Forces and Causes in Evolutionary Theory

Christopher Stephens†

The traditional view of evolutionary theory asserts that we can usefully understand natural selection, drift, mutation, migration, and the system of mating as forces that cause evolutionary change. Recently, Denis Walsh and Robert Brandon have objected to this view. Walsh argues that the traditional view faces a fatal dilemma and that the force analogy must be rejected altogether. Brandon accepts the force analogy but argues that drift, rather than the Hardy-Weinberg law, is the best candidate for a zero-force law. Here I defend the traditional view against these objections.

1. Introduction. Contemporary evolutionary theory invokes a number of factors—natural selection, random genetic drift, mutation, migration, and the system of mating—to explain evolutionary change. How should we understand these factors in evolutionary theory? Are natural selection and drift causes of evolutionary change or merely explanatory of such changes? Do these factors act in force-like ways, and if so, is there a "zero-force" law in evolutionary theory? In his book *The Nature of Selection* (1984), Elliott Sober defends what I will call the *traditional view*: "All possible causes of evolution may be characterized in terms of their 'biasing effects.' Selection may transform gene frequencies, but so may mutation and migration. And just as each possible evolutionary force may be described in terms of its impact on gene frequencies, so it is possible for a cause of evolution to be present without producing changes in gene frequencies. . . . All this is merely to locate evolutionary theory in familiar territory: it is a theory of forces" (31).

The traditional view is committed to five main theses:

†To contact the author, please write to: Philosophy Department, University of British Columbia, 1866 Main Mall, E370, Vancouver, BC V6T 1Z1, Canada; e-mail: cstephen@interchange.ubc.ca.

Philosophy of Science, 77 (December 2010) pp. 716–727. 0031-8248/2010/7705-0026$10.00
Copyright 2010 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/656821 Published online by Cambridge University Press

1. Natural selection, mutation, migration, drift, and so on are *conceptually* distinct.
2. Each of these factors—whether alone or in combination—is a possible *explanation* of evolutionary change.
3. Each of these factors should further be understood as separate *causes*.
4. The way in which these factors combine and interact to produce change in gene and/or genotype frequencies can be usefully thought of as a theory of forces. We can describe how each of these factors singly and in various combinations will combine to lead to evolutionary changes.
5. The Hardy-Weinberg law is the best candidate for a *zero-force law*—a statement about what will happen to a system when no forces act on it.

The main goal of this paper is to defend the traditional view against recent criticisms. Some of the recent criticisms are mistaken, while others are consistent with the traditional view, properly understood. Theses 1–5 are stated in approximate order of increasing controversy. The traditional view has come under attack from a number of different directions.[1] Critics such as Walsh (2007) argue that even though theses 1 and 2 are true, theses 3, 4, and 5 are false. Walsh's worry is that these factors such as natural selection and drift are not causes, and if they are not causes, then they are not forces. What are selection, drift, and so on if not causes of evolutionary change? Walsh argues that we should instead think of selection, drift, and these other factors as merely statistical explanations of evolutionary change. In an earlier paper (Stephens 2004), I criticized Matthen and Ariew's (2002) and Walsh et al.'s (2002) objections to the traditional view, and since then Walsh (2007) has replied with new arguments and objections. One of my goals in this paper is to assess his new objections. I argue here that they fail to undermine the traditional interpretation of evolutionary theory.[2]

1. It is a delicate exegetical question the extent to which Beatty (1984), Matthen and Ariew (2002), and/or Walsh, Lewens, and Ariew (2002) should be understood as challenging thesis 1 as well as theses 3, 4, and 5. All sides in these debates seem to accept some version of thesis 2. It is also possible that some critics of the traditional view would accept thesis 4 while rejecting thesis 3—as long as thesis 4 is given a noncausal reading. For replies to this earlier work, see Reisman and Forber (2005), Abrams (2007), Forber and Reisman (2007), Haug (2007), and Shapiro and Sober (forthcoming).

