

2002

The Trials of Life: Natural Selection and Random Drift

Denis M. Walsh

Tim Lewens

André Ariew

University of Rhode Island

Follow this and additional works at: https://digitalcommons.uri.edu/phl_facpubs

Terms of Use

All rights reserved under copyright.

Citation/Publisher Attribution

Denis Walsh, Tim Lewens, and André Ariew, "The Trials of Life: Natural Selection and Random Drift," *Philosophy of Science* 69, no. 3 (September 2002): 429-446.

Available at: <https://doi.org/10.1086/342454>

This Article is brought to you for free and open access by the Philosophy at DigitalCommons@URI. It has been accepted for inclusion in Philosophy Faculty Publications by an authorized administrator of DigitalCommons@URI. For more information, please contact digitalcommons@etal.uri.edu.

The Trials of Life: Natural Selection and Random Drift*

Denis M. Walsh^{†‡}
University of Edinburgh

Tim Lewens
University of Cambridge

André Ariew
University of Rhode Island

We distinguish dynamical and statistical interpretations of evolutionary theory. We argue that only the statistical interpretation preserves the presumed relation between natural selection and drift. On these grounds we claim that the dynamical conception of evolutionary theory as a theory of forces is mistaken. Selection and drift are not forces. Nor do selection and drift explanations appeal to the (sub-population-level) causes of population level change. Instead they explain by appeal to the statistical structure of populations. We briefly discuss the implications of the statistical interpretation of selection for various debates within the philosophy of biology—the ‘explananda of selection’ debate and the ‘units of selection’ debate.

*Received August 2001; revised March 2002.

†Send requests for reprints to the authors. Walsh: Department of Philosophy, University of Edinburgh, David Hume Tower, George Square, Edinburgh, Scotland EH26 9NJ; Lewens: Department of History and Philosophy of Science, University of Cambridge, Free School Lane, Cambridge, UK CB2 3RH; Ariew: Department of Philosophy, 170 Chaffee Building, University of Rhode Island, Kingston, RI USA, 02881.

‡We wish to thank the following for help and discussion. Alexander Bird, Anjan Chakravartty, Kent Holsinger, Deborah Kohn, Peter Lipton, Joel Pust, Alex Rosenberg, various members of the Philosophy Workshop, HPS, University of Cambridge, attendees at the ISHPSSB conference in Quinnipiac. We would like to acknowledge the invaluable comments offered by Elliott Sober. We would also like to thank the Seven Stars Bakery in Providence, Rhode Island for all the coffee and for not minding all the noise. We reserve especial thanks to R. C. Lewontin, for inspiration, hospitality, and support. Some of the ideas in this paper converge on those expressed in Matthen and Ariew (2002). These were developed largely independently.

Philosophy of Science, 69 (September 2002) pp. 452–473. 0031-8248/2002/6903-0000\$10.00
Copyright 2002 by the Philosophy of Science Association. All rights reserved.

1. Introduction. Evolutionary theory is commonly described as a theory of forces. Evolutionary forces—natural selection, random drift, mutation, recombination, and migration—are said to act on populations to bring about changes of trait frequencies. When a population changes, it is supposed that some combination of evolutionary forces may be acting concurrently. Evolutionary theory allows us to decompose the disparate, independent forces acting on a population and to distinguish their respective causal contributions. Hence one can ask how much of some observed change owes itself to each force. This *dynamical* conception of evolutionary theory is neatly summarized by Sober:

In evolutionary theory, the forces of mutation, migration, selection, and drift constitute causes that propel a population through a sequence of gene frequencies. To identify the causes of the current state . . . requires describing which evolutionary forces impinged. (1984, 141)

At the same time, biologists frequently talk about the various elements of evolutionary theory in strictly statistical terms. One of the most commonly encountered analogies for the process of evolution is that of a blindfolded selector drawing balls from an urn. The metaphor is thought to illuminate the irreducibly probabilistic nature of evolutionary processes. Other statistical metaphors abound too: selection is spoken of as “discriminate sampling” (Beatty 1984). Drift is spoken of variously as “indiscriminate sampling,” or “sampling error.” Natural selection, by this way of thinking, is a mere consequence of a statistical property of a population—its variation in fitness (Endler 1992). We shall call this interpretation, in contrast, the *statistical* conception.

The dynamical and statistical descriptions of evolutionary theory are often used interchangeably, or admixed indiscriminately. Perhaps the thought is that these are simply alternative, yet equivalent, ways of talking about the same theory. But, it seems to us that there are important differences between these conceptions. If evolutionary theory is a theory of forces it isn't a theory about the statistical structure of populations (and vice versa). Our task in this paper is to articulate and contrast the statistical and dynamical conceptions of evolutionary theory, and to make a case for the statistical conception. We restrict our attention to natural selection and drift, in the hope that the lessons learned there will generalize. Selection and drift are not forces acting on populations; they are statistical properties of an assemblage of “trial” events: births, deaths and reproduction. The only genuine forces going on in evolution are those taking place at the level of individuals (or lower) and none of these (and no aggregate of these) can be identified with either selection or drift.

2. Feathers and Coins. In order to fix the distinction between statistical and dynamical theories more vividly, compare two experimental set-ups. In the first a feather is dropped from a height of 1m. In the second, ten coins are drawn at random from an array of 1000 coins: 500 with heads up, 500 tails up. In the first experiment the feather falls to the ground. In the second the sequence of trials (draws) eventuates in a distribution of heads and tails. In each of the cases we might make a prediction concerning the outcome. We might predict that the feather will fall somewhere not too far from a point where the surface of the earth intersects a line drawn from the feather to the center of the earth (allowing for air currents and so forth). In the coin case, we predict a distribution of heads and tails: perhaps 5 heads and 5 tails. In each case we acknowledge a certain likelihood of the outcome diverging from that predicted. In each case, then, there is an expected outcome and also an expectation of some divergence from that outcome, or *error*. There is a significant difference, though, between the respective notions of expectation and error in each of these cases.

