Natural Selection, Mechanism, and the Statistical Interpretation

Fermín C. Fulda*†

What is natural selection? I address this question by exploring the relation between two debates: Is natural selection a mechanism? Is natural selection a causal or a statistical theory? I argue that the first can be assessed only relative to a model and that, following the second, there are two fundamentally different and independent kinds of models, Modern-Synthesis and Darwinian models. MS-models, I argue, are not mechanistic even if they are causal. D-models, in contrast, are mechanistic. A causal-mechanistic interpretation of D-models is thus compatible with a statistical interpretation of MS-models. Natural selection, I conclude, lacks a single, unifying nature.

1. Introduction. Natural selection (NS) is the central concept of evolutionary theory. Yet, its nature and role remain highly contested issues in the philosophy of evolutionary biology. Two recent debates address this problem. One concerns whether NS is a mechanism. The other concerns whether the theory of NS is a causal or a statistical theory. Perhaps because it seems obvious that a mechanistic account of NS implies a causal interpretation, there has been no in-depth analysis of how these two debates are related. This article aims to fill this lacuna. I believe that we can gain important insights about the nature and role of NS by clarifying this relation.

*To contact the author, please write to: The Rotman Institute of Philosophy, Department of Philosophy, Stevenson Hall 2150, Western University, London ON N6A 5B8, Canada; e-mail: ffulda@uwo.ca.

†I wish to thank audiences at the meeting in Montreal of the International Society for the History, Philosophy, and Social Studies of Biology, especially Eric Saidel for organizing the session on natural selection and mechanism; the Institute for the History and Philosophy of Science and Technology (Paris) workshop on statistical and causal approaches to natural selection, particularly Philippe Huneman for putting it together; the Philosophy of Science Association meeting in Atlanta; and the Rotman Institute of Philosophy at Western University. I would also like to thank Denis Walsh and André Ariew for insightful comments. This research was supported by the Social Sciences and Humanities Research Council of Canada.

Philosophy of Science, 84 (December 2017) pp. 1080–1092. 0031-8248/2017/8405-0023$10.00

Section 2 introduces these debates. Section 3 argues that whether NS is a mechanism can only be assessed relative to a model (sec. 3.1) and that there are two fundamentally different and mutually independent kinds of models of NS, Darwinian (D) and Modern-Synthesis (MS) models (sec. 3.2). Section 4 argues that since MS-models represent the statistical properties (growth rates) of abstract entities (trait types), they are not mechanistic even if they are causal (sec. 4.1). In turn, D-models of selection-for represent the activities of concrete entities that produce an increase in individual fitness and hence constitute a distinctive kind of historical-ecological, outward-looking mechanistic model (sec. 4.2). Section 5 concludes that a causal-mechanistic interpretation of D-models of selection-for is after all compatible with a statistical interpretation of MS-models and that since NS plays two distinct and largely independent roles in evolutionary theory, it lacks a single, unifying nature.

2. Two Debates about Natural Selection

2.1. Is Natural Selection a Mechanism? One way in which the nature of NS has been addressed is by asking whether it is a mechanism in the sense of the "new mechanistic philosophy of science." This is the view that a distinctive form of scientific explanation consists in identifying and describing the mechanisms that causally produce a phenomenon (Glennan 2002; Woodward 2002; Bechtel and Abrahamsen 2005). According to a prominent definition, "Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up conditions to finish or termination conditions" (Machamer, Darden, and Craver 2000, 3). Mechanisms typically have a hierarchical, multilevel organization, exhibit regular but not exceptionalness behavior, are spatially and temporally extended, have a function, and are causally productive and counterfactual supporting. A crucial feature of mechanisms is their ontic character; that is, they are physical entities engaging in concrete activities. Mechanistic explanations, in turn, subsume a given phenomenon under a general model or "schema." The model explains how some phenomenon works by representing how the relevant entities are spatially and temporally organized so as to engage in the activities that produce the phenomenon.

This philosophy has proven to be an insightful framework for understanding explanatory practices across many biological sciences. This suggests a potential for unification (Craver and Darden 2013). But given the centrality of the theory of evolution by NS in biology, this unification is contingent on a mechanistic account of NS. After all, NS is typically glossed in scientific, philosophical, and informal contexts as a mechanism. For example, according to a recent textbook on evolutionary biology, "The main purpose of evolutionary biology is to provide a rational explanation for the extraordinarily

complex and intricate organization of living things. To explain means to identify a mechanism that causes evolution and to demonstrate the consequences of its operation" (Bell 2008, 1).[1] And for the philosophers Skipper and Millstein (2005, 328–29), "There is no question that contemporary evolutionary biology exemplifies the view that natural selection is a mechanism."

