# Why Selection and Drift Might Be Distinct

Jessica Pfeifer†‡

In this paper, it is argued that selection and drift might be distinct. This contradicts recent arguments by Brandon (2005) and Matthen and Ariew (2002) that such a distinction "violates sound probabilistic thinking" (Matthen and Ariew 2002, 62). While their arguments might be valid under certain assumptions, they overlook a possible way to make sense of the distinction. Whether selection and drift can be thought of as causally distinct, I argue, depends on the source of probabilities in natural selection. In particular, if the probabilities used in defining fitness values are at least partly a result of abstracting from or ignoring certain features of the environment, then selection and drift might in fact be causally distinct.

**1. Introduction.** Sober and Millstein have recently been criticized for treating natural selection and random genetic drift as distinct forces or processes. Matthen and Ariew argue that Sober's treatment of selection and drift as distinct "violates sound probabilistic thinking" (2002, 62). Brandon (2005) makes similar arguments in response to Millstein (2002). However, as I will show, their arguments are valid only if they make certain assumptions about the source of probabilities in natural selection. I will establish this by considering a commonly held view about the source of probabilities in natural selection—that they are at least partly a result of abstracting from or ignoring certain features of the environment. I will argue that such a view is reasonable, and, given this view, it is not conceptually confused to treat selection and drift as causally distinct.

**2. Why Such a Distinction Seems Conceptually Confused.** Since fitness is

†To contact the author, please write to: Department of Philosophy, University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore, MD 21250; e-mail: pfeifer@umbc.edu.

‡Thanks to Jim Thomas, Roberta Millstein, and the members of the D.C. History and Philosophy of Biology Group, especially Lindley Darden, Eric Saidel, Ilya Farber, and Jason Baker.

Philosophy of Science, 72 (December 2005) pp. 1135–1145. 0031-8248/2005/7205-0040$10.00
Copyright 2005 by the Philosophy of Science Association. All rights reserved.



https://doi.org/10.1086/508122 Published online by Cambridge University Press

defined probabilistically, an organism's or trait's actual offspring contribution can vary from that expected from the organism's or trait's fitness. In its broadest construal, random genetic drift occurs whenever this happens. To be more precise, "drift is any deviation from the expected result due to sampling error" (Brandon 2005). (I will discuss below a narrower definition of 'drift' that Millstein appears to defend and why such a narrower definition might be less desirable.) Since fitness is defined probabilistically, when an outcome differs from what is expected, it is not due to an additional causal factor—drift. Exactly the same thing explains the cases where the expected outcome happens and the cases where the actual outcome varies from that expected.

> Consider this analogy. You toss a coin four times. What would explain the outcome *two heads*? Answer: the physical setup of the coin-tossing trials. What would explain the outcome *four heads*? Answer: the same thing. (Matthen and Ariew 2002, 60–61)

Therefore, it does not make sense to treat selection and drift as distinct causal factors.

This argument alone, however, will not work against Millstein, since she, unlike Sober, seems to restrict the phrase 'random genetic drift' to cases where the organisms or traits involved are selectively neutral and, therefore, natural selection is precluded. In contrasting her view with Brandon's, she states, "I, on the other hand, would label the deviations from expectations as drift if and only if there were no physical differences that were causally relevant to the change in frequency" (2005, 172). Therefore, there seem to be two distinct types of cases, according to Millstein: those where the organisms or traits are equally fit (and therefore sampling is indiscriminate) and those where they are not (and therefore sampling is discriminate). Random drift occurs when, and only when, they are equally fit (when there are no physical differences that are causally relevant to differential success). When they are not equally fit, selection becomes the operating causal factor. In both cases, the actual outcome can vary from the expected outcome; however, 'random genetic drift' refers only to cases of sampling error when the organisms or traits are equally fit. Therefore, Millstein is not guilty of violating "sound probabilistic thinking." Random drift and selection are distinct because they are mutually exclusive.[1]

1. In a personal communication, Millstein has said that her view is compatible with the view I defend at the end of this paper, and therefore she is not committed to drift being limited to the neutral case. However, according to the view I defend, drift can occur even when there are physical differences that are causally relevant to changes in frequency, which is inconsistent with what Millstein explicitly states in the citation

https://doi.org/10.1086/508122 Published online by Cambridge University Press

The question that arises for such a view is whether there is good reason to restrict the definition of 'drift' in this way. Brandon (2005) argues that there is not. Not only does the narrower definition not fit well with the way many biologists define 'drift,' but, more importantly, there is no qualitative difference between the way the actual outcomes vary from the expected result in the two cases; in both cases, it is due to sampling error. This latter argument, though not directed at Millstein, is also developed in Brandon and Carson 1996.