2. Matthen and Ariew (2009) have also replied (in a different way) to my earlier criticisms. Space prohibits dealing with their arguments here, though I plan to respond to them elsewhere.

https://doi.org/10.1086/656821 Published online by Cambridge University Press

However, in his recent work, Robert Brandon (2006) accepts theses 1–4 but objects only to thesis 5. His view is that we can and should think of evolutionary theory as a theory of forces (and causes) but that the traditional view is mistaken about what the zero-force law is. I will both respond to his criticisms of my 2004 defense of the traditional view and raise some difficulties with his alternative.

I will proceed as follows. In section 2, I elaborate the traditional view and provide reasons for thinking that it is correct. Since theses 1 and 2 are common ground with the critics that I am focusing on in this paper, I will not spend time defending these theses.[3] Instead, I will focus on why theses 3, 4, and 5 are attractive. In section 3, I then consider and respond to Brandon's reasons for rejecting thesis 5 of the traditional view and raise a limitation for his alternative approach. In section 4, I consider Walsh's more radical skepticism and examine the arguments that he gives in favor of rejecting theses 3, 4, and 5. I argue that his objections to the traditional view fail: he has not given us sufficient reason to doubt theses 3, 4, and 5. In section 5, I conclude by defending a modest pluralism about how to understand the structure of evolutionary theory: one that admits pros and cons to both the traditional interpretation and those such as Brandon that accept the first four theses of the traditional view but reject the fifth.

2. Goals and Attractions of the Traditional Approach. The traditional approach to understanding evolutionary theory has as one of its main goals representing the structure of the possible factors such as selection or drift in such a way that we can tell what will happen if none of these forces is at work: the zero-force law tells us about what *does not change* in such a case. In particular, in the traditional view, the zero-force law can be characterized as saying that the genotype (or merely gene) frequencies do not change if no forces are at work.[4] The second, related, attraction of the Hardy-Weinberg zero-force law is that it tells us *when a force has played a role*. The Hardy-Weinberg law tells us how the frequencies of traits in the gametes of the parents are related to the frequencies of traits in the offspring that develop from those gametes. The fact that genotype frequencies do not change is not sufficient to claim that no force was at work, but the fact that the frequencies do change is sufficient to claim that some force must have been at work.

3. For a defense of the claim that selection and drift are conceptually distinct, see Millstein (2002) and Stephens (2004).

4. The precise formulation of the zero-force law is a delicate issue in part because there are factors that can affect genotype frequencies without affecting gene frequencies (see Stephens 2004).

https://doi.org/10.1086/656821 Published online by Cambridge University Press

Besides providing a zero-force law that tells us when a force must have been at work, the force analogy also describes each of the possible factors and tells us how and when they can combine and/or cancel one another out.[5] Besides the fact that these forces do not have to combine additively, the traditional view of evolutionary theory is disanalogous with Newtonian physics in a couple of other ways: it is described against a background of a substantial biological assumption—namely, Mendelism. Furthermore, as will be discussed below, drift is a distinct kind of force from the other factors such as selection, mutation, and migration.

Despite these limitations, there are important respects in which the various factors in evolutionary theory can combine and interact in force-like ways. We can, for example, think about composing and decomposing causes within a particular force. Migration can be broken down into two causes: immigration and emigration. Together, these combine to give us the net effect of migration from one population to another. Similarly, mutation can be decomposed into two parts: if A and a are the alternate alleles at a locus, we consider both the rate at which A mutates into a and the rate at which a mutates into A.

Furthermore, there is composing and decomposing between different forces: the net mutation rates might favor A becoming a more than the reverse, but selection might favor A over a. So these different forces can combine to work in opposition to one another or in the same direction. Drift is perhaps a special case, and I will say more about it in the next section since it is the focus of Brandon's skepticism about the traditional view and his alternative way of representing the structure of evolutionary theory.

3. Brandon's Zero-Force Evolutionary Law (ZFEL). In his paper "The Principle of Drift: Biology's First Law," Brandon (2006) proposes an alternate zero-force law to the traditional Hardy-Weinberg one. His alternative is what he calls the evolutionary principle of drift, which has two parts:

(A) A population at equilibrium will tend to drift from that equilibrium unless acted on by an evolutionary force. (A population at rest will tend to start moving unless acted on by an external force.)

