Causal Composition in Evolutionary Theory

Charles H. Pence

Preprint of August 16, 2020

Abstract

I argue here that the debate over the causal status of natural selection, genetic drift, and fitness has foundered on an unrecognized issue in the metaphysics of science. Many parts of this debate rely, tacitly or explicitly, on a solution to what I will call "the causal composition question" – defining the circumstances in which micro-level causes compose to produce a macro-level causal process. Further, evolutionary biology constitutes a particularly poor environment in which to study this general question regarding causation. I thus claim that this is a place where we must foster engagement between philosophers of biology and metaphysicians of science.

One of the most heated debates in the literature on the foundations of evolutionary theory over the last fifteen years has concerned whether natural selection and genetic drift are *causal processes.* This issue, in turn, has touched on several other core features of evolutionary theory. Are we right to describe selection, drift, mutation, migration, and so on as analogous to Newtonian forces, each of which drives a population in a given direction? If we are, which of these should be considered part of the "inertial" state (analogous to Newton's first law), and which should be considered "special" forces (analogous to Newtonian gravitation and described by the second law)? If selection is causal, is fitness causal as well? Or is fitness merely a non-causal property of organisms, tallied for biological convenience?

Broadly, two positions in this debate have solidified. The "causalist" picture, canonically stated by Sober (1984), considers selection and drift to be causal processes, which in many cases can profitably be compared with Newtonian forces. The "statisticalist" interpretation, which was inaugurated by Walsh et al. (2002) and Matthen and Ariew (2002), disagrees. On this interpretation, selection, drift, and fitness are all non-causal. Rather, they are simply useful statistical summaries of the genuinely causal events that occur at other (individual, genetic, biochemical) ontological levels.[1] The force analogy, the statisticalists argue, breaks down, particularly in the case of genetic drift.

A useful way to characterize these two positions in more detail is, in fact, to begin with what *isn't* at stake. Both sides grant an identical conception of the individual-level

1. I am cognizant of the trouble with ontological "levels" talk (see, e.g., Eronen, 2015, pp. 39–40 for a helpful survey of these worries). I don't, however, think the idea that natural selection, if it is a genuine causal process, is in some sense composed of causal events involving individual organisms or their biochemistry is especially controversial. I also must note that I am forced by considerations of length to smooth over a variety of differences within each of the statisticalist and causalist views, which are by no means monolithic.

causal account of the biological world. Individuals live and die, they are predated, they mutate, they give birth to other individuals. All of these are causal events, all are known to occur, and all are, in some sense or another that remains to be specified, "responsible for" the higher-level phenomena that we call natural selection and genetic drift.

This commonality is best illustrated by a brief comparison. Matthen and Ariew, from the statisticalist camp, state explicitly that at the individual level, "we do not want to debate … whether differences in fitness-relevant traits actually played a role in evolution" (Matthen and Ariew, 2009, p. 203). Shapiro and Sober, on the causalist side, hold that the appropriate supervenience base for natural selection includes some subset of "the causal processes that impinge on individual organisms" (Shapiro and Sober, 2007, p. 251). It is clear, then (modulo some superficial differences in terminology), that both sides suppose that happenings in the lives and deaths of individual organisms are causal, and that these in some way result in evolutionary change.

Second, it is granted that natural selection and genetic drift (and, hence, evolutionary change in general) are phenomena that demand probabilistic explanations, and they are expressible mathematically as a set of statistical trends that exist within populations of organisms. This is a foundational part of the statisticalist picture. Of the causalists, Hodge's (1987) position perhaps makes this most clear – selection and drift *just are* varieties of sampling for Hodge, and this sampling is best described by the statistical formulation of traditional population genetics.

This, it would seem, is a fairly robust picture of the biological world. We have agreement on both the underlying causes that (in some way or another) constitute the evolutionary process, and on the broad structure of the statistical framework that must be used to describe population change over time. What, then, remains at stake, after this description of the world is agreed upon?

I will consider here one significant issue which persists – and argue that it forms a major part of the intractability of this debate.[2] The causalist/statisticalist debate depends on a resolution to the following general worry: in which circumstances do the organism- and biochemical-level causes agreed upon by both sides *compose* to produce causal processes of natural selection, genetic drift, and so on, active at "higher" ontological levels?[3] This, in turn, is just a specification of a general question that arises in the metaphysics of science. When do causes at a lower level combine in such a way as to produce a causal process active at a higher level? I will call this the *causal composition question* (or CCQ).[4]

Before we begin, I should quickly justify why I take this to be a novel and important contribution to the debate. First, as we will see below, the metaphysics of science is not frequently discussed directly in the literature in the philosophy of biology – these questions are normally approached only obliquely, via generalization from biological cases. Second, while there is some literature that has come close

2. One further question must, unfortunately, fall outside the scope of this paper: differences over the very definitions of natural selection and genetic drift themselves. I hope to explore this in future work.

3. This is not to say that the causalist/statisticalist debate *reduces* to one over causal composition: this would mean that there are no peculiarities at all introduced by the biological facts of the matter, a claim that it will be clear I reject in the following.

4. It has also been termed an instance of "epiphenomenalism" by Shapiro and Sober (2007), about which more later.

to discussing CCQ, albeit in different terms (about which more later), considering this question explicitly will, I argue, give us the opportunity to cultivate a novel and productive relationship between the philosophy of biology and the metaphysics of science. To start evaluting the role of CCQ, then, we should consider how the participants in this debate define the more fundamental notion of a causal process.

1. Identifying Causal Processes

A number of authors have offered accounts of how we might identify causal processes – or, to put the question more clearly, how we might differentiate genuinely causal processes from mere sequences of events that do not cohere sufficiently to qualify as a causal process. Salmon (1984, 1994) and Dowe (1992, 1995) take the notion of 'process' to be an undefined primitive, and separate causal processes from non-causal processes by considering the transmission and conservation of quantities through these processes over time. Chakravartty (2007, p. 108) offers a more robustly realist account, on which causal processes consist of the continuous interaction of manifesting causal dispositional properties.

