



CHICAGO JOURNALS



Of Mice and Metaphysics: Natural Selection and Realized Population-Level Properties

Author(s): Matthew C. Haug

Reviewed work(s):

Source: *Philosophy of Science*, Vol. 74, No. 4 (October 2007), pp. 431-451

Published by: [The University of Chicago Press](#) on behalf of the [Philosophy of Science Association](#)

Stable URL: <http://www.jstor.org/stable/10.1086/522900>

Accessed: 04/01/2013 19:20

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



The University of Chicago Press and Philosophy of Science Association are collaborating with JSTOR to digitize, preserve and extend access to *Philosophy of Science*.

<http://www.jstor.org>

Of Mice and Metaphysics: Natural Selection and Realized Population-Level Properties*

Matthew C. Haug^{†‡}

In this paper, I answer a fundamental question facing any view according to which natural selection is a population-level causal process—namely, how is the causal process of natural selection related to, yet not preempted by, causal processes that occur at the level of individual organisms? Without an answer to this grounding question, the population-level causal view appears unstable—collapsing into either an individual-level causal interpretation or the claim that selection is a purely formal, statistical phenomenon. I argue that a causal account of realization provides an answer to the grounding question. By applying this account of realization to the natural selection of melanism in rock pocket mice, I show how an alternative, formal account of realization, favored by proponents of the statistical interpretation, misses biologically important features. More generally, this paper shows how metaphysical issues about realization normally discussed in the philosophy of mind apply to debates in philosophy of biology. Thus, it is a first step toward fleshing out the oft-noted similarities between debates in these areas.

1. Introduction. Recently, several philosophers have argued that an adequate understanding of evolutionary theory after the Modern Synthesis calls for “a radical revision of received conceptions of causal relations in evolution” (Matthen and Ariew 2002, 57). Contrary to the traditional

*Received July 2006; revised March 2007.

[†]To contact the author, please write to: The College of William and Mary, Box 8795, Williamsburg, VA 23187-8795; e-mail: mchaug@wm.edu.

[‡]Previous versions of this paper were presented at the Cornell Philosophy Department Workshop, Brandeis University, the University of Iowa, Southern Methodist University, the College of William and Mary, and the 2006 Central Division APA meeting. Thanks to the members of those audiences for their comments and questions and especially to Robert Skipper, my commentator at the APA. I am also grateful to Richard Boyd, André Ariew, Mohan Matthen, Eric Swanson, and two anonymous referees for helpful conversations, comments, and/or correspondence. This material is based upon work supported under a National Science Foundation Graduate Research Fellowship.

Philosophy of Science, 74 (October 2007): 431–451. 0031-8248/2007/7404-0001\$10.00
Copyright 2007 by the Philosophy of Science Association. All rights reserved.

view held by both biologists and philosophers of biology, natural selection is not a cause of evolutionary change. Rather, it is a statistical trend, a “*formally* characterized phenomenon, a statistical property of physical substrates that possess certain metrical properties” (Matthen and Ariew 2002, 68–69, emphasis in original). According to this purely *statistical interpretation*, the only causal processes in evolution are at the level of biological individuals (or below), and none of these can be identified with metrical properties like selection and drift (Walsh, Lewins, and Ariew 2002, 453; cf. Matthen and Ariew 2002, 69).

These authors single out for critique what they call the *dynamical interpretation* of natural selection, according to which natural selection theory is understood as a theory of forces analogous to Newtonian mechanics. Some responses to this critique have defended parts of the Newtonian analogy (Stephens 2004). Others have argued that natural selection is an individual-level causal process, according to which “selection [is] a contingent causal process in which individual fitnesses are the causes and subsequent population differences are the effects” (Bouchard and Rosenberg 2004, 710; see also Rosenberg and Bouchard 2005). Still others have argued that natural selection is a population-level causal process in which variation in heritable abilities to survive and reproduce causes differential reproductive success and, when a population is not in equilibrium, changes in genotype and/or phenotype frequencies (Millstein 2006; see also Reisman and Forber 2005).

I think that the dynamical interpretation is flawed; there are too many important differences between forces like gravity and phenomena like natural selection and drift for the analogy between Newtonian mechanics and evolutionary theory to be particularly helpful. For instance, distinct evolutionary factors do not in general decompose and recombine like vectors. Of course, even if natural selection is not a force, it may still be a cause (cf. Stephens 2004). Further, there are powerful arguments that natural selection must be population-level causal process, if it is a causal process at all (see Millstein 2006; Glymour 1999; Sober 1984). After the modern synthesis, evolutionary theory is a probabilistic theory that abstracts away from the contingencies and vicissitudes of individual organisms’ life histories.

Thus, the view of natural selection as a population-level causal process is attractive. However, there is a fundamental question that faces any such view: namely, what I call the *grounding question*.

Grounding Question. How are population-level causal processes related to the causal processes that occur at the level of individual organisms, and how is this relation compatible with the *causal* nature of population-level processes?

All sides of the debate agree that there are causal processes that (exclusively) involve properties of individual organisms. Further, as I discuss below, proponents of the population-level causal process view acknowledge that the causal processes involving individual organisms are somehow responsible for, determine, or are the basis for the process of natural selection.

Given this consensus, the grounding question is structurally similar to the so-called “causal exclusion problem” in the philosophy of mind (see, e.g., Kim 1989, 1998). Indeed, some authors have argued that the individual-level causal processes (events or properties) preempt or screen off the population-level processes (events or properties) that they determine, in the same way that the physical realizers of mental properties allegedly preempt the causal efficacy of those mental properties (if they are held to be distinct from their realizers) (see Walsh 2000, especially p. 139 and p. 150; Walsh 2004).¹ Given that individual-level events are causally sufficient for the evolutionary fate of a population, it seems that one must either abandon the causal efficacy of population-level phenomena like selection (i.e., adopt the statistical view) or identify selection with an individual-level causal process. Thus, it looks like the population-level causal view is unstable, just as Kim has alleged that nonreductive physicalism about the mind is unstable.

