Can There Be Stochastic Evolutionary Causes?

Patrick Forber and Kenneth Reisman†‡

Do evolutionary processes such as selection and random drift *cause* evolutionary change, or are they merely convenient ways of describing or summarizing it? Philosophers have lined up on both sides of this question. One recent defense (Reisman and Forber 2005) of the causal status of selection and drift appeals to a manipulability theory of causation. Yet, even if one accepts manipulability, there are still reasons to doubt that genetic drift, in particular, is genuinely causal. We will address two challenges to treating drift as causal within a manipulation framework. We will argue that both challenges ultimately fail, but that they raise interesting and subtle issues about the nature of causation and the differences between selection and drift.

1. Introduction. Do evolutionary processes such as selection and random drift *cause* evolutionary change, or are they merely convenient ways of describing or summarizing it? Philosophers have lined up on both sides of this question (Sober 1984; Rosenberg 1994; Glymour 1999; Matthen and Ariew 2002; Walsh, Lewens, and Ariew 2002; Stephens 2004; Reisman and Forber 2005; Shapiro and Sober 2007). One recent defense of the causal status of selection and drift appeals to a manipulability theory of causation. The core of manipulability theories states, roughly, that if an appropriately controlled manipulation of variable A results in a systematic change in variable B, then A is a cause of B. We have argued that biological variables such as the strength of selection and initial population size can be empirically manipulated to produce systematic changes in populations; hence, if one accepts a manipulability theory, then one should regard selection and the founder effect (a type of drift) as causes of population change.

The most direct way to dismiss this argument is to reject manipulability

†To contact the authors, please write to: Patrick Forber, Tufts University, Philosophy Department, Miner Hall, Medford, MA 02155; e-mail: patrick.forber@tufts.edu.

‡Thanks to audiences at the 'PBDB1' Conference and the 2006 PSA Meeting for valuable feedback.

Philosophy of Science, 74 (December 2007) pp. 616–627. 0031-8248/2007/7405-0006$10.00

as a criterion for identifying causal relations. We will not defend the manipulability approach here; excellent defenses already exist (Pearl 2000; Woodward 2003). Yet, even if one accepts manipulability, there are still reasons to doubt that genetic drift, in particular, is genuinely causal. We will address two challenges to treating drift as causal within a manipulation framework. The first challenge accepts the basic manipulability criterion of causation, but denies that drift is the kind of process that can be manipulated. The second challenge accepts that drift is amenable to manipulation, but denies that manipulations of drift produce *systematic* changes in a population. Manipulations of drift produce *random* changes, whereas manipulations of selection produce systematic changes in a population. These arguments would put a wedge between the causal reality of 'deterministic' selection and that of 'stochastic' drift; while selection can be manipulated to produce population change, drift cannot.

In this paper, we articulate these challenges to the causal status of drift and examine their implications. We will argue that both challenges ultimately fail, but that they raise interesting and subtle issues about the nature of causation and the differences between selection and drift.

2. Mathematical versus Causal Relations. We claim that drift is an evolutionary cause because intervening on population size, the key variable in the drift model, brings about a change in the distribution of outcomes predicted by the model. To support our claim we pointed to a study done by Dobzhansky and Pavlovsky (1957) as an example of a manipulation of drift. The study shows that decreasing the number of founding members in replicate populations produces an increase in the variability of evolutionary outcomes across those replicate populations. Hence, we concluded that drift, in the form of the founder effect, is a cause of this kind of evolutionary change.

Did Dobzhansky and Pavlovsky perform a genuine manipulation of drift? Arguably, they did not; drift processes are not amenable to manipulation because the models of drift identify *mathematical* or *conceptual* relations rather than causal ones.[1] The relationship between population size and the distribution of outcomes due to drift should be interpreted as mathematical one, described by the model, rather than a causal one in the evolutionary process. Changing population size does not *cause* a change in the variability of outcomes, but is rather a mathematical consequence of the probability calculus, or so the objection goes.