The expectation for the trajectory of the feather is generated by summing those forces known to be acting on the feather. The feather is affected not only by the force of gravity but also by attractive forces from other bodies, electromagnetic forces, forces imparted by random movements of the air molecules, etc. Those forces whose effects can be predicted get factored into the expected outcome. Those that can't are ignored, or at least assumed to cancel one another out. Things are quite different with the random drawing of coins. The expected outcome is *not* generated by attending to the forces acting on the coins, but by taking into account the structure of the population being sampled. Every set-up in which coins are sampled at random faces a probability distribution around each of the possible outcomes (sequence of heads and tails). The probability distribution is given by the laws of probability. In this case, the likelihood of r heads occurring in N trials is given by:

$$P(r) = (N! / r!(N-r)!) \pi^r (1 - \pi)^{N-r}$$

where π is the chance of a single random draw coming up heads. The distribution entails that the most likely outcome is 5 heads and 5 tails. So, whereas the expected outcome in the first experiment is generated by attending the forces acting on the feather, in the second it is generated by attending to the statistical structure of the population of coins.

The concept of "error" differs between the two set-ups as well. Error in the first experimental set-up comes from the failure to take properly into account all the forces acting on the feather. There are two important features of this *dynamical* conception of error. The first is that the net

forces can be decomposed into *component forces* and some of these identified as the sources of “error.” The second is that to the extent that our predictions and explanations contain an error term, this is a reflection of our *ignorance* of the forces involved. Once we know about the forces causing the error, these can either be factored into the expected outcome, or eliminated from the experimental set-up.

Error in the coin case is different; it is a built-in feature of the laws of probability. The expected outcome and the error distribution around it are consequences of the fact that this experimental set-up consists of an assemblage of independent trial events. There is no force or single process that tends to direct the final distribution of trials toward 5:5 while other forces tend to deflect the outcome away from that result. In the coin drawing case, we cannot distinguish between the *causes* of the expected outcome and the *causes* of the error. This *statistical* error is ineliminable and not a reflection of our ignorance of the forces at work.

These two experimental outcomes—feathers and coins—require different sorts of explanation. The explanation of the trajectory of a falling feather requires a dynamical theory, i.e., a theory of forces. The explanation of the outcome of a series of trials requires a statistical theory, i.e., a theory that deals with the statistical structure of the population in question and the probabilistic nature of sampling. One may well be able to explain, in dynamical terms, why *each* draw of a coin finished the way it did. By extension of this procedure one could explain dynamically the outcome of the series of draws, by explaining each severally. But, the availability of such a dynamical explanation does not remove the need for a distinct, statistical explanation of the same outcome. The dynamical account does not explain why, given the distribution of heads and tails in the population being sampled, the expected outcome is to be expected—(why *is* 5 heads and 5 tails the best bet?). Nor does the dynamical view account for the distribution of errors (why is 6 heads and 4 tails more likely than 9 heads and one tail?) Nor does it explain the way the probability distribution in this series of ten trials is systematically related to the distribution in trials of 20, 40, or 100. (Why is 9 heads and 1 tail more likely than 99 heads and 1 tail?).

In general, statistical and dynamical theories differ significantly in their structure, their explanatory apparatus and in the explananda they are applied to. There is little reason to believe that dynamical and statistical interpretations could be equivalent descriptions of the same theory. This raises the question: “is evolution a statistical theory or a dynamical one?” Is selection a theory of forces, such as might be invoked in an explanation of the feather case, or is it a statistical theory concerning a sequence of trial events (births, deaths and reproduction)? We in-

tend to argue—against the prevailing orthodoxy—for the statistical interpretation.

We proceed in the following way. We first discuss the interpretation of random drift. It is easier to make the case for the statistical interpretation here. We then discuss natural selection. Despite the prevalence of dynamical talk about selection, the statistical interpretation makes more sense on its own terms. In addition, once it is conceded that drift is to be interpreted statistically, only the statistical conception of natural selection can accommodate the presumed relation between selection and drift. We end with some general considerations concerning the statistical interpretation of evolutionary theory.

3. Random Drift. The term drift is routinely applied to a disparate array of phenomena. These include: (i) *Sewall Wright Effect*: Large populations are often subdivided into smaller sub-populations. As a consequence, the changes in trait frequencies within a population, when summed over the population as a whole are more likely to diverge from the predicted outcome than if the population is not subdivided. (ii) *Hagedoorn Effect*: Mendel's law of segregation predicts equal numbers of each of a parent's alleles in the gamete pool. But, given that each individual produces only a small number of offspring from a huge number of gametes, there is always a considerable chance that the gametes which make it into the subsequent generation will misrepresent the structure of the population of gametes as a whole. (iii) *Indiscriminate Sampling*: Some processes remove individuals from populations indiscriminately. For example, lightning, when it does strike, does not usually discriminate between fit and less fit individuals (or traits), so fitness differences cannot predict or explain the consequences to a population of lightning strikes. (iv) *The Finiteness of Natural Populations*: The Hardy-Weinberg Law says that in infinite populations (of diploid organisms) there is no change in gene frequencies when there is no variation in gene fitnesses. But natural populations are finite in size; often they are small. In finite populations there will always be some non-negligible chance that trait frequencies will diverge from expectation.

What unites these disparate phenomena under the rubric "drift"? Clearly the notion of drift imports some notion of chance. This is more often than not contrasted with the apparent determinism of selection. Sober (1984) and Dobzhansky et al. (1977) offer vivid, and not unrepresentative, examples:

In evolutionary theory, mutation and selection are treated as *deterministic* forces of evolution. *Chance*, as the term is used, is not a property of the selection and mutation processes but stands in contradis-

inction to them. *Random genetic drift*, on the other hand, is the source of the stochastic element in evolution (Sober 1984, 110)¹

. . . the directional force of evolution, *natural selection*, acts on the basis of conditions existing at the broad level of the environment as it affects populations. . . . With respect to population differentiation, the effects of *chance* [random drift] have considerable influence when populations are small and there is little or no counteracting pressure of selection. (Dobzhansky et al. 1977, 6–7) . . .

The most interesting stochastic process is random genetic drift due to accidents of sampling, especially in small populations. (Dobzhansky et al. 1977, 157)

It isn't clear what sense of 'chance' is being invoked here. It could be that drift comprises a class of stochastic or unpredictable forces, like those that might act upon the feather, or that drift is just statistical error of the sort we encountered in the coin sampling set-up. We take these in turn.