The grounding question has not been addressed in the literature on natural selection, and the main goal of this paper is to redress this. I show how a metaphysical framework normally discussed in debates in the philosophy of mind can be fruitfully employed in philosophy of biology. Thus, this paper is a first step toward fleshing out the similarities between metaphysical debates in these areas that have been noted in passing by others (e.g., Fodor 1997, 162–163, n. 7; see also Shapiro and Sober, forthcoming). In particular, I use a version of the realization relation, which has been put forth as the relation between mental and physical properties, to address issues about the causal status of natural selection. The realization relation is commonly defined as follows: an entity (property, state, object) *P* realizes another entity (property, state, kind, predicate/formula) *Q* if and only if *P* satisfies the condition (fills the role) that is definitive of *Q*. *P* is a realizer of the realized entity *Q*. There are, of course, various ways of making this account more precise by specifying the nature of the condition or role (e.g., whether it is causal, computational, mathematical,

1. Matthen and Ariew explicitly state that they are not relying on an exclusion-based argument to support the statistical interpretation (2002, 82). They support an epiphenomenal view of selection on distinct grounds. They also appeal to considerations of irreversibility and discontinuity to argue that natural selection is not a causal process, but rather merely a ‘stochastic cause’. See n. 17.

or folk-theoretic) and what one takes ‘filling’ that role to consist in. Below, I discuss two such ways; one involves causal relations, while the other is purely formal.

After discussing how the grounding question is related to the exclusion problem in Section 2, I sketch a causal account of realization in Section 3. I argue that a causal account of realization can be used to provide an adequate answer to the grounding question. If the population-level properties involved in the process of selection are realized in this way, then they are grounded in individual-level causal properties so that the exclusion problem does not arise. In Section 4, I consider in detail the case of adaptive melanism in rock pocket mice and show how the causal account of realization applies to it. I use this example to argue that the *formal* account of realization sketched by Matthen and Ariew cannot completely capture the phenomenon of natural selection. I conclude in Sections 5 and 6 by presenting and responding to an argument for the claim that natural selection cannot provide unified causal explanations.

2. The Grounding Question and the Exclusion Problem. According to Millstein, the proponent of the population-level causal view does not “deny that populations are composed of individuals or that causes are acting on those individuals. Nor [does she] deny that individual-level causal events (living, dying, and reproducing) are in some sense responsible for selection. What [she denies] is that such individual-level causal events constitute the selection process itself” (2006, 650). In claiming that individual-level causal events are “in some sense responsible for selection,” Millstein rightly rejects a strong emergentist view according to which natural selection is a novel force that appears in populations of organisms. Yet this appears to sit uneasily with the claim that these individual-level causal events do not ‘constitute’ the selection process. The apparent tension in this passage suggests a crude version of the exclusion problem; if individual-level causal events are responsible for selection, what ‘causal work’ is left for the distinct population-level process of selection to do?

In order to state the problem more precisely, note that individual-level causal events (as long as they are individuated partially in terms of the environment in which they take place) *must* constitute (determine or ground) the selection process. Otherwise, the population-level causal view would violate plausible completeness and exclusion principles.² In this

2. Thus I disagree with the following passage from Reisman and Forber: “First off, we want to be clear that we readily accept that there are individual-level causes that operate and that can explain the results of the Dobzhansky and Pavlovsky experiment. However, we think that *in addition* to these individual-level causes, there are also population-level causes and that is what drift and selection are” (2005, 1120, my emphasis).

context it is not the completeness of physics that is salient but what might be called the *completeness of Darwinian biology*.

Completeness of Darwinian Biology. Every evolutionary outcome is completely determined by a causal process that involves only births, deaths, mutations, genetic recombination events, etc.³

Yet this completeness principle alone does not rule out a strong emergentist view of selection, for it is compatible with the existence of a novel, completely distinct force of selection that acts independently of Darwinian events in the production of evolutionary outcomes. So we also need an *exclusion principle*.

Exclusion Principle. Barring cases of overdetermination, there cannot be two or more causal processes responsible for a given evolutionary outcome.

If these two theses are accepted, then it appears that one cannot consistently claim that natural selection is a population-level causal process that is not identical to any individual-level causal process. Just as the exclusion problem in the philosophy of mind is supposed to force a choice between the reduction or causal impotence of mental properties, we seem to be forced to claim that natural selection is either an individual-level causal process or an epiphenomenal population-level process.

Upon reflection, it is apparent that the reasoning behind the exclusion problem plays two roles in the metaphysics of evolution. The first is laudable and sound: to rule out strong emergentism regarding selection and drift—for instance, it rules out the view that individual-level causal events are ‘responsible for’ selection in that they *cause* selection which in turn causes evolutionary change. The second role is pernicious and unsound: to force a choice between epiphenomenalism and reduction to individual-level processes—thus leading to the claim that the grounding question rests on a mistake and should be given a deflationary response by denying that there is a population-level causal process that is distinct from the individual-level causal processes that are taking place in the population.

In the next section, I spell out in what sense individual-level causal properties/relations are ‘responsible for’ and ‘constitute’ the properties involved in population-level causal processes like selection. I argue that natural selection is a causal process that is realized by, but not identical

3. This completeness principle must be formulated in terms of a fairly ‘egalitarian’ or coarse-grained notion of causation—one that is not inextricably tied to explanation (see Hall 2004; Shoemaker 2003, n. 3). For this reason, it might be clearer to formulate completeness in terms of ‘causal sufficiency’, as Kim does.

to, the individual-level causal processes occurring in a biological population. This allows me to endorse the first role of exclusionary reasoning while rejecting the second. Thus, I preserve the grain of truth in the statistical interpretation while rejecting radical, anticausal claims that are an integral part of it.

3. A Causal Account of Realization Provides an Answer to the Grounding Question. First of all, I believe that supervenience does not provide substantive answer to the grounding question for reasons similar to those that have convinced many philosophers to reject supervenience-based formulations of physicalism (see, e.g., Horgan 1993; Wilson 2005). Supervenience requires only covariation between a set of supervenient properties and a set of base properties. Thus, it is compatible with a variety of metaphysical positions that fail to answer the grounding question (and thus run afoul of the exclusion problem), e.g., emergentism and dualist occasionalism. Supervenience is also compatible with a deflationary treatment of grounding question since ‘metrical properties’ (mere averages, variances, etc.) supervene on individual-level properties and events. These compatibilities obtain because supervenience is silent about the metaphysical ground of the covariation relations it implies. In particular, there is no reason to think that the supervenient properties need even be eligible for causal efficacy.