These views have, to some degree, been explicitly considered in the philosophy of biology. Millstein (2013), for example, has endorsed the Salmon account. Matthen and Ariew, on the other hand, seem to have a more stringent criterion in mind, on which natural selection fails to be a causal process because "it is not a process in which the earlier events cause the latter" (2002, p. 79). Although this view is minimally spelled out, Matthen has argued elsewhere (2011) that a causal process must have either "an overarching cause that is responsible for each event in the sequence, and the order in which they occur," or be such that each event in the sequence is causally responsible for future events in the sequence.[5] (Whether this is partial causal responsibility or complete causal responsibility is left unspecified.)

It is clear that this distinction makes a significant difference. Consider a particularly nebulous process, such as the urbanization of twentieth-century America. On a more minimal account like that of Salmon or Dowe, this may well qualify as a causal process, if we can establish the entities involved and the quantities preserved across time. On the more demanding definition of Matthen and Ariew, however, this may well not be a causal process, if it is the result of a constellation of diffuse causes, none of which can integrate the various states in a sufficiently tightly-knit manner.

When we apply this worry about causal processes to the biological world in particular, this vagueness in definition remains problematic. A lower threshold for what counts as a causal process may well incline one to be a causalist (as more processes, perhaps with selection and drift among them, will qualify as causal), while a higher threshold may support a statisticalist position.

One might worry that this lack of clarity on the definition of a causal process threatens the entire discussion of CCQ that I hope to pursue here. Without a stable and sufficiently shared definition of a causal process, there may not even be a sense in

5. Matthen has suggested to me (pers. comm.) that the definition of causal process intended in the earlier Matthen and Ariew (2002) is Salmon's, but it is not clear that Salmon's view is consistent with the quote above. More recent statisticalist work (Walsh et al., 2017) has argued that whatever notion of causal process is in play, the analytic, statistical inference from trait fitnesses to population change is incompatible with it. This latter claim rests, as well, on unspecified features of the analysis of causal processes.

which we can have a debate over whether or not selection and drift are in fact causal. I want to respond to this objection in two different ways. First, part of the debate about whether or not events qualify as a casual process just *is* an argument about CCQ. The example that I briefly sketched above concerning urbanization is in part a question of whether or not a higher-level process of urbanization emerges from lower-level events (movements of people, construction of cities, etc.). In that sense, debate over the definition of a causal process feeds into debate about CCQ, and vice versa.

Second, insofar as I don't have the space (and, as I will argue below, appropriate theoretical resources) here to resolve CCQ, the kind of interplay that I will consider below between the metaphysics of science and the philosophy of biology will offer us equally powerful ways in which we might help resolve the question of the nature of causal processes. My choice to emphasize CCQ rather than this definitional concern arises, in part, because CCQ is broadly unconsidered in the philosophy of biology literature, and will remain even after this definitional question is resolved.

In what follows, then, I will take a loosely defined, common-sense notion of a causal process rather than adopting a more formal notion like that of Salmon and Dowe. As it turns out, this will suffice for introducing and elucidating the stakes of CCQ.

2. Multi-Level Causal Systems

This much should already lead us to think that issues in the metaphysics of science have a very significant impact upon the causalist/statisticalist debate. But we can make the worry here significantly more precise. When we look at many of the specific arguments offered by both sides, we will see that they primarily concern not the specifics of the case of drift and selection, but the circumstances, in general, in which a set of lower-level causal events constitutes a higher-level causal process, which I have termed the *causal composition question.*[6] Let's consider three such arguments in turn.

Individual-level causal selection is a category mistake. Several causalist readings of natural selection place the causal action at the level of individual organisms (e.g., Hodge, 1987; Bouchard and Rosenberg, 2004). Walsh et al. (2002) offer an argument to the effect that this move is a category mistake. Selection, they argue, is a sorting process. And sorting "is not a force that causes a coin to fall head or tail" – rather, the sorting outcome "is explained and predicted by appeal to some statistical property [of the sequence], an *average* of individual propensities" (Walsh et al., 2002, p. 463). We therefore commit a category mistake if we attempt to "point to one or other of these individual-level causal processes and say *this* is a cause of the error and *this* one is a cause of sorting" (Walsh et al., 2002, p. 465).

The causal interpretation fails to understand population subdivisions. Consider a coin-tossing example in which two coins are tossed 50 times each (Walsh, 2007, p. 293).[7] The tosses are performed ten at a time by two experimenters – Walsh provides

6. I will consider arguments grounded more directly in biological facts below.

7. It is interesting to note (in connection with my claim below that these arguments are non-biological) that while several initial expressions of this argument were phrased in biological terms (Walsh et al., 2002; Walsh, 2004; Pigliucci and Kaplan, 2006), the most recent (and clearest) version comes from

simulated data for such an experiment. Now, Walsh writes, "[t]here are at least three different, equally legitimate ways to describe this process" (2007, p. 293): (1) a single series of coin tosses, 49 heads and 51 tails; (2) two series of 50 tosses, 20/30 and 29/21; and (3) 10 series of 10 tosses, ranging from 5/5 to 7/3. If drift is just the analogue of sampling error in this experiment, then, Walsh argues, the causal interpretation "is committed to the claim that drift-the-cause is strong in the aggregate of 10 sequences of 10 tosses," as well as the claim that "[i]n the single sequence of 100 tosses, however, drift is not very strong at all. But these are not two populations; *they are different ways of describing the same population*" (Walsh, 2007, p. 296, original emphasis). Thus a causal notion of drift is both strong and not-strong in the same population, a supposed contradiction.