While this objection does track an important difference between stan-

1. Woodward (personal communication), based on Hausman and Woodward (1999), mentions this as a potential problem with manipulating drift. Rosenberg's (1994) argument that drift is a fiction also coheres with this objection.

dard models of selection and drift, a difference we will discuss below, it overlooks the role mathematical models play in science. Evolutionary biologists often use mathematical relationships in models to represent causal relationships in the evolutionary process. The *models* may describe mathematical relations, but the interesting question concerns whether manipulations of *biological systems* yield evolutionary outcomes predicted by the model used to represent the system. The instantiation of the model relationship in a real biological system, confirmed by manipulation of the relevant variables in that system, makes it causal. To convey the force of this reply, and the intuition behind the original objection, we need to delve into the mathematical details of drift and selection.

Random drift includes a heterogeneous group of processes that posit stochastic fluctuations rather than fitness differences to explain evolutionary change. Beatty (1984; 1992) analogizes the family of drift hypotheses to 'indiscriminate sampling' processes. The standard model of drift, the Wright-Fisher model (Gillespie 1998), makes this analogy clear. The Wright-Fisher model assumes that a finite number of individuals in the parent generation contribute genes to an infinite pool of gametes. The genes for the offspring then are sampled randomly from the pool. Beatty calls the sampling indiscriminate because it occurs without regard for any phenotypic features of individuals. This process of 'sampling' can change the gene frequency in the offspring generation according to probabilities specified by the following *binomial equation* from probability theory, where $2N$ is the total number of alleles in a generation (twice the population size N in a diploid system), p is the frequency of A_1 allele and i is the number of A_1 alleles in the next generation. The function $C(2N, i)$ gives the binomial coefficient which specifies the number of possible ways for iA_1 alleles to be transmitted in a diploid population of size N:

$$P(X = i) = C(2N, i)p^i(1 - p)^{2N-i}.$$

The strength of drift depends on population size (N). Smaller populations experience stronger drift; decreasing N increases the chance that stochastic fluctuations will change allele frequencies. Figure 1 shows a single-step drift process for a population of size 4. The family of drift models also includes *founder effects*. Founder effects occur when a small subpopulation colonizes a new niche or locale (Dobzhansky and Pavlovsky 1957; Mayr 1963). That small subpopulation represents a random subset of the genetic variation found in the overall population, analogous to the sampling of alleles for an offspring generation from the much larger pool of gametes, and hence counts a form of drift.

The binomial equation provides a model of how stochastic fluctuations can lead to evolutionary change. The probabilities of each possible outcome *are* a consequence of the mathematics. Yet notice that the mathe-