But if supervenience fails to provide an answer, why think that there is a substantive answer to the grounding question at all? According to the statistical interpretation, there is not. The core claim of the statistical interpretation is that natural selection is a mere ‘trend’, a change in ‘stochastic properties’—ones that are mathematical (not nomic or causal) consequences of the individual-level properties of organisms in that population (Matthen and Ariew 2002, 81). Stochastic properties belong to statistical populations and are not causally efficacious. For example, it is perfectly mathematically or formally legitimate to calculate stochastic/metrical properties of the statistical population of all the moths alive in 1945 or 1951 in Australia or England. If natural selection is merely a stochastic trend, then the grounding question can be dismissed as misguided. For, there would be no causal process of natural selection to ground in the first place.

However, there is good reason to think that such ‘stochastic’ properties are not the kind of ‘population-level’ properties that are involved in the process of natural selection. Evolutionary processes involve causal properties of a *biological* population of organisms that are *realized* by the properties of individual organisms and their relations to other organisms and their environment. Examples of such properties of biological populations include a population’s size, rate of migration/gene flow, and var-

iation in phenotypic traits (which result in variation in ability to survive/reproduce) within a population. In the next section, I show how some of these properties are realized in a particular case and argue that they are not ‘mere averages’; they are not stochastic properties of a statistical population of organisms.⁴

If stochastic properties alone fail to capture certain features of a population undergoing natural selection (cf. Millstein 2006, which makes similar points), then it is clear why the statistical interpretation misses this. It conflates two notions of ‘population’: (1) a *mathematical* notion of population—an abstract, mathematical object used in statistics whose membership conditions can be biologically arbitrary, and (2) an empirical, *biological* notion of population, “a group of conspecific organisms that occupy a more or less well defined geographic region and exhibit reproductive continuity from generation to generation” (Futuyma 1986, 554–555).⁵ Importantly, organisms in a biological population are bound together or unified by *causal* processes such as mating, competition, niche construction, and being subject to common ecological circumstances (e.g., floods, fires, droughts). Further, biological populations plausibly fit into the levels of organization in nature that are ordered by the part/whole relation (e.g., subatomic particles, atoms, molecules, cells, *individual* organisms, *populations* of organisms).

It is important to note there can be statistical populations at *any* level of organization within the latter hierarchy, e.g., the average mass of a kind of cell (a cellular-level property). Hence, even if the statistical inter-

4. Of course, scientists may use sampling and replication to estimate population-level properties and their effects. Statistical methods may be used to measure or estimate a property even if the property is causally efficacious. Cf. the point below about fitness and the causally efficacious ‘components’ that underlie it. Genetic variation can be quantified in various ways, e.g., by heterozygosity or, at the DNA level, by nucleotide diversity (the average proportion of nucleotides that differ between any two copies of a gene). Some might question whether the population-level property *variation* is ever causally efficacious (which might be motivated by skepticism about population-level causation in general). In response, I would point to examples that can plausibly be understood only by taking *variation in a property* to be a cause, e.g., when the temperature gradient over the surface of a large mirror causes deformation in the mirror. See also Millstein’s discussion of frequency-dependent selection (2006, 629ff.)

5. Shapiro and Sober (forthcoming) make a similar point: “[The proponents of the statistical interpretation] repeatedly assert that selection exists when two sets of individuals (call them the x-individuals and the y-individuals) differ in fitness. But this can occur when x has nothing causally to do with y; x and y might just as well be at opposite ends of the universe. Our reply is that no biologist would treat two individuals as part of the same (token) selection process if they were at opposite ends of the universe” The reason that no biologist would do so is that such individuals are not part of the same biological population.

pretation were correct, it would imply nothing about the level of organization at which causal processes occur. The core claim of the statistical view does not show that all realized properties of a biological population are causally inefficacious, any more than it shows that all realized properties of a cell are causally inefficacious, *assuming that there is no reason to claim that realized properties, in general, are causally inefficacious.*

Yet Matthen and Ariew sketch and utilize a formal account of realization, according to which realized properties *are* causally inefficacious.⁶ According to this view of realization, (1) realizers of more abstract selection formulae are produced by the logical operation of conjunction, (2) it is solely the ‘formal character’ of natural selection theory that accounts for the fact that natural selection is multiple realizable, and (3) what unifies all of the realizers of a single realized property is their formal, not *causal*, structure (see Matthen and Ariew 2002, 74–77, 71–72, 81). If this account of realization were the only one available, then the realized properties of a biological population of organisms would be merely metrical or ‘stochastic’. Thus, the *only* causally efficacious properties would be the basic, unrealized properties at the level of individual organisms (if any; perhaps such properties would only be found at the fundamental physical level).

However, there is a notion of realization that is not purely formal and which does not lead to the conclusion that selection is not a causal process. Roughly, according to this kind of realization, a property X of entity o is realized by properties Y_1, \dots, Y_n of o 's parts and their environment if and only if X bestows causal powers to o in virtue of the fact that Y_1, \dots, Y_n bestow causal powers to o 's parts, and not vice versa (see Gillett 2002).⁷ I think that it is at best a programmatic sketch of a complete account of realization. But the details of such an account will not matter for the debate about selection or for answering the grounding question. Alternatively, one could adopt something like Shoemaker's account (2003) of microrealization, except that in the biological case the realizers will not be states of affairs involving microphysical entities propertied and related in certain ways but rather states of affairs involving individual organisms (and their environment) propertied and related in certain ways.

6. Or at least aren't part of causal processes. See n. 17.

7. Mohan Matthen (personal communication) claims that this account of realization is included as a special case of the formal realization relation. It's true that the formal 'hierarchical' model of realization can *model* every instance of causal realization, i.e., for every instance of causal realization there is some instance of formal realization that is isomorphic to it. However, this isomorphism does not preserve causal structure. The formal image of causally realized property is not itself causally efficacious. (E.g., a supercomputer may be able to perfectly model a tornado, but the simulation will not have all of the causal powers that the tornado does.)

Given this account of how properties realize other properties, one can provide a derivative account of how causal processes realize other causal processes (where a causal process is a sequence of causally related property instantiations). Suppose we have a causal process ψ in which the instantiation of property C is causally sufficient for the instantiation of property E . If ϕ is a causal process in which the instantiation of Y_1, \dots, Y_n is causally sufficient for the instantiation of Z_1, \dots, Z_n , then ϕ realizes _{p} ψ if and only if C is realized by Y_1, \dots, Y_n and E is realized by Z_1, \dots, Z_n (in the sense given in the previous paragraph). The claim that natural selection is a realized _{p} causal process is thus shorthand for the claim that the population-level properties, C and E , which are the cause and effect in the process of natural selection are realized by properties of and relations between the individuals in the population, and the properties/relations that realize C are causally sufficient for those that realize E . In what follows, I drop the subscript from process-realization.