However, Skipper and Millstein are skeptical that NS is a mechanism in the above sense. First, is not clear what the relevant entities (e.g., genes, organisms, traits, populations) and activities (e.g., nourishment, growth, survival, reproduction) are or whether they are arranged into the kind of stable spatial and temporal organization characteristic of mechanisms. But without the appropriate organization, there is no causal productivity and hence no explanatory power. In addition, the theory of NS is probabilistic. Evolutionary outcomes are made more or less likely by the factors that explain them, but they are not guaranteed by these factors. So, if NS is a mechanism, it must be a stochastic mechanism. However, the leading accounts of mechanism are deterministic. They will always produce (ceteris paribus) the same outcome. NS thus also fails to exhibit the regularity of behavior characteristic of mechanisms.

In response, Barros (2008, 318) proposed a more inclusive conception of mechanism in which NS is "a two-level, multistage stochastic mechanism." The "two-level" condition refers to the individual and population levels. The "multistage" condition refers to the temporally extended character of NS operating across generations. Finally, the "stochastic" condition refers to its probabilistic character, its being a "biased stochastic mechanism." Barros's strategy is first to represent diagrammatically this mechanism in a specific case of selection in a particular population and then abstract away the specific details to represent the general mechanism of NS.

But Havstad (2011) objected that Barros's general account is descriptively inadequate. It represents the abstract roles of selectees, selectors, and selective pressures irrespective of the actual, concrete biological realizers of these roles. Hence, it applies to any selective process, biological or not. The problem is that specific episodes of selection are too variable to be generalized from one particular case. The size of the beaks of finches that varies with seed size was selected for ease of foraging. The coloration of peppered moths that varies with the local background was selected for camouflage. The bacterial strains whose genes mutate in response to antibiotics were selected for their resistance. And the shell shape of intertidal snails that varies with crab predation was selected for decreasing the likelihood of being crushed by crabs, their predators. Each of these cases involves a very different set of entities and ac-

1. The quotation is taken from McKay and Williamson (2010).

tivities to be instances of a single mechanism. Furthermore, the idea of a general mechanism whose entities and activities are specified abstractly is incompatible with the ontic character of a mechanism. There seems to be no coherent notion of a general mechanism of NS.

2.2. Is Natural Selection a Causal or a Statistical Theory? The nature of NS has also been the focus of the debate about the interpretation of evolutionary theory. The neo-Darwinian theory of evolution by NS emanated from MS consists of models of population dynamics that represent changes in the structure of a population in terms of the concepts of selection, drift, and fitness. The degree and direction of evolutionary change are measured in terms of variation in fitness. When the population varies in fitness, it is said to undergo selection. The debate concerns the nature of the reality that these models aim to represent.

According to the standard causal or dynamical interpretation, models of population dynamics represent the causes of changes in trait frequencies in a population (Sober 1984; Bouchard and Rosenberg 2004; Brandon and Ramsey 2007; Shapiro and Sober 2007; Stephens 2010). NS is a discrete, independent force or causal process that "propels" the population through these changes. For some, this implies that in addition to the individual-level causes, NS is a further population-level cause (Millstein 2006). This interpretation conceives the theory of NS in analogy to Newtonian mechanics.

In contrast, according to the statistical interpretation, models of population dynamics represent statistical parameters that correlate but do not cause changes in trait frequencies (Ariew and Matthen 2002, 2005, 2009; Walsh, Lewens, and Areiw 2002; Walsh 2007, 2010, 2013; Ariew, Rice, and Rohwer 2015; Walsh, Ariew, and Matthen 2017). NS is a higher-order effect that results from the causal interactions at the individual level. So, in addition to the individual-level causes, no further causes need to be added at the population level. This interpretation conceives the theory of NS in analogy to an investment portfolio or a life-expectancy analysis, both of which represent general trends in idealized aggregates of individual-level events.

