speciation at the Diploid Level." *American Naturalist* 109 (1975): 83 – 92.

[21] Carson, Hampton L. and Alan R. Templeton. "Genetic Revolutions in Relation to Speciation Phenomena: The Founding of New Populations." *Annual Review of Ecology and Systematics* 15 (1984): 97 – 131.

[22] Cheetham, Alan H. "Tempo of Evolution in a Neogene Bryozoan: Rates of Morphologic Change Within and Across Species Boundaries." *Paleobiology* 12 (1986): 190 – 202.

[23] ———. "Tempo of Evolution in a Neogene Bryozoan: Are Trends in Single Morphologic Characters Misleading?" *Paleobiology* 13 (1987): 286 – 296.

[24] Collins, Aurthur. "The Use of Statistics in Explanation." *British Journal for the Philosophy of Science* 17 (1966): 127 – 140.

[25] Coyne, J. A., N. H. Barton, and M. Turelli. "Is Wright's Shifting Balance Process Important in Evolution?" *Evolution* 54 (2000): 306 – 317.

[26] Craver, Carl F. "Role Functions, Mechanisms, and Hierarchy." *Philosophy of Science* 68 (2001): 53 – 74.

[27] Crow, James. "Genes that Violate Mendel's Rules." *Scientific American* 240 (1979): 134 – 143.

[28] Cummins, Robert. "Functional Analysis." In Allen *et al.* [5] (1998).

[29] Dawkins, Richard. *The Extended Phenotype.* New York: Oxford University Press (1982).

[30] ———. *The Selfish Gene.* Oxford: Oxford University Press, new ed. (1989).

[31] ———. *The Blind Watchmaker.* New York: W. W. Norton and Company, reissued, new ed. (1996).

[32] ———. "Universal Darwinism." In Hull and Ruse [71] (1998).

[33] Dennett, Daniel. "Taking the Intentional Stance Seriously." *Behavioral and Brain Sciences* 6 (1983): 379 – 390.

[34] ———. *Darwin's Dangerous Idea.* New York: Touchstone (1995).

[35] Dietrich, Michael. "The Origins of the Neutral Theory of Molecular Evolution." *Journal of the History of Biology* 27 (1994): 21 – 59.

[36] Dikmen, Sureyya, Charles H. Bombardier, *et al.* "Natural History of Depression in Traumatic Brain Injury." *Archives of Physical Medicine and Rehabilitation* 85 (2004): 1457 – 1464.

[37] Dobzhansky, Theodosius G. *Genetics and the Origins of Species.* New York: Columbia University Press, 3rd ed. (1951).

[38] Dray, William. *Laws and Explanation in History.* Oxford: Oxford University Press (1957).

[39] Ducasse, C. J. "Critique of Hume's Conception of Causality." *Journal of Philosophy* 63 (1966): 141 – 148.

[40] Earman, John. *A Primer on Determinism.* Boston: D. Reidel Publishing Company (1986).

[41] El-Serag, H. B., P. Pilgrim, and P. Schoenfeld. "Systematic Review: Natural History of Irritable Bowel Syndrome." *Alimentary Pharmacology and Theraputics* 19 (2004): 861 – 870.

[42] Eldredge, Niles and Stephen Jay Gould. "Punctuated Equilibria: An Alternative to Phyletic Gradualism." In Thomas Schopf, ed., *Models in Paleobiology.* San Francisco: Freeman, Cooper and Company (1972).

[43] Endler, John. *Natural Selection in the Wild.* Princeton: Princeton University Press (1984).

[44] Ereshefsky, Marc, ed. *The Units of Evolution: Essays on the Nature of Species.* Cambridge: MIT Press (1992).

[45] Falconer, Douglas S. *Introduction to Quantitative Genetics.* New York: Wiley (1989).

[46] Fetzer, James A. "A Single Case Propensity Theory of Explanation." *Synthese* 28 (1974): 171 – 198.