3.1. Dynamical Drift. If drift and selection are distinct forces (or causal processes) we should be able to distinguish the effects of one from the other just as we can distinguish the effects of gravity from the effects of other forces in the trajectory of a feather. Rosenberg (1994) has recently argued that there is no discernible causal process deserving of the label "drift." His attack on drift has two prongs. The first argues that drift cannot be a stochastic process. Organisms die, survive and reproduce; these processes are the subject matter of evolutionary theory and they are, one and all, wholly deterministic. So if drift is a process (force) it must be a deterministic one. The second line of attack claims that the events that are labelled "drift" events, like lightning strikes, etc. are no different in kind from selection events. If all the facts were known, any process that causes a change in trait frequencies would be counted as a selection process. Consequently, there is no such process or force as "random drift."

In support of this nihilism about drift, Rosenberg offers a parable. We are asked to suppose that a population of giraffes comprises short- and long-necked forms. Because the long-necked forms are better able to reach acacia thorns, population biologists assign them higher fitnesses and predict that the frequency of the long-necked forms will increase. Contrary to biologists' predictions the long-necked forms decrease in prevalence. Lacking an explanation for this change biologists attribute it to drift. Unbeknownst to the biologists, however, poachers are culling the tallest giraffes. What, to the investigator, looked like an unexplainable random

1. Sober (1993) makes the point that the determinism of selection doesn't entail that it is non-chancy.

variation from the predicted outcome of selection turns out to be a reflection of nothing more than the investigator's ignorance of the relevant selective forces. There is actually in this population unobserved selection *against* long necks. So, Rosenberg concludes that once we know all the facts we no longer need the category "drift." Invoking the probabilistic notion of drift merely reflects our *ignorance* of the relevant causes of evolutionary change.² The position is a venerable one, originating with A. J. Cain:

This is the real basis for every postulate of . . . genetical drift. The investigator finds that he, personally, cannot see any [evidence of selection], and concludes that, therefore, there is none. (1951)

Rosenberg's argument is effective against the view that drift is a distinct sort of force. However, his contention that drift explanations merely reflect our ignorance of the causes of evolutionary change does not follow. Rosenberg's interpretation of drift matches our account of "dynamical error" above. As in the case of the feather, "dynamical error" really is a reflection of ignorance concerning the causes of a particular outcome. But as we saw with the example of coins, when we are dealing with an ensemble of events whose collective outcome is predicted or explained on the basis of some statistical property of the population, there is still an explanatory role for the irreducibly probabilistic concept of *statistical* error. Statistical error is not a reflection of ignorance of causes, but a consequence of sampling. It is not found in the cause of an individual draw of a coin, but in a sequence of trials in which coins are drawn. Analogously, even though we do not find instances of drift by looking at individual events of birth, death, and reproduction, it does not follow that drift is a mere reflection of our ignorance of the causes of evolutionary change. For all Rosenberg's argument establishes, drift might be statistical error.³

3.2. *Statistical Drift.* The statistical interpretation of drift is widespread; drift is often described as sampling error (Beatty 1984). Clearly some drift phenomena are best accounted for in this way, for example the Sewall Wright effect and the Hagedoorn effect. Here the frequency of traits in a small subset of a population stands a reasonable chance of diverging from

2. One of the obvious flaws in Rosenberg's scenario is that there really is a selective explanation to be had. But this isn't obviously so for all cases of drift. The Wright effect, the Hagedoorn effect (etc.) don't appear to be susceptible to this kind of selective explanation. Millstein (1996) effectively makes the point that only some of the presumed instances of drift can be argued away in this manner.

3. We thank Alexander Rosenberg for corroborating this interpretation of his (former) position. It is worth noting that Rosenberg's revised position (2001) is much more congenial to our "statistical interpretation" of both selection and drift.

the frequencies in the population as a whole, simply as a function of sampling. Similarly, a lightning strike might take a very small random sample of the population, one that misrepresents the frequencies of traits in the population as a whole. This really is “sampling error.” But other putative forms of drift don’t have this feature. When drift occurs as a consequence of small overall population size there is no larger population literally being sampled. However, in these cases what happens is that the distribution of fitnesses in the population yields a prediction concerning the way in which a population will change. Drift is manifested as a difference from the outcome predicted by the fitnesses in the population. The law of large numbers tells us that the likelihood of significant divergence from these predictions is an inverse function of the size of the population. The small size of a population increases the chances of error of the sort we encountered in the coin example (Crow and Kimura 1970). This is just a fact about the statistical structure of populations.⁴ All these presumed cases of drift, then, share a common feature. Drift, in any of its forms, is a statistical property of an ensemble of trials or events: *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.

It is hardly surprising then that, as Rosenberg points out, we do not find instances of drift when we look at individual events of births and deaths. In just the same way we do not find “sampling error” when we look at an individual draw of a coin. A *sequence* of such draws may exhibit error—divergence from the expected outcome—but it is simply a category mistake to suppose that the error is to be found in one or other of the individual draws. By analogy, it is a category mistake to suppose that drift is to be found in an individual instance of a birth or death. But if drift *were* a force that causes individual births, deaths, and reproductions we should expect to find its effects in such singleton events.

This raises a question: if random drift is best interpreted statistically, how should selection be construed—dynamically or statistically? The answer to this question must be one that does justice to the explanatory role of natural selection and captures the presumed relation between drift and selection.

4. Natural Selection. There are dynamic and statistical conceptions of selection, just as there are of drift. The dynamical interpretation is pretty

4. The lightning strike case can also be understood in this second way. Suppose two types of organism in a population have equal fitnesses. Events like lightning strikes are more likely to affect the types differently when the population is small than when the population is large.

pervasive: selection is almost universally spoken of a deterministic force which causes individual births, deaths, and reproductions (Waters 1991). Elliott Sober offers the clearest articulation of the dynamical conception.

Selection *for* is the causal concept *par excellence*. Selection for properties causes differences in survival and reproductive successes. . . . 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. (Sober 1984, 100; emphasis in original)

Compelling as this dynamical interpretation of selection is, there are reasons to resist it. We offer three. The first is that it misrepresents the explanatory role of fitness in natural selection theory; only the statistical conception gives an adequate account. The second is that the dynamical theory of selection fails to make room for the relation of selection to drift. The third is that though there are causal processes, even forces, which cause changes in trait frequencies these cannot be *identified* with natural selection because these forces also cause drift. We take these in turn.