If we can show that properties of biological populations are realized in this way by individual-level properties, then we will have provided a substantive answer to the grounding question. First, this account of realization clearly allows causal properties (and causal processes in which they are involved) to be realized by other causal properties (and processes). Second, these population-level properties are more abstract and more modally flexible than their realizers—they may be instantiated when their realizers are not (i.e., they may be multiply realizable). Hence, they are arguably not identical to their realizers.⁸ Third, and most importantly, realized population-level properties are compatible with the completeness of Darwinian biology. They do not compete with their realizers for causal sufficiency because the properties they involve are abstract ‘logical parts’ of the properties that realize them (in the same way that *being red* is part of *being scarlet* or *being translucent* is part of *being water*).⁹ In this way, they are not something *in addition to* the causal processes involving Darwinian events. Further, since they are not independent of the individual-level properties that realize them, they do not violate any exclusion principle that would rule out strong emergentism or dualism regarding selection. There are not two independent causal processes that result in

8. Some think that ‘local reductions’ are helpful here (e.g., Kim 1998). Although I can’t fully respond to this strategy here, I think that it results in missed objective biological generalizations (see Section 4 below).

9. Cf. Yablo’s claim (1992) that determinates and determinables don’t compete for causal sufficiency. I’m not certain that realizers and realized properties, in general, are related as determinates to determinables, but I am certain that they don’t compete for causal sufficiency because they are not independent of one another. Cf. Yablo (1992, 259).

a single outcome. That is, the exclusion principle, as stated above, is too strong. It should be reformulated so that it allows for cases in which simultaneous causal processes are related by the causal realization relation.

Of course, if natural selection is captured equally well by a formal realization relation, then all of this is neither here nor there. In the next section, I spell out in detail how the causal realization framework applies to a particular case of natural selection: the evolution of melanism in rock pocket mice. And I argue that a formal realization relation cannot completely capture biologically important features of this case.

4. Natural Selection for Melanism in Rock Pocket Mice. The rock pocket mouse, *Chaetodipus intermedius*, lives in rocky areas of the desert southwest. Most of their habitat consists of light-colored rocks, and the mice typically have light-colored pelage. However, on several dark lava flows separated from each other by hundreds of kilometers, the overwhelming majority of mice have dark fur. It is believed that melanism in these mice is an adaptation against predation. Owls are one of their major predators, and earlier experiments have shown that owls prey differentially on light and dark mice.

Coat color in mice, and in mammals generally, is determined by the amount of two forms of melanin: eumelanin, which is brown or black, and pheomelanin, which is yellow to dull red. The relative amount of these pigments is controlled by the interaction of the melanocortin-1 receptor (MC1R), a seven transmembrane G-protein-coupled receptor expressed at high levels in melanocytes, and the agouti signaling protein which is an antagonist of MC1R.

Nachman, Hoekstra, and D'Agostino (2003) found that four, nonsynonymous single-nucleotide polymorphisms (at nucleotide positions 52, 325, 478, and 699) were responsible for charge-changing amino acid polymorphisms (Arg-18-Cys, Arg-109-Trp, Arg-160-Trp, and Gln-233-His). These polymorphisms were found only in the population of dark mice at Pinacate, Arizona and were in complete linkage disequilibrium with one another.¹⁰ Of the 17 dark mice caught at this site, 11 were homozygous for the dominant, dark (*D*) allele, while six were heterozygous, resulting in an allele frequency of 28/34, or approximately 82%. Nachman, Hoekstra, and D'Agostino claim that one or more of these amino acid mutations is responsible for the melanic phenotype in the Pinacate population. Hence, they conclude that this locus is under disruptive natural selection across the two habitats (i.e., the dark lava flow and surrounding

10. That is, these four polymorphisms were always found together.

light rocky areas).¹¹ This conclusion is supported the fact that the perfect association found between phenotype and genotype is extremely unlikely in a panmictic population with no assortative mating unless the given gene (or a tightly linked gene) is responsible for the phenotype (Nachman, Hoekstra, and D'Agostino 2003, 5271).¹² Further, the possibility that a linked gene is responsible for the melanic trait is unlikely given (1) that no other genes in pigmentation pathways are linked to the Mc1r locus in humans or mice, and (2) that other amino acid substitutions in Mc1r have been shown to be responsible for melanic phenotypes in mice and other organisms, including geese, bananaquits, and dogs.

However, this evidence is not conclusive, for population structure (e.g., isolation by distance) could be responsible for the observed phenotypic variation. That is, the perfect association between melanic phenotype and the *D* allele at the Mc1r locus could be spurious; the Mc1r gene could not be involved in the production of the phenotype at all. Further, even if the Mc1r locus is responsible for the melanic phenotype, this does not imply that *selection* was responsible for the phenotypic and genotypic polymorphism. Although “the reduced variability seen among the dark Mc1r alleles [as compared to the light haplotypes] is the expected pattern if selection has recently fixed an adaptive substitution” (Nachman, Hoekstra, and D'Agostino 2003, 5272), “two standard tests of neutrality [the McDonald-Kreitman test and the HKA test] fail to detect selection on Mc1r in the Pinacate population, either in the total sample or in the subpopulations on light and dark substrate” (5273).

If demographic history or population structure caused the observed variation, one would expect all loci in the genome to have been equally affected. Thus, if population history or structure were responsible for the habitat-restricted variation, then variation at neutral loci should be expected to correlate with variation in phenotype. This hypothesis was tested by sequencing two neutral mitochondrial DNA (mtDNA) loci, COIII and ND3 and using these data to construct a phylogeny (Nachman, Hoekstra, and D'Agostino 2003). Haplotypes from light and dark mice were found to be haphazardly distributed on this phylogeny; there was no correlation between variation at these neutral sites and phenotypic variation, sup-

11. Disruptive natural selection occurs when individuals at phenotypic extremes are favored. Here, the light mice were favored on the light sand while the dark mice were favored on the lava flow.