While I think Brandon is correct, this does not entail that we *must* broaden the definition. Millstein is correct to point out that there might be multiple senses of 'drift' being used in biology. However, only the broader definition will do the work Millstein wants the concept of drift to do—in particular to make sense of the neutralist-selectionist debate. For the simplest version of this debate, her definition works perfectly well. The neutral theory, in its simplest form, asserts that most mutations are either neutral or deleterious. Selection will act to weed out most deleterious mutations, so these will rarely have much evolutionary impact. Therefore, most evolutionary change is due to drift acting on neutral mutations. The neutralist still requires the concept of sampling error to be applicable when selection is occurring, but the narrower definition does not preclude this. It simply is not considered random drift. The neutralist-selectionist debate can then be characterized, as it often is, as a debate about how much of evolution is due to drift—how often is evolutionary change due to selection and how often is it due to drift. Since these are mutually exclusive, there is no problem making sense of the debate. If the neutralist position were really this simple, this would provide a good reason for adopting the narrower definition I have attributed to Millstein. It would make sense of why the debate is often framed in terms of how much evolution is due to drift.

However, the neutralist position is not (or at least is no longer) so simple. As Beatty has noted (1984, 200), even Kimura's strong position has always had some "slack." Even in Kimura's 1968 paper, he does not claim that mutations are strictly neutral, but that they are "neutral (or nearly neutral)" (1968, 625). If this is how the neutralist position is framed, the narrower definition will be unhelpful. In cases of near neutrality, selection is operative, and, therefore, according to the narrower definition drift is not occurring. Consequently, the debate could no longer be framed as a debate about how much of evolution is due to drift rather than selection, since the neutralist would no longer be claiming that most of

above. My interpretation of Millstein might be incorrect. Nevertheless, the view I attribute to her is a possible way to treat selection and drift as distinct and is important to recognize even if it is not Millstein's.

https://doi.org/10.1086/508122 Published online by Cambridge University Press

evolution is due to drift. According to the narrower definition, the neutralist position could still hold even if drift never occurred. Consequently, statements like the following would not make sense:

> Finally, if my chief conclusion is correct, and if the neutral or nearly neutral mutation is being produced in each generation at a much higher rate than has been considered before, then we must recognize the great importance of random genetic drift due to finite population number in forming the genetic structure of biological populations. (Kimura 1968, 626)

If we define 'drift' narrowly, it would make drift irrelevant to the neutralist-selectionist debate. It is precisely for this reason that Beatty (1984) chooses not to define 'drift' in this narrow way.

If we are interested in making sense of the neutralist-selectionist debate, then the broader definition ought to be used. And, if the broader definition is correct, it seems from the above argument that treating natural selection and drift as distinct causes (or forces or processes) does not make sense. However, as I will show, sense can be made of the distinction, given certain assumptions about the source of probabilities in natural selection. This is not to say that Brandon's or Matthen's and Ariew's arguments are incorrect; rather, they overlook a consistent way to treat selection and drift as distinct. Moreover, I am not attempting to show that we ought to hold such a position, but rather that such a position does not, given certain assumptions, violate "sound probabilistic thinking."

**3. The Source of Probabilities in Natural Selection.** What I will argue is that, assuming that the probabilities used to define fitness are at least partly a result of ignoring or abstracting away from some features of the environment, it makes sense to treat selection and drift as causally distinct. The goal of this section is not to argue that we ought to accept this assumption about the source of probabilities in natural selection, but merely to outline some of the reasons for accepting it. My intent is simply to establish that the assumption is reasonable. I am primarily interested in what this assumption would imply about the possibility of treating selection and drift as distinct, which I will pursue in the next section. If it is a reasonable assumption, and if it allows for the possibility of treating selection and drift as causally distinct, then treating them as distinct does not involve a conceptual confusion.