[47] Fisher, Sir Ronald A. and E. B. Ford. "The Spread of a Gene in Natural Conditions in a Colony of the Moth *Panaxia dominula L.*" *Heredity* 1 (1947): 143 – 174.

[48] ———. "The 'Sewall Wright' Effect." *Heredity* 4 (1950): 117 – 119.

[49] Floch, Martin H. and Iona Bina. "The Natural History of Diverticulitis: Fact and Theory." *Journal of Clinical Gastroenterology* 38, supplement 1 (2004): S1 – S7.

[50] Freud, Sigmund. "Three Essays on Sexuality." In Peter Gay, ed., *The Freud Reader*. New York: W. W. Norton and Company (1989).

[51] Friedman, Michael. "Explanation and Scientific Understanding." In Pitt [115] (1988).

[52] Gage, Kenneth and Michael Y. Kosoy. "Natural History of Plague: Perspectives from More Than a Century of Research." *Annual Review of Entomology* 50 (2005): 505 – 528.

[53] Garland, James A., Hubert E. Jones, and Ralph L. Kolodny. "A Model for Stages of Development in Social Work Groups." In Saul Bernstein, ed., *Explorations in Groupwork*. Boston: Milford House (1973).

[54] Gayon, Jean. *Darwinism's Struggle for Survival: Heredity and the Hypothesis of Natural Selection*. Cambridge: Cambridge University Press (1998).

[55] Gillespie, John. *The Causes of Molecular Evolution*. New York: Oxford University Press (1991).

[56] ———. *Population Genetics: A Concise Guide.* Baltimore: The Johns Hopkins University Press (1998).

[57] Gilligan, Carol. *In a Different Voice.* Cambridge: Harvard University Press (1982).

[58] Glennan, Stuart. "Mechanisms and the Nature of Causation." *Erkenntnis* (1996): 49 – 71.

[59] ———. "Rethinking Mechanistic Explanation." *Philosophy of Science* 69 (2002): S342 – S353.

[60] Goodnight, C. J. and M. J. Wade. "The Ongoing Synthesis: A Reply to Coyne, Barton, and Turelli." *Evolution* 54 (2000): 317 – 324.

[61] Gould, Stephen Jay and Niles Eldredge. "Punctuated Equilibrium: The Tempo and Mode of Evolution Reconsidered." *Paleobiology* 3 (1977): 115 – 151.

[62] ———. "Punctuated Equilibrium Comes of Age." *Nature* 366 (1993): 223 – 227.

[63] Griffiths, Paul. "Functional Analysis and Proper Functions." In Allen *et al.* [5] (1998).

[64] Hempel, Carl G. "Aspects of Scientific Explanation." In *Aspects of Scientific Explanation* [65] (1965).

[65] ———. *Aspects of Scientific Explanation.* New York: The Free Press (1965).

[66] ———. "The Function of General Laws in History." In *Aspects of Scientific Explanation* [65] (1965).

[67] ———. "The Logic of Functional Analysis." In *Aspects of Scientific Explanation* [65] (1965).

[68] ———. "Studies in the Logic of Explanation." In *Aspects of Scientific Explanation* [65] (1965).

[69] Hubby, J. L. and Richard C. Lewontin. "A Molecular Approach to the Study of Genic Heterozygosity in Natural Populations. I. The Number of Alleles at Different Loci in *Drosophila Pseudoobscura*." *Genetics* 54 (1966): 577 – 594.

[70] Hull, David. "The Particular-Circumstance Model of Scientific Explanation." In Matthew H. Nitecki and Doris V. Nitecki, eds., *History and Evolution.* Albany: SUNY Press (1992).

[71] Hull, David L. and Michael Ruse, eds. *The Philosophy of Biology.* New York: Oxford University Press (1998).