12. Panmictic populations are those in which all individuals are potential partners. Assortative mating occurs when organisms are more likely to mate with phenotypically similar individuals than with ones that differ phenotypically.

porting the claim that selection, not population structure and history, was the cause of the spread of the dark allele and phenotype on the lava flow.¹³

Using a larger sample size, Hoekstra, Drumm, and Nachman (2004) confirmed the absence of population structure in the Pinacate population; there was no correlation between mtDNA haplotype and either geography or color phenotype (see also Hoekstra, Krenz, and Nachman 2005). Using estimates of migration rates and Mc1r allele frequencies and a simple model of migration-selection balance, Hoekstra, Drumm, and Nachman (2004) estimated selection coefficients of 0.013–0.390 against light mice/alleles on dark rocks and of 0.0002–0.0200 against dark mice/alleles on light rocks. “Together these results suggest that natural selection acts to match pocket mouse coat color to substrate color, despite high levels of gene flow between light and melanic populations [at Pinacate]” (Hoekstra, Drumm, and Nachman 2004, 1329).¹⁴

Several population-level causal processes are occurring in the Pinacate population: selection against the *d* allele on the dark lava flow, somewhat weaker selection against the *D* allele on the light rocks surrounding the lava, and asymmetric migration/gene flow between the subpopulations on the lava flow and on the surrounding light rock. These processes are distinct, and gene flow is opposing selection in this population. If selection were not occurring (or were weaker), then the population would likely become more homogenous; if gene flow were not occurring (or were weaker), then the differences between the subpopulations in the two habitats would likely become greater.

Further, the properties involved in these processes are realized by different properties of and relations between individual mice and their environment. In selection, the property *variation in camouflage ability (X)* of the population of mice (which, in turn, in part realizes *variation in the mice's abilities to survive/reproduce*) is causally efficacious. This property is realized by the coat colors of the individual mice in that population, together with their relation to the substrate on which they live, their relations to other mice on that substrate (with respect to coat color/reflectance), and their relations to predators like owls and other mammals (Y_1, \dots, Y_n). By contrast, other properties of the individual mice and their environment will realize the properties involved in the process of gene flow, such as the speed and agility of individual mice, their ‘attrac-

13. In these alternative scenarios, the traits of the mice would be causally irrelevant to the habitat-specific variation. Thus, on one interpretation, the variation would be produced by (a variety of) the process of drift.

14. Note that a number of strong assumptions go into the model used, some of which may not be biologically realistic.

tiveness' to potential mates, and the availability of food resources on the different substrates.

Assuming that the camouflage ability of individual mice can be quantified (relative to a substrate), the variance of a randomly selected sample of mice will be an estimate of *variation in camouflage ability* within the population. However, neither this estimate, nor the mathematical property of the true population variance, can be identified with *variation in camouflage ability* in the mice population. To do so would be to conflate a mathematical population with a biological one. The same variance could be possessed by a spatiotemporally and causally diffuse set of mice, but in such a case, it would not contribute the causal powers that *variation in camouflage ability* contributes to the Pinacate population, e.g., the power to adapt to a dark substrate. (See the discussion of temperature below.) *Variation in camouflage ability* contributes this power to the population of mice in virtue of the fact that the light coat color of some mice in the population enables them to blend into the sand while the dark coat color of other mice enables them to blend into the lava. Similar claims will be true for any other powers contributed by *variation in camouflage ability*. Thus, *variation in camouflage ability* is causally realized by the individual-level properties and relations in the mice population.

These sketches of the realization relations that obtain in the mice population reveal that the formal account of realization is unable to capture the phenomenon of natural selection. In order for natural selection to be realized, it is crucial that the population of mice be spatiotemporally and causally unified. Since the relation between stochastic realized entities and their realizers is purely formal or mathematical, it leaves out these interactions between mice, predators, and their environment which are required for the realization of natural selection.

To see this more clearly, it may help to consider a somewhat simpler example. A particular sample of water can be thought of as a *causally unified* population of molecules and ions. Population-level properties of the sample, like translucence, electrical conductance, temperature, and viscosity, are realized by properties of and relations between the molecules and ions. But the formal account of realization is unable to distinguish between these causally efficacious population-level properties and mere 'stochastic' population-level properties. For instance, the same average kinetic energy, and thus the same 'temperature', could be *formally* realized by a sample of water in the glass on my table and by a motley assortment of H₂O molecules scattered across the solar system.

5. Causal Explanation and Natural Selection. According to the causal view of realization, a number of properties are simultaneously realized in the Pinacate population: *variation in genotype*, *variation in coat color*,

variation in camouflage, and *variation in ability to survive/reproduce*. If the discussion in Sections 2 and 3 is correct, these properties do not compete for causal sufficiency since each is causally realized by the properties of and relations between individual mice (and their environment). However, even if this discussion provides an answer the metaphysical issue raised by the grounding question, these population-level properties arguably still compete in the context of causal explanations (cf. Yablo 1992, 273ff.). That is, if the above discussion is correct, it shows that an exclusion principle regarding causal processes (or causal sufficiency) is too strong. However, one might still accept an exclusion principle regarding causal *explanation*, to the effect that there cannot be two or more independent causal explanations of a given evolutionary outcome. Given such a principle, which of these properties should be used in explanations cited in a biologically general theory of natural selection? In the last section, I left this unclear, as have previous defenses of the population-level causal view.¹⁵

In fact, there is a powerful dilemmatic argument the upshot of which is that, no matter what properties appear in it, a suitably *general* natural selection theory cannot offer *causal* explanations. On the one hand, fairly specific, fine-grained population-level properties either will not be sufficient for selection or will miss objective biological similarities between populations. On the other hand, abstract, coarse-grained properties will be either too causally heterogeneous (too disjunctive) or too ‘causally thin’, and hence will be unprojectible and useless for empirical, causal explanations.

