

Selection, Drift, and the “Forces” of Evolution*

Christopher Stephens^{†‡}

Recently, several philosophers have challenged the view that evolutionary theory is usefully understood by way of an analogy with Newtonian mechanics. Instead, they argue that evolutionary theory is merely a statistical theory. According to this alternate approach, natural selection and random genetic drift are not even causes, much less forces. I argue that, properly understood, the Newtonian analogy is unproblematic and illuminating. I defend the view that selection and drift are causes in part by attending to a pair of important distinctions—that between *process* and *product* and that between *natural selection* and *fitness*.

1. Introduction. Philosophers and evolutionary biologists often appeal to an analogy with Newtonian mechanics to understand the interaction of natural selection, mutation, migration, and random genetic drift.¹ According to this traditional approach, factors such as selection and drift are viewed as “forces” that are the possible causes of evolutionary change. Perhaps the most developed use of the Newtonian analogy to discuss evolutionary theory occurs in Elliott Sober’s book, *The Nature of Selection* (1984):

All possible causes of evolution may be characterized in terms of

*Received June 2003; revised December 2003.

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

[‡]Thanks to Robin Andreasen, André Ariew, John Beatty, Vincent Bergeron, Bruce Glymour, Daniel Hausman, Selman Halabi, Katriina Ilves, Tim Lewens, Mohan Matthen, Anya Plutynski, Alex Rosenberg, Steve Savitt, Elliott Sober and Denis Walsh for helpful comments and discussion. Thanks also to Lisa Lloyd for her comments on an earlier version of this paper at the 2003 Central Division meetings of the American Philosophical Association.

1. Biologists who make use of the force analogy include Roughgarden (1979, 18) and Gillespie (1998, 19). Biologists who criticize the force analogy include Endler (1986).

Philosophy of Science, 71 (October 2004) pp. 550–570. 0031-8248/2004/7104-0009\$10.00
Copyright 2004 by the Philosophy of Science Association. All rights reserved.

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 to locate evolutionary theory in familiar territory: it is a theory of forces. (34)

There are four important Newtonian features to what I'll call the *traditional view* of evolutionary theory.

Forces as causes. Natural selection, mutation, migration and random genetic drift are forces (causes) that can result in changes in trait frequencies (evolution).

Zero-force law. Evolutionary theory has a zero-force law that states what will happen if no forces impinge on the system.

Singleton force models. Evolutionary theory provides models to represent how each force—selection, drift, and so on, acts alone—and also an account of how to combine them.

Resolution of forces. Different possible evolutionary forces such as selection and mutation can interact and combine in a Newtonian fashion. Net forces can be decomposed into component forces.²

Recently, there have been a growing number of challenges to these claims. Walsh (2000, 2004) and Walsh, Lewens, and Ariew (2002) argue that the forces analogy is misguided and defend what they call a "statistical" interpretation of evolutionary theory. Matthen and Ariew (2002) argue, for somewhat different reasons, that the force analogy is misleading and should be rejected as a way to understand the various possible factors involved in evolutionary theory.

Although there are some differences in detail, all of these critics want to defend a different interpretation of evolutionary theory. Following Walsh, Lewens, and Ariew (2002), I will call this alternative the *statistical interpretation* of evolutionary theory.³ In addition to denying the theses

2. It is unclear to what extent the traditional view accepts the resolution of forces condition. Matthen and Ariew (2002, 59) describe Sober as accepting this condition; however, Sober (1984, 115–117) denies that drift is decomposable in this way, and points out that forces do not have to combine additively as they do in Newtonian physics (32). Matthen and Ariew misrepresent Sober's view on other details as well. I return to these issues later.

3. Sterelny and Kitcher (1988, 344–345), in an influential article on the units of selection problem, also defend a version of the statistical interpretation. They claim that the aim of evolutionary theory is just to make clear the "central tendencies" in evolving populations. Rosenberg and Bouchard (2004) develop a critique of the central tendencies approach that is in some ways complementary with my approach here.

of the force analogy, defenders of the statistical interpretation defend the following positive claim.⁴

Evolutionary forces are pseudo-processes. Natural selection, mutation, migration, and drift are not real forces acting on populations; rather, they are merely statistical byproducts of a population that result from a sequence of events (births, deaths, and reproduction). There are no genuine forces or causes in evolution at the population level.

Here is an analogy: think of the actuarial notion of “overall life expectancy.” Overall life expectancy is not a *cause* of how long one lives; rather, one’s life expectancy is a statistical summary of information about all (or a large number of) the possible causes (each weighted by its chance of occurring) that might affect your survival. Similarly, on the statistical interpretation of evolution, selection, drift, mutation, and migration are not forces that act on populations; rather, they are statistical properties of a collection of single trial events (Walsh, Lewens, and Ariew 2002, 453; Matthen and Ariew 2002, 59).

Understanding these issues properly is important, according to these critics. In addition to its own intrinsic interest, these authors argue that some of the most heated debates in evolutionary theory are illuminated by the statistical interpretation—e.g., what natural selection explains, adaptationism, and the units of selection problem. I have my doubts. My contention is that these criticisms of the traditional view (properly understood) fail. Throughout the paper, I clarify the sense in which the Newtonian analogy is—and is not—appropriate. In addition, I defend the view that selection, drift, and so on are causes of evolutionary change.

2. Combining Different Components of Fitness. Part of the problem in thinking about the causes of evolution is that there are *several* distinct issues about how to harmonize various causes—one set of issues concerns how to make sense of various components of natural selection, and another concerns how to combine natural selection with other factors such as drift, mutation, and migration. It is important to keep these issues distinct. I will argue that some of these problems have a more “Newtonian” resolution than others do.