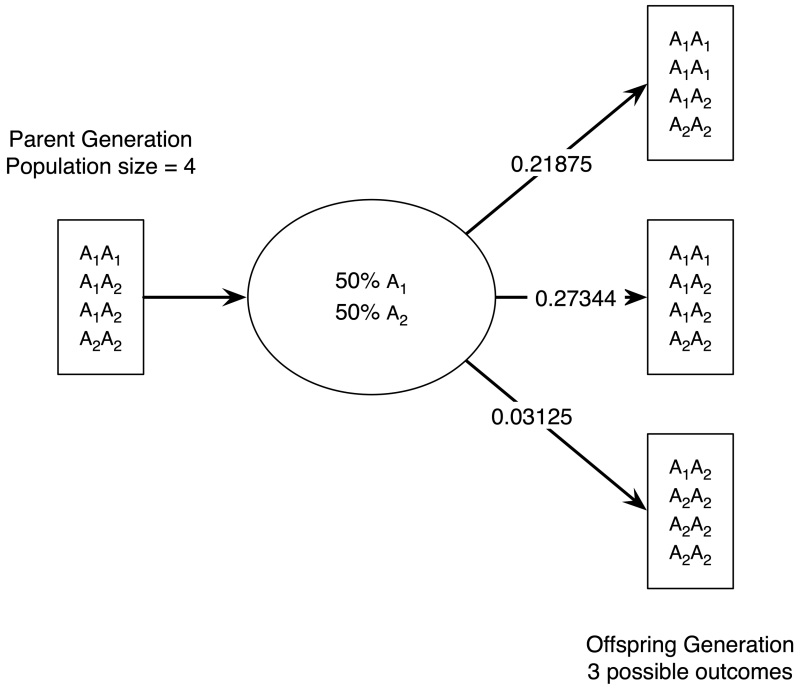

Figure 1. A representation of the Wright-Fisher model of random genetic drift. The frequencies of A_1 and A_2 alleles in the infinite gamete pool are determined by the frequencies found in the parent generation. Of the 9 possible outcomes for the offspring generation, 3 are given above. The probability of each outcome is determined according to the binomial equation. No change (middle) or small deviations (top) are more likely than large deviations (bottom), but all outcomes, including the extinction of A_1 or A_2, are possible.

matical binomial equation, to explain evolutionary change, must have a biological interpretation. In the drift model the parameters N, p, and i represent biologically meaningful variables: population size, allele frequency and the number of alleles in the next generation. This biological gloss adds something to the binomial equation; it shows how the mathematical model can represent how biological systems can change over time. Experimental manipulation, guided by mathematical models, is a good strategy for identifying causal relations in biological systems. Thus, manipulating the population size *in a biological system* helps identify causal relations between the variable and potential evolutionary outcomes.

Moreover, if the objection were right, if the relation between population size and evolutionary outcomes described by the drift model counts as a

https://doi.org/10.1086/525608 Published online by Cambridge University Press

mathematical one, then it seems that relations between fitness values and evolutionary outcomes described by selection models should be, too. That is, the objection and our reply should apply equally to mathematical models of selection *and* drift. Their causal status stands or falls together. A closer look at the standard selection model makes this clear.

Consider selection in an infinite, randomly mating diploid population with two alleles, A_1 at frequency p and A_2 at frequency $q = 1 - p$. Suppose the relative fitnesses are: $w(A_1A_1) = 1$, $w(A_1A_2) = 1 - s$, and $w(A_2A_2) = 1 - 2s$, where s is the selection coefficient.[2] The change in the allele frequency over the next generation (Δp) is given by the following equation:

$$\Delta p = (pqs)/(1 - 2pqs - 2q^2s).$$

Here we have another mathematical relationship described by the selection model. The change in frequency (Δp) is entailed by the initial frequencies (p, q) and the selection coefficient (s), given the assumptions of the model. Manipulating selection in some biological system, by (say) intervening to set $s = 0$ and nullifying selection, reveals that selection causes changes in the allele frequency. Without selection no change occurs.

If manipulating the selection coefficient (s) and observing the predicted change in some biological system is a reliable strategy for showing that selection causes evolutionary change, then manipulating population size (N) should be a similarly reliable strategy to show that drift causes evolutionary change. Population size can be manipulated in experimental populations. The drift model, drawing on the binomial equation, makes probabilistic predictions about how allele frequencies in the population may fluctuate. If the outcomes of the manipulations in replicate experimental populations change according to the predictions of the model then this justifies attributing a causal relationship between population size and the pattern of outcomes across replicates. Representing a relationship mathematically in a model does not preclude that the relationship in a biological system may be causal. Thus, the first objection ultimately fails; we can manipulate the relevant drift variables in biological systems.

3. Levels of Analysis. We suspect that the resistance some have to treating drift as a cause trades on an important difference between the standard model of drift and the standard model of selection. The selection model makes a definite prediction about how the population *will* change given initial conditions, whereas the drift model makes a probabilistic prediction about how the population *may* change. The perceived difference between

2. This represents an additive fitness structure; having two deleterious A_2 alleles is twice as bad as having one. For the purposes of the model let $0 \leq s \leq 0.5$.

selection and drift reflects a difference in the *level of analysis* required to explore the respective evolutionary dynamics of the standard models. Evolutionary events and causes may be discerned at multiple levels of analysis. While there is no uniquely correct way of partitioning the various levels, a partition relevant to much of evolutionary biology distinguishes between the *individual*, *population*, and *ensemble* levels.