Consider Scriven's (1959) example of two identical twins, one of whom (Josh) gets struck by lightning before reaching sexual maturity and dies, while the other (Joe) does not. Assume they are phenotypically identical. Given that fitness is always relative to an environment, how fine-grained ought our description of their environment be? Ought we say that Josh

https://doi.org/10.1086/508122 Published online by Cambridge University Press

and Joe were in different environments—that Josh was in an environment containing a lightning strike, whereas Joe was not—and that Josh's fitness in this environment is 0, or close to 0? If we do not want to say this, we must abstract from the lightning strike when describing the environment. Then we can say that they are in the same environment and that they are equally fit, since they are phenotypically identical; one of them, by "chance," just happens to do better.

In specifying an organism's (or trait's) fitness, we could fill in all the gory details of the environment. The question the twins example raises is whether we ought to do so. There are a number of arguments one can give for ignoring some of the details. First, it is argued that the use of probabilities in defining fitness does not rely on the assumption that the underlying (or micro) processes are indeterministic. As Sober argues, "If we assign the two twins identical fitness values, this implies no commitment that their difference in survivorship must be an irreducible matter of chance" (1984, 130). However, if the world is deterministic, and we filled in all the details, there would be no room left for probabilities in defining fitness. Moreover, even if the underlying micro-processes were not deterministic and such indeterminism "percolated up" to make fitness probabilistic (as Sober 1984, Brandon and Carson 1996, Stamos 2001, and Glymour 2001 have argued is possible), this would not yield the probabilities of interest to evolutionary biologists. As Graves, Horan, and Rosenberg have argued, even if probabilities can percolate up, the macro-level processes would "asymptotically approach determinism" (1999, 145). If so, all the probabilities used to define an organism's fitness would be close to 1 or 0. In addition, if we were to fully describe all the details of the environment, it is likely that few, if any, individuals would occupy the same environment. As Sober (1984) and Sterelny and Kitcher (1988) argue, this would prevent biologists from making clear the general patterns in evolution. Moreover, as Brandon argues (1990), selection requires that competing organisms be in the same environment. Therefore, if few (or no) individuals ever occupied the same environment, selection would rarely (if ever) occur. Instead, evolution would primarily be a function of how individuals are distributed across different environments.

Given these arguments, it is reasonable (and perhaps correct) to think that the probabilities used in defining fitness values are at least partly a result of ignoring or abstracting from some features of an organism's or type's environment. What remains to be discussed is what precisely this might mean and why this allows for the possibility of treating selection and drift as causally distinct. To this I now turn.

**4. Why Selection and Drift Might Be Causally Distinct.** What might it mean to ignore or abstract from certain features of the environment? Let

https://doi.org/10.1086/508122 Published online by Cambridge University Press

**E** be the conjunction of environmental factors that are not being ignored or abstracted from and $\mathbf{I}_{1,\ldots,n}$ be those features of the environment that are ignored. By ignoring or abstracting from the **I**'s, we can define an organism's or type's fitness relative to **E**, rather than relative to **E** and some combination of **I**'s. Suppose there are two environmental factors being abstracted from, $\mathbf{I}_1$ & $\mathbf{I}_2$. What we want is to focus on the *p*(Individual or type has *x* offspring/**E**) and not the *p*(Individual or type has *x* offspring/**E** & $\mathbf{I}_1$ & $\mathbf{I}_2$). This is formally the position outlined by Sober (1984), though as I will discuss there are various ways we might make sense of it.[2] However, as Sklar (1970) has argued and Sober acknowledges (1984, 131), in order to get a determinate value for *p*(Individual or type has *x* offspring/**E**), we would need to make it relative to, or conditional on, how the **I**'s are distributed.[3] The question, though, is how this might be understood such that we are actually ignoring or abstracting from the **I**'s and thereby yielding the probabilities we want in the process.

It cannot be conditional on how the **I**'s are *actually* distributed in any particular case. Consider first what this would mean for individuals. For an individual, say Joe, it would mean making the probability that Joe has a certain number of offspring conditional on the actual combination of **I**'s to which Joe is exposed. In such a case, no causally relevant factors are being ignored. For types, such as dark winged moths, it would mean making the probability that dark winged moths have a certain number of offspring conditional on how the ignored factors are actually distributed across moths with dark wings—what percentage of dark winged moths are exposed to which combinations of **I**'s. For types, however, there is an additional complication. Having dark wings is not the only phenotypic (or genotypic) factor affecting reproductive success. There are two general ways to deal with this complication: we can fully specify the type (phenotype or genotype) or we could define the fitness of the type as the average expected offspring contribution of individuals of that type. For example, we might fully specify the type as those with dark wings, who fly fast, and so on. Alternatively, we might want to make claims about the fitness of having dark wings, without fully specifying the phenotype or genotype, since this would allow for greater generality. Then we would need to make the value either relative to a particular population or conditional on how the other phenotypic (or genotypic) properties are dis-

2. In fact, in some places Sober seems to defend a view incompatible with the position I ultimately defend in this paper.