The first harmonization problem is how to combine different components of fitness. That is, *within* the concept of natural selection, how do the different “forces” add up? Take the trait, resistance to malaria. Suppose that organisms in the population are either resistant to malaria or not, and are either quick or slow. The fitness of these combinations is represented in Table 1.

4. These critics don’t deny the Singleton Force Models condition, provided that talk of “forces” is expunged.

TABLE 1

SPEED	RESISTANCE TO MALARIA	
	Yes	No
Quick	<i>W</i>	<i>X</i>
Slow	<i>Y</i>	<i>Z</i>

We can then sample individual organisms from the population to estimate the four fitnesses. Suppose we get the estimates shown in Table 2.

We can then conclude that being both quick and having resistance to malaria is the fittest combination of traits, being slow and resistant to malaria is the next fittest, and so on. We can also say (in this case) that resistance to malaria makes a *bigger* difference to an organism’s fitness than being quick (rather than slow).⁵ Finally, we can define the fitness values of *singleton* traits. For instance the fitness of running fast will be somewhere between the values in the top row (*W* and *X*), depending on how often fast organisms are resistant to malaria.

We could make a three dimensional $2 \times 2 \times 2$ table if we wanted to include the trait of being strong or weak. If we wanted to consider continuous, rather than dichotomous traits, the table could be easily expanded. Matthen and Ariew (2002, 67) ask “How do these fitness factors add up?” as if this were unanswerable. In fact, empirical observation is how one determines how these factors “add up.” If we sample the members of the population, the estimated fitnesses may tell us that the traits interact—the relationship between these various possible traits may be additive or non-additive.

Notice various *Newtonian* features of this harmonization problem. Suppose that one population of genetically identical organisms is fairly quick but somewhat weak while another population of organisms is slow but fairly strong. It is possible that members of both populations of organisms might have the same *overall* chance of survival. Overall fitness is like a net result of component vectors. A billiard ball can be stationary because two equal and opposite forces are pushing on it, where the strength of these forces may have any number of values provided they are opposite and equal in strength. Different combinations of component forces on the billiard ball can lead to the same net force just as different combi-

5. The numbers in this table happen to be additive—that is, resistance to malaria provides a .6 boost to the organism’s fitness *independent* of whether the organism is quick or slow. This is simply an artifact of my example; nothing about evolutionary theory requires that different components of fitnesses must combine additively.

TABLE 2

SPEED	RESISTANCE TO MALARIA	
	Yes	No
Quick	1	.4
Slow	.8	.2

nations of traits can result in different ways of having the same overall fitness.⁶

3. Combining Selection and Mutation. In addition to arguing that there is a harmonization problem *internal* to the notion of fitness, some critics argue that there is no way to make sense of drift, mutation, migration, recombination, and selection as causes. In this section I explain how selection and mutation can be combined, with an eye to showing how Newtonian this harmonization is. Although these critics don't discuss mutation in any detail, it is worth showing how mutation and selection combine, since it is importantly different from the way that drift is combined with the other forces of population genetics.⁷

Ridley (1996, 45–46) describes a simple case where there is genetic variation at a locus with two alleles, A and a . Suppose further that there is selection against the dominant allele (A), so that the fitnesses of the three genotypes AA , Aa and aa are $(1 - s)$, $(1 - s)$ and 1 , respectively. Imagine further that mutation opposes selection. Let v = probability that a mutates into A .⁸ What will the equilibrium frequency (p) of A be in this case? Here $p = v/s$. Since mutation rates are generally small (typically, $v \approx 10^{-6}$ or 10^{-7}), even a modest selection pressure of $s = .01$ means that the equilibrium frequency of A will be very small. Notice the sense in which there is a *direction* to the force of mutation, and notice how (in this case) it *opposes* selection. We can also talk about cases where the force of mutation *more or less strongly* opposes selection, and cases where mutation operates in the *same direction* as selection. It makes perfect sense

6. Another way of thinking about “components of fitness” involves separating and then combining viability and fertility. For details, see Sober (2000, 60).

7. Walsh, Lewens, and Ariew (2002) write “[w]e restrict our attention to natural selection and drift, in the hope that the lessons learned there will generalize” (453). Their neglect of these other factors produces a misleading picture of evolutionary theory because the interaction of selection with mutation and migration is especially analogous to Newtonian mechanics.

8. The rate that A mutates into a is also relevant; for simplicity, suppose that this value is equal to zero. For a more detailed model that takes into account nonzero mutation rates in both directions, see Roughgarden 1979, chapter 3.

to say that as the force of selection gets stronger relative to an opposing mutation rate (if e.g., $s = .1$ instead of $s = .01$), the equilibrium frequency of A gets smaller. There are (in principle) an infinite number of distinct combinations of mutation rates and selection coefficients that will yield the same equilibrium frequency for A . This is analogous to Newtonian mechanics where any number of different component forces can combine to yield a particular net force. The point is that it makes perfect sense to add selection and mutation in a way analogous to the addition of Newtonian forces.

It is worth noting that similar methods for resolving the strength of biological forces occurs in the units of selection debate—one can figure out what group selection, acting alone, would predict about trait frequencies, what individual selection would predict, and what various compromises there might be between these two forces. For instance, in his famous critique of group selection, G. C. Williams (1966) points out that female biased sex ratios would be evidence for group selection. One can also ask, what would the force of individual selection predict (by itself)? Williams suggested that in this case the sex ratio will be approximately 1:1.⁹ On the other hand, if the force of group selection were strong, one would expect an extreme female bias in sex ratio. One can understand these as two distinct forces, and if the observed sex ratio is, for instance, slightly biased toward females, this is evidence that group selection played some role, and if it is more strongly female biased, this is evidence that the force of group selection was stronger in that population.