The individual level of analysis is concerned with particular organisms. Events such as births and deaths of individual organisms, and properties such as phenotypes, are visible at this level of analysis. The population level of analysis is concerned with particular populations. Events such as changes in the frequencies of types within a population, and properties such as the size of a population, are all visible at this level of analysis. Notice that the population level of analysis also picks out information about the organisms in a population, but this is information about types or aggregates of individuals rather than token individuals. Finally, the ensemble level of analysis is concerned with events and properties pertaining to metapopulations (i.e., populations of populations). For example, an ecologist studying a type of population in the laboratory might set up many replicates of that type. While each replicate is a token population in itself, the population of replicates is an ensemble. Likewise, the frequency of a given type of population within the ensemble is an ensemble-level property. Any time we are interested in types of populations rather than token populations then we have ascended to an ensemble level of analysis.

How does our observation about levels of analysis connect up with the manipulability conception of causation? On this conception, a causal relationship is something that holds between *variables*, and a given variable will always be visible at some level of analysis. A variable is any property or event in a system which has several possible values and for which it possible to intervene so to change the value. In an important sense, it is up to us which variables in a biological system we identify to manipulate.

We can choose variables situated at different levels of analysis, and observe the patterns that occur as a result at different levels as well. Variables whose values are associated with states and properties of particular organisms are visible at the individual level. Variables whose values are associated with states and properties of particular populations (or types of organisms) are visible at the population level, and so on. There is no one privileged level of analysis where all causal relations are situated; causal relations can be identified and manipulated at all of these levels, from individual to ensemble.

To say that a variable X causes change at some level of analysis is just to say that certain manipulations to X would have systematic effects on some variable Y at that level of analysis. Regardless of the relevant level for X, if manipulation of X bring out about systematic changes to an

individual-level variable, then we may conclude that X causes individual-level change. If manipulations to X also bring about systematic changes to population or ensemble-level variables, then we may conclude that X causes change at those levels as well.

Manipulations to X could have systematic effects on variables at one level of analysis, and yet not at others. Thus, if we are interested in whether variable X is a cause, the answer may depend upon the level of analysis for the relevant Y variables. For example, suppose that X variable is the initial size of a type of population. It is feasible that manipulations to this variable would have no systematic effect on relevant variables at the individual level (e.g., manipulations to X could be uncorrelated with the fertility or lifespan of particular organisms), no systematic effect on relevant variables at the population level (e.g., manipulations to X could be uncorrelated with the equilibrium gene frequencies of token populations), and some systematic effect on variables at the ensemble level (e.g., manipulations to X could be correlated with *variance* in equilibrium gene frequencies across many replicate populations). Dobzhanksy and Pavlovsky (1957) conducted a manipulation that had the latter result: manipulating founding population size correlated with a change in the variability of equilibrium gene frequencies observed for the ensemble of experimental populations.

The difference between the standard models of selection and drift reflects a difference in the level of analysis necessary to guide manipulations of biological systems. Manipulating the selection coefficient determines the outcome in a single population, whereas manipulating population size requires an ensemble of populations to observe the pattern of population-level outcomes. Resolving these levels of analysis also helps explain why drift often gets described as 'sampling error' or 'noise'.[3] Stochastic noise is often introduced into formal models to investigate the strength or robustness of the signal. The signal determines the expectation and the noise increases the variance of the outcome distribution without changing the expectation. This is true when comparing standard models of selection and drift *at the population level of analysis*. The selection coefficient affects the expectation of future gene frequency distributions (signal), while drift only affects the variance of those distributions (noise). Yet at the ensemble level the *expectation* of future population-type frequency distributions changes in the drift model—at this level drift is not simply noise. In the simple system discussed above (Figure 1) the expected frequency of the 50%-A_1/50%-A_2 type population will decline faster with smaller population size N. Hartl and Clark (1997, 273–278, esp. Figures 7.4 and 7.5) give a

3. Thanks to Christian List for making this point clear.

textbook example of this kind of ensemble-level analysis of drift using the Wright-Fisher model and data from Buri (1965).