Northcott (2010, pp. 459–460) provides a causalist rebuttal to this argument. He has us imagine 100 slaves pushing a large rock in service of the pharaoh. The pharaoh is impressed by uniformity of effort and angered by lack of uniformity, so the higher the variance of the pushes applied by the laborers, the more the pharaoh will become angry. We may plausibly say, in such a case, that the variance of the 100 pushes *causes* the pharaoh's anger. This cause has exactly the same strange property with respect to subdivisions as genetic drift did in the example described above – the pharaoh should be angrier (on the average, at least) at smaller groups of workers than he should be at the group as a whole. But this doesn't necessarily entail that the variance of the group as a whole somehow *fails to cause* the pharaoh's anger.

Selection is a spontaneous statistical tendency, not a force. Matthen and Ariew (2009) lay out and critique a particular picture of the causal interpretation, one which they claim derives from Sober (1984). On this view, natural selection exists as an intervening step in the causal diagram of any evolutionary process – variation in traits causes selection, which in turn causes the births and deaths of animals responsible for evolutionary change. Such a "tertium quid," they argue, is unnecessary. Consider a series of tosses of a coin biased toward heads. They claim that "[t]he bias implies that it is probable, in a series of tosses of this coin, that heads will come up more often than tails. No process of 'toss selection' is needed for this result"(Matthen and Ariew, 2009, p. 206). The addition of a causally potent natural selection (on this reading of the causal interpretation) is therefore superfluous.

What's the point in discussing these three arguments? One feature of them, I think, stands out – or should. None of them are phrased with any reference to actual biological instances of drift and selection. In this and many other arguments relating to the causal efficacy of selection and drift, we see apple carts (Walsh, 2007), Newtonian gravitation, centers of mass (Matthen and Ariew, 2009), pharaoh's laborers (Northcott, 2010), scatter plots, smoking, heart disease, painkillers (Walsh, 2010), race cars (Ramsey, 2016), and (many, many) coin flips.

In and of itself, of course, this is not necessarily a problem. Philosophers often draw analogies with various kinds of generalized cases in order to bring important considerations or morals extracted from them back to the original question at issue. In this case, however, I believe something more problematic is in fact taking place. For all of the varied analogies that are deployed, attempting to harness intuitions arising in a number of other domains, the debate over causation in evolutionary theory has

Walsh (2007), in which he uses the example of a coin-tossing experiment.

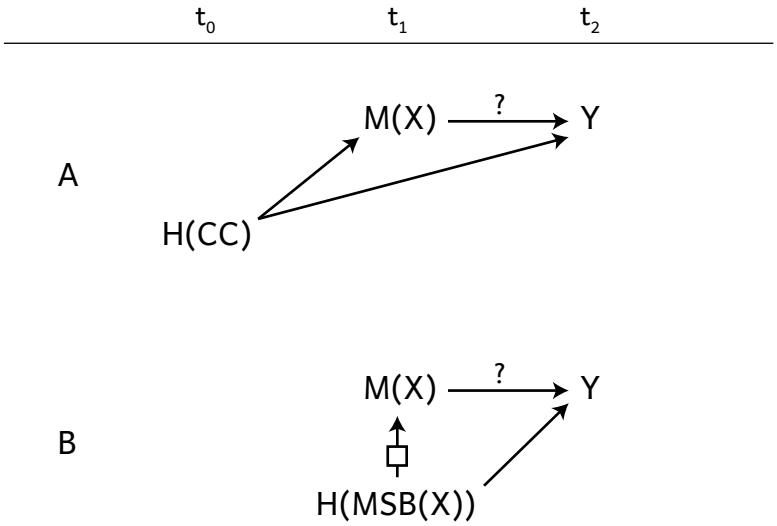

Figure 1: Two ways to test for the presence of a cause. $H(X)$ represents holding condition X fixed, $M(X)$ represents manipulating X. CC are the common causes of X and the later event Y, and $MSB(X)$ is the micro-level supervenience base of the macro-level phenomenon X. The boxed arrow holds with nomological necessity, and the question-marked arrow is the causal relation to be studied. After figure 5 from Shapiro and Sober (2007).

remained persistently intractable. And the targets of these various analogies seem to be clustered around a particular, general question in the metaphysics of multi-level causation, rather than a problem in the philosophy of biology. The debate, that is, seems to have been circling around a central question in the metaphysics of science without, in general, having engaged with metaphysics directly.

The clearest connection between the biological debate and a broader question about causation comes from Shapiro and Sober (2007). They consider both the causalist/statisticalist debate and that over epiphenomenalism of mental properties in the philosophy of mind – arguing that the question of mental epiphenomenalism and CCQ are two sides of the same coin. In essence, they claim, both of these boil down to a dispute over how to appropriately perform a manipulation to test whether or not a macro-level cause is present (see figure 1).

In both cases, we want to know whether or not manipulating the macro-level event X (at time t_1) will result in a change in the macro-level event Y (at time t_2). The question, Shapiro and Sober argue, is what we hold fixed in order to bring this manipulation about in the appropriate way. Two obvious possibilities present themselves. First, we could hold fixed the prior (at some earlier time t_0) common causes of X and Y (called CC in the figure) – this is "option A." Alternatively, we could hold fixed the micro-level supervenience base of X (called $MSB(X)$ in the figure) – this is "option B." They argue that adopting option B amounts to considering something nomologically impossible: considering whether or not Y would occur if a nomologically-connected sufficient condition for X occurred (namely, $MSB(X)$), but X did not. Option A is thus the correct choice. The statisticalists, they hope to show, have argued on the basis of option-B analyses of selection and drift, while causalists have argued (correctly)

from option-A analyses.

Whatever the merits of this particular argument about manipulability turn out to be,[8] Sober and Shapiro have correctly identified what *type* of case we are dealing with. CCQ is a general concern within the domain of the metaphysics of science, to be analyzed using our best theories about causation and supervenience in multi-level systems.