[72] Humphreys, Paul. *The Chances of Explanation: Causal Explanation in the Social, Medical, and Physical Sciences.* Princeton: Princeton University Press (1989).

[73] Jackson, Jeremy B. C. and Alan H. Cheetham. "Phylogeny Reconstruction and the Tempo of Speciation in Cheilostome Bryozoa." *Paleobiology* 20 (1994): 407 – 423.

[74] Jaffar, Shabbar, Alison D. Grant, *et al.* "Systematic Review: Natural History of HIV-1 and HIV-2 Infections in Adults in Africa: A Literature Review." *Bulletin of the World Health Organization* 82 (2004): 462 – 469.

[75] Jeffrey, Richard. "Statistical Explanation vs. Statistical Inference." In Salmon [127] (1971).

[76] Kauffman, Stuart. *Origins of Order.* Oxford: Oxford University Press (1993).

[77] Keller, Joseph. "The Probability of Heads." *American Mathematical Monthly* 93 (1986): 191 – 197.

[78] Kimura, Motoo. "Evolutionary Rate at the Molecular Level." *Nature* 217 (1968): 624 – 626.

[79] ———. "The Neutral Theory of Molecular Evolution." *Scientific American* 241 (1979): 98 – 126.

[80] ———. *The Neutral Theory of Molecular Evolution.* Cambridge: Cambridge University Press (1983).

[81] King, Jack Lester and Thomas H. Jukes. "Non-Darwinian Evolution." *Science* 164 (1969): 788 – 798.

[82] Kitcher, Philip. "Explanatory Unification." In Pitt [115] (1988).

[83] Kitcher, Philip and Kim Sterelny. "The Return of the Gene." *Journal of Philosophy* 85 (1988): 335 – 358.

[84] Kohlberg, Lawrence. "The Claim to Moral Adequacy of a Highest Stage of Moral Development." *Journal of Philosophy* 70 (1973): 630 – 646.

[85] Kwan, P. and J. W. Sander. "The Natural History of Epilepsy: An Epidemiological Perspective." *Journal of Neurology, Neurosurgery, and Psychiatry* 75 (2004): 1376 – 1381.

[86] Lande, Russell. "Natural Selection and Random Drift in Phenotypic Evolution." *Evolution* 30 (1976): 314 – 334.

[87] Lee, K. S. "Natural History of Chronic Subdural Haemotoma." *Brain Injury* 18 (2004): 351 – 358.

[88] Lewens, Tim. "Adaptationism and Engineering." *Biology and Philosophy* 17 (2002): 1 – 31.

[89] Lewontin, Richard C. *The Genetic Basis of Evolutionary Change.* New York: Columbia University Press (1974).

[90] Lewontin, Richard C. and J. L. Hubby. "A Molecular Approach to the Study of Genic Heterozygosity in Natural Populations. II. Amount of Variation and Degree of Heterozygosity in Natural Populations of *Drosophila Pseudoobscura.*" *Genetics* 54 (1966): 595 – 609.

[91] Machamer, Peter, Lindley Darden, and Carl F. Craver. "Thinking About Mechanisms." *Philosophy of Science* 67 (2000): 1 – 25.

[92] Mallet, J. and M. Joron. "Evolution of Diversity in Warning Colour and Mimicry: Polymorphisms, Shifting Balance and Speciation." *Annual Review of Ecology and Systematics* 30 (1999): 201 – 233.

[93] Mallet, J. and M. C. Singer. "Individual Selection, Kin Selection, and the Shifting Balance in the Evolution of Warning Colours: The Evidence of Butterflies." *Biological Journal of the Linnean Society* 32 (1987): 337 – 350.

[94] Mallet, J. and J. R. G. Turner. "Did Forest Islands Drive the Diversity of Warningly Coloured Butterflies? Biotic Drift and the Shifting Balance Theory." *Philosophical Transactions of the Royal Society of London B* 351 (1996): 835 – 845.