4. Process versus Product Notions of Drift. Critics of the traditional view also raise substantive objections to the standard treatment of random genetic drift. Some argue for a kind of blurring of the distinction between selection and evolution, while others argue that a proper (statistical) understanding of drift and natural selection undermines the traditional view's claim that drift can be a cause of evolution. In other cases, they argue that a proper understanding of drift shows that there cannot be a zero-force law in evolutionary biology. I argue below that many of these criticisms can be rebutted if proper attention is paid to the distinction between process and product.

4.1. Is Drift a Cause? Several reasons are given for thinking that drift cannot be a cause. At the end of Section II of their paper, Matthen and

9. Only "approximately" a 1:1 sex ratio because, as R. A. Fisher ([1930] 1999) argued, we should expect an *equal investment* in the two sexes if there is random mating. If the average son costs less than the average daughter, for instance, we should expect an excess in the number of males.

Ariew (2002) claim that “. . . *it is incoherent to think of the component factors contributing to evolutionary change by separate action. As a consequence, the analogy with Newtonian forces collapses*” (59–60).

Notice that even if they are right about the combination of selection and drift, the Newtonian analogy hardly *collapses*. As we have seen, it works just fine in the problem of combining components of fitness and in the problem of combining selection and mutation. Let us now turn their reasons for thinking that selection and drift cannot be combined.

One objection these critics raise is the *causal decomposition problem*. Drift should not be viewed as a component force that can be added to selection for the same reason that it doesn't make sense to say how much of a coin's landing heads is due to chance and how much is due to the propensity of the coin to land heads. Matthen and Ariew (2002, 62) illustrate this with a pair of cases in which one organism has good eyesight and another has bad eyesight. In the first case, the organism with bad eyesight falls off a cliff and dies while the good-sighted one survives. In the second case, lightning kills the organism with the better eyesight, while the one with poorer eyesight survives. Matthen and Ariew maintain that it doesn't make sense to say that selection operates in the first case while drift operates in the second.

Defenders of the traditional view, however, *agree* that one cannot say how much drift and selection each contribute in an individual case.¹⁰ The point is that the effect of drift is only properly understood at the population level. It is a *population level* cause. One sees the differential causal impact of drift *only* by comparing populations of different sizes. Drift plays a larger role in flipping a fair coin 10 times than it does in flipping a coin 10,000 times. Notice, however, the nature of the force: it must be described as an *expectation*. We can say, for instance, that there is a *greater chance* that a fair coin flipped only 10 times will result in more than 60% heads than a fair coin flipped 10,000 times will result in more than 60% heads. Suppose we flip a fair coin 10 times and get exactly 5 heads and then we flip a fair coin 10,000 times and we get 6,281 heads. Does this show that drift played a greater role in the case with the larger number of tosses since it deviated farther from “expectation”? *No*. Since claims about drift are only probabilistic expectations, they are not falsified by the existence of an improbable event. Drift *still* plays a larger role in the smaller number of coin flips. But what this means is given to us by the laws of probability. As Gillespie (1998, 158) notes, if we have n indepen-

10. Sober (1984) makes a similar analogy when discussing drift: “If a fair coin lands heads six out of ten times it is tossed, there is no saying how much its probability of landing heads contributed to this result as compared with the fact that it was tossed only ten times” (117).

dent trials and the probability of success in any one trial is p , the chance of getting i successes is:

$$Pr\{X = i\} = (n!)/(i!(n - i)!)p^i(1 - p)^{n-i}.$$

It is this equation that allows us to define drift as a *process*.

The traditional view is defensible if we attend to an important distinction between drift as a *process* and drift as a *product*.¹¹ It is the process notion of drift that is needed to understand its role as a cause.¹² In a population of a given size, drift as a process of indiscriminate sampling *always* has the same force. It is part of the definition of drift that it is stronger when the population is smaller. However, it is not part of the definition of drift that the effect (product) of drift must be the same in all cases where the process of drift is equally powerful. Because drift is a probabilistic cause, the *same* causal force can have two *different* outcomes. Walsh, Lewens, and Ariew focus on describing drift as accounting for certain kinds of *error*—but “error” in their sense is merely a *product*—it is not a process, so one has to be careful not to define drift as whatever amount of deviation there is from expectation.

I am accusing these authors of (implicitly) committing a kind of *operationalist* fallacy vis-à-vis drift. Operationalism, as the claim that theoretical terms should have testable consequences, may be good *methodological* advice. But operationalism as the view that one should *define* a theoretical term by whatever is evidence for it is a mistake. We don’t want to define temperature as whatever a thermometer indicates—if we did, then the thermometer couldn’t be mistaken about the temperature.

Do these critics blatantly commit such a fallacy? Perhaps not. But it

11. Millstein (2002) argues convincingly that Beatty (1984) fails to consistently distinguish between the process and product notions of drift. Although I am in agreement with much of Millstein’s essay, she mistakenly argues that one can only make sense of drift and selection as processes (as opposed to outcomes or products) if one models natural selection probabilistically. But there is nothing about viewing selection deterministically (*viz.*, that the future frequencies of traits are logically implied by the relevant fitness values and their starting frequencies) that prevents one from distinguishing process and product notions of selection and drift. In her discussion of Brandon and Carson (1996), she accurately pinpoints the confusion in their urn example; however, she fails to realize that there *must* be drift (in the process sense) in their urn example, since it involves a finite number of draws. She asks, “But, regardless of outcome, why should we consider this example to model drift in any way?” (49) My point is that it does (and must) include drift as a process unless one is assuming that the number of draws is infinite.