3. It might be objected that we could make the fitness values relative to the assumption that none of the ignored factors are present. However, this would also be a claim about how the ignored factors are distributed.

https://doi.org/10.1086/508122 Published online by Cambridge University Press

tributed among dark winged moths. The latter would allow us to abstract from a particular set of individuals and from assumptions about which individuals have which combinations of properties. I will focus on this last option, since it in many ways allows for the greatest generality. The points I will make, however, follow for the other versions as well.

Let $\mathbf{P}_{1,\ldots,n}$ be the various other phenotypic (or genotypic) properties. The fitness of dark moths would be equivalent to the average expected number of offspring of all dark moths, given **E** and given how the various combinations of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$ are distributed across moths with dark wings—what percentage of dark moths have or are exposed to which combinations of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$ (e.g., 10% are exposed to $\mathbf{I}_1$ and $\mathbf{P}_1$, 30% are exposed to $\mathbf{I}_1$ and $\mathbf{P}_2$, and so on). This would allow us to abstract from a particular set of individuals and from which individuals have which combinations of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$. We are thereby abstracting from information relevant to what happens to individuals. However, we are not abstracting from any information relevant to what happens to the type. As a result, we would undercut the very reasons we had for ignoring environmental factors in the first place. Making the values relative to the distribution in the way outlined would not only fail to give us the probability values of interest to biologists, it would also make it difficult, if not impossible, to get the sort of generality it was argued we need.

Consider first the former. If the number of offspring produced by a dark moth is a deterministic result of **E** and the combination of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$ to which the moth is exposed, then $p$(dark moths have two offspring/**E**) will be equivalent to the frequency of dark moths with two offspring, given **E** and given the percentages of dark moths that are exposed to each combination of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$. In other words, what happens to the type would not be probabilistic, and there would be no room for drift. If, on the other hand, the number of offspring produced by a dark moth is a probabilistic result of **E** and the combination of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$ to which the moth is exposed, but the probabilities approach 1 or 0, as per the arguments in the previous section, we would still not get the values we want. Suppose for simplicity that there are only two possible combinations of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$. Suppose that, given one of these combinations along with **E**, the probability that a dark moth has two offspring approaches 1, and given the other combination it approaches 0. We would also need to make the value for the type relative to the percentage exposed to each combination. Suppose that 60% of the moths are exposed to the first combination, and 40% to the latter. This would mean that 60% are virtually guaranteed two offspring and 40% are virtually precluded from having two. While there would be room for drift, it would be negligible. This would yield probabilities, but, assuming

https://doi.org/10.1086/508122 Published online by Cambridge University Press

the arguments of the previous section are correct, they are not the probability values of interest to biologists.

Additionally, if we only abstracted from which particulars have which combinations of $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$, we would also fail to get to the generality we need. As argued, we need generality for two reasons: so that individuals of different types can be in the same environment and to make the general patterns of evolution clear. As Brandon argues, in order for selection to occur, organisms of competing types would need to be in the same environment. However, by making fitness values conditional on how ignored factors are distributed, what counts as the same environment would be defined in part by this distribution. Any difference in the distribution of ignored factors would count as a difference in the environment. Therefore, competing types in the same population would need to have the same percentages exposed to the various combinations of ignored factors in order to be in the same environment. Since this is very unlikely, it would make selection nearly impossible. This would also make it difficult to draw generalizations about more than one population. There would be facts about particular populations abstracted from—in particular which individuals are exposed to which combinations of the $\mathbf{P}_{1,\ldots,n}$ and $\mathbf{I}_{1,\ldots,n}$. Consequently, there might be more than one population where the various combinations of **I**'s and **P**'s are distributed in exactly the same way. However, it would be very unlikely, and it certainly would not be likely that many populations would have the same distribution.