[95] ———. "Biotic Drift or the Shifting Balance: did Forest Islands Drive the Diversity of Warningly Coloured Butterflies?" In P. R. Grant, ed., *Evolution on Islands.* Oxford: Oxford University Press (1998).

[96] Martin, Ronald E. *Taphonomy: A Process Approach.* Cambridge: Cambridge University Press (1999).

[97] Matthen, Mohan and André Ariew. "Two Ways of Thinking about Fitness and Natural Selection." *Journal of Philosophy* 99 (2002): 55 – 83.

[98] ———. "How to Understand Casual Relations in Natural Selection: Reply to Rosenberg and Bouchard." *Biology and Philosophy* 20 (2005): 355 – 364.

[99] Maynard-Smith, John. "Optimization Theory in Evolution." *Annual Review of Ecology and Systematics* 9 (1978): 31 – 56.

[100] Mayr, Ernst. "Change of Genetic Environment and Evolution." In J. Huxley, A. C. Hardy, and E. B. Ford, eds., *Evolution as a Process.* London: Allen and Unwin (1954).

[101] ———. *Animal Species and Evolution.* Cambridge: Harvard University Press (1963).

[102] ———. *Systematics and the Origin of Species From the Point of View of a Zoologist.* Cambridge: Harvard University Press, reprint ed. (1999).

[103] Michod, Richard E. *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality.* Princeton: Princeton University Press (1999).

[104] Millstein, Roberta. *The Chances of Evolution: An Analysis of the Roles of Chance in Microevolution and Macroevolution.* Phd dissertation, University of Minnesota (1997).

[105] ———. "Are Random Drift and Natural Selection Conceptually Distinct?" *Biology and Philosophy* 17 (2002): 33 – 53.

[106] ———. "Selection vs. Drift: A Response to Brandon's Reply." *Biology and Philosophy* 20 (2005): 171 – 175.

[107] Nei, Masotashi. *Molecular Evolutionary Genetics*. New York: Columbia University Press (1987).

[108] Newton, Cathryn R. and Léo Laporte. *Ancient Environments*. Englewood Cliffs. Prentice-Hall, 3rd ed. (1989).

[109] Northen, Helen and Roselle Kurland. *Social Work with Groups*. New York: Columbia University Press, 3rd ed. (2001).

[110] Ohta, Tomoko. "The Nearly Neutral Theory of Molecular Evolution." *Annual Review of Ecology and Systematics* 23 (1992): 263 – 286.

[111] ———. "Mechanisms of Molecular Evolution." *Philosophical Transactions of the Royal Society of London, B* 355 (2000): 1623 – 1626.

[112] ———. "Near-Neutrality in Evolution of Genes and Gene Regulation." *Proceedings of the National Academy of Sciences, USA* 99 (2002): 16134 – 16137.

[113] Oster, G. and P. Alberch. "Evolution and Bifurcation of Developmental Programs." *Evolution* 36 (1982): 444 – 459.

[114] Peebles, P. James, David N. Schramm, *et al.* "The Evolution of the Universe." *Scientific American* 271 (1994): 52 – 57.

[115] Pitt, Joseph, ed. *Theories of Explanation.* New York: Oxford University Press (1988).

[116] Popper, Karl. *The Logic of Scientific Discovery.* London: Routledge (1992).

[117] Press, Frank, Raymond Siever, *et al. Understanding Earth.* New York: W. H. Freeman and Company, 4th ed. (2003).

[118] Provine, William B. *Sewall Wright and Evolutionary Biology.* Chicago: University of Chicago Press (1986).

[119] Railton, Peter. "A Deductive-Nomological Model of Probabilistic Explanation." In Pitt [115] (1988).

[120] Reeve, Hudson Kern and Paul W. Sherman. "Adaptation and the Goals of Evolutionary Research." *Quarterly Review of Biology* 68 (1993): 1 – 32.