12. Even if drift is understood as a process, one might object that this does not, by itself, settle the issue of whether it is a cause. In this section I am concerned to rebut objections to the claim that drift could be a cause. For a more positive argument for thinking of drift as a cause, see Section 6.2.

is *implicit* in some of the remarks that they make. At least, they seem to commit the fallacy when they are describing what the traditional view says about drift. For instance, Walsh, Lewens, and Ariew (2002) write:

. . . *drift is statistical error*. A series of births, survivals, deaths, and reproductions manifests drift just if the outcome—measured as changes in trait frequencies—diverges from that predicted by differences in fitness. (459, emphasis theirs)

They seem to be equating drift with “actual sampling error.” In another case they describe a simple urn example in which the sample drawn perfectly matches the overall frequency of different kinds of balls in the urn. They claim that in a similar situation in which “the outcome is precisely that predicted by differences in trait fitnesses; there is selection but no drift” (464).

If we interpret these quotes as referring to drift as a *product*, then they avoid the operationalist fallacy and can be understood as making an innocuous epistemic point. If the population does not deviate from what is expected by fitness differences, we may not have evidence that drift is at work. However, my point is that the problems that Walsh, Lewens, and Ariew raise for the standard view disappear if you think about the *criterion* for drift and not simply about what is *evidence* of drift. On the process notion of drift, it is false to say that “drift is statistical error.” It is consequently wrong for them to conclude that “there is selection but no drift.”¹³

4.2. *Is There a Zero-force Law in Evolutionary Biology?* Matthen and Ariew (2002) raise three objections to the possibility of a zero-force law in evolutionary biology. First, they ask us to imagine a case where there are two populations, each with the same number of individuals that have trait T_1 (the fitter trait) and trait T_2 (which is less fit). In one of the populations, the fitter trait T_1 goes to fixation, but in the other population, the less fit trait T_2 goes to fixation. They point out, correctly, that there is no *extra* factor present in the second population not present in the first that explains why in the second case the less fit trait goes to fixation (after all, by hypothesis, the two populations were exactly the same initially). But they then draw the following conclusion “So one cannot say that ‘if genotype frequencies depart from equilibrium, some force must have been at work’. Because the causes here are probabilistic, change might have the same cause as equilibrium” (61).

13. Notice, even if one accepts a purely statistical interpretation of selection and drift, one ought to take seriously the difference between the criterion for drift and what is evidence for it.

Matthen and Ariew seem to think that the traditional view is committed to saying that drift only operates in the second case (or is stronger in the second population). But this is to confuse process and product. Drift operates with *equal* strength in both cases, and is not the sort of thing that is in the business of explaining why in one case the fitter trait swept through the population while in the other case the less fit trait reached equilibrium—at least not when the two populations are identical in size.¹⁴ Notice that that on the traditional view, the point is just that departure from Hardy-Weinberg equilibrium is *sufficient* (rather than necessary) for thinking that some force must have been at work. Just because the first population does *not* deviate from Hardy-Weinberg equilibrium does not mean that no other force *is* at work.

The situation is the same in Newtonian physics. Deviation from the zero-force law provides *sufficient*, but not necessary conditions for concluding that some force must have been at work. Two forces can act in equal and opposite strength and so the system remains at equilibrium—change is not *necessary* for some force to be at work. This is true in both Newtonian physics and evolutionary theory.

These critics also argue against the idea of a zero-force law in evolutionary biology because they say that “there is no common currency in which to compare the contributions of different evolutionary theories” (Matthen and Ariew 2002, 68). This is true if one means no common currency for their *causal* powers, but false if one is talking about their *effects*. On the traditional view, the impact of all the so-called forces (such as selection, migration, mutation, and drift) can be cashed out in the language of genotype (or gene) frequencies. Still, here is an important disanalogy with Newtonian physics. In the evolutionary case, the *effects*, but not necessarily the *causes*, can be understood in a common currency. This is enough, however, to provide us with a zero-force law.

Finally, Matthen and Ariew (2002) argue that “[it] makes no sense to think of natural selection acting ‘on its own,’ in the way that gravitation acts alone on an uncharged particle” (68). Throughout their paper, they make a point of saying that natural selection must operate in a “substrate.” Whoever thought otherwise? Not defenders of the traditional zero-force approach. Here is Sober, once again, from the *Nature of Selection*:

Some aspects of the Mendelian process—e.g., the “linkage” of genes that are located on the same chromosome—are treated as evolutionary forces. *The point is that a substantive Mendelian mechanism is*

14. If one is bothered by thinking of a process that operates with more or less strength, one can simply say that the sampling is equally (in)discriminate in both cases.

assumed to be at work even when all evolutionary forces are said to be absent. (35, emphasis mine)

So the zero-force law in evolutionary biology is *disanalogous* to Newtonian theory in this way. The zero-force law in evolutionary biology is formulated *against the background* of a substantive biological assumption. But no one ever said the analogy was perfect. Still, *given* an assumption about Mendelian inheritance, the zero-force law can be formulated and works remarkably well, as we have seen.¹⁵

This reveals why there is no problem here with the traditional view's expression of the adaptationism debate. Matthen and Ariew (2002) think that the traditional approach tries to ask: "What would happen if selection were to act *by itself*?" (66) and point out correctly that it doesn't make sense to answer this.