It is important to note that I do not intend to claim that there is a sharp dividing line between questions that are metaphysical and those that lie within the domain of the philosophy of science. (I am not even certain that either category can be clearly defined.) Rather, the contrast here is one of methods and general approach. When we ask for a detailed conceptual analysis of a concept arising in evolutionary theory, we turn, in general, to those who work in the philosophy of biology. When we ask for the expertise relevant to resolving general questions about relationships of multi-level causation, on the other hand, we turn, in general, to those who work in metaphysics – as, I will argue, we should do more often in the case at issue.

2.1. Resolving the Debate Without CCQ

One might object here that a number of authors have claimed to have offered complete, causal reconstructions of an approach to this debate, which have not involved directly addressing CCQ. In particular, these authors have attempted to present a causalist picture that operates only at the population level, thus rendering the question of composition from the individual level superfluous. If that's true, then one might argue that CCQ is at best an interesting but unnecessary side-note to the broader causalist/statisticalist debate.[9]

I want to close this section, then, by briefly sketching the view of Reisman and Forber, as a representative example. I will argue that while one might consider a portion of the debate resolved by population-only interpretations, the question of the relationship between the individual and population levels will remain, and as long as it does, a response of some sort to CCQ will be required.

In short, Reisman and Forber (2005) argue that a manipulationist view of causation, when applied to population-level evolutionary phenomena, will – at least for certain kinds of examples, which they describe in some detail – give us the possibility of clean "interventions," in Woodward's sense (2003). The existence of such interventions – systematically alter the founder effect, as Dobzhansky and Pavlovsky did in a real population, and you will control the strength of drift, which in turn will alter future population composition – indicates that genetic drift is a causal process (and, by analogy, natural selection is as well). Indeed, one strength of these manipulations is taken to be that they never "attempted to identify how selection and drift cause individual births or deaths" (Reisman and Forber, 2005, p. 1119).

The question is whether this solely population-level story can really remain so under increased scrutiny. I don't think it can. Consider their approach to natural selection. Imagine, they write, that we were to "change the fitness structure … such

8. In particular, one might be worried that the analogy between epiphenomenalism in philosophy of mind and selection and drift is defective. If the philosophy of mind example concerns whether micro-level determinates instantiate (or not) macro-level determinables, rather than causation, then it, unlike the biological case, may well have an *a priori* answer. Thanks to Anjan Chakravartty and Mohan Matthen for pointing out this objection.

9. I thank an anonymous reviewer for pushing this objection.

that one homozygous type had the greatest fitness" (Reisman and Forber, 2005, p. 1119). How is this type fitness to be understood? Reisman and Forber don't say. Traditionally, type fitnesses in causal interpretations of evolutionary theory are taken to be derived from individual fitnesses, which in turn estimate the ability of individual organisms to survive and reproduce in particular contexts. Intervening on population fitnesses without intervening on these individual abilities to survive and reproduce is a tall order; if this is what Reisman and Forber really mean here, we need an account of how it is possible. Otherwise, interventions on natural selection will implicate, whether we like it or not, changes at the individual level – and thus we need a story about how those individual changes are related to population changes, which will entail understanding how those individual-level and population-level causes are related.

In short, given that evolving populations are composed of individuals, and it is individuals that engage in classic behaviors like eating, mating, and dying, I find it unlikely that we will be able to construct a "free-floating," population-level understanding of the causal structure for evolution that thereby avoids embroiling itself in CCQ.

3. Analyzing Causal Composition

The mere fact that we have identified a metaphysical question within a debate in the philosophy of biology is, in itself, rather unsurprising. Biological systems are causally complex, and hence we should expect that their structures could offer up novel challenges to general work on causation. The question we must turn to, then, is one of the relationship envisioned between CCQ and the question of the causal structure of evolution in particular. I can see three ways in which we might envision such a relationship. On the first, work in the philosophy of biology as it exists now could be used to inform the debate on CCQ in the metaphysics of science. Second, work in current metaphysics of science on CCQ could be used to directly inform the debate over causation in evolution. And third, we might envision a new kind of interaction, where the two fields work collaboratively to resolve both CCQ in general and the biological case in particular. I will argue for the third, after dispensing with the first two.

3.1. Applying Extant Biology to Metaphysics

We will not, I claim, be able to directly apply work in the philosophy of biology as it is currently found to the general problem of CCQ. Consider two possible ways in which we could do so. We could choose to invoke either complex biological case studies – "reading off" the metaphysics from the biology – or we could simplify them in a bid to get the biological cases to map more clearly onto metaphysical questions. Neither avenue is promising.

First, consider a turn toward biological practice, extracting the correct answers to our metaphysical worries from the use of biological theories in real-world examples from the scientific literature. One strand within the debate over the causal status of evolutionary theory has in fact attempted to work in this manner, remaining biologically informed and working out the way in which selection and drift are distinguished (and, it is argued, used *as causes*) in evolutionary biology. Most prominent here is the work of Millstein, who has used examples from the montane willow beetle (Millstein,

2006), land snails (Millstein, 2008), and the early development of the theory of genetic drift by Fisher and Wright (Millstein et al., 2009). The effort to bring the debate over causation back to particular biological contexts is long overdue. But when applied to CCQ, however, I am skeptical of its chances for success. Millstein herself is all too cognizant of the limitations of this approach when the questions at issue are as abstract as these. She notes that her examples can only show that viewing "natural selection [as] a causal process *is consistent with at least some biological practice*" (Millstein, 2006, p. 637, emphasis added), or that "Wright and Fisher had good reason to think that there were biological processes in nature that *needed to be represented in drift models* in order to better track evolutionary changes in a population" (Millstein et al., 2009, p. 8, emphasis added). In both cases, we see a clear recognition of the difficulty of making metaphysical inferences from the practice of working biologists, whether by limiting the scope of the inference or speaking merely about the representation of factors in modeling.