The result of all this emphasis on levels of analysis is that whether processes such as selection and drift are causes of evolution depends closely on what variables, and what levels of analysis, one privileges. Someone interested exclusively in causal analysis of individual-level variables may arrive at a different conclusion than someone interested in population or ensemble-level variable. Once we see this, the matter of whether selection and drift are evolutionary causes ceases to be a binary question. Rather, the question becomes: are there levels of analysis for which selection or drift causes evolutionary change, and, if so, what levels are they?

4. Selection and Drift: New Developments. Clarifying the different levels of analysis helps show that the contrast between selection and drift is not as sharp as it is traditionally conceived. Manipulations of both selection and drift produce the kind of systematic changes indicative of a causal relationship between the manipulated variables and evolutionary outcomes. Moreover, identifying these systematic changes often requires an analysis at the ensemble level.

The traditional picture of natural selection and random drift treats selection as a *deterministic* model and drift as a *stochastic* model. Because stochastic models make probabilistic predictions about the evolutionary trajectory of a population they require ensemble level analysis to explore their dynamics. Recent theoretical work shows that this way of construing the distinction does not generalize to more complex models. Two relatively new developments suggest that models of the evolutionary process require, as a general rule, analysis at the ensemble level; the deterministic model of selection is the exception. The first concerns a hitchhiking model in molecular evolution that describes a process dubbed *genetic draft* (Gillespie 2000a, 2000b, 2001). The second concerns a family of models that incorporate temporal or spatial *variation* into evolutionary fitness (Frank and Slatkin 1990).

Sober (1984, 110) makes the definitive statement of the traditional contrast in the philosophical literature:

> In evolutionary theory, mutation and selection are treated as *deterministic* forces of evolution. . . . *Random genetic drift*, on the other hand, is the source of the stochastic element in evolution. . . . Natural selection is typically claimed to be a deterministic process in this sense: *When it acts alone, the future frequencies of traits in a population are logically implied by their starting frequencies and the fitness values of the various genotypes*.

For population genetic selection models of infinite populations this is

true. In these models fitness differences do *determine* the evolutionary trajectory of a population. Population level of analysis suffices for these models. Observing the effect of a manipulation of fitness parameters in one token population provides the necessary information to attribute causal status to selection. The difference between deterministic selection models and stochastic drift models is predicated on a difference in idealizations. Drift models must assume a finite population and are stochastic because they make *probabilistic* predictions about future trait frequencies in a target population. Such models require an ensemble of populations to completely explore the evolutionary dynamics. At the ensemble level of analysis drift models exhaustively *determine* the outcome for the ensemble. Genetic draft and variation in fitness show that complex models of selection also require an ensemble of populations to completely explore evolutionary dynamics. Thus, the contrast between deterministic models of selection and stochastic models of drift depends on the difference between population and ensemble levels of analysis. Let us explain.

Gillespie describes the theoretical model of *genetic draft*.[4] The model of draft builds upon the basic model of genetic hitchhiking proposed by Maynard Smith and Haigh (1974). Hitchhiking works like this. Consider a situation where no recombination occurs; the fate of all the variants on a chromosome or region of DNA are forever linked. Some adaptive mutation can, by chance, occur in a chromosome or region of DNA. Selection will, if the mutation overcomes the risk of initial loss due to drift, sweep this adaptive mutation to fixation, along with all linked neutral variants. Neutral alleles 'draft' alongside the evolutionary trajectory of the adaptive mutation much like professional cyclists draft behind the leaders of the peleton. When there is free recombination each site has an independent evolutionary fate; no draft occurs. Draft differs from drift in that selection for a linked adaptive mutation drives stochastic changes in neutral allele frequencies. Draft models, unlike drift models, need not assume a finite population; draft can be effective in both large and small populations, whereas drift requires smaller population sizes to be effective (Gillespie 2001). Through linkage draft causes what Sober (1984, 97–102) calls *selection of* the neutral variant.