[121] Rosenberg, Alexander. "The Supervenience of Biological Concepts." *Philosophy of Science* 45 (1978): 368 – 386.

[122] ———. "Fitness." *Journal of Philosophy* 83 (1980): 457 – 473.

[123] Rosenberg, Alexander and Frédéric Bouchard. "Matthen and Ariew's Obituary for Fitness: Reports of its Death Have Been Greatly Exaggerated." *Biology and Philosophy* 20 (2005): 343 – 353.

[124] Roughgarden, Jonathan. *Theory of Population Genetics and Evolutionary Ecology: An Introduction.* New York: Macmillan (1979).

[125] Rudwick, M. J. S. "The Inference of Function from Structure in Fossils." In Allen *et al.* [5] (1998).

[126] Salmon, Wesley. "Statistical Explanation." In *Statistical Explanation and Statistical Relevance* [127] (1971).

[127] Salmon, Wesley, ed. *Statistical Explanation and Statistical Relevance*. Pittsburgh: University of Pittsburgh Press (1971).

[128] Salmon, Wesley. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press (1984).

[129] Schiller, L. Y. "Stages of Development in Women's Groups: A Relational Model." In Roselle Kurland and Robert Salmon, eds., *Group Work Practice in a Troubled Society: Problems and Opportunities*. Binghamton, NY: The Haworth Press (1995).

[130] Scriven, Michael. "Explananation and Prediction in Evolutionary Theory." *Science* 130 (1959): 477 – 482.

[131] ———. "Explanations, Predictions, and Laws." In Pitt [115] (1988).

[132] Shanahan, Timothy. "Selection, Drift, and the Aims of Evolutionary Theory." In Paul Griffiths, ed., *Trees of Life*. Dordrecht: Kluwer Academic Publishers (1992).

[133] Simpson, George Gaylord. "The Problem of Plan and Purpose in Nature." *The Scientific Monthly* 64 (1947): 481 – 495.

[134] ———. *Tempo and Mode in Evolution*. New York: Columbia University Press (1984).

[135] Singh, Simon. *Fermat's Enigma: The Epic Quest to Solve the World's Greatest Mathematical Problem*. New York: Walker and Company (1997).

[136] Skipper, Robert. *The R. A. Fisher-Sewall Wright Controversy in Philosophical Focus: Theory Evaluation in Population Genetics*. Phd dissertation, University of Maryland, College Park (2000).

[137] ———. "The Persistence of the R. A. Fisher-Sewall Wright Controversy." *Biology and Philosophy* 17 (2002): 341 – 367.

[138] Sober, Elliot. "Common Cause Explanation." *Philosophy of Science* 51 (1984): 212 – 241.

[139] ———. *The Nature of Selection*. Chicago: University of Chicago Press (1993).

[140] ———. "Six Sayings About Adaptationism." In Hull and Ruse [71] (1998).

[141] Stanley, Steven. "A Theory of Evolution Above the Species Level." *Proceedings of the National Academy of Sciences, USA* (1975): 646 – 650.

[142] ———. *Macroevolution*. Baltimore: The Johns Hopkins University Press (1998).

[143] Stanley, Steven and Xiangning Yang. "Approximate Evolutionary Stasis for Bivalve Morphology Over Millions of Years: A Multivariate, Multilineage Study." *Paleobiology* 13 (1987): 113 – 139.

[144] Stoppard, Tom. *Rosencrantz and Guildenstern are Dead*. New York: Grove Press (1967).

[145] Wade, M. J. and C. J. Goodnight. "Wright's Shifting Balance Theory: An Experimental Study." *Science* 253 (1991): 1015 – 1018.

[146] Walsh, Denis M., Tim Lewens, and André Ariew. "The Trials of Life: Natural Selection and Random Drift." *Philosophy of Science* 69 (2002): 452 – 473.

[147] Williams, George C. *The Pony Fish's Glow, and Other Clues to Plan and Purpose in Nature.* New York: Basic Books (1997).