And, I claim, there is a well-grounded philosophical reason for the existence of this difficulty. Several authors, perhaps most persuasively Waters (2011), have made the point that many practicing biologists are what we might call "toolbox theorists." As he puts it:

> Perhaps, given the messiness of the world, the ideal theory turns out to be more like a toolbox than a fundamental framework. Perhaps what is sometimes thought to be a single concept, such as fitness, consists of a family of concepts, each useful for theorizing about different aspects, parts, or scales of entangled causal processes. (Waters, 2011, p. 240)

Setting aside the accuracy of this view as a broader prescription for science or philosophy as a whole, we have ample evidence that practicing biologists do often view their theoretical enterprise in this way. Picking up Waters's invocation of fitness, for example, the biologist de Jong has reviewed the use of 'fitness' (in only the restricted numerical sense familiar from population genetics) in the biological literature, and covers dozens of fitness measures which can be categorized in several ways, and which feature in several different models of the selective process. In her conclusion, however, she asks the following provocative question:

> What is fitness supposed to measure? Some sort of overall performance or quality of design or aptness for life, or general adaptedness? *This seems to be a discussion that is standing outside the practice of evolutionary biology.* What we usually ask is how such measures … are interrelated, and how they relate to the change in phenotypic traits. Given knowledge of the life history of the population, the causes and values of the fitness components, *these relations can be spelled out in specific models. We need not ask whether expected time to extinction or genotypic weight is the proper fitness measure.* What we would like to know is how they are related in a mechanistic model for a specified situation. (de Jong, 1994, p. 18, emphasis added)

I have quoted extensively from de Jong because her survey provides a particularly clear demonstration of toolbox theorizing. The question of the "correct" model of fitness, she argues, lies entirely outside the province of evolutionary biology. Rather, we take a particular system, with a particular model which specifies the components of fitness of interest. This model and system, then, provide a context in which we study

the relationship of the various fitness measures we find in our toolbox. If this is how practicing biologists view fundamental theoretical concepts like fitness, selection, and drift, then it is a mistake to turn to the biological literature in an attempt to resolve metaphysical questions such as these factors' status as causal processes. There are simply no fundamental concepts there to analyze.

Perhaps, then, this is not the direction to look in order to apply biological concerns to CCQ. Consider, rather, how we might reduce the biological cases to their simpler fundamentals. This is, indeed, common practice as well within the debate. Philosophers of biology who work at the general or foundational level at which the causalist/statisticalist debate is situated often – of necessity – employ particularly simple models of selection and drift. In the context of these austere biological models, we can most clearly partition the influences of selection and drift, forming tractable thought experiments on which the contributions of the various components of the evolutionary process can be considered.

This effort, however, suffers from a significant problem of its own. It's not clear that we really are still exploring – even approximate – biological cases when we simplify the biological world to this extent. The precise interplay of selection, drift, mutation, and migration required to offer a plausible explanation of even a relatively uncomplicated Mendelian trait like blood type in humans is incredibly hard to describe in detail.[10] As another example, the standard recurrence equations used to provide predictions for how trait fitness affects a population when selection alone is at work are derived in the context of Hardy-Weinberg equilibrium. But populations at Hardy-Weinberg equilibrium are highly idiosyncratic – in fact, the standard response to evolutionary stasis is to invoke powerful stabilizing selection, as noted by McShea and Brandon (2010, p. 120). Simplifying the biological case enough that it could bring insights to CCQ therefore runs the risk of removing from it any relevant biological details – all the while without acknowledging that there is a general metaphysical problem at issue in the first place.

Thus, whether we attempt to do so by way of complicating or simplifying the biological examples here, it seems as though there is no straightforward way in which the questions raised in the study of the causal structure of evolutionary theory can be made to speak directly to CCQ.

The problem can be made all the more acute when we consider that it seems as though there are areas throughout biology that seem to successfully "work around" CCQ – that is, they offer us understanding and control of multi-level causal systems, whether we have a clear philosophical account of CCQ or not. Neuroscience is an important example (and hence the connection that Shapiro and Sober (2007) make with the philosophy of mind), but there are others, including in areas as diverse as the role of random search in foraging theory (Glymour, 2001). The prospects for creating a generalized theory that begins with biological practice in areas tat are this profoundly disparate also seem grim.

3.2. Applying Extant Metaphysics to Biology

What about the other direction? Perhaps we can approach our best theories of contemporary causation and "read off" the correct solution to the biological case from

10. I thank Mark Jordan for the example as well as for pushing me to consider the impact of biological complexity on the debate.

them?

In general, the case here is not much better. There is, indeed, one small branch of the literature that has considered questions similar to CCQ, arising from work in causal graph theory, which builds on the classic work of Spirtes et al. (2001). As Spirtes and Scheines note, the application of this formalism is rendered problematic when "two or more variables are deterministic functions of one another, not deliberately, but because of redundant measurements, or underlying lawlike connections" (2004, p. 834). Manipulating a biological population would be, in their sense, an "ambiguous manipulation" – because we have no way of knowing which of the multiply realizable individual-level structures is responsible for giving rise to the population variables that we see before us. Unfortunately they do not give us a general algorithm for resolving these questions, and the proofs available in their work do not seem to include any that map onto the case of evolutionary theory.

Glymour takes up the same theory to approach the question of mental causation, or, more precisely, "time series of macroscopic quantities that are aggregates of microscopic quantities" (2007, p. 330). But the key move with respect to CCQ is performed rather by picking out the causal *roles* that macro-level variables play:

> The causal role of a system of macroscopic properties is the conditional independence graph, or diagram, of an aggregation of microscopic properties, together with the values of any causally relevant parameters; each macroscopic property is an unknown function of the collection of microscopic properties. (Glymour, 2007, p. 342).

Of course, this does not mean that just any graph we might draw connecting macroscopic and microscopic properties will automatically describe a causal role for those macroscopic properties. As Glymour puts it, "not every way of aggregating to form macroscopic variables will yield screening off relations that reflect a causal structure; indeed, most ways will not" Glymour (2007, p. 342). This approach thus might be able to offer us a way in which we could resolve CCQ, by examining in which cases these macro-level causal roles actually exist.