Draft models of the evolutionary process introduce two stochastic elements. First, whether an adaptive mutation occurs on a linked site in the chromosome or DNA region is a matter of chance. What linked neutral variants will draft on the adaptive mutation depends on the probabilistic rate of adaptive mutations for that set of linked sites. Second, whether a recombination event occurs to uncouple the fate of some linked

4. Skipper (2005) discusses draft models and the interaction between draft and drift.

neutral allele from the fate of the adaptive allele is also a matter of chance. This stochastic element is introduced by making a more realistic assumption about the rate of recombination. In natural populations recombination is neither complete nor free ($1 < r < 0$). Recombination can occur frequently or rarely and there are empirically estimated rates of recombination (r) for specific lineages. When $1 < r < 0$ it is a matter of chance whether the linkage between neutral and adaptive alleles is broken due to some recombination event. Increasing the rate of recombination decreases the effectiveness of draft.

On the traditional picture, increasing population size makes the evolutionary dynamics at the population level more deterministic because it decreases the stochastic element introduced by drift. The model of genetic draft, by adding new stochastic elements, shows that the evolutionary process will continue to have stochastic dynamics even when the population size increases. Like drift, the draft model requires an ensemble of populations to explore all the possible evolutionary trajectories. Incorporating recombination and linkage leads to stochastic draft dynamics due to selection for linked adaptive mutations. Draft makes *selection of* stochastic at the population level, and thus requires analysis at the ensemble level.

Yet draft models do not invalidate the view that some simple direct *selection for* models are deterministic. If there is no drift then the future frequency of the adaptive variant in the draft case will still be entailed by the fitness values. While strictly true, what matters is whether the case of selection in an infinite population is generally representative and merits the traditional distinction between deterministic selection and stochastic drift. At the very least draft models show that drift is not the only source of stochastic elements, and that a kind of selection (*selection of* linked sites) requires analysis at the ensemble level.

Models that incorporate variation in fitness show that *selection for* can require ensemble-level analysis. Frank and Slatkin (1990) review different models that incorporate some kind of variation in fitness and locate them in a unified mathematical framework. Mathematically these models are quite complex. Yet conceptually they are straight forward. These models aim to capture biological cases where some factor in the environment or in development causes variability in the fitness values. The environment may regularly oscillate between wet and dry years, individuals may have developmental mechanisms that cause the reproductive success of a single genotype to vary, or there may be spatial heterogeneity in the environment. All these factors introduce variation into fitness and thus require a *variance* term to adequately predict the evolutionary dynamics. Also, all variation in fitness models assume a finite population size; in the infinite case the central limit theorem eliminates the need for a variance term. Without

needing to dwell on the mathematical details, these models have a clear conceptual consequence regarding what sort of variables are necessary to capture the evolutionary dynamics for different classes of selection models.

Suppose we want to predict how the frequency p of some type A_1 will change in a population of A_1 and A_2 types over one generation of selection. As with the standard selection model, we want to predict the difference in frequency Δp based upon the specific kind of model. The standard selection model counts as deterministic due to the assumption of an infinite population. The predictive variable Δp specifies the evolutionary outcome for all populations that share the initial conditions. The population-level analysis suffices because the model entails that the frequency will change by Δp. But this is not true for all selection models. As Frank and Slatkin show, the variation in fitness cases require both an expectation *and* a variance term to predict the *expected* outcome for a populations with the same initial conditions: $E(\Delta p)/\mathrm{Var}(\Delta p)$. In this case, the predictive variable $E(\Delta p)/\mathrm{Var}(\Delta p)$ provides a *probabilistic* prediction for the evolutionary trajectory of a population (Frank and Slatkin 1990, 245). Thus, the variation in fitness cases require ensemble-level analysis.

The traditional picture that treats selection as deterministic only fits when the population level of analysis suffices to explore the evolutionary dynamics of selection models. The draft and variation in fitness models show that there are robust stochastic elements in evolutionary models incorporating selection that, like drift, require ensemble-level analysis. The predictions of ensemble-level models are probabilistic, and so they make counterfactual commitments that require an ensemble of populations to explore. The level of analysis guides the construction of experiments capable of detecting causal relations. For natural populations in which the idealization of an infinite population size is legitimate, manipulating a single population suffices to detect the causal impact of selection. Detecting draft, drift, or variation in fitness in natural populations requires different manipulations. Respecting the difference in levels of analysis, manipulations to detect these stochastic evolutionary causes must be replicated across an ensemble of populations.