4.1. Selection and Fitness. The central explanatory concept in natural selection theory is *fitness*. In fact there are two distinct concepts of fitness in play in evolutionary theory and it is important to distinguish their respective roles. *Individual fitness* is a causal/dispositional property of an individual organism, its propensity to survive and reproduce in a given environment (Mills and Beatty 1979). *Trait fitness* is a statistical property of a *trait type*. It is a function of the *mean* and *variance* of the fitnesses of individuals within a population (Sober 1993; Maynard Smith 1986; Sober 2000). It is easy to overlook the differences between these two varieties of fitness, yet the differences are crucial.

The objective of natural selection theory is to explain and predict changes in the relative frequencies of heritable traits within a population. The change that selection explains is a consequence of *variation in fitness* (Lewontin 1970, 1974; Brandon 1990). But which *kind* of fitness? Certainly variation in individual fitness is required. If individuals do not vary in their individual fitnesses, one expects no differential survival and reproduction, and hence no changes in the relative frequencies of heritable traits. But mere differences in individual fitnesses are *not sufficient* to bring about changes in trait frequencies.⁵ For that, variation in *trait* fitnesses is required. The models (Figure 1, below) demonstrate this. Table A in each model shows the fitnesses and starting and finishing frequencies of the





5. Sober and Lewontin (1982) suggest that variation in individual fitness is both necessary and sufficient. Endler's (1986) characterization of selection also entails that variation in *individual* fitness is both necessary and sufficient.

Model 1: Variation in individual fitness/no variation in trait fitnesses

A. Individual Genotype Fitnesses and Frequencies









Individual Genotype	Individual Fitness	Freq. F ₀	Freq. F ₁
1 	1.6	10	16
2 	0.4	10	4
3 	1.2	10	12
4 	0.8	10	8
5 	0.8	10	8
6 	1.2	10	12
7 	0.4	10	4
8 	1.6	10	16
Total		80	80

B: Trait Frequencies





Trait Fitness	Traits					
	Big	Little				
Trait Fitness	1.0	1.0	1.0	1.0	1.0	1.0
Freq. F ₀	40	40	40	40	40	40
Freq. F ₁	40	40	40	40	40	40

Model 2: Variation in individual fitness/variation in trait fitnesses

A. Individual Genotype Fitnesses and Frequencies

Individual Genotype	Individual Fitness	Freq. F ₀ Gen	Freq. F ₁ Gen
1 	1.6	10	16
2 	1.2	10	12
3 	0.8	10	8
4 	1.6	10	16
5 	0.4	10	4
6 	0.8	10	8
7 	0.4	10	4
8 	1.2	10	12
Total		80	80

B: Trait Frequencies

Trait Fitness	Traits					
	Big	Little				
Trait Fitness	0.8	1.2	1.1	0.9	1.3	0.7
Freq. F ₀	40	40	40	40	40	40
Freq. F ₁	32	48	44	36	52	28

Haploid population with 3 traits, 2 alleles each: 1) Big: Little, 2) Triangle: Square, 3) Black: Grey

Figure 1. The models depict the individual fitnesses, trait fitnesses and changes in trait frequencies for a population of haploid individuals with 3 traits, 2 alleles each, 8 total genotypes. *Model 1* There is variation in individual fitnesses (table A) but no variation in trait fitnesses (table B). There is no change in trait frequencies (table B). *Model 2*. The population exhibits the same amount of variation in individual fitness (Table A), with variation in trait fitnesses (Table B). Here there is change in trait frequencies. The models demonstrate i) that heritable variation in *individual* fitnesses is not sufficient for changes in trait frequencies under selection, whereas heritable variation in *trait* fitnesses is and (ii) that the degree and direction of change in trait frequencies is predicted and explained by differences in *trait* fitnesses. Trait fitnesses are calculated as the average of individual fitnesses. The models assume no drift.

individual genotypes. Table B shows the *trait* fitnesses (calculated, for each trait, as the average individual fitness of those with that trait). In Model 1, there is variation in individual fitnesses (Table A), yet there are no changes in *trait* frequencies (Table B). The reason is that there is no variation in trait fitnesses (Table B). In Model 2 the degree of variation in individual fitnesses is held constant (Table A) and the trait fitnesses made to vary (Table B). Here there are changes in trait frequency (Table B). So while variation in individual fitnesses is necessary for changes in relative trait frequencies it is not sufficient; variation in *trait* fitnesses is, however, both necessary and sufficient.

There is a further question of which of these parameters predicts and explains changes in relative trait frequencies. Here again, the models demonstrate the importance of *trait* fitnesses. If one wants to explain (or predict) the relative success of, say, BIG as compared to Little in either model, one abstracts away from individual fitnesses by averaging across the fitnesses of those individuals with the BIG trait and those individuals with the Little trait. By this method one obtains the *trait* fitnesses for each trait. In Model 1 the equivalence of the trait fitnesses predicts and explains the lack of change in the relative frequencies of BIG and Little in the F_1 generation. In Model 2 the variation in trait fitnesses predicts (and explains) the preponderance of Little over BIG in the F_1 generation. This is not to deny that differences in trait fitnesses are a consequence of differences in individual fitnesses. Nor is it to deny that differences in individual fitnesses *cause* the changes in trait frequencies that natural selection theory explains. It is simply to assert that if you want to explain how a population will change in its trait frequencies, you have to know its *trait* fitnesses. The point to be made here is just that natural selection explanations appeal to a set of statistical properties of populations, viz. the mean (and variance) of fitnesses between trait types. Explanations of this sort do not advert to forces.