I think that the first horn is incontrovertible. For instance, in three other melanic populations of mice living on dark lava flows in New Mexico, none of the four amino acid substitutions in Mc1r were observed (Hoekstra and Nachman 2003). Further, the polymorphisms that were observed at the Mc1r locus in these populations were not associated with coat color. This suggests that: “Adaptive melanism has evolved independently in different populations of *Chaetopidus intermedius* and has done so through changes at different genes” (Hoekstra and Nachman 2003, 1191). If these other mutations had subsequently occurred in a few mice in a population that was already fixed for the *D* allele at the Mc1r locus, there would be no further selection for melanism because the mice would

15. E.g., Millstein claims that “heritable differences in physical survival and/or reproductive abilities” (2006, 631) are the cause in the process of natural selection. However, she also claims that both “variation in genotypes” and “variation in the abilities of genotypes” are the “causal engine of selection” (645) and that “we have selection when heritable differences in physical characteristics (yielding differences in survival and/or reproductive abilities) are causally relevant to differences in reproductive success” (645).

be (virtually) indistinguishable to predators. Thus, not just any specific *variation in genotype* is causally sufficient for selection.¹⁶

Further, the same is true of *variation in phenotype*. Suppose that the *D* allele has pleiotropic effects so that all mice with this allele are slightly more sensitive to pain and suppose that this trait has no consequences for survival and reproduction. Then, *variation in sensitivity to pain* will be perfectly correlated with *variation in coat color*, but although there will be selection of mice who are sensitive to pain there will not be selection for sensitivity to pain (see Sober 1984). Thus, not just any heritable *variation in phenotype* is causally sufficient for selection.

Finally, such properties like *variation in genotype* and *variation in phenotype* are unable to capture certain biological similarities. Consider the rock pocket mice again. Hoekstra and Nachman (2003) showed that there were no significant differences between rock color among four lava flows in New Mexico and Arizona. And, although there were minor phenotypic differences between populations—melanic mice in the New Mexico populations were slightly, but significantly, darker than those in the Arizona population—it is clear that similar selection pressures existed in these different populations (Hoekstra and Nachman 2003). Despite differences in causal processes occurring in these populations at the genetic level, it is biologically important to recognize that the *same type* of relatively abstract, causal natural selection process acted in each population. *Variation in camouflage ability* (which must include relations to the environment and not just the coat color of mice) caused changes in phenotypic and genotype frequencies over time. Finer-grained population-level properties (like *variation in genotype*) are unable to capture this important biological similarity. Further, this similarity with respect to selection processes is distinct from *other* causal similarities that might obtain between populations: e.g., similar amounts of gene flow, similar population structure/history.

For similar reasons, one should not claim that the various selection pressures acting on a population are identical to an instance of the causal process of natural selection. Rather, the abstract causal selection process is realized by the various selection pressures acting on different populations. As Walsh, Lewens, and Ariew (2002, 467) argue, in what I call the *identity* argument, identifying selection pressures such as predation and sunlight exposure with natural selection would lead to the untenable conclusion that natural selection is identical to drift (or that selection constitutes drift), since in certain cases these selection pressures lead to drift.

16. Cf. Millstein: “It does not matter if a particular individual has a superb ability to survive and reproduce—there will be no selection at all if other individuals who are physically different have the same ability” (2006, 644, emphasis in original).

Thus, on this proposal the causal interpretation allegedly “short-circuits itself”; there is no “causal (or explanatory) work for drift to do” (Walsh, Lewens, and Ariew 2002, 467).

Using the mice example again, we can avoid this conclusion by claiming that what one might call the ‘maximal realizer’ (i.e., the individual-level properties of all of the mice in the population and the entire habitat) simultaneously *realizes* both selection and drift. (Just as the microstructure of a sample of water simultaneously realizes its transparency, conductivity and viscosity.) Realized processes, like selection and drift, are different ways of abstracting away from the details in the maximal realizer by focusing on different *causally* relevant aspects of that realizer (heritable variation in a trait that is causally relevant reproductive success, and, perhaps, population size, respectively).

These considerations raise a problem for Millstein’s (2006) defense of the population-level causal view. In her argument, she focuses on the causal basis of a particular selection pressure (and corresponding adaptation). Hence, it is not clear that her discussion of how biologists established the ‘causal basis of selection’ at a certain locus completely succeeds in showing that natural selection is a causal process. Rather, it seems that it instead shows *how* differences in willow leaf beetles’ ability to withstand extreme temperatures are *realized* (partially) in a polymorphism at the PGI locus. That is, it spells out the details of a particular instance of the general realization framework presented here. I do not, of course, object to Millstein’s naturalistic use of real biological examples to make a metaphysical point. And I agree that the physiological work done to uncover the ‘underlying mechanism’ supports the ecological claim that the ability to withstand extreme temperatures is the relevant adaptation for these beetles. However, this does not show how natural selection as a population-level causal process is related to the individual-level causal processes that occur (the account of realization given here is needed to do that). Without providing this—without answering the grounding question—the population-level causal view is undermined. Further, Millstein’s discussion does it tell us whether the specific *variation in ability to withstand extreme temperatures* or the more abstract *variation in ability to survive/reproduce* is the property that should be cited in a causal, selectional explanation in this case.

However, if it is the more abstract property of *heritable variation in the ability to survive and/or reproduce* that is employed in natural selection explanations, then the second horn of the dilemma outlined above comes into play. For, variation in the *ability to survive and/or reproduce* looks suspiciously similar to variation in *fitness*. Elliott Sober (1984) has argued that fitness is not causally explanatory because it is a (heterogeneous) disjunctive property—that is, because fitness is multiply realizable and

there is no causal process that unites all of the various realizers. Different realizers of fitness produce differential reproductive success/failure via different causal processes. Even if one argues that *ability to survive and/or reproduce* is not identical to *fitness* (as suggested by some passages in Millstein [2006]), the same argument applies to the former as to the latter.

Another way of pushing this horn of the dilemma arises out of Matthen and Ariew's understanding of multiple realizability (2002). They take Lawrence Sklar's understanding (1999) of what makes quantities in statistical thermodynamics, like entropy and temperature, multiply realizable and suggest that it applies to the properties that appear in natural selection theory as well. According to this picture, it is the formal character of statistical thermodynamics that accounts for its 'universal applicability' to heterogeneous domains. A quantity like temperature is not a (sparse) physical property but "a way of characterizing ensembles" (Sklar 1999, 195). Statistical thermodynamics "mathematically demonstrates a phenomenon that occurs regardless of [the] physical details of the system" (Matthen and Ariew 2002, 80). Similarly, "natural selection is the consequence of heterogeneous processes in substrates *sharing only a formal structure*" (81, my emphasis).¹⁷

According to this view, natural selection theory does not provide causal explanations because the predicates that appear in it do not pick out sparse properties. Thus, because the theory of natural selection only captures formal unity, one should not expect all of the satisfiers of a selection formula to have any properties in common other than those that are

17. While discussing statistical thermodynamics, Matthen and Ariew suggest that natural selection is a kind of cause (a 'stochastic cause') that differs in important ways from genuine causal processes (in particular, fundamental physical processes). This is reminiscent of Skinner's claim that: "Darwin discovered the role of selection, a kind of causality very different from the push-pull mechanisms of science up to that time" (quoted in Hull, Langman, and Glenn 2001, 511) and also brings to mind Ned Hall's distinction between dependence and production. As Matthen and Ariew (2002) discuss, natural selection is apparently discontinuous, and changes in individual-level properties affect the population-level property of trait fitness instantaneously, thus violating a plausible locality condition on causal production. If natural selection were not a causal process, but merely a stochastic cause, a "running tally of lower level events" (Matthen and Ariew 2002, 79), then this would provide for another deflationary response to the grounding question. There would be no distinct causal process of natural selection to ground in processes involving individual-level events and properties.