The details, however, are murky, and I lack the space to pursue such a reconstruction here. For the moment, I will note only that Otsuka has worked extensively on the effort to interpret evolutionary theory in terms of such directed causal graphs (Otsuka, 2016a,b, 2019), with the caveat that this approach has yet to include the connections between the individual and population levels that would be necessary in order for Glymour's understanding of macroscopic causal roles to apply. Much more work remains to be done here.

To sum up, then, there is surprisingly little literature for us to read off of in the first place, excepting some results in causal graph theory (and a literature on the composition of Newtonian forces, about which more in a moment). It will thus be a challenge to link extant metaphysical work with CCQ in the evolutionary context.

But further, and forming perhaps a more significant worry, we find the inverse of the problem dealt with at the end of the last subsection. It is so difficult to apply contemporary theories of causation to an arena as complex as evolution that it is unlikely, as a result, that any genuine progress on CCQ could be made by appealing to a few "intuitive" biological examples. It is not "intuitively clear," for example, even in these simple cases, whether natural selection satisfies the criteria of a manipulationist, counterfactual, or a mechanism-based account of causation, nor should we expect it

to be.[11] And further, intuition can be an unsteady guide in many evolutionary cases. As mentioned above in the contexts of Hardy-Weinberg equilibrium and blood type, cases which seem simple on one axis (population stasis or simple Mendelian heredity, respectively) can frequently be profoundly complex along other axes, and complex in ways that produce a very opaque causal structure, making our attempts to apply an approach to causation perilous.

3.3. A New Form of Engagement

Readers would be forgiven at this point for despairing of the possibility of ever connecting results concerning a general question in metaphysics to the particular details of a case in a complex special science.[12] But I believe there is a model to which we may turn, by which metaphysical and practical questions can engage in a profitable, dialectical growth, a reciprocal engagement which has thus far been entirely absent from the literature in the philosophy of biology.

As an instance of this possible method, consider the discussion of Newtonian forces. In taking only a few samples, we can find a full gamut of papers, ranging from involvement directly with the science to comparatively "pure" metaphysics. We see direct discussion of the history of the force concept (Jammer, 1957), or the way in which we might flesh out that concept in great technical detail (Earman and Friedman, 1973). But we also see conceptual analyses at a more general level (Wilson, 2007), along with efforts to categorize forces at a fairly high level of abstraction (Ellis, 1963). We then move toward the metaphysics of forces themselves (Massin, 2009), whether or not certain additive forces can be said to really exist (Creary, 1981), or even Wilson's connection of the force literature to Horgan's "superdupervenience" (Wilson, 2002). Note as well that the publication dates of these articles are effectively random as we move across this spectrum – we have here a genuine case where issues are simultaneously being worked out at all levels of explanatory grain, often by the same authors, and often in dialog with one another.

Such an approach, when applied to the case here, would push us to consider a class of questions – the kinds of generalizations that take up the center of the spectrum in the force literature – which have been more or less ignored in the debate so far in the philosophy of biology. For example, where else might the causal composition question crop up in other sciences? Are there ways in which theories of causation need to be modified or extended, given that we have a genuine need of an answer to CCQ in the biological world?

Of course, any particular debate might become unbalanced in either one way or the other when such a relationship is at issue. This debate, I have claimed, has suffered in virtue of having effectively no connections to the metaphysics of science. Another arena (one commentator has mentioned the debate over the concept of species as a possible candidate) might suffer on the contrary from too much metaphysics. The intent, rather, is balance: engagement such as this will cease to be fruitful should it become too one-sided.

11. For example, see Ariew et al. (2015) for an account of the debate over the applicability of a manipulationist definition of causation in biology, or Pérez-González and Luque (2019) for worries about the use of mechanism.

12. Readers who do not believe such connections are possible or fruitful in the first place are encouraged to stop reading the paper now.

If such an outcome can be avoided, a reciprocal movement between the fine details of scientific cases and metaphysical reasoning has the potential, at least, to offer new perspectives of precisely the sort that might reinvigorate a debate otherwise desperately in need of a path forward.

4. Objections and Connections

Before I conclude, I should pause to raise one important objection, and draw a number of connections between the program that I propose here and the work of a number of other authors.

Other mentions of CCQ. I am not, of course, the only author to have noted the existence of CCQ in this literature. In a paper which teases apart the debate over the causal nature of genetic drift into a number of further questions, Plutynski argues that one such component of the debate is "the metaphysical question: is drift (and, for that matter, natural selection) an epiphenomenal, or a genuinely supervenient, 'population-level' causal process?" (Plutynski, 2007, pp. 161–162). She notes the existence of debate over the metaphysics of causation, but does not (as it lies outside her scope) deeply engage with any literature on the question.

For their part, the statisticalists assert that the composition question is unsolvable, without further argument. Matthen and Ariew, for example, write that "there is no straightforward translation between levels with regard to causal statements" (2002, p. 63). Matthen notes that an element of the debate concerns "intuitions about the compositionality of causes" (2009, p. 465), but does not follow up on the details of anything like CCQ. There has thus been a general understanding that CCQ is a relevant question within this literature, but no real progress toward its solution.

Connecting causation and probability. That said, several of the authors involved in this debate have also been working in neighboring areas of the metaphysics of science, particularly on the metaphysics of probability, in ways that could perhaps bear interestingly on CCQ. Following upon the growing recognition that there is a consistent way to think about objective probability at the macroscopic level, even in a deterministic world (Ismael, 2009; Lyon, 2011; Ismael, 2011), several authors – most importantly Abrams (2012; 2015) and Strevens (2011; 2013; 2016) – have offered compelling pictures of the ways in which causal structures at one level of organization can give rise to repeatable, objective probabilities at higher levels of organization.