5. Critics of the force analogy (such as Matthen and Ariew 2002) sometimes mistakenly characterize the traditional view as committed to thinking that the forces must combine additively (see also Matthen and Ariew 2009). As Sober says, "Newtonian mechanics has made vector addition a familiar paradigm for computing the net effect of forces acting in concert. But it is only one example, and other more complex interactions are certainly possible. . . . Each theory of forces must solve this compositional problem for itself" (1984, 31–32).

https://doi.org/10.1086/656821 Published online by Cambridge University Press

(B) A population on evolutionary trajectory t, caused by some net evolutionary force F, will tend to depart from the extrapolated path predicted based on F alone (in either direction and magnitude or both) even if no other evolutionary force intervenes, unless F continues to act. (A population in motion will tend to stay in motion, but change its trajectory, unless continually acted on by an external force.) (328)

Brandon has three criticisms of the traditional view's approach to the zero-force law:

1. It depends on a contingent biological foundation (Mendelism).
2. Drift is not a force.
3. There is a better alternative—namely, ZFEL.

I have already commented on criticism 1: it is an important disanalogy with some theories of forces. But in the absence of an alternative that itself lacks any disanalogies with traditional theories of forces, it is an inconclusive objection. I will return to this issue when I consider Brandon's alternative zero-force law.

Brandon gives two reasons for thinking that drift is not a force. First, he claims that drift has only a magnitude, not a direction. If the Hardy-Weinberg law is to be a zero-force law, as the traditional view asserts, then Brandon says that "drift needs to be a force *just like* the other evolutionary processes" (2006, 324; emphasis mine). The traditional view, of course, acknowledges drift's anomalous status—"It [drift] is a force of a different color" (Sober 1984, 117)—but thinks that it is still illuminating to consider drift as one force among the others.

It is not hard to see why: although drift is a different kind of force from selection, mutation, migration, and so on, population genetics models still tell us how to combine drift with these other factors. Each of these factors can be given a common language in terms of either effects on gene or genotype frequencies. We can also talk about drift coming in different strengths, and we can talk about how its effects can be in the same direction as, or opposed to, other forces.

How is this possible if, as Brandon emphasizes, drift does not have a direction at the level of genic (allele) frequencies; by definition, isn't drift supposed to be nondirectional? Indeed it is. But remember that population biologists are interested in two types of evolutionary change: change in gene frequencies and change in genotype frequencies. In response to my earlier defense (Stephens 2004) of the force analogy, Brandon writes, "He [Stephens] suggests that it [drift] predictively leads to loss of heterozygosity and increase of homozygosity. That is a prediction, but a prediction without a direction. It does not say which of A_1A_1 or A_2A_2 will increase in

https://doi.org/10.1086/656821 Published online by Cambridge University Press

frequency. In physics that would be like saying a 20-Newton force is acting on object *A*. Such a statement either makes no sense (the magnitude, but not direction has been specified) or is incomplete. . . . Notice that we were able to specify the direction of the other evolutionary processes" (2006, 325). But the mere fact that there is an important disanalogy here with Newtonian physics is not enough to sink the traditional view—given that the traditional view acknowledges such disanalogies and, more important, that Brandon's own favored alternative has disanalogies of its own.

It is worth emphasizing that there is an important sense in which drift does have a direction. *Drift tends to remove variation from natural populations*. This is an important role that drift plays, in part because there are other forces that can increase variation in a population (e.g., mutation). So mutation can oppose the direction of drift in this sense. Furthermore, selection can oppose drift as well (e.g., selection can favor a new advantageous mutation that drift opposes). Population genetics is full of models and equations that tell us how to balance the direction of drift (removing variation) with forces that maintain variation. One of the main obsessions of much of population genetics over the last century has been how much variation is in nature and what its possible causes are. This is no small matter.