5. Conclusion. Both drift and selection models describe mathematical relations between variables and evolutionary outcomes. Testing whether these mathematical models represent causal relations in biological systems requires manipulating the relevant variable in the system and observing the effects. On the manipulability conception of causation there is no reason why higher level variables cannot be causal. Manipulating the relevant variables for selection and drift, especially in complex cases that include genetic draft or variance in fitness, leads to systematic ensemble-level consequences. So both selection and drift should count as evolu-

tionary causes. Yet these evolutionary causes are 'high level' causes, discernible at the population or ensemble levels of analysis. This implies that often the causal signature of evolutionary processes cannot be resolved at the individual level or below.

REFERENCES

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

——— (1992), "Random Drift", in E. F. Keller and E. A. Lloyd (eds.), *Keywords in Evolutionary Biology*. Cambridge, MA: Harvard University Press, 273–281.

Buri, Peter (1965), "Gene Frequency in Small Populations of Mutant *Drosophila*", *Evolution* 10: 367–402.

Dobzhansky, Theodosius, and Olga Pavlovsky (1957), "An Experimental Study of the Interaction Between Genetic Drift and Natural Selection", *Evolution* 11: 311–319.

Frank, Steven A., and Montgomery Slatkin (1990), "Evolution in a Variable Environment", *American Naturalist* 136: 244–260.

Gillespie, John H. (1998), *Population Genetics: A Concise Guide*. Baltimore: Johns Hopkins University Press.

——— (2000a), "Genetic Drift in an Infinite Population: The Pseudohitchhiking Model", *Genetics* 155: 909–919.

——— (2000b), "The Neutral Theory in an Infinite Population", *Gene* 261: 11–18.

——— (2001), "Is the Population Size of a Species Relevant to its Evolution?", *Evolution* 55: 2161–2169.

Glymour, Bruce (1999), "Population Level Causation and a Unified Theory of Natural Selection", *Biology and Philosophy* 14: 521–536.

Hartl, Daniel L., and Andrew G. Clark (1997), *Principles of Population Genetics*. Sunderland: Sinauer Associates.

Hausman, Daniel M., and James Woodward (1999), "Independence, Invariance and the Causal Markov Condition", *British Journal for the Philosophy of Science* 50: 521–583.

Matthen, Mohan and André Ariew (2002), "Two Ways of Thinking about Fitness and Natural Selection", *Journal of Philosophy* 99: 55–83.

Mayr, Ernst (1963), *Animal Species and Evolution*. Cambridge, MA: Harvard University Press.

Pearl, Judea (2000), *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press.

Reisman, Kenneth and Patrick Forber (2005), "Manipulation and the Causes of Evolution", *Philosophy of Science* 72: 1113–1123.

Rosenberg, Alexander (1994), *Instrumental Biology, or the Disunity of Science*. Chicago: University of Chicago Press.

Shapiro, Larry, and Elliott Sober (2007), "Epiphenomenalism: The Do's and Don'ts", in G. Wolters and P. Machamer (eds.), *Thinking about Causes: From Greek Philosophy to Modern Physics*. Pittsburgh: Pittsburgh University Press, 235–264.

Skipper, Robert A. (2006), "Stochastic Evolutionary Dynamics: Drift versus Draft", *Philosophy of Science* 72: 655–665.

Smith, John M., and John Haigh (1974), "The Hitch-hiking Effect of a Favourable Gene", *Genetical Research* 23: 23–35.

Sober, Elliott (1984), *The Nature of Selection*. Chicago: Chicago University Press.

Stephens, Chris (2004), "Selection, Drift, and the "Forces" of Evolution", *Philosophy of Science* 71: 550–570.

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

Woodward, James (2003), *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.