Here again, an analogy with sequences of coins suggests itself. Consider a process we might call a “sorting” process in which the frequency of the heads-to-tails changes in a series of trials. U.S. pennies have a systematic bias which can be demonstrated by standing them up on their edges on a flat surface. If you thump the surface with your fist the pennies will tumble with the tails side up, on average, 60% of the time. Suppose we begin with 10 U.S. pennies lined up on a flat surface with a distribution of 5 heads and 5 tails facing the observer. Then we bang the table. After the pennies land we put them on their edges with the upper surfaces facing the observer. We expect that the frequency of tails to heads facing the observer will change from 5:5 to 6:4. The sorting process—the change in relative frequencies of heads and tails—is explained by appeal to the systematic bias. It is important to note here that the systematic bias is not a property

of any individual coin. It is an average. It is a fair bet that each coin has a different bias. The difference between the *average* propensity of coins to fall tails and the *average* propensity of coins to fall heads explains and predicts the change in relative frequencies of heads and tails. Sorting of this kind offers a close analogy to natural selection. The bias of each individual coin is the analogue of individual fitness and the bias of heads versus tails within the population of coins is the analogue of trait fitness. In each case a change in the structure of a population is explained and predicted by appeal to some statistical property, an *average* of individual propensities.

The analogy between sorting and selection is instructive in other ways. It demonstrates the error of the dynamical interpretation. Sorting is not a force. Sorting is an ensemble-level phenomenon—a property of a *sequence of trials*. Consequently, it is inappropriate to point to an individual trial and ask if it is a consequence of sorting. If sorting is an appropriate analogue for selection then selection is not a property of individual events but a property of a sequence of trials of births and deaths. Just as sorting is not a force that causes a coin to fall head or tail, selection isn't a force that causes an individual to live or to die. One cannot point to a single event of, say, a birth or death and ask whether it is a consequence of the process of natural selection. The passage from Sober quoted above (p. 3) suggests that on the dynamical conception of selection one ought to be able to do so.

In both sorting and selection the explanations are indifferent to the nature of the individual-level forces involved. So long as the forces are such as to preserve the statistical structure of the population it doesn't really matter what they are for the purposes of statistical explanation. It should not be alarming that there is a theory that explains by appeal to population-level statistical properties as opposed to individual causal properties. The kinetic theory of gases is an obvious example of such a theory. The properties described by the phenomenological gas laws—temperature, pressure, volume—are, like trait frequencies, population-level properties. The nature and relation of these ensemble-level properties is explained in terms of a statistical property, *mean kinetic energy*. The *forces* involved may simply be those given by Newtonian mechanics, but it is not these forces that the kinetic theory appeals to. It is the *statistical* property of mean kinetic energy that explains the changes in the ensemble-level properties of the gas.⁶

6. Maxwell is credited with being the first to appreciate the distinctive explanatory role of statistical properties of ensembles: "In studying the relations of quantities of this kind, we meet with a new kind of regularity, the regularity of averages . . ." (1873, 373–374 quoted in Gigerenzer et al. 1989, 62). There remains, of course, the problem of demonstrating how these statistical properties of ensembles can be described in terms of the non-statistical properties of individuals (see Sklar 1999).

In short, natural selection occurs only when the relative frequency of trait types changes in a population as a consequence of differences in the *average* fitnesses of individuals in different trait-classes.⁷ This is what we call the statistical interpretation of natural selection.

4.2. *Selection and Drift.* One motivation for the dynamical model of evolutionary theory is that it makes sense of the presumed relation of drift and selection. Drift and selection can occur together in a population or alone. We can, for example, ask what sorts of changes a population might go through if it manifested drift without selection or selection without drift. We can also sensibly ask how much of the change in a population is attributable to selection and how much to drift. It is helpful in doing this to appeal to the imagery of decomposing Newtonian forces in factoring out the distinctive effects of drift and selection (see Sober 1984, chapters 1 and 4). But this force-talk is by no means obligatory. The statistical interpretation also allows one to distinguish selection and drift. For example, in Model 2 the trait fitnesses predict 52 Black and 28 Grey in the F_1 generation. If the result is 54 Black and 26 Grey there is selection for Black, because the relative frequencies of the traits change in a way predicted quite accurately by differences in trait fitnesses. But because the outcome diverges somewhat from expectation, there is drift too. Were the result to be 56 Black and 24 Grey, there would be comparatively more drift. It is even possible on the statistical interpretation to conceive of either selection or drift occurring in the absence of the other. For example in Model 2 there are 52 Black and 28 Grey in the F_1 generation. The outcome is precisely that predicted by differences in trait fitnesses; there is selection but no drift. If in Model 1 there are 41 Black and 39 Grey, in the F_1 generation, then there is a change in trait frequencies where trait fitnesses predict no change; there is drift, but no selection. The point is simply that the conceptual distinction between drift and selection can be drawn without requiring the metaphysical distinction between *forces* that cause drift and those that cause selection.

One reason to suppose that we cannot decompose drift and selection into component forces is that the same processes that cause natural selection cause drift. Again consider the sorting analogy. The average bias of U.S. pennies predicts a 6:4 ratio of tails to heads in a series of trials. If 1000 coins produce an outcome of 641 tails and 359 heads there is sorting, and there is error too. The causal factors which determine the outcome of the series of trials include the *individual* biases of the coins (not the average bias) and the impetus from the pounding on the table. These factors together cause the *distribution* of heads and tails by causing each coin, sev-

7. "Trait class" here refers to a class of individuals possessing the same heritable trait.

erally, to fall either heads or tails. We cannot point to one or other of these individual-level causal processes and say *this* is a cause of the error and *this* one is a cause of sorting.⁸ The same forces cause both. Similarly selection and drift, on the statistical construal, are properties of the outcome of a series of trials (births, deaths, and reproductions). The forces that cause the births, deaths, and reproductions also cause the outcome. If the outcome constitutes both selection and drift, then taken together the same forces cause both selection and drift, without being decomposable into those that cause the selection and those that cause the drift.⁹

A further reason to adopt the statistical conception of selection is that it preserves the presumed *explanatory* relation between selection and drift. As we mentioned, drift explains what natural selection cannot in the sense that drift accounts for those differences between the actual outcome of a series of births, deaths, and reproductions and the outcome predicted by differences in trait fitness. The dynamic interpretation of selection distorts this explanatory relation.

John Beatty's (1984) discussion of selection and drift illustrates the point. Beatty offers the following scenario. A population comprises dark and light moths in an environment of dark and light trees. There are more dark trees offering camouflage for more dark moths. So dark moths have, on average, higher fitnesses than light moths. However, it just so happens that over the course of a few generations dark moths, by chance, land more commonly on light trees than we would predict. The upshot is that the light trait increases with respect to the dark trait contrary to the predictions made on the basis of trait fitnesses. Beatty asks: "is this unexpected change in relative frequencies of light and dark attributable to selection or drift?" The answer ought to be clear. In so far as the trait fitnesses fail to predict or explain the changes in trait frequencies, there is error. We are even told where the error occurs; dark moths land more frequently on light trees than would be predicted by their relative frequencies. So, the change in trait frequencies is a consequence of sampling error; this is a case of drift.