Brandon also objects that "drift is not a 'special' force in evolution; it is the default position. By that I mean that it is part and parcel of a constitutive process of any evolutionary system" (2006, 325). Brandon is pointing out that in any real population (because finite), drift will always be at work. Other forces such as selection are not always present in a population. So drift has a special place in evolutionary theory and should be the zero-force law. My reply to this objection is that we should not mix how pervasive drift is with whether it is a force or a cause. We would not rule out gravity as a force in a Newtonian universe that contained objects with mass simply because gravity would be part of any such system.

What about Brandon's own alternative? One of the nice features of the traditional view with Hardy-Weinberg as a zero-force law is that if there is deviation from what Hardy-Weinberg predicts, some force must have been at work. The zero-force law provides one with a condition such that if it is violated, then some force—drift, selection, mutation, and so on—must be working. The best that Brandon can say with his alternative framework is that if there is deviation from the zero-force law, then we can conclude that probably some force was at work. So the traditional view has a kind of heuristic epistemic advantage in this regard. Whether one thinks of the structure of evolutionary theory in Brandon's way or in the traditional way, one has to admit certain disanalogies with the

https://doi.org/10.1086/656821 Published online by Cambridge University Press

Newtonian picture. On the traditional view, drift must be treated as a different kind of force from selection, mutation, and so on. On Brandon's view, the zero-force law no longer provides a condition for what will not change under certain circumstances. Both of these features—the commensurability of forces and the ability to state a condition for what will not change under certain circumstances—were attractive features of the zero-force approach in physics. One lesson here is that it appears that we cannot have both of these features in biology.

I think that the right solution here is a kind of modest pluralism about the structure of evolutionary theory. Each of these ways of thinking about the structure of evolutionary theory is an acceptable one—each has its pros and cons as a theory of forces. On neither approach does evolutionary theory turn out to be just like Newtonian physics. On both approaches we can make sense of the way in which components of a force (e.g., immigration and emigration) can combine in forcelike ways, as well as how different components (e.g., selection and drift) can work for or against one another. Each interpretation may have its own useful heuristic value, but it is hard to see that there is anything that would decide that one interpretation is better than another *tout court*.

4. Walsh's New Argument for Statisticalism. In his earlier papers (e.g., Walsh et al. 2002), Denis Walsh argued against the traditional (he calls it "dynamical") view and in favor of a purely statistical interpretation of evolutionary theory. In Walsh (2007), he raises a new objection to the traditional, dynamical interpretation. His main argument is in the form of a dilemma for the traditional view: it must either accept a contradiction or accept a loss of explanatory power (compared to the statistical view, which faces neither problem). My aim in this section is to argue that Walsh's dilemma is a false one: the traditional, dynamical view faces neither a contradiction nor a loss of explanatory power.

Walsh's alternative interpretation of evolutionary theory is that modern synthesis explanations do not account for evolutionary change by citing the actions of specific causal processes of drift, selection, and so on. Rather, they explain simply by citing the statistical properties of populations. In Walsh's view, "what it is for a change in relative trait frequencies to constitute selection (or drift) is merely for it to be susceptible to a certain kind of statistical description" (2007, 282).

Walsh develops his dilemma by first assuming that "*causal relations are description-independent*. By this I mean that if x causes y, then this relation holds no matter how x and y are described" (2007, 292–93). I have some

https://doi.org/10.1086/656821 Published online by Cambridge University Press

questions about this assumption and how it is to be interpreted.[6] At any rate, I agree that if selection and drift have rank order effects like those described in his apple and coin examples (discussed below), this would undermine the dynamic interpretation. Even if we accept his assumption D1 ("If selection and drift occur within a population, they do so *no matter how the population is described*" [293]), I think that he presents a false dilemma.

Walsh develops his dilemma by describing a coin-tossing analogy. In various coin-flipping scenarios that I describe below, he argues that there is no fact about how much drift or selection there is; only the statistical interpretation can properly account for this. The coin-flipping experiment goes like this. Two fair coins (coin 1 and coin 2) are each tossed 50 times. The first experimenter tosses coin 1 10 times, and the second experimenter tosses coin 2 10 times. Then each experimenter hands her coin off to another experimenter, who tosses the coin 10 more times, until each coin has been tossed a total of 50 times (10 times each by five different experimenters). We are supposed to think of the probability of the coin landing heads or tails as analogous to the fitnesses of two alleles at a locus and the number of tosses as analogous to drift.