But no one who formulates the debate about adaptationism thinks that natural selection operates without *any* substantive biological background assumptions. An adaptationist model is often one in which drift and mutation rates are ignored, the population is assumed to reproduce asexually, and so on. An adaptationist hypothesis asserts that making these idealizations won't affect the predictive accuracy of the model in any substantial way. The disagreement takes place against a common background of agreement about certain sorts of constraints. The issue is what *further* constraints may be required. The point of an optimality model is to represent *local optima*, that is, optima *given* a set of constraints.

5. Natural Selection and Fitness. In addition to the difference between process and product notions of drift, there is another important distinction that is sometimes neglected. This is the distinction between *selection* and *fitness*. An organism's overall fitness is like the actuarial idea of life expectancy. Life expectancy is a kind of summary of a number of different possible factors that affect how many more years someone will survive. Although life expectancy reflects a number of possible causal factors that

15. Matthen and Ariew (2002) also criticize Sober's formulation of the zero-force law in terms of Hardy-Weinberg because it applies only to sexually reproducing organisms. They suggest that we should understand the traditional view by thinking of the steady state in terms of gene, rather than genotype, frequencies remaining the same (58). A couple of points are worth making here. First, Sober (1984, 36) himself makes the same point, and notes that there are really *two* ways to formulate the zero-force law. Second, Matthen and Ariew's suggestion to use no change in *gene* frequencies as the zero-force law is inadequate for sexually reproducing organisms because genotype frequencies can change without change in gene frequencies. For a simple example, see Sober (2000, 4).

might shorten or lengthen one's life span, many (if not most) of these factors will never play an *actual* causal role.

An organism's overall fitness is similar in that it is a reflection of different kinds of *possible* causes that might affect an organism's survival and reproductive success. Thus it could include the chance of encountering a certain disease and the chance of being killed by the disease if encountered. It could include the chance of encountering a certain sort of predator and the chance of being detected and then eaten by the predator, and so on. As with the case of life expectancy, many of these causal possibilities may never come to pass. The organism in question may never encounter a particular disease or predator. It is therefore reasonable to think that the overall fitness does not cause anything; rather, it is merely a reflection of causal possibilities.

Selection, on the other hand, is supposed to be a *causal* notion on the traditional view. Overall fitness might inform one that a certain trait is likely to spread, but it does not tell us *why* the trait is fitter. To get at the causes, we must look at which traits are favored by selection, and which traits are not. This distinction is important since many of the objections to the traditional view depend on showing that overall fitness is causally inert—a claim these critics have in common with views such as Sober's. Still, they do not always seem to be aware of this fact, and it leads them astray. Or so I shall argue.

One reason that these critics are skeptical of the traditional view is that they do not think that selection is a force that can cause individual organisms to live or die. Recall the coin analogy that was discussed previously—one cannot ask of an individual coin flip whether its coming up heads was due to drift or to the bias of the coin. I argued that the traditional view *agrees* with this claim—one cannot say of any given flip how much is due to the bias of the coin and how much is due to drift. If selection is identified with overall fitness, selection is not a force that can kill an individual organism. So what is the dispute about?

While it is true that *overall* fitness—either individual or trait—does not cause anything, much less evolution, the traditional view maintains that *selection* does do causal work, whereas these critics deny this. For example, Walsh, Lewens, and Ariew (2002) appear to miss the importance of this distinction for the traditional view. First, they quote the following passage from Sober (1984):

Selection *for* is the causal concept *par excellence*. Selection for properties causes differences in survival and reproductive success An organism's overall fitness does not cause it to live or die, but the fact that [for example] there is selection against vulnerability to predators may do so. (100)

Notice that this passage from Sober invokes *both* the causal notion of selection *and* the non-causal notion of overall fitness (and that these two notions are separated by the ellipsis in their quote). However, in response to this above quote, they claim that the traditional view “misrepresents the explanatory role of *fitness* in natural selection theory” (460, emphasis mine). Why is that? Well, they go on to point out that it is not *individual fitness*—a summary of an individual’s dispositional properties to survive and reproduce in a given environment—that is necessary and sufficient to explain changes in trait frequencies. Instead, it is *trait fitness*—a statistical summary of a particular trait type—that is necessary and sufficient to explain changes in trait frequencies. The details of why this is the case do not need to concern us here. They are correct about this point. It is, however, irrelevant to the issue at hand. They think that this fact about trait fitness shows that the *forces as causes* condition is mistaken and that we must accept the statistical interpretation of evolutionary theory.

The traditional view can take on board Walsh, Lewens, and Ariew’s point about it being variation in trait fitness, rather than individual fitness as being necessary and sufficient to bring about changes in trait frequencies.¹⁶ Still, on the traditional view, natural selection should not be identified with *either* individual fitness or trait fitness. On the traditional view, *neither* overall individual fitness *nor* overall trait fitness is causally efficacious. Each reflects certain possible causal facts, but each is itself a statistical summary. However, the traditional view maintains that *natural selection* can be a cause—it is selection that causes differential survival and reproduction.

Walsh, Lewens, and Ariew (2002) raise another objection to the traditional view that seems to depend on missing this distinction. They say point out that natural selection theory generally explains by appeal to statistical properties. While it is true that we can explain how a population can be expected to change by citing trait fitnesses (which are statistical properties). If we want to know *why* the trait fitnesses have the values they have, however, we need to appeal to a causal notion of selection.