How are these debates related? Perhaps surprisingly, the view that NS is a mechanism rather than a force or causal process has not been explicitly put forward by defenders of the causal interpretation. However, if NS is a mechanism, then it follows that the causal interpretation is true. Mechanisms after all provide causal explanations. So, we seem to have a straightforward inference from the mechanistic status of NS to the causal interpretation of the theory. A causal interpretation, however, does not imply a mechanistic account of NS, for not all causal explanations cite mechanisms. Some cite laws, regularities, counterfactual relations, or dispositions. For example, Glennan (2008) and Huneman (2012) appeal to a counterfactual concept of causal rel-

evance to argue that models of population genetics may be interpreted causally even if the population-level factors these models cite are not causally productive—as individual-level interactions are—of the outcomes they explain. So, even if NS is not a mechanism, the causal interpretation might still be true.

Conversely, if the statistical interpretation is true, then it seems to follow that NS cannot be a mechanism. This interpretation denies that variation in fitness causes population change. Therefore, if models of population dynamics describe only statistical correlations between the fitness of trait types and their frequency in the population, then these models do not specify the mechanisms that underpin the values of these parameters. The mechanistic account of NS, it seems, is incompatible with the statistical interpretation. This, I take it, is the intuitive picture of the relation between these debates.

3. Two Ambiguities about Natural Selection

3.1. The Ontological and the Semantic Mode. Despite its intuitive plausibility, this picture of the relation between these debates is predicated on two important ambiguities, each corresponding to one of the debates. The first concerns what it 'is' for NS to be a mechanism and hence how the question must be formulated. The standard approach asks whether the process of NS itself can be identified and hypothesized as a distinctive mechanism over and above the mechanisms involved in specific instances. Call this way of posing the question 'the ontological mode'. I take the rationale for the ontological mode to be as follows. Each particular episode of selection involves a distinctive mechanism. Yet, there must be something in common among these cases that makes them all instances of NS. To preserve the popular conviction that NS is a mechanism, we must hence hypothesize a general mechanism of NS whose entities and activities are abstractly specified. We can then fill in the concrete entities and activities that occupy that role in particular instances (Barros 2008).

But we have seen the difficulties of this approach (Havstad 2011). The problem is how to reconcile the ontic character of the actual mechanisms involved in particular cases with the generality of the theory.[2] This problem, I claim, is a consequence of framing the question in the ontological mode. But even if NS itself is not a mechanism, it does not follow that there is no role for mechanism in the theory of NS. Selectionist *explanations*, rather than NS *itself*, may be mechanistic. The new mechanistic philosophy is after all an ac-

2. Besides, just because there are causes of evolution in particular cases of selection, it does not follow that there are mechanisms for NS. As mentioned before, a cause need not be a mechanism.

https://doi.org/10.1086/694008 Published online by Cambridge University Press

count of scientific (causal) explanation, among other things. As Barros (2008, 307) puts it, "The mechanistic approach is a strategy for explaining natural phenomena." Mechanistic explanations consist of a model or "schema" that represents how a given phenomenon works (Glennan 2005; Craver 2006). So, whether NS is a mechanism can only be assessed relative to a model. The proper question should then be whether *models* of NS are mechanistic. Call this way of posing the question 'the semantic mode'.

The semantic mode allows us to reconcile the ontic character of the actual mechanisms involved in particular instances of selection with the generality of the theory. Each particular instance of selection involves a different mechanism, that is, a specific set of concrete entities and activities. But these different mechanisms figure in models that have the same representational form, namely, that such-and-such set of entities and activities are productive, in the relevant context, of an increase in the relative fitness of organisms with that trait, thereby changing the structure of the population. The specific sets of concrete entities and activities vary across individual cases. But they all are productive of the same evolutionarily salient effect. So, although each model cites a different mechanism, they are all mechanistic models. The generality is thus a feature of the form of the model, not of any putative abstract mechanism. The ontic character, in turn, is a feature of the particular mechanism mentioned in the content of the model. So, we can preserve the ontic character of mechanisms by locating generality in the form of the model, not in the mechanism. Barros is thus right about the need for a general account but wrong about the prospects of a coherent notion of a general mechanism involving abstract activities and entities.

Once we make the semantic ascent from processes to models, the problem of identifying the entities and activities that constitute the general mechanism of NS dissolves. And so does the motivation for reifying NS as a distinctive (general) mechanism over and above individual-level mechanisms. So, even if there is no single general mechanism that can be identified with NS, it is possible to interpret NS models mechanistically. The semantic mode allows us to vindicate a coherent role for mechanism in the theory of NS without ontological addition.