In Walsh's experiment, the results of the coin tossing are as follows: in the 10 series of 10 tosses of the coins, we get the following numbers of heads/tails: 4/6, 7/3, 3/7, 5/5, 4/6, 5/5, 5/5, 5/5, 4/6, 7/3. We can, as Walsh notes, describe these results in at least three ways. We can describe them as the result of 10 experiments (each experimenter is a treatment), or we can summarize the results for each coin: coin 1 20H/30T and coin 2 29H/21T. Finally, we can describe the overall results: 49H and 51T. Which of these is the correct description? Walsh argues that each of these three descriptions is equally legitimate. He accepts Millstein's (2002) and Stephens's (2004) earlier emphasis on the distinction between process and product (outcome) notions of selection and drift.[7] Drift-the-process varies with the effective population size, so that a smaller population size means that drift is stronger. According to Walsh, there are three different ways of describing the strength of drift-the-product in this experiment: drift being the strongest according to the first description (where there are only

6. Does it, e.g., allow that causation may be contrastive so that x rather than z might cause y but x *simpliciter* does not? Some analyses of causation are contrastive in this way.

7. The process-product distinction is one of the two distinctions that Stephens (2004) argued that critics of the traditional view missed. Walsh (2007) now accepts this distinction. However, it is unclear the extent to which he (or Matthen and Ariew 2009) understands the importance of the other major distinction that I emphasized in my earlier paper—that between *natural selection* and *fitness*.

https://doi.org/10.1086/656821 Published online by Cambridge University Press

10 trials) and drift being the weakest in the third description (where each coin is described as being tossed 50 times). Each of these descriptions is explanatory, according to Walsh. We need the first description in order to explain why there is so much observed deviation from 50% H and 50% T in each of these 10 flip experiments. However, we need the third description to explain why the overall outcome (49% T and 51% H) is so close to the expectation of a fair coin. In particular, as Walsh notes, we can appeal to the central limit theorem to explain this because it "entails that the means of samples of the same size drawn from a normal distribution will themselves be normally distributed about the population mean" (2007, 296).

So now Walsh presents the dynamical view with the dilemma. It must accept either of two horns: (1) Drift is both objectively strong (if we describe the experiment as 10 cases of 10 flips) and objectively weak (if we think of it as an overall population of 100 flips). (2) One of these accounts—saying that the experiment is really 10 trials of 10 tosses, or saying that it is two trials of 50 tosses, or saying that it is 100 tosses—must be the "canonical description."

Horn 1 is a contradiction, and so that leaves only horn 2, according to Walsh. He raises two problems with the second horn. First, it is arbitrary to pick one of these descriptions as the correct one: there is no reason in these coin-tossing experiments to think that one of these three alternative descriptions is the objectively right one. Second, even if one does pick one of these accounts as canonical, one will forfeit explanatory power. Why is this? If the experiment is viewed as 10 series of 10 flips, then one cannot explain why overall (49H and 51T) there is not much drift. Conversely, if one understands the experiment as a single series of 100 tosses, then there is no explanation for why there is often so much more deviation from 50% H and T in the subsamples of 10 flips. Each explanation, Walsh argues, is explanatorily indispensable, and only the statistical interpretation can recognize this.

How might a defender of the traditional interpretation respond to Walsh's arguments? There are, I think, at least two responses open to the defender of the traditional, dynamical interpretation. Both of these responses result from a common source: it is not clear that Walsh's coin-tossing analogies are sufficiently analogous to real cases of drift and selection. Allow me to explain.