Although a complete response to this suggestion is a subject for another paper, I think that the distinction between dependence and production is orthogonal to that between individual-level and population-level causation. Matthen and Ariew may have shown that natural selection is not a fundamental (physical) causal process, but no one ever claimed that it was. It is a causal process in the same sense that heating, perception, photosynthesis, erosion, and sexual discrimination are causal processes (for the two last examples, see Shapiro and Sober, forthcoming).

mathematically definitive of the natural selection formula. Any generalizations of natural selection theory are a priori or conceptual once an adequate selection formula is determined.¹⁸

6. Causal Processes, Disjunctive Properties, and Empirical Generalizations. In this section I argue that the second horn of the dilemma can be blunted. It is not *heritable variation in ability to survive and/or reproduce* that is distinctive of causal selection explanations, but the more specific property of *heritable variation in ability to survive/reproduce due to the interaction of a particular trait and the environment*.¹⁹ The focus on a particular trait ensures that this population-level property is not ‘wildly’ disjunctive at the level of biological individuals. It is the causally efficacious (population-level) property that biologists attempt to measure by calculating trait fitness. This is reflected in the fact that the simplest models in population genetics simply *assume* that all of the fitness differences are determined by a single locus. Also, when calculating trait fitness by averaging over the survival/reproductive success of all organisms with a given trait, one attempts to determine the mean ability or advantage bestowed by that single trait. That is, it is assumed that with a large enough sample the advantages and disadvantages (abilities and disabilities) bestowed by other traits will ‘average out’ yielding a normal distribution. Of course, this process can always be confounded by perfect correlations between traits (e.g., genetic hitchhiking). But in the absence of such problems, trait fitness will be a good measure in expectation of the ability conferred by that particular trait, which is a component of the organism’s overall ability

18. There are interesting similarities between Matthen and Ariew’s account of realization and the account of realization that is part of the Ramsey-Lewis framework for defining theoretical terms, especially as it is used by Kim (1998). According to this view of realization, multiply realized ‘properties’ (viz., property designators) are “causally and nomologically heterogeneous kinds, and this at bottom is the reason for their inductive unprojectibility and ineligibility as causes. . . . What lends unity to the talk of dormitivity and such is conceptual unity, not the unity of some underlying property” (Kim 1998, 110). One difference is that for Lewis at least, the ‘conceptual unity’ is bestowed by platitudes or folk theory, while in Matthen and Ariew’s framework it is a matter of nontrivial mathematics.

19. An anonymous referee remarked that the same proposal has been put forth as a solution to the ‘tautology problem’—whether the ‘principle of natural selection’ has empirical content. I agree with much of the motivation behind this solution. For instance, Brandon and Beatty claim that the propensity interpretation was put forth to break the ‘explanatory circularity’ of claims involving fitness—to develop an account of fitness that would allow for “genuine explanatory accounts of differential reproduction in terms of differential fitness” (1984, 343). However, in this paper, I remain neutral on the nature of fitness, whether the ‘tautology problem’ is really a problem, and whether the ‘propensity interpretation’ solves it.

to survive and reproduce. And of course, *variation* in this component will be a population-level property.

This view of things nicely fits the practice and assumptions of biologists. Given that “elucidating the causes of population divergence is a central goal of evolutionary biology” (Hoekstra, Krenz, and Nachman 2005, 217), a complete understanding of those causes requires that “we identify traits that are ecologically important and have some understanding of how these traits affect fitness in different environments” (Nachman, Hoekstra, D’Agostino 2003, 5268).

This suggestion also handles “multiply beneficial traits”—traits that cause survival and reproductive success (or failure) via two or more distinct causal processes—perhaps as tusks and feathers do (for scientific examples see the references in Glymour [1999]). Since such traits contribute to survival and reproductive success (or failure), they are part of a selection process. Further, while for some purposes it is obviously useful to distinguish between the distinct causal processes that lead to this outcome (e.g., defense against predators and attracting mates), there is no reason why there cannot be a more abstract causal process that incorporates each of these more fine-grained processes. In some cases, this kind of process will be hard to describe, but that does not make this an ad hoc maneuver. In fact, such an abstract, unifying process will be scientifically important whenever there are trade-offs between the various fine-grained processes. Even when each ‘component’ is beneficial, one will not always be able to find their net effect by adding the effects each would have in isolation. The more abstract causal process will include any interaction effects that obtain when the individual processes are co-instantiated.

It is, of course, an empirical matter what causal generalizations and processes are there to be found. For instance, there may be a causal process that is common to every case covered by the very general account of selection given by Hull, Langman, and Glenn (2001). But such a causal process would be so abstract that it is doubtful that there are any novel, empirical generalizations to discover about it. That is, it is doubtful that there is an empirical science that covers all of the instances of the selection process sketched by Hull, Langman, and Glenn. Of course, the debate at hand is concerned solely with Darwinian natural selection in populations of biological organisms (what Hull, Langman, and Glenn [2001] characterize as ‘gene-based selection’), and there is good reason to think that the trait-specific population-level properties that are involved in the causal process of selection are projectible. That they are is one reason why evolutionary biology is a unified subdiscipline. For instance, one can make predictions about how populations of different organisms will evolve based on abstract evolutionary similarities between them (e.g., if they are

affected by similar selection pressures), even when the organisms are phylogenetically and ecologically quite different. Instances of convergent evolution are examples of this.