Furthermore, assessing the mechanistic status of NS relative to a model has important implications for the debate about the interpretation of evolutionary theory. We saw that if NS is a mechanism in the ontological sense, then the causal interpretation is vindicated. Conversely, if NS is a statistical theory, then NS cannot be a mechanism. But even if there is no population-level causal process of selection as the statistical interpretation claims, it remains an open possibility that some models of NS could be mechanistic. So, a mechanistic account of NS models does not require the truth of the causal interpretation. Hence, it is in principle compatible with the statistical interpretation. The question is whether there are any NS models that can be in-

terpreted mechanistically. To answer this question, we must know what kinds of models of NS there are. This takes us to the second ambiguity.

3.2. Modern-Synthesis and Darwinian Selection. According to Grene (1961, 31), "We must . . . distinguish between 'genetical selection,' which is purely statistical, and Darwinian selection which is environment-based and causal. They remain two distinct concepts with a common name."[3] Defenders of the statistical interpretation have articulated this difference in terms of two kinds of models, MS-models and D-models (Walsh 2013; Walsh et al. 2017).

MS-models are familiar from population genetics studies of general trends in a population. They represent changes in the trait structure of a population as a function of heritable variation in trait fitness. 'Trait fitness' is a statistical property, namely, an estimate of the growth rate of an abstract (idealized) entity, a trait type. These models predict that when growth rates of trait types vary, populations change in their trait structure. They specify the strength and direction of change in trait frequencies irrespective of what causes the values of these parameters. Insofar as these models are explanatory, they explain changes in trait structure as an analytic consequence of the distribution of probabilities of trait types.

In contrast, D-models are familiar from ecological studies of specific selective forces acting on populations in the wild. They represent changes in the structure of an actual population as the causal consequence of the activities among individual organisms and their environments, "the struggle for existence." These models tell us that when individuals vary in fitness, some lineages become more prevalent in the population than others. 'Individual fitness' is a causal property, a propensity to survive and reproduce of a concrete entity, an individual organism. D-models explain changes in lineage structure as a causal consequence of the activities of individual organisms.

These models are fundamentally different. They have different representational contents. MS-models explicitly represent the relative growth rates of trait types. D-models explicitly represent the causal properties of organisms. They have different conditions of application. MS-models are substrate neutral and hence require any ensemble of abstract entities (a general population). D-models require a particular biological population causally interacting in a shared environment. They have different explananda. MS-models explain changes in trait structure: as some abstract trait types grow at a greater rate than others, some trait types increase in frequency relative to others. D-models explain changes in lineage structure (Walsh et al. 2017): as some individuals survive and reproduce more frequently, some lineages become more prevalent than others. They have different ontologies. MS-models quan-

3. The quotation is taken from Walsh et al. (2017).

tify over abstract entities, namely, trait types or classes. D-models quantify over concrete entities, individual organisms. Each trait class will typically have many individuals as members. And each individual organism will be a member of many trait classes. Since these models carve up the population in different ways, their respective ontologies are incommensurable (Walsh et al. 2017). Finally, MS- and D-models play different roles in evolutionary biology. MS-models predict and, according to statisticalists, also explain population-level regularities in terms of population-level statistical parameters. D-models explain and predict population-level regularities in terms of individual-level ecological interactions.

The distinction between these two kinds of models is not an artifact of the statistical interpretation. It is more or less explicit in what some defenders of the causal interpretation have said. Sober (1984, 59) tells us that "whereas it is mainly ecology that tries to provide source laws for natural selection, the consequence laws concerning natural selection are preeminently part of the province of population genetics. It doesn't matter to the equations in population genetics why a given population is characterized by a set of selection coefficients. . . . These values may just as well have dropped out of the sky." I take it that "source laws" correspond to D-models while "consequence laws" correspond to MS-models. Notice that the content of MS-models is insensitive to, and hence largely independent from, the actual causal conditions that determine the values of the selection coefficients. Okasha (2006) concurs: "Population genetics models are (deliberately) silent about the causes of the fitness differences between genotypes whose consequences they model. . . . To fully understand evolution, the ecological factors that lead to these fitness differences must also be understood." And so does Lewens (2010, 331): "The biologist can attend to the specific ecological episodes that result in the demise of some individuals, and the successes of others. Alternatively, she can move away from this focus on causal interactions between individuals and their environments, with a view to understanding trends in trait frequencies in populations."