Rather than coin tossing, let us consider a more realistic case in which drift exists: in gamete production and selection. Suppose we have a parent organism that has either *A* or *a* allele at some locus with 50%/50% representation in the gametes (e.g., it is a male organism, and half of its sperm have *A* and half have *a*). If no other evolutionary factors are at work (no selection, etc.) and if we ignore drift, then 50% of the offspring

https://doi.org/10.1086/656821 Published online by Cambridge University Press

of such parents will receive *A* and 50% will receive *a* at that locus. Now, imagine 10 such individuals each producing 10 (offspring) gametes, so that the case is structurally similar to Walsh's coin-tossing scenarios. That is, there are some individuals who have quite skewed gamete production (e.g., 70% of one or the other), but the overall population has close to 50% *A* and 50% *a*. We can then ask two questions about this population: why is there so much deviation from 50/50 in so many of the 10 individuals? And why in the overall population of 100 individuals is there less deviation?

The dynamical interpretation can answer these questions by thinking of drift as a two-place relation: we can say that relative to an individual parent, drift is strong, but relative to the overall population, drift is weak(er). One is not saying, of the overall population, that drift is both strong and not-strong, *simpliciter*. Instead of asserting *X* and not-*X*, one is instead saying *X* with respect to *Y* but not-*X* with respect to *Z*. So there is no contradiction. Drift is strong in each subpopulation but weak overall. In this case, we can think of the 10 individual parent organisms as each providing a microenvironment against which drift is operating. If these 10 individuals were isolated on 10 different islands, then population geneticists would not be inclined to think of them as one population of 100. In such a case, *facts about the biology and environment* make it the case that we would treat the effective population size as 10 for purposes of computing drift. We could of course still appeal to statistics to explain why the overall average across these 10 populations was closer to 50/50, but a biologist would not need to invoke drift to explain this.

However, if these 10 individual parents were in the same geographic area (i.e., their offspring had some chance of interacting with the other parents' offspring in some kind of population), then most population biologists would be inclined to use 100 as the effective population size number for purposes of computing drift. But the reason is that they are usually interested in such larger-scale trends. They want to know why the gene or genotype frequencies are the way that they are in the overall population. But this would not prohibit them from asking the question about the particular individuals if they had wanted to. Theoretically, we can say that drift is strong locally, in the short run, and so on but weak in the overall population.

According to Walsh, we can run an analogous set of experiments to illustrate the description-dependence of selection. If we imagine that coin 1 has a .6 chance of landing heads and coin 2 a .4 chance of landing heads, we can again have each of the 10 experimenters toss a coin 10 times (so each coin is tossed a total of 50 times in groups of 10). This time, however, before each toss an experimenter chooses a coin (either coin 1 or coin 2) at random. So in the overall population, there is the

https://doi.org/10.1086/656821 Published online by Cambridge University Press

same probability of getting heads or tails—namely, .5. But within a group of 10 flips, there may be cases in which, by chance, an experimenter chooses coin 1 more times than coin 2. The defender of the traditional view can make the same reply to this case, mutatis mutandis.

Walsh concludes by claiming that "there is no objective fact of the matter how much selection or drift a population is experiencing" (2007, 301). This would come as quite a surprise to biologists involved in debates and discussions about the relative strength of these factors. A defender of the statistical interpretation might be tempted to say: it is precisely because there is no fact of the matter that there is so much unresolved controversy about the neutralism debate, for instance. But this, I think, is a highly implausible claim. For one thing, there are obvious cases in which one knows—on broadly theoretical grounds—that drift is the cause of the evolution of a trait.[8]

But perhaps it remains impossible to disentangle drift and selection in cases in which the focus is on phenotypic characteristics rather than molecular ones. But if so, then skeptics need to examine realistic cases and explain how such debates rest on confusion. Millstein (2008) argues that the so-called Great Snail Debate between biologists of the 1950s is a case in which biologists did successfully distinguish between the effects of selection and those of drift and were able to clarify the kind of observations that would likely count in favor of an important role for drift. These snails live in a number of relatively isolated populations, with some populations dominated by banded shells and others by unbanded ones. Furthermore, the frequencies of various shell colors also vary from population to population. Some biologists, such as Cain and Sheppard (1950), sought to give a nearly exclusive role to selection in explaining the patterns of variation found in the shells. Others, such as Lamotte (1952), argued that in addition to selection, drift played an important role, in large part because there is a greater divergence in shell properties among the smaller populations than among the larger ones. As Millstein points out, it is hard to make sense of either side in these debates if there are not really objective ways to determine, for instance, what the effective population size is for these groups of snails. I encourage critics of the traditional, dynamical interpretation to discuss more realistic examples such as these.