Beatty's answer is less straightforward. On the one hand, he says, it looks like drift: the outcome is a substantial deviation from what is predicted by fitness differences. On the other, he says, the change in trait frequencies is caused by natural selection—in the form of predation pressure. So the change in trait frequencies is wholly explained by selection. He ends up conceding that we have to say "both." But this is unsatisfactory. If the divergence from expectation is drift then it *cannot* be explained

8. Matthen and Ariew (2002) make a similar point on independent grounds.

9. See below for an argument that one cannot identify selection (or drift) with the forces that cause the outcome.

by selection and if the outcome is wholly explained by selection then it cannot an instance of drift.¹⁰

As we diagnose it, it is the dynamic view of selection that causes the problem. Where selection is seen as a force that causally determines individual death or survival, it completely explains the changes in the structure of this population. Consequently there is no explanatory role left for drift. So, when Beatty cites predation pressure as the cause of differential mortality and identifies predation pressure with “selection,” it seems clear that “selection” (so construed) explains the changes in population structure—and drift doesn’t.¹¹ But as we have seen, natural selection theory is not in the business of explaining changes in the structure of a population by citing the forces that cause individual-level events. Instead, it explains by appeal to the statistical properties of populations—differences in trait fitnesses. Where differences in trait fitnesses account for the observed change, the change is attributed to selection. Where differences in trait fitnesses fail to account for the observed change because of some statistical error, the outcome is explained by appeal to drift. The statistical interpretation preserves the presumed explanatory relation between selection and drift, whereas the dynamical interpretation obscures it.

Yet another reason to think of both selection and drift statistically is that doing so preserves the idea that whether an outcome is to be attributed to drift or selection is sometimes a matter of explanatory context. Suppose, to vary Beatty’s scenario somewhat, there is by chance a disproportionate number of light trees in one patch of the moths’ habitat. One would find, then, that within this patch light moths do better than would be predicted on the basis of trait fitnesses calculated across the entire population. Is *this* outcome selection or drift? The answer depends upon the explanatory context. If the objective of the study is to explain changes in trait frequencies across the entire population, then what goes on in the aberrant patch looks like drift. Differences in trait fitnesses between light and dark morphs across the population as a whole do not reflect the relative successes of light and dark morphs within the patch. On the other hand, if the objective is to explain why the light morph systematically does better *within* the patch than across the population as a whole, we would point to the fact that the trait fitness of the light morph is higher *within* the patch than across the population taken as a whole. The patch and the population as a whole constitute two distinct selective regimes. There is selection *for* the light trait within the patch but selection *against* it across the population. So the outcome within the patch is either drift or selection depending upon how the trait fitnesses are calculated.

10. Beatty expresses a certain discomfiture concerning his proposed solution.

11. These are precisely the grounds on which Rosenberg (1994) argues against drift as a genuine evolutionary phenomenon.

On the statistical view, some sequence can be understood as either drift or selection, depending how one calculates trait fitnesses. There may be no single correct way to calculate the trait fitnesses; often enough it is simply a matter of our explanatory project. In contrast, if selection and drift are distinct forces it cannot be a matter of explanatory context whether an outcome is correctly ascribable to drift or selection. There ought to be an objective fact of the matter which force is operating and when. But this position, as demonstrated by both Beatty's and Rosenberg's scenarios, fails to preserve the relation between selection and drift. For this reason the statistical interpretation of selection and drift is preferable.

4.3. *Selection and its Causes.* Selection and drift may not be *forces*, but there are *selective forces* acting on a population.¹² Predation, sunlight, and competition, for example, are all selective pressures that operate upon a population causing differential survival and reproduction of individuals. Why not identify predation pressure, sunlight, and competition as *the* forces of selection? This would reinstate the dynamical conception of selection, without reifying a distinct force of selection.¹³ But it is a mistake to identify natural selection with the forces causing differential survival and reproduction for reasons we have already 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. This is unsatisfactory for a number of reasons. Not only does it obscure the distinctive explanatory role of drift as we saw in the preceding section, it also eliminates the so-called "force of drift." The advantage of the dynamical conception is supposed to be that it allows us to identify *both* selection and drift as distinct forces and to partition their respective effects. But identifying the force of selection with the underlying individual-level forces causing changes in trait frequencies leaves no causal (or explanatory) work for drift to do. The dynamical interpretation short-circuits itself. In general, it is a mistake to identify natural selection with the individual-level processes that cause it. But doing precisely this is the error of the dynamical interpretation.

4.4. *Pragmatic Dynamics.* Elliott Sober (1984) offers yet another attempt to justify the dynamical view of both selection and drift, this time on pragmatic grounds. We need a division of explanatory labor—drift

12. The distinction between selection and selective forces is made in Lewens (2001).

13. This suggestion has been made to DMW a number of times. Thanks to Kent Holsinger and Jon Hodge for pressing this line of enquiry.

and selection explain different things—and the dynamical view gives us one. Sober is sensitive to the tension between the dynamical and the statistical readings of the central evolutionary concepts. He says:

It is not controversial that two “factors” influence whether the percentage of heads one gets on a run of independent tosses will fall in a given interval. These are (1) the coin’s probability of landing heads and (2) the number of times the coin is tossed. But it would be bizarre to treat these two factors as constituting separate processes or forces. Why, then, is drift separated from selection, mutation, and other deterministic factors in evolutionary theory?

The reason is to facilitate comparison among different populations. Two populations may be characterized by the same set of selective coefficients and yet differ in size. Two other populations may have the same effective size and yet differ as to the selective forces that impinge on them. Separating selection and drift yields concepts that are needed to mark important biological distinctions. (Sober 1984, 115)

The distinction is illustrated, by Sober, with the case of coins.