If making the fitness values conditional on the distribution of **I**'s is not the same as abstracting from or ignoring the **I**'s, then what could the latter mean? What we really want is to average over different ways the distribution might be. We want to specify what to expect on average if we looked at enough populations or if we looked at a large enough population. We can specify this using probabilities rather than specifying the actual distribution. We need not assume this probability represents any indeterminism in the world, nor need it reflect our ignorance. As Sober has argued, probabilities allow us to carve out generalizations (1984, 127). What the probability represents is what on average the distribution of ignored factors would be if we looked at enough populations or if we looked at a large enough population. We can then weight the various combinations of **I**'s according to their frequency, and we could represent this as a probability. I will focus on what this would mean for types, since this is in many ways less problematic. I will also be assuming for simplicity that the type is fully specified, and therefore there are no other relevant phenotypic or genotypic properties. Assume also that there are only two ignored environmental factors. The $p$(Type **A** has 2 offspring/**E**) would be equivalent to

https://doi.org/10.1086/508122 Published online by Cambridge University Press

$$[p(\text{Type } \mathbf{A} \text{ has 2 offspring}/\mathbf{E} \ \&\ \mathbf{I}_1 \ \&\ \mathbf{I}_2)p(\mathbf{I}_1 \ \&\ \mathbf{I}_2)$$

$$+ p(\text{Type } \mathbf{A} \text{ has 2 offspring}/\mathbf{E} \ \&\ \mathbf{I}_1 \ \&\ \sim\mathbf{I}_2)p(\mathbf{I}_1 \ \&\ \sim\mathbf{I}_2) + \cdots].$$

It is important to note that $p(\mathbf{I}_1 \ \&\ \mathbf{I}_2)$ does not represent the propensity of $(\mathbf{I}_1 \ \&\ \mathbf{I}_2)$ occurring. It represents the frequency $(\mathbf{I}_1 \ \&\ \mathbf{I}_2)$ would be expected in a large enough number of populations or in a large enough population. It is also important to note that what would be expected in a large number of populations might not be equivalent to what would be expected in a very large population. Which of these one chooses in specifying fitness values does not matter for the argument I am making. However, it might make a difference to what sorts of generalizations can be drawn and whether the fitness values can be used to make predictions about what happens as the population size varies.

If we abstract from environmental features in this way, then it makes sense to think of drift and selection as causally distinct. The fitness value is conditional on how we would expect the **I**'s to be distributed in a large enough number of populations or in a large enough population. However, how the **I**'s are actually distributed in a particular population may vary from this. When it does, this results in drift. Therefore, drift is caused by the distribution of **I**'s, in particular by their being distributed in a way different from what would be expected in a large enough number of populations or in a large enough population. Selection, on the other hand, occurs whenever competing types are exposed to **E**. Exposure to **E** causes selection, and the distribution of **I**'s causes drift. This will become clearer with an example.

Suppose you want to specify the fitness of dark and light moths in an environment where 60% of the trees are dark and 40% are light. Assume for simplicity that whenever a dark moth lands on a light tree it is eaten, and when a light moth lands on a dark tree it is eaten. There may be features of the environment we are ignoring that make it deterministic whether a particular moth will land on a particular tree. Let **E** be an environment with 60% dark trees, and $\mathbf{I}_{1,\ldots,n}$ be all the ignored factors, including those that affect whether a moth lands on a particular tree. Assume that what tree a moth lands on is not a function of the moth's color. If it were, then we ought not ignore this factor. If we were to look at enough populations (or at a large enough population) with 60% dark trees, we might then expect the **I**'s to be distributed so that both dark and light moths land on dark trees about 60% of the time. We would weight the various combinations of **I**'s according to their expected frequency, and we could represent these as probabilities.

There might also be other ignored environmental features that affect moths' reproductive success without affecting whether they land on dark

https://doi.org/10.1086/508122 Published online by Cambridge University Press

or light trees. Assume, for example, that disease on average strikes and kills 40% of moths. Assume, though, that whether an organism contracts the disease is not related to its color. Therefore, we might expect the disease to be equally distributed among dark and light moths in a large number of populations (or in a large enough population). Consequently, we can assume that on average 40% of both dark and light moths will die of disease. We can represent this as each moth having a probability of 0.4 of contracting disease. What we have effectively done is made the fitness values relative, not to the actual distribution of **I**'s, but to what on average this distribution would be if we looked at enough populations (or if we looked at a large enough population), and we specify this using probabilities. Since disease on average strikes 40% of moths, and since wing color is not linked to whether an organism contracts the disease, we would expect that, if we looked at enough populations (or a large enough population), disease would be distributed across approximately 40% of both dark and light moths.