5. Conclusions. I have argued that Brandon's objections to the traditional (dynamical) view fail. Furthermore, on either the traditional approach or Brandon's alternative, one has to accept certain disanalogies with New-

8. For instance, Ridley (1996, 181) notes that most biologists accept that the rapid evolution of pseudogenes is due to drift; the reason is that some pseudogenes cannot be transcribed.

https://doi.org/10.1086/656821 Published online by Cambridge University Press

tonian physics. At the same time, using Brandon's framework may have different heuristic value than the traditional approach, and so we should accept a moderate pluralism between these two approaches to thinking about the structure of evolutionary theory.

Walsh (2007) raises a new objection to the dynamical approach: that it either is contradictory or must face a loss of explanatory power. I have argued that this is a false dilemma and that the traditional view faces neither problem. The coin cases he describes are not sufficiently analogous to real cases of drift and selection. There may of course be reasons to doubt that drift and selection are forces or causes, but they are not the reasons given in Walsh's paper.

REFERENCES

Abrams, Marshall. 2007. "How Do Natural Selection and Random Drift Interact?" *Philosophy of Science* 74:666–79.

Beatty, John. 1984. "Chance and Natural Selection." *Philosophy of Science* 51:183–211.

Brandon, Robert. 2006. "The Principle of Drift: Biology's First Law." *Journal of Philosophy* 103:319–35.

Cain, Arthur, and Philip Sheppard. 1950. "Selection in the Polymorphic Land Snail *Cepaea nemoralis*." *Heredity* 4:275–94.

Forber, Patrick, and Kenneth Reisman. 2007. "Can There Be Stochastic Evolutionary Causes?" *Philosophy of Science* 74:616–27.

Haug, Matthew. 2007. "Of Mice and Metaphysics: Natural Selection and Realized Population–Level Properties." *Philosophy of Science* 74:431–51.

Lamotte, Maxime. 1952. "Le rôle des fluctuations fortuites dans la diversité des populations naturelles de *Cepaea nemoralis* (L.)." *Heredity* 6:333–43.

Matthen, Mohan, and André Ariew. 2002. "Two Ways of Thinking about Fitness and Natural Selection." *Journal of Philosophy* 119:55–83.

———. 2009. "Selection and Causation." *Philosophy of Science* 76:201–24.

Millstein, Roberta. 2002. "Are Random Drift and Natural Selection Conceptually Distinct?" *Biology and Philosophy* 17:33–53.

———. 2008. "Distinguishing Drift and Selection Empirically: 'The Great Snail Debate' of the 1950s." *Journal of the History of Biology* 41:339–67.

Reisman, Kenneth, and Patrick Forber. 2005. "Manipulation and the Causes of Evolution." *Philosophy of Science* 72:1113–23.

Ridley, Mark. 1996. *Evolution*. 2nd ed. Cambridge: Blackwell Science.

Shapiro, Larry, and Elliott Sober. Forthcoming. "Epiphenomenalism—the Do's and the Don'ts." In *Studies in Causality: Historical and Contemporary*, ed. G. Wolters and Peter Machamer. Pittsburgh: University of Pittsburgh Press.

Sober, Elliott. 1984. *The Nature of Selection*. Cambridge, MA: MIT Press.

Stephens, Christopher. 2004. "Selection, Drift, and the 'Forces' of Evolution." *Philosophy of Science* 71:550–70.

Walsh, Denis. 2007. "The Pomp of Superfluous Causes: The Interpretation of Evolutionary Theory." *Philosophy of Science* 74:281–303.

Walsh, Denis, Tim Lewens, and André Ariew. 2002. "The Trials of Life: Natural Selection and Random Drift." *Philosophy of Science* 69:452–73.

https://doi.org/10.1086/656821 Published online by Cambridge University Press