Two coins may have the same chance of landing heads but may be tossed a different number of times. Two other coins may be tossed the same number of times, but may differ in their chances of landing heads. The first two coins have something in common, and so do the second. Probability theory provides us with the conceptual machinery needed to describe what these points of similarity amount to. (Sober 1984, 115).

Ironically, Sober’s reason for thinking that selection and drift are forces is our reason for thinking that selection and drift are appropriately statistical—viz. that probability theory provides us with the conceptual machinery to make the distinction. The distinction between those effects that are due to selective co-efficients (trait fitnesses) and those that are merely the consequences of population size is indeed an important one. But the distinction can be made *within* the ambit of the statistical reading. One is explained by the differences in average propensities among organisms of different trait classes. The other is explained as the error in a sequence of trials (births, deaths, and reproductions).

5. Implications of Statistical Selection. The distinction between the dynamical and statistical interpretations of evolutionary theory has implications for a number of current debates in the philosophy of biology. We restrict our attention to two.

5.1. What Natural Selection Explains. There is an ongoing dispute within the philosophy of biology concerning what precisely natural selec-

tion explains. Specifically, some authors claim that natural selection can explain for each individual organism why it has the traits it has (this is the so-called “Positive View”) and others deny it (this is the so-called “Negative View”). Those who hold the Negative View (led by Lewontin (1983) and Sober (1984, 1995)) point to a distinction between variational and developmental explanations. A variational explanation explains changes in the structure of a population by appeal to the amount and kind of variation within that population. A developmental explanation accounts for those features which are consequences of the developmental processes occurring within individuals. According to the Negative View, natural selection theory provides variational explanations of the changes in structure of a population *without* explaining the possession of traits by individuals. Those who support the Positive View maintain that *by* changing the structure of a population, selection becomes part of the causal history of an individual’s possession of its particular traits.

The dynamic conception of selection appears to support the Positive View and undermine the Negative View. If natural selection is a force that causes individuals to die, survive, and reproduce, then for any individual F , that is born with heritable trait a , if selection for a caused F ’s birth, then selection is part of the causal background of F ’s possessing a . Under these circumstances, it seems entirely appropriate to say that natural selection partially explains the possession of a trait by an individual. Some proponents of the Negative View respond that the processes of heredity and development which cause an individual to have the traits it has *pre-empt* selection. Hence selection is not part of the explanation of why an individual has the traits it has.¹⁴ But this is unsatisfactory. In general, distal causes of an event are pre-empted by more proximal causes, but they are none the less causes, and consequently no less part of the complete causal explanation of the event.¹⁵ So, the appeal to pre-emption cannot save the Negative View.

In contrast, the statistical conception of selection favors the Negative View over the Positive. As we saw in the preceding section, natural selection theory explains changes in the structure of a population, but not by appeal to the individual-level causes of births, deaths, and reproductions. As natural selection theory does not deal in the causes of these individual-level phenomena, it cannot yield causal explanations of them. If the statistical conception is correct, then natural selection theory can no more explain why a particular individual has the traits it has than the kinetic

14. There are many variants on the pre-emption argument and responses to it. (See Neander 1995; Walsh 1998; Matthen 1999; Lewens 2001; Pust 2001.)

15. A similar point is made for mental causation by Yablo (1992).

theory of gases can explain why a particular molecule has the velocity it has. It seems to us, then, that the statistical versus the dynamical interpretation of selection is crucial to the resolution of this debate over what natural selection explains.

5.2. The Units of Selection. If the statistical interpretation of selection practically decides one issue in current philosophy of biology, it renders another more or less nugatory. The units of selection debate is construed as a question concerning the causes of changes in trait frequencies in a population. According to genic selectionists, natural selection operates at the level of the gene only. This claim is interpreted in either of two ways: (i) differences in gene fitnesses are adequate to explain all changes in the structure of a population that are attributable to selection (Dawkins 1976, 1982; Williams 1966), or (ii) the causes of those changes in trait frequencies in a population are to be found exclusively in the context-sensitive causal powers of genes (Sterelny and Kitcher 1988). Pluralists, those who maintain that genes, organisms, and groups are all genuine units of selection, argue against (i) that explaining changes in trait frequencies by calculating gene fitnesses may be a handy “bookkeeping device” but it fails to expose the real sub-population-level *causes* of evolutionary change. They argue against (ii) by claiming that the causes of evolutionary change can only be captured by citing the causal properties of all these levels of organization.

It is clear that some pluralists object to genic selectionism on dynamic grounds. Sober and Lewontin (1982) aver that to ascribe changes in the structure of a population to the force of impinging on the level of genes alone is a distortion. Genic selectionism is wrong in general because it fails to recognize the multiple levels at which the forces of selection act.¹⁶ The changes in gene frequencies explained by natural selection theory are very often merely the resultant of selection forces acting at different levels of organization. The bookkeeping device of calculating gene fitnesses misleads because it fails to identify the causally relevant properties that bring about population change. Of course this objection holds only if natural selection theory explains by appeal to forces and to the sub-population-level causes of population change.

From the perspective of the statistical interpretation, the units of selection debate appears deeply uninteresting. Natural selection explains changes in the structure of a population, but not by citing or identifying forces or sub-population-level causes. So the objection that pluralism iden-

16. Kenneth Waters (1991) argues effectively that Sober and Lewontin’s objection to genic selectionism is predicated on what he calls realism concerning the force of selection.

tifies the relevant forces and that gene selectionism doesn't—or vice versa—completely misses the point. On the statistical interpretation, natural selection theory identifies classes (e.g., those sharing the same allele, or those sharing the same allelic combination), such that differences between their respective statistical properties (calculated as the mean and variance of individual fitnesses) adequately predict and explain changes in the structure of the population. This is merely bookkeeping and so long as the books are kept properly natural selection theory is able to do its work.