In any particular population, it may be completely deterministic how the ignored factors are actually distributed, and it may be completely deterministic how the dark and light moths fare, given this distribution. In any particular population, there might be causes for why the moths do not land on just any tree and for why disease is not equally distributed in both dark and light moths. These features cause the actual offspring contribution of dark and light moths to vary from what would be expected if we looked at enough populations (or at a large enough population). In other words, how the ignored features are actually distributed can result in drift. They may not cause drift, because the actual distribution may be representative of what we would expect in multiple populations (or in a large enough population). These ignored features, however, are not the cause of selection. The cause of selection for having dark wings is that predators cannot see dark moths on dark trees as well as light moths. So long as dark and light moths sometimes land on dark trees, there is selection *for* dark wings, and selection for *light* wings if both land on light trees. Moreover, these selective factors are what make it the case that in a large number of populations (or in a large enough population) with both dark and light moths, when 60% of the trees in the environment are dark, dark moths will on average have greater reproductive success.

It has been argued that selection and drift are caused by the very same factors, but if this analysis is correct that might not be true. Drift is caused by the distribution of ignored factors; selection for the trait in question is caused by those features that are not ignored.

**5. Conclusion.** Brandon (2005) and Matthen and Ariew (2002) have argued that, assuming we define drift broadly so that selection and drift

https://doi.org/10.1086/508122 Published online by Cambridge University Press

can co-occur, it makes no sense to think of drift and selection as either being or having distinct causes. I have shown that, if the probabilities involved in fitness values are (at least partly) a result of ignoring features of the environment, then it does not violate "sound probabilistic thinking" to assert that drift and selection are causally distinct. I have not argued that we should accept such an assumption. Therefore, I have not shown that drift and selection ought to be thought of in this way. I have only argued that it is reasonable to do so.

## REFERENCES

Beatty, John (1984), "Chance and Natural Selection", *Philosophy of Science* 51: 183–211.
Brandon, Robert (1990), *Adaptation and Environment*. Princeton, NJ: Princeton University Press.
——— (2005), "The Difference between Drift and Selection: A Reply to Millstein", *Biology and Philosophy* 20: 153–170.
Brandon, Robert, and Scott Carson (1996), "The Indeterministic Character of Evolutionary Theory: No 'No Hidden Variables Proof' but No Room for Determinism Either", *Philosophy of Science* 63: 315–337.
Glymour, Bruce (2001), "Selection, Indeterminism, and Evolutionary Theory", *Philosophy of Science* 68: 518–535.
Graves, Leslie, Barbara Horan, and Alex Rosenberg (1999), "Is Indeterminism the Source of the Statistical Character of Evolutionary Theory?", *Philosophy of Science* 66: 140–157.
Kimura, Motoo (1968), "Evolutionary Rate at the Molecular Level", *Nature* 217: 624–626.
Matthen, Mahon, and Andre Ariew (2002), "Two Ways of Thinking about Fitness and Natural Selection", *Journal of Philosophy* 99: 55–83.
Millstein, Roberta (2002), "Are Random Drift and Natural Selection Conceptually Distinct?", *Biology and Philosophy* 17: 33–53.
——— (2005), "Selection vs. Drift: A Response to Brandon's Reply", *Biology and Philosophy* 20: 171–175.
Scriven, Michael (1959), "Explanation and Prediction in Evolutionary Theory", *Science* 130: 477–482.
Sklar, Lawrence (1970), "Is Probability a Dispositional Property?", *Journal of Philosophy* 67: 355–366.
Sober Elliott (1984), *The Nature of Selection*. Cambridge, MA: MIT Press.
Stamos, David (2001), "Quantum Indeterminism and Evolutionary Biology", *Philosophy of Science* 68: 164–184.
Sterelny, Kim, and Philip Kitcher (1988), "The Return of the Gene", *Journal of Philosophy* 85: 339–361.

https://doi.org/10.1086/508122 Published online by Cambridge University Press