[148] Wilson, R. A. *Species: New Interdisciplinary Essays.* Cambridge: MIT Press (1999).

[149] Winklerprins, Vince, Lotte van den Nieuwenhoof, *et al.* "The Natural History of Asthma in a Primary Care Cohort." *Annals of Family Medicine* 2 (2004): 110 – 115.

[150] Wittgenstein, Ludwig. *Philosophical Investigations.* New York: Macmillan, 3rd ed. (1968).

[151] Wright, Larry. "Functions." In Allen *et al.* [5] (1998).

[152] Wright, Sewall. "On the Roles of Directed and Random Changes in Gene Frequency in the Genetics of Populations." *Evolution* 2 (1948): 279 – 294.

[153] ———. "Character Change, Speciation, and the Higher Taxa." In *Evolution* [155] (1986).

[154] ———. "Coefficients of Inbreeding and Relationship." In *Evolution* [155] (1986).

[155] ———. *Evolution.* Chicago: University of Chicago Press (1986).

[156] ———. "Evolution in Mendelian Populations." In *Evolution* [155] (1986).

[157] ———. "Fisher and Ford on 'The Sewall Wright Effect'." In *Evolution* [155] (1986).

[158] ———. "Mendelian Analysis of the Pure Breeds of Livestock I: The Measurement of Inbreeding and Relationship." In *Evolution* [155] (1986).

[159] ———. "Mendelian Analysis of the Pure Breeds of Livestock II: The Duchess Family of Shorthorns as Bred by Thomas Bates." In *Evolution* [155] (1986).

[160] ———. "Mendelian Analysis of the Pure Breeds of Livestock III: The Shorthorns." In *Evolution* [155] (1986).

[161] ———. "On the Roles of Directed and Random Changes in Gene Frequency in the Genetics of Populations." In *Evolution* [155] (1986).

[162] ———. "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution." In *Evolution* [155] (1986).

[163] Zuckerandl, E. and L. Pauling. "Evolutionary Divergence and Convergence in Proteins." In V. Bryson and H. J. Vogel, eds., *Evolving Genes and Proteins.* New York: Academic Press (1965).

Vita

Personal Information

- Date of Birth: February 27, 1972.

Education

- Ph.D., Philosophy, May 2006, The Johns Hopkins University, Baltimore, MD.

- M.A., Philosophy, May 1999, Johns Hopkins.

- B.A., Biology and Philosophy, April 1995, The Univ. of Calif. at Santa Cruz, Santa Cruz, CA.

Honors and Awards

- David Sachs Memorial Fellowship. The Johns Hopkins University. 2001 – 2002.

- Dean's Teaching Fellowship. Johns Hopkins. Fall 2000.

- Honors, Philosophy. The Univ. of Calif. at Santa Cruz. April 1995.

- Honors, Senior Essay, Biology. The Univ. of Calif. at Santa Cruz. April 1995.

Published Reviews

- *Function, Selection, and Design,* and *Nature's Purposes: Analyses of Function and Design in Biology,* in *The Australasian Journal of Philosophy,* 80(2002): 126 – 8.

- *Darwinism's Struggle for Survival,* in *Mind* 110(2001): 204 – 7.

- *Teleological Language in the Life Sciences,* in *The Australasian Journal of Philosophy* 77(1999): 385 – 7.

Teaching Experience

Drew University, Madison, New Jersey, 2002 - 2005

- Philosophy of Science, Philosophy of Language, Seminar in Contemporary Philosophy, Analytic Philosophy, and Theory of Knowledge.

Johns Hopkins, 1997 - 2000

- Theory of Evolution and Philosophic Classics.

- Courses assisted with: Introduction to Deductive and Inductive Logic, Philosophic Classics, and Philosophical Problems.

The George Washington University, Washington, D.C., 1999

- Introduction to Logic.