Of course, this consideration doesn't adjudicate in favor of gene selectionism or against it. Natural selection theory doesn't privilege one way of keeping the books over another. For instance, if you want to explain why, in the case of a heterotic trait at equilibrium, the frequencies of diploid genotypes (aa , AA , and Aa) change over the course of a generation, you calculate the fitnesses of trait classes picked out by those diploid traits. If at the same time you want to explain why the gene frequencies (a and A) don't change you simply calculate the trait fitnesses of classes demarcated by possession of a and A .¹⁷ Similarly, if you want to explain why the altruist trait is increasing in frequency in the population as a whole despite declining within each group, it is best not to calculate the within-group fitness of altruism, but to use the fitness of the altruistic trait across the population as a whole. On the other hand, if you want to explain the tendency of altruism to decrease in each group, you had better calculate the within-group fitnesses. On the assumption that these changes in relative trait frequencies are attributable to selection, the objective is simply to find the statistical properties of populations that show those changes to be likely. It is not to uncover the level of organization at which the force of selection operates. In short, the statistical conception of selection shows the units of selection debate to be simply the “units of bookkeeping” debate.

6. Conclusion. We have attempted to draw a distinction between the dynamical and statistical interpretations of evolutionary theory and to argue for the latter over the former. The widely held view that evolutionary theory is a theory of forces is mistaken. It is, instead, a statistical theory; it explains the changes in the statistical structure of a population by appeal to statistical phenomena—*trait* fitnesses and “sampling” error.

REFERENCES

Ariew, André (1998), “Are Probabilities Necessary for Evolutionary Explanations?”, *Biology and Philosophy* 13: 245–253.

17. Sober and Lewontin (1982) use this very example to demonstrate that gene selectionism commits the so-called “bookkeeping fallacy.”

- Beatty, John (1984), "Chance and Natural Selection", *Philosophy of Science* 51: 183–211.
- (1992), "Random Drift", in E. F. Keller and E. Lloyd (eds.), *Keywords in Evolutionary Biology*. Cambridge, Mass.: Harvard University Press, 273–281.
- Brandon, Robert (1990), *Adaptation and Environment*. Cambridge, Mass.: MIT Press.
- Cain, Arthur J. (1951), "So-called Non-adaptive or Neutral Characters in Evolution", *Nature* 168: 424.
- Crow, James and Mooto Kimura (1970), *An Introduction to Population Genetics Theory*. Minneapolis: Burgess Press.
- Dawkins, Richard. (1976), *The Selfish Gene*. Oxford: Oxford University Press.
- (1982), *The Extended Phenotype*. Oxford: Oxford University Press.
- Dobzhansky, Theodosius, Francisco Ayala, and G. Ledyard Stebbins (1977), *Evolution*. San Francisco: Freeman.
- Endler, John A. (1986), "Natural Selection in the Wild", *Princeton Ecological Monographs*. Princeton: Princeton University Press.
- Gigerenzer, Gerd et al. (1989), *The Empire of Chance: How Probability Changed Science and Life*. Cambridge: Cambridge University Press.
- Lewens, Tim (2001), "Sex and Selection: A Reply to Matthen", *British Journal for the Philosophy of Science*. 52: 589–598.
- (2001), *Organisms and Artefacts: Design in Nature and Elsewhere*. PhD. Dissertation. Cambridge: University of Cambridge.
- Lewontin, Richard C. (1970), "The Units of Selection", *Annual Review of Ecology and Systematics* 1: 1–14.
- (1983), "Darwin's Revolution", *New York Review of Books* 30: 21–27.
- (1974), "The Genetic Basis of Evolutionary Change", New York: Columbia University Press.
- Matthen, Mohan (1999), "Evolution, Wisconsin Style: Selection and the Explanation of Individual Traits", *British Journal for the Philosophy of Science* 50: 143–150.
- Matthen, Mohan and André Ariew (2002), "Two Ways of Thinking about Fitness", *Journal of Philosophy* 49(2): 55–83.
- Maynard Smith, John (1986), *Evolutionary Genetics*. Oxford: Oxford University Press.
- Mills, Susan and John Beatty (1979), "The Propensity Interpretation of Fitness", *Philosophy of Science* 46: 263–286.
- Millstein, Roberta (1996), "Random Drift and the Omniscient Viewpoint", *Philosophy of Science* 63 (suppl.): S10–S18.
- Neander, Karen (1995), "Explaining Complex Adaptations: A Reply to Sober's Reply to Neander", *British Journal for the Philosophy of Science* 46: 583–587.
- Pust, Joel (2001), "Natural Selection Explanation and Origin Essentialism", *Canadian Journal of Philosophy* 31: 201–220.
- Rosenberg, Alexander (1994), *Instrumental Biology or the Disunity of Science*. Chicago: University of Chicago Press.
- (2001), "Discussion Note: Indeterminism, Probability, and Randomness in Evolutionary Theory", *Journal of Philosophy*: 68: 536–544.
- Sklar, Lawrence (1999), "The Reduction(?) of Thermodynamics to Statistical Mechanics", *Philosophical Studies* 95: 187–202.
- Sober, Elliott (1984), *The Nature of Selection*. Cambridge, Mass.: MIT Press.
- (1993), *Philosophy of Biology*. Boulder, Colo.: Westview Press.
- (1995), "Natural Selection and Distributive Explanation: A Reply to Neander", *British Journal for the Philosophy of Science* 46: 384–397.
- (2001), "The Two Faces of Fitness", in R. Singh, Diane Paul, Costas Crimbas, and John Beatty (eds.), *Thinking about Evolution: Historical, Philosophical and Political Perspectives*. Cambridge: Cambridge University Press, 309–321.
- Sober, Elliott and Richard C. Lewontin (1982), "Artifact, Cause and Genic Selection", *Philosophy of Science* 49: 157–180.
- Sober, Elliott and David Sloan Wilson (1994), "A Critical Review of Philosophical Work on the Units of Selection Problem", *Philosophy of Science* 61: 534–555.
- Sterelny, Kim and Philip Kitcher (1988), "The Return of the Gene", *Journal of Philosophy* 85: 339–360.

- Walsh, Denis (1998), "The Scope of Selection: Neander and Sober on What Natural Selection Explains", *Australasian Journal of Philosophy* 76: 250–264.
- (2000), "Chasing Shadows: Natural Selection and Adaptation", *Studies in History and Philosophy of Biological and Biomedical Sciences* 31: 135–153.
- Waters, C. Kenneth (1991), "Tempered Realism about the Forces of Selection", *Philosophy of Science* 58: 553–573.
- Williams, George C. (1966), *Adaptation and Natural Selection*. Princeton, N. J.: Princeton University Press.
- Yablo, Steven (1992), "Mental Causation", *Philosophical Review* 101: 245–280.