They consider a related objection (in section 4.3 of their paper). They imagine that someone will say that predation, sunlight, and so on are all various selective pressures that can act on a population causing differential survival. However, their response to this is another place where they seem to commit an operationalist fallacy vis-à-vis drift. They say that

it is a mistake to identify natural selection with the forces causing differential survival and reproduction for the reasons we have already

16. Assuming we “hold fixed” all other possible factors that could influence trait frequencies.

seen. Predation pressure, sunlight, and competition cause changes in trait frequencies. Sometimes, as we saw in Beatty's scenario, *the changes constitute drift*. So if predation, sunlight, and competition (etc.) are identified with the force of selection, the force of selection causes drift. (467, emphasis mine)

Here, when they write "the changes constitute drift," they're not simply claiming that such changes are *evidence* of drift—they are saying that they *constitute* it. This is incorrect—it is an example of failing to distinguish between process and product notions of drift. Once we understand the process notion of drift, we no longer have to say that the same facts *constitute* both selection and drift—though of course selection and drift can each (or both) *be responsible for* the same outcomes.

6. Can Drift Be a Cause? So far, we have seen that many of the objections to the traditional view rest on failing to pay attention to two different distinctions—one between process and product and another between selection and fitness. What reason is there for thinking that drift is a cause?

6.1. Does Drift Have a Direction? Matthen and Ariew (2002) argue that drift cannot be a force or cause because "drift is not the sort of thing that can play the role of a force—it does not have predictable and constant direction" (61). This is false. Drift does have a direction—it serves to eliminate heterozygosity. Over the long run, pure random drift causes a population to become homozygous at a locus because in every generation, some genes will, by chance, fail to reproduce and others will, by chance, leave "extra" copies. Eventually, if there is no interference by other forces (e.g., mutation), the population will become homozygous at that locus. This is easy to see with a simple example: suppose a population starts off with 3 individuals who are all heterozygous (Aa) at a certain locus. When the genes are randomly sampled to produce the next generation, there are, suppose, 4 copies of A and 2 copies of a , instead of the original 3 copies of each. At some later point, chance may lead to 5 A and 1 a , and so on. Given enough time, and no intervening forces, the population will drift to either all A or all a .

Perhaps what they have in mind by saying that drift does not have predictable direction is that with respect to *allelic* frequencies, drift does not have a direction. There is a sense in which this is true, but notice that we *can* speak of drift as tending to reduce allelic variation. It is natural to think of drift as a cause because *drift can make a difference*. Suppose we have a given population with two kinds of organisms, X and Y , and

the only forces at work are selection and drift.¹⁷ If X organisms are fitter than Y organisms, we should expect trait X to increase, and Y organisms to decrease in frequency. But *how confident* should we be that trait X will increase in frequency? That depends on the population size. If the population is very small so drift plays a large role, we should not be very confident of the expectation that X organisms increase in frequency. If, on the other hand, the population is large, then drift is insignificant and we should have more confidence that X will increase. So drift plays an important role in telling us about the *probability* of evolutionary change; it is natural to treat it as a cause because it makes such a difference.

6.2. More Reasons to Think That Drift Is a Cause. Here is another example that illustrates why it is natural to talk about drift as a cause. Suppose there are two sets of populations, each with two kinds of otherwise identical individuals. One kind of individual has a trait (T_1) is fitter than an alternative (T_2). In one population set, there are 20 (isolated) groups each with 6 individuals, whereas in the other set there are 20 (isolated) groups each with 1,000 individuals. Suppose that every group in each population set begins with exactly 50% T_1 and 50% T_2 . Imagine that the population evolves, and, as expected, the number of groups in which T_1 goes to fixation in the second population set is much higher than the number of groups in which T_1 goes to fixation in the first population. We now ask: *Why* did the trait with the higher fitness (T_1) sweep to fixation in all the groups in the second set but not in all the groups in the first set? Answer: drift. The fact that the effective population size of each group in the first population set is much smaller means that drift is a much greater force.

One might object that drift cannot be a force or cause because we cannot specify the degree to which drift is involved in a given case in which both selection and drift are at work. For instance, Matthen and Ariew (2002) object that “[t]he proposition that drift was involved to degree p in this history generally has no truth value” (65). Although they are correct that one cannot say that drift was involved to degree p , we can still talk about drift as a cause and about the relative importance of drift in a population. The way this works is to compare the process in question to a hypothetical one in which the population size is much bigger or smaller. Consider the debate about neutralism. Neutralism is usually formulated as saying that most evolutionary changes at the molecular level are a result of drift; selectionists claim, on the contrary, that most evolutionary changes at the molecular level are driven by natural selection.

17. Imagine that the organisms reproduce asexually and that like always produces like. This example is from Sober (1984, 116).

To say that the evolution at the molecular level is mostly the result of drift amounts to a claim that the selective coefficients of the variants in question are all zero (i.e., no fitness variation exists between the alternatives).¹⁸ To say that the changes are mostly a result of selection means that there is a selective difference between the variants in question. This latter claim does *not* rule out the possibility that drift was involved to some extent, as it must be in any finite population.

The debate therefore turns on how likely mutations with a certain sort of selection coefficient arise. Neutralists argue that most mutations are neutral or "nearly" neutral, and very few are selectively favored. Selectionists argue that exactly neutral mutations are very unusual and that enough favorable mutations occur to account for all (or nearly all) molecular evolution. Whether drift plays a large role depends on whether $Ns \gg 1$, where N is the effective population size and s is the selection coefficient. Even if trait X has a selective advantage over Y and trait Y prevails anyway, this does not mean we get to say that neutralism is true or that drift is *the* explanation in this case. The theories still make different predictions about the *frequency distributions* for rates of mutation and selection coefficients. One of the reasons that the debate has persisted is that it is often difficult to determine in practice whether or not alternative molecular traits are functionally equivalent. It is widely accepted, however, that neutralism is true for at least some kinds of molecular evolution. *Pseudogenes* are bits of DNA for which there are good theoretical reasons for thinking that they have no function. Ridley (1996) notes that "[s]ome pseudogenes, for example, cannot be transcribed, because they lack promoters and introns" (181). Currently, most biologists accept the fact that the best explanation of the rapid evolution of pseudogenes is drift.