Because of their fundamental difference, the assessment and interpretation of one kind of model is largely independent of the assessment and interpretation of the other kind. Thus, they are mutually independent. This has important implications for both debates. First, whether the theory of NS is a causal or statistical theory is really two different questions: Are MS-models causal or statistical? And are D-models causal or statistical? Defenders of the statistical interpretation are explicit that the debate is about MS-models (Walsh et al. 2017). Besides, it would be preposterous to claim that D-models are merely statistical. So this debate concerns only MS-models. Similarly, whether NS is a mechanism is really two different questions: Are MS-models mechanistic? And are D-models mechanistic? Let us consider each question in turn.

4. Modern-Synthesis Patterns and Darwinian Mechanisms

4.1. MS-Models Are Not Mechanistic. MS-models represent the growth rates of abstract entities, trait types, not the entities and activities of concrete entities, individual organisms. They are also substrate neutral and hence lack the ontic character of mechanisms. The maximal generality of their content and the abstract character of their ontology render MS-models inapt for mechanistic interpretation. As Havstad (2011, 522) puts it, "How natural selection works can get filled in with a mechanistic account only by moving to the individual case. At the general level, natural selection is an abstract process without the entities/parts and activities/interactions characteristic of mechanism." So, even if MS-models are causal, they are not mechanistic. The inherent inaptness of MS-models to be mechanically interpreted is presumably what underlies the tendency in the debate to implicitly ask only whether D-models are mechanistic. Skipper and Millstein's (2005, 330) schema, for example, is meant to capture "the causal crux of selection," namely, "Darwin's 'struggle for existence.'" And for Barros (2008, 312), "Natural selection occurs when a particular trait gives an individual organism a selective advantage over other individuals who do not possess the trait."

4.2. D-Models Are Mechanistic. D-models, in contrast, represent the causes of individual survival, reproduction, and death and quantify over concrete entities, namely, organisms that engage in a variety of activities that are relevant in their struggle for existence. The specificity of their content and the ontic character of their ontology indicate that these models are apt for mechanistic interpretation. I believe that close attention to the role that the notion of 'selection for' (Sober 1984) plays in D-models shows that these models are in fact mechanistic.

Suppose hearts became prevalent in a population by NS. To explain *why* they were selected, we need to look at what hearts did that contributed to individual fitness in the relevant context. Hearts can pump blood and make pulse sounds, among other activities. These activities are reliably correlated in the relevant context. This means that in that context all hearts that pump blood make pulse sounds and vice versa. But while pumping blood contributed to individual fitness, making pulse sounds did not. So there was selection of pumping blood and making pulse sounds, but only selection for pumping blood. Hearts became prevalent in the population then *because* they pumped blood. More generally, a trait T in population P was selected for doing X in historical/ecological context E if and only if past tokens of T causally contributed to the fitness of individuals who possessed T in E by doing X and those individuals who did not have T were on average less fit than those individuals who did. In contrast, to say that there was selection

of trait T in P is to say that individuals with T have a higher average fitness than individuals who lack T. 'Selection for' denotes the activities that caused differences in individual fitness and hence denotes the causes of selection. In contrast, 'selection of' denotes the effects of selection. Knowing that there was selection of a trait is silent about the activities of that trait that caused differences in individual fitness. So 'selection for' entails 'selection of' but not vice versa. This allows us to distinguish those traits that became prevalent in a population because they enhanced the fitness of their bearers from those that became prevalent because they were reliably correlated or locally coextensive with fitness-enhancing traits but made no distinctive contribution to individual fitness.

This model of the evolution of hearts explicitly represents the causes of individual survival and reproduction in an actual population. So, it is an instance of a D-model. It also explains the causes of individual survival and reproduction by explicitly representing a specific activity, pumping blood, of a concrete entity, the heart. So, the model is mechanistic. D-models that represent what a trait was selected for are thus mechanistic (Fulda 2015). This is not surprising. Although Sober (1984) originally described the selection for/of distinction in terms of "source laws" and "consequence laws" respectively, he introduced the distinction using a mechanism: a set of balls of different sizes and colors and a multilevel structure with holes of different sizes at different levels. These entities are organized such that the balls engage in the activity of falling through holes of the appropriate size under the force of gravity. Whether the balls can go through the holes depends on their respective sizes, not their colors. So the balls are selected for their size, not for their color, although as a side effect there is selection of color. We can explain the distribution of size and color at the appropriate levels by citing the way these entities are organized and engage in the relevant activities. So, although the distinction was baptized nomologically, it was conceived mechanistically.