Finally, the discussion in this paper assumes that causal processes do not involve only interactions or quantities from fundamental physics. In particular, it is assumed that there can be causal processes involving chemical and biological properties (cf. n. 7). Of course, this is compatible with all such processes in the actual world being ultimately realized by interactions between fundamental physical properties. (Indeed, it is plausible that physicalists are committed to every causal process being so realized.) However, spelling out exactly how that realization would go and giving a general account of causal processes are tasks for another occasion.²⁰

7. Conclusion. If natural selection is a population-level causal process, then it is a process whereby *variation in heritable ability to survive/reproduce due to the interaction of a particular trait and the environment* (e.g., *variation in camouflage ability* in the rock pocket mouse example) causes differential reproductive success. This causal process is compatible with the existence of causal processes at the level of individual biological organisms because the population-level properties involved in natural selection are causally realized by properties of and relations between individual organisms: *variation in camouflage ability* confers causal powers on the population of mice (e.g., the power to adapt to a dark substrate) in virtue of the fact that coat color confers causal powers on individual mice (together with relations to their environment). The formal account of realization proposed by defenders of a purely statistical interpretation of natural selection is unable to capture the phenomenon of natural selection and its relation to other evolutionary processes like gene flow and population subdivision. The causal realization framework captures these phenomena and also shows how natural selection theory is able to provide causal explanations and empirical generalizations and predictions.

REFERENCES

- Bouchard, Frédéric, and Alex Rosenberg (2004), "Fitness, Probability and the Principles of Natural Selection", *British Journal for the Philosophy of Science* 55: 693–712.
 Brandon, Robert, and John Beatty (1984), "The Propensity Interpretation of 'Fitness'—No Interpretation Is No Substitute", *Philosophy of Science* 51: 342–347.

20. Although I agree with much of what Schaffer (2003) says about causal processes, I don't think that his suggestion that "the best general understanding of a complete process is in terms of lawful subsumption across sequences of events" (34), at least if laws are taken to be exceptionless. I think a better strategy would appeal to generally stable mechanisms or perhaps Schaffer's other suggestion of 'intrinsically intact' sequences.

- Fodor, Jerry (1997), "Special Sciences: Still Autonomous after All These Years", *Philosophical Perspectives* 11: 149–163.
- Futuyma, Douglas J. (1986), *Evolutionary Biology*. Sunderland, MA: Sinauer Associates.
- Gillett, Carl (2002), "The Dimensions of Realization: A Critique of the Standard View", *Analysis* 62: 316–323.
- Glymour, Bruce (1999), "Population Level Causation and a Unified Theory of Natural Selection", *Biology and Philosophy* 14: 521–536.
- Hall, Ned (2004), "Two Concepts of Causation", in John Collins, Ned Hall, and L. A. Paul (eds.), *Causation and Counterfactuals*. Cambridge, MA: MIT Press, 225–276.
- Hoekstra, Hopi E., Kristen E. Drumm, and Michael W. Nachman (2004), "Ecological Genetics of Adaptive Color Polymorphism in Pocket Mice: Geographic Variation in Selected and Neutral Genes", *Evolution* 58: 1329–1341.
- Hoekstra, H. E., J. G. Krenz, and M. W. Nachman (2005), "Local Adaptation in the Rock Pocket Mouse (*Chaetodipus intermedius*): Natural Selection and Phylogenetic History of Populations", *Heredity* 94: 217–228.
- Hoekstra, H. E., and M. W. Nachman (2003), "Different Genes Underlie Adaptive Melanism in Different Populations of Rock Pocket Mice", *Molecular Ecology* 12: 1185–1194.
- Horgan, Terence (1993), "From Supervenience to Superdupervenience", *Mind* 102: 555–586.
- Hull, David L., Rodney E. Langman, and Sigrid S. Glenn (2001), "A General Account of Selection: Biology, Immunology, and Behavior", *Behavioral and Brain Sciences* 24: 511–573.
- Kim, Jaegwon (1989), "Mechanism, Purpose and Explanatory Exclusion", *Philosophical Perspectives* 3: 77–108.
- (1998), *Mind in a Physical World*. Cambridge, MA: MIT Press.
- Matthen, Mohan, and André Ariew (2002), "Two Ways of Thinking about Fitness and Natural Selection", *Journal of Philosophy* 99: 55–83.
- Millstein, Roberta L. (2006), "Natural Selection as a Population-Level Causal Process", *British Journal for the Philosophy of Science* 57: 627–653.
- Nachman, Michael W., Hopi E. Hoekstra, and Susan L. D'Agostino (2003), "The Genetic Basis of Adaptive Melanism in Pocket Mice", *Proceedings of the National Academy of Sciences* 100: 5268–5273.
- Reisman, Ken, and Patrick Forber (2005), "Manipulation and the Causes of Evolution", *Philosophy of Science* 72: 1113–1123.
- Rosenberg, Alex, and Frederic Bouchard (2005), "Matthen and Ariew's Obituary for Fitness: Reports of Its Death Have Been Greatly Exaggerated", *Biology and Philosophy* 20: 343–353.
- Schaffer, Jonathan (2003), "Overdetermining Causes", *Philosophical Studies* 114: 23–45.
- Shapiro, Larry, and Elliott Sober (forthcoming), "Epiphenomenalism—the Do's and the Don'ts", in G. Wolters and P. Machamer (eds), *Studies in Causality: Historical and Contemporary*. Pittsburgh: University of Pittsburgh Press.
- Shoemaker, Sydney (2003), "Realization, Micro-realization, and Coincidence", *Philosophy and Phenomenological Research* 67: 1–22.
- Sklar, Lawrence (1999), "The Reduction(?) of Thermodynamics to Statistical Mechanics", *Philosophical Studies* 95: 198–202.
- Sober, Elliott (1984), *The Nature of Selection*. Cambridge, MA: MIT Press.
- Stephens, Christopher (2004), "Selection, Drift, and the 'Forces' of Evolution", *Philosophy of Science* 71: 550–570.
- Walsh, D. M. (2000), "Chasing Shadows: Natural Selection and Adaptation", *Studies in the History and Philosophy of Biology and the Biomedical Sciences* 31: 135–153.
- (2004) "Bookkeeping or Metaphysics? The Units of Selection Debate", *Synthese* 138: 337–361.
- Walsh, D. M., Tim Lewens, and André Ariew (2002), "The Trials of Life: Natural Selection and Random Drift", *Philosophy of Science* 69: 452–473.
- Wilson, Jessica (2005), "Supervenience-Based Formulations of Physicalism", *Nous* 39: 426–459.
- Yablo, Steven (1992), "Mental Causation", *Philosophical Review* 101: 245–280.