This is like a case in which you know, through a careful physical examination, that a coin is perfectly balanced. You then have good theoretical reasons to think that the coin is fair. Biologists have discovered something analogous about pseudogenes. Since you have a strong reason to think the coin is unbiased, this means that there is no "selective" advantage to heads. This fact reveals an important asymmetry that is often overlooked. In real populations, drift *always* exists to some extent as a force—although when the effective population size is large the force is very weak or insignificant compared to other factors.

18. Sometimes neutralism is characterized by saying that the neutral variants have *no* fitness difference ($s = 0$) (Ridley 1996). In other cases, biologists describe neutralism so that s is very close to zero, but the effect of selection is overpowered by that of drift, due to the small population size. For instance, Kimura characterizes the issue of whether drift plays a significant role by whether $Ns \gg 1$.

6.3. *Two Questions about the Causal Role of Drift.* The lesson is that one must keep distinct two questions:

1. *How Much—Individual:* How much do drift and selection each contribute in an individual case where both “forces” are present?
2. *How Much—Population:* In a population of organisms or (traits that are evolving in a population), how does the force of drift compare in strength to the force of selection?

Notice that even though there is no determinate answer to question (1), there can still be cases where we can say that a trait’s evolution is simply a result of drift. This is because there are populations where drift is the *only* factor at work and hence there is no problem in apportioning the degree of casual influence in an individual event to *drift alone*. We can also talk about drift’s comparative strength with selection and other forces at the population level.

So there are two senses in which we can say that a trait is a result of drift. On the one hand, there are the special cases where we think, on general theoretical grounds, that the alternative phenotypes are selectively equivalent and so *any* change must be due to drift. At the same time, in cases where *both* selection and drift are at work, we can talk about which one is stronger in a *population level* sense. Again there is a sense in which defenders of neutralist hypotheses at the molecular level will maintain that most of the population level effects are due to drift as opposed to selection. Here, one can determine how strong the force of drift is if one can determine the effective population size.

7. Why Selection and Drift Are Conceptually Distinct. Critics such as Matthen and Ariew (2002) sometimes argue that selection cannot be a cause because it is just a measure of some effect. They write that “selection is not a cause of growth . . . in this conception; it is the mathematical aggregate of growth taking place at different rates” (74). Just as we earlier distinguished between process and product notions of drift, it is important to distinguish between *selection as a product* and *selection as a process*. Selection as a *product* is the fact that growth of a certain kind occurs; selection as a process, on the other hand, is what biologists refer to when talking about a *cause* of evolution. These critics have failed to show that there is anything wrong with this traditional approach to thinking of selection as a cause.

Matthen and Ariew also argue that “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” (78). They do hedge this claim a bit, when they acknowledge that “it is legitimate to ask, *in a statistical sense*, how much

TABLE 3

POPULATION SIZE	FITNESS VALUES	
	$w_1 = w_2$	$w_1 \neq w_2$
Finite	Drift only	Both selection and drift
Infinite	No drift or selection	Selection only

the causation of *B* is due to competitive advantage” (78). Ultimately, however, they reject this approach because these factors cannot be understood “in an ontologically separate way” (78).

It is unclear, however, what they mean by “ontologically separate” and how it is supposed to be distinct from “statistically separate.” Consider two traits with fitness values w_1 and w_2 as shown in Table 3.

Table 3 tells us the conditions under which selection and drift are relevant factors. Since selection can be relevant without drift and vice versa, selection and drift are *conceptually* distinct. Why not also say that they are “ontologically” distinct? We can, after all, model selection and drift separately. These critics need to provide an account of what “ontological separation” means that explains why drift and selection aren’t ontological separated in the relevant way.¹⁹

Matthen and Ariew do argue for a distinction between what they call “probabilistic” or “stochastic” causation on the one hand and “fundamental” causation on the other. This distinction is central to their skepticism about attributing causes to various factors such as natural selection, drift, mutation, and so on. They argue for this distinction on the basis of what they call “discontinuity” and “irreversibility.”

They state two criteria about what it takes for something to be a “fundamental physical process” that is “strictly law governed” (79). Such a process must be both *continuous* and *time-symmetric*. They then argue that since natural selection fails these two criteria, it cannot be a fundamental physical process. But it is far from clear that these are requirements on fundamental physical processes. My contention is that these are ad hoc requirements. What if one of our fundamental physical theories turns out not to meet these criteria? Quantum phenomena raises doubts

19. Incidentally, in their discussion of individual and population levels, they claim that “any death, from whatever cause, results in evolutionary change since it results in a change of gene frequencies” (2002, 63, emphasis theirs). This is false. Here is a simple counterexample: suppose a population has three individuals, one *AA*, one *Aa* and one *aa*. If the second organism dies, there is no change in gene frequencies. Counterexamples can also be given to the claim that any death results in a change in *genotype* frequencies (as well as the claim that any death results in *either* a change in gene or genotype frequencies). Some deaths don’t change *any* of these frequencies. I leave the formulation of these counterexamples as an exercise to the reader.

about both of these criteria. For example, while it is true that quantum mechanics is usually understood as time symmetric, this is far from clear in the case of quantum field theory, which is plausibly just as (if not more) fundamental. Since, as Sklar (1992, 130) notes, there are certain phenomena in quantum field theory that do seem to require time-asymmetric laws, does this mean that quantum field theory cannot count as a theory that describes “fundamental physical processes”? In addition, quantum mechanics throws doubt on Matthen and Ariew’s other criterion, since there are important senses in which quantum phenomena is discontinuous.