Such work is not, of course, directly targeted at the resolution of CCQ. The aim, rather, is to offer a clear theoretical and metaphysical basis for the inference to objective macro-scale probability, which seems to occur in circumstances as diverse as games of chance, biological populations, and complex systems theory. But it does seem that this is the trend within the current literature most applicable to the kinds of general questions that I raise here. Placing this research in dialogue with more contemporary work on causation might serve as a way in which to advance the debate.

Metaphysics or explanation? Finally, one major objection should be considered. A number of recent articles, particularly from the statisticalist tradition, have attempted to shift the focus of the debate. Rather than being a question of the metaphysical structure of evolutionary theory, we might think about the debate as being instead over

the nature of evolutionary explanations. Matthen (2009), for instance, has described a theory of "statistically abstractive explanation," on which evolutionary theory deploys a specific and different sort of scientific explanation, by abstracting away from the details of particular populations. Ariew et al. (2015) have taken a similar approach, as has Walsh (2013).

On this basis, one might well argue for the rejection of my entire approach here. The objection would go something like the following. The metaphysics of evolutionary theory is not complex in the slightest. There are individual organisms (setting aside, of course, the difficulty of questions about biological individuality), which bear traits (setting aside, of course, the difficulty of questions about delineating and distinguishing traits), and which form populations that change over time. To the extent that causes impinge on these populations, they are just the everyday sorts of causes that lead to death and mating and eating and so forth. What is in fact interesting is how we offer explanations for that population change. What we therefore need is (at the most general) work in general philosophy of science on the structure of explanations in statistical theories, or even (at the least general) simply analyses of the in-practice explanations that biologists offer for population change.

I would respond to this critique in two ways. First, it is not clear that, in fact, these papers that claim to shift the focus to explanation have indeed done so, and other works which have been published since have returned the focus to the underlying metaphysical structure of evolutionary theory. When Matthen puts forth his account of statistically abstractive explanation, this is supposed to be an explanation which is only applicable in some sorts of circumstances – in other words, the system to be explained must be structured in such a way that "the reference classes display variability that is traced to the omission of theoretically inadmissible features" (Matthen, 2009, p. 477), where these inadmissible features are those that are are deemed irrelevant to the theoretically contextualized study of the phenomena at issue. But when he turns to the justification of this model, he argues from the causal structure of a particular sort of coin-flipping processes (about which more in a moment; Matthen, 2009, pp. 483–484), which seems to return our focus to metaphysics.

Further, in their most recent "manifesto" description of the statisticalist position, Walsh et al. (2017) are quite equivocal on this question. On the one hand, they often frame their view in terms of, for instance, the difference between "explaining changes in trait distribution and explaining the ecological causes of population change" (Walsh et al., 2017, p. 5). On the other hand, the explicit statement of their position comes in terms of properties of models of natural selection, a decidedly less explanation-focused and more metaphysical approach.

As a second avenue of response, even if a shift to explanation was fully executed by some of the relevant authors, I do not believe that it actually avoids the necessity of engaging with the kinds of questions that I consider here.[13] Matthen is undoubtedly correct that some kinds of causal influences are not taken to be theoretically relevant by practicing population geneticists. But what could be the grounds for a justified exclusion of such claims from evolutionary explanations? What might give rise to the different approaches that authors have taken to analyzing these explanations, with Matthen arguing that they exclude the individual-level causes of drift or selection, while Millstein (2008) argues that a different case study foregrounds such causal links? It seems inescapable that the justification for such exclusions, particularly if one (as

13. I thank Daniel Swaim for encouraging me to think about this line of response.

Matthen does) wants to draw an entirely general claim about all evolving populations from them, will be grounded in the causal structure of the evolving systems at issue – hence, the return of CCQ.

To put the point differently, and perhaps more provocatively, one may propose any number of accounts of scientific explanation, and these may more or less well track the nature of biological practice. But if one intends to argue in turn not only that these explanations are present in practice, but that their use is justified, one must have a sufficiently clear grasp of the metaphysics of the system being explained to support such a justification. Such a grasp will require engagement at the sort of general level that I have discussed here.

5. Conclusions

I hope, therefore, to have argued for the following. An important and underappreciated facet of the debate over the causal status of natural selection and genetic drift is its reliance on a question from the metaphysics of science: when do causes at one level compose to produce a causal process at a higher level? We see this most clearly from the highly generalized examples often used in discussion of the causalist/statisticalist debate. This would not, in and of itself, constitute a problem, were it not that – primarily due to the distance between our understanding of CCQ in abstract cases and the real biological world – selection and drift constitute a particularly poor environment in which to study CCQ, and vice versa. Following on the literature on classical mechanics, I then pointed to a model of interaction between philosophers of science and metaphysical concerns that is, I claim, a possible way out of the morass in which the debate finds itself.

What, then, does this mean overall? I hope that this constitutes an argument for increased engagement between philosophers of biology and metaphysicians of science.[14] The debate over the fundamental causal structure of evolution can provide just one of what I suspect are manifold fruitful places for philosophers of biology to interact with the work being done in the metaphysics of science. If my arguments here are correct, our prospects for understanding that structure are otherwise rather bleak.

Acknowledgments

This project has had a particularly long development time. Many thanks to those who read and offered comments on its various ancestors, including several anonymous reviewers, along with Anjan Chakravartty, Mohan Matthen, Roberta Millstein, Elliott Sober, Daniel Swaim, and Grant Ramsey. Special thanks to Mohan Matthen for a helpful discussion about his views on the definition of causal processes, as well as access to some unpublished material. Thanks to a group at the Aarhus Center for Science Studies for reading a draft version, especially Samuel Schindler, Line Edslev Andersen and Pierre Saint-Germier. Thanks as well to audiences at VUB, especially Steffen Ducheyne and Erik Weber; at the KU Leuven Center for Logic and Philosophy of Science, especially Bendik Aaby, Stijn Conix, and Hugh Desmond; at PhiloSTEM-3, especially Mark Jordan, Timothy Fuller, and Bernd Buldt; and at the Notre Dame

14. The story of the estrangement of the philosophy of biology from metaphysics, in stark contrast with the philosophy of physics, deserves to be told, though I am unqualified to tell it.