Consider the case of Darwin's finches. The length of the beak of finches was selected for its contribution to foraging. A mechanistic D-model of why longer beaks were selected first identifies the relevant entities, beaks of different lengths, and the relevant activities, picking seeds inside trees. Then it describes how this foraging activity causally produced an increase in the individual fitness of finches with beaks with lengths that correlate with the depth of the holes in the trees. Because finches with long beaks had a higher average fitness than finches with shorter beaks, they became prevalent in the population, thus changing the structure of the finch lineage. This explanation is mechanistic. Now consider the case of the peppered moths. The coloration of peppered moths was selected for its contribution to camouflage and hence for avoiding predators. A mechanistic D-model of why dark coloration was selected identifies as the relevant entities the wings of different colors and

placing themselves in a dark background as the relevant activity.[4] Then it describes how the activity of placing dark-colored wings in a dark background provided camouflage against predators. Because moths with dark wings had a higher average fitness than moths with light-colored wings, they became prevalent in the population, thus changing the structure of the moth lineage. This explanation is mechanistic.

Paradigmatic cases of mechanistic explanations explain how some phenomenon presently works by citing the relevant current entities and activities. But selection-for explanations explain why some phenomenon occurs in terms of the past activities of past entities, so they are historical. Furthermore, they are context sensitive in an ecological sense. The description of how the entities and activities produced the relevant effect—and hence how selection-for mechanisms are individuated—necessarily includes information about the relevant aspect of the environment. This context sensitivity is thus not just a background condition for the proper performance of these mechanisms; it is constitutive of its identity. Furthermore, mechanistic explanations can be inward looking into its inner components or outward looking into the wider system of which they are part (Craver 2001, 2012; Bechtel and Abrahamsen 2005). The explanation of how the heart pumps blood is inward looking. It represents how the activities of ventricles and atria produce the pumping effect. This is a lower-level physiological explanation. The explanation of why hearts were selected, in turn, is outward looking. The explanation represents how the activity of pumping blood of an entity, the heart, causally contributes to the organism's overall capacity to survive and reproduce in the relevant environment, thus changing the structure of the population. This is a higher-level evolutionary explanation. D-models of selection-for are thus historical-ecological instances of outward-looking mechanistic explanations.

5. Conclusion. I have argued that whether NS is a mechanism concerns only D-models and that D-models of selection-for are indeed mechanistic. I have also argued that whether NS is a causal or a statistical theory concerns only MS-models and that these models are not mechanistic even if they are causal. We can now return to the relation between these debates. Since these two kinds of models are fundamentally different and mutually independent, it follows that a causal-mechanistic interpretation of D-models is after all compatible with a statistical interpretation of MS-models. I conclude that as far as contemporary evolutionary biology goes, NS lacks a single, unifying nature.

4. Skipper and Millstein (2005, 341) claim that being camouflaged is not an activity on the grounds that it is a passive defense mechanism. But the moths have to *place* themselves in the appropriate background in order to be camouflaged. Placing is an activity.