Of course it could turn out that our understanding of these fundamental physical processes is mistaken, but it is odd that Matthen and Ariew would want to commit themselves to such a controversial position. Even if our best current theories turn out to be wrong, the theories that replace them might turn out to describe time-asymmetric or discontinuous processes at the most fundamental level of nature.

8. Conclusion. I have set out to defend the use of the Newtonian analogy by the traditional view. The new wave of criticisms has done nothing to undermine the view that evolutionary theory is analogous to Newtonian mechanics in many ways. In particular, it makes perfect sense to think of selection, mutation, migration, and drift as causes since they are factors that *make a difference*. All of them can make a difference in the frequency of genes and genotypes. Furthermore, these causal factors can often combine in Newtonian ways, with one factor canceling out or augmenting the effect of another. For instance, heritable variation in fitness is not sufficient for evolution to occur, since factors such as drift or mutation can counteract the effects of such variation.

There are of course important disanalogies between evolutionary theory and Newtonian physics. For one thing, the zero-force law in evolutionary theory is formulated against a background biological assumption. Also, there is nothing like the notion of drift in Newtonian physics. Here, drift is a different kind of force, and it does not make sense to ask, in a particular case, how much effect drift had as compared to selection. Some might be reluctant to call it a force because of this. Providing one is clear about the sense in which drift is a different (non-Newtonian) ‘force’, I don’t see the harm. At any rate, they have not undermined the idea that drift is a kind of population level cause. If we keep the distinction between process and product in mind, we can still make sense of drift as a cause—it does not have to be simply delegated to a “statistical” idea, as Walsh, Lewens, and Ariew suggest—nor does it collapse with selection, as Matthen and Ariew sometimes suggest. Critics of traditional view have also missed the important distinction between selection and fitness. The

traditional view *agrees* that an organism's overall fitness (or trait fitness) is causally inert.

Of course, I have not dealt with all the objections that one might raise to the appeal to causes in population biology. Those with reductionist tendencies might argue that the real causal action goes on only at some more fundamental level. Still, we should not lose sight of the enormous pragmatic advantages to using causal talk at higher levels. It is perfectly natural to appeal to causes in scientific modeling in order to discuss which factors make a difference, which are counterfactual supporting, and to distinguish cause from correlation.

Walsh, Lewens, and Ariew (2002) motivate their paper by asking a simple question "is evolutionary theory a statistical theory or a dynamical theory?" (455) where a statistical theory is supposed to be analogous to thermodynamics and a dynamical theory is one like Newtonian physics. They argue that evolutionary theory is a statistical, rather than dynamical theory. My claim is that this is a *false dichotomy*—I hope you can see that evolutionary theory has elements of both statistical and dynamical theories. In this sense it is neither exactly like Newtonian physics nor exactly like statistical mechanics. Its unique status in this regard is one of the reasons why the theory continues to engage our philosophical interest.

REFERENCES

- Beatty, John (1984), "Chance and Natural Selection", *Philosophy of Science* 51: 183–211.
- 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.
- Endler, John (1986), *Natural Selection in the Wild*. Princeton, NJ: Princeton University Press.
- Fisher, Ronald A. ([1930] 1999), *The Genetical Theory of Natural Selection*. New York: Oxford University Press.
- Gillespie, John (1998), *Population Genetics: A Concise Guide*. Baltimore: Johns Hopkins University Press.
- Matthen, Mohan, and André Ariew (2002), "Two Ways of Thinking about Fitness and Natural Selection", *Journal of Philosophy* 119(2): 55–83.
- Millstein, Roberta (2002), "Are Random Drift and Natural Selection Conceptually Distinct?", *Biology and Philosophy* 17: 33–53.
- Ridley, Mark (1996), *Evolution*, 2d ed. Cambridge: Blackwell Science.
- Rosenberg, Alex, and Frederic Bouchard (2004), "Fitness, Probability, and the Principles of Natural Selection", *Biology and Philosophy*, forthcoming.
- Roughgarden, Jonathan (1979), *Theory of Population Genetics and Evolutionary Ecology*. New York: Macmillan.
- Sklar, Lawrence (1992), *Philosophy of Physics*. Boulder, CO: Westview Press.
- Sober, Elliott (1984), *The Nature of Selection*. Cambridge, MA: MIT Press.
- (2000), *Philosophy of Biology*. Boulder, CO: Westview Press.
- Sterelny, Kim, and Philip Kitcher (1988), "The Return of the Gene", *Journal of Philosophy* 85 (7): 339–61.

- Walsh, Denis (2000), "Chasing Shadows: Natural Selection and Adaptation", *Studies in History and Philosophy of Science* 31C (1): 135–153.
- (2004), "Bookkeeping or Metaphysics? The Units of Selection Debate", *Synthese*, forthcoming.
- Walsh, Denis, Tim Lewens, and André Ariew (2002), "The Trials of Life: Natural Selection and Random Drift", *Philosophy of Science* 69: 452–473.
- Williams, George C. (1966), *Adaptation and Natural Selection*. Princeton, NJ: University Press.