https://doi.org/10.1086/694008 Published online by Cambridge University Press

REFERENCES

Ariew, André, and Mohan Matthen. 2002. "Two Ways of Thinking about Fitness and Natural Selection." *Journal of Philosophy* 99:55–83.
———. 2005. "How to Understand Causal Relations in Natural Selection: Reply to Rosenberg and Bouchard." *Biology and Philosophy* 20:355–64.
———. 2009. "Selection and Causation." *Philosophy of Science* 76:201–24.
Ariew, André, Collin Rice, and Yasha Rohwer. 2015. "Autonomous-Statistical Explanations and Natural Selection." *British Journal for the Philosophy of Science* 66:635–58.
Barros, Benjamin D. 2008. "Natural Selection as a Mechanism." *Philosophy of Science* 75:306–22.
Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanist Alternative." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:421–41.
Bell, Graham. 2008. *Selection: The Mechanism of Evolution*. New York: Oxford University Press.
Bouchard, Frédéric, and Alexander Rosenberg. 2004. "Fitness, Probability, and the Principles of Natural Selection." *British Journal for the Philosophy of Science* 55:693–712.
Brandon, Robert, and Grant Ramsey. 2007. "What's Wrong with the Emergentist Statistical Interpretation of Natural Selection and Random Drift?" In *Cambridge Companion to the Philosophy of Biology*, ed. David Hull and Michael Ruse, 66–84. Cambridge: Cambridge University Press.
Craver, Carl F. 2001. "Role Functions, Mechanisms, and Hierarchy." *Philosophy of Science* 68:53–74.
———. 2006. "When Mechanistic Models Explain." *Synthese* 153:355–76.
———. 2012. "Functions and Mechanisms: A Perspectivalist Account." In *Functions: Selection and Mechanisms*, ed. Philippe Huneman, 133–58. Dordrecht: Springer.
Craver, Carl F., and Lindley Darden. 2013. *In Search of Biological Mechanisms: Discoveries across the Life Sciences*. Chicago: University of Chicago Press.
Fulda, Fermín C. 2015. "A Mechanistic Framework for Darwinism; or, Why Fodor's Objection Fails." *Synthese* 192:163–83.
Glennan, Stuart. 2002. "Rethinking Mechanistic Explanation." *Philosophy of Science* 69 (Proceedings): S342–S353.
———. 2005. "Modeling Mechanisms." *Studies in History and Philosophy of Biological and Biomedical Sciences* C 36:443–64.
———. 2008. "Productivity, Relevance, and Natural Selection." *Biology and Philosophy* 24:325–39.
Grene, Marjorie. 1961. "Statistics and Selection." *British Journal for the Philosophy of Science* 12:25–42.
Havstad, Joyce C. 2011. "Problems for Mechanism and Natural Selection." *Philosophy of Science* 78:512–23.
Huneman, Philippe. 2012. "Natural Selection: A Case for the Counterfactual Approach." *Erkenntnis* 76:171–94.
Lewens, Timothy M. 2010. "The Natures of Selection." *British Journal for the Philosophy of Science* 61:313–33.
Machamer, Peter, Lindley Darden, and Carl F. Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science* 67:1–25.
McKay, Phyllis, and Jon Williamson. 2010. "Function and Organization: Comparing the Mechanisms of Protein Synthesis and Natural Selection." *Studies in History and Philosophy of Biological and Biomedical Sciences* 42:279–91.
Millstein, Roberta L. 2006. "Natural Selection as a Population-Level Causal Process." *British Journal for the Philosophy of Science* 57:627–53.
Okasha, Samir. 2006. "Population Genetics." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University. https://plato.stanford.edu/archives/win2016/entries/population-genetics/.
Shapiro, Larry A., and Elliott Sober. 2007. "Epiphenomenalism—the Do's and the Don'ts." In *Studies in Causality: Historical and Contemporary*, ed. Gereon Wolters and Peter Machamer, 235–64. Pittsburgh: University of Pittsburgh Press.
Skipper, Robert A., and Roberta L. Millstein. 2005. "Thinking about Evolutionary Mechanisms: Natural Selection." *Studies in History and Philosophy of Biological and Biomedical Sciences* 36:327–47.

Sober, Elliott. 1984. *The Nature of Selection*. Chicago: University of Chicago Press.
Stephens, Christopher. 2010. "Forces and Causes in Evolutionary Theory." *Philosophy of Science* 77:716–27.
Walsh, Denis M. 2007. "The Pomp of Superfluous Causes: The Interpretation of Evolutionary Theory." *Philosophy of Science* 74:281–303.
———. 2010. "Not a Sure Thing." *Philosophy of Science* 77:147–71.
———. 2013. "Descriptions and Models: Some Responses to Abrams." *Studies in History and Philosophy of Biological and Biomedical Sciences* 44:302–8.
Walsh, Denis M., André Ariew, and Mohan Matthen. 2017. "Four Pillars of Statisticalism." In *Philosophy, Theory, and Practice in Biology*, forthcoming.
Walsh, Denis M., Timothy Lewens, and André Ariew. 2002. "Trials of Life: Natural Selection and Random Drift." *Philosophy of Science* 69:58–83.
Woodward, James. 2002. "What Is a Mechanism? A Counterfactual Account." *Philosophy of Science* 69 (Proceedings): S366–S377.