Dissertation Writing Seminar, especially Dan Hicks, Robert Audi, and Martin Sticker.

Bibliography

Abrams, M. 2012. *Mechanistic Probability*. **Synthese**, 187(2):343–375. doi: 10.1007/s11229-010-9830-3.

———. 2015. *Probability and Manipulation: Evolution and Simulation in Applied Population Genetics*. **Erkenntnis**, 80(S3):519–549. doi: 10.1007/s10670-015-9784-4.

Ariew, A., C. Rice, and Y. Rohwer. 2015. *Autonomous-Statistical Explanations and Natural Selection*. **British Journal for the Philosophy of Science**, 66(3):635–658. doi: 10.1093/bjps/axt054.

Bouchard, F. and A. Rosenberg. 2004. *Fitness, Probability and the Principles of Natural Selection*. **British Journal for the Philosophy of Science**, 55(4):693–712. doi: 10.1093/bjps/55.4.693.

Chakravartty, A. 2007. *A Metaphysics for Scientific Realism: Knowing the Unobservable*. Cambridge University Press, Cambridge.

Creary, L.G. 1981. *Causal Explanation and the Reality of Natural Component Forces*. **Pacific Philosophical Quarterly**, 62:148–157.

de Jong, G. 1994. *The Fitness of Fitness Concepts and the Description of Natural Selection*. **Quarterly Review of Biology**, 69(1):3–29. doi: 10.1086/418431.

Dowe, P. 1992. *Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory*. **Philosophy of Science**, 59(2):195–216. doi: 10.1086/289662.

———. 1995. *Causality and Conserved Quantities: A Reply to Salmon*. **Philosophy of Science**, 62(2):321–333. doi: 10.1086/289859.

Earman, J. and M. Friedman. 1973. *The Meaning and Status of Newton's Law of Inertia and the Nature of Gravitational Forces*. **Philosophy of Science**, 40(3): 329–359.

Ellis, B. 1963. *Universal and Differential Forces*. **British Journal for the Philosophy of Science**, 14(55):177–194. doi: 10.1093/bjps/XIV.55.177.

Eronen, M.I. 2015. *Levels of Organization: A Deflationary Account*. **Biology & Philosophy**, 30(1):39–58. doi: 10.1007/s10539-014-9461-z.

Glymour, B. 2001. *Selection, Indeterminism, and Evolutionary Theory*. **Philosophy of Science**, 68(4):518–535. doi: 10.1086/392940.

Glymour, C. 2007. *When Is a Brain like the Planet?* **Philosophy of Science**, 74(3): 330–347. doi: 10.1086/521968.

Hodge, M.J.S. 1987. *Natural Selection as a Causal, Empirical, and Probabilistic Theory*. In Krüger, L., G. Gigerenzer, and M.S. Morgan, editors, **The Probabilistic Revolution, Volume 2: Ideas in the Sciences**, pages 233–270. Bradford Books, Cambridge, MA.

Ismael, J. 2009. *Probability in Deterministic Physics*. **Journal of Philosophy**, 106(2): 89–108.

———. 2011. *A Modest Proposal about Chance*. **Journal of Philosophy**, 108(8): 416–442.

Jammer, M. 1957. *Concepts of Force: A Study in the Foundations of Dynamics*. Harvard University Press, Cambridge, MA.

Lyon, A. 2011. *Deterministic Probability: Neither Chance nor Credence*. **Synthese**, 182: 413–432. doi: 10.1007/s11229-010-9750-2.

Massin, O. 2009. *The Metaphysics of Forces*. **Dialectica**, 63(4):555–589. doi: 10.1111/j.1746-8361.2009.01213.x.

Matthen, M. 2009. *Drift and "Statistically Abstractive Explanation"*. **Philosophy of Science**, 76(4):464–487. doi: 10.1086/648063.

———. 2011. *Commentary on Grant Ramsey, 'Drift as an Individual-Level Process'*.

Matthen, M. and A. Ariew. 2002. *Two Ways of Thinking about Fitness and Natural Selection*. **Journal of Philosophy**, 99(2):55–83.

———. 2009. *Selection and Causation*. **Philosophy of Science**, 76(2):201–224. doi: 10.1086/648102.

McShea, D.W. and R.N. Brandon. 2010. *Biology's First Law: The Tendency for Diversity and Complexity to Increase in Evolutionary Systems*. University of Chicago Press, Chicago.

Millstein, R.L. 2006. *Natural Selection as a Population-Level Causal Process*. **British Journal for the Philosophy of Science**, 57(4):627–653. doi: 10.1093/bjps/axl025.

———. 2008. *Distinguishing Drift and Selection Empirically: "The Great Snail Debate" of the 1950s*. **Journal of the History of Biology**, 41(2):339–367. doi: 10.1007/s10739-007-9145-5.

———. 2013. *Natural Selection and Causal Productivity*. In Chao, H.K., S.T. Chen, and R.L. Millstein, editors, **Mechanism and Causality in Biology and Economics**, pages 147–163. Springer, New York.

Millstein, R.L., R.A. Skipper, and M.R. Dietrich. 2009. *(Mis)Interpreting Mathematical Models: Drift as a Physical Process*. **Philosophy and Theory in Biology**, 1: e002.

Northcott, R. 2010. *Walsh on Causes and Evolution*. **Philosophy of Science**, 77(3): 457–467. doi: 10.1086/652965.

Otsuka, J. 2016a. *Causal Foundations of Evolutionary Genetics*. **British Journal for the Philosophy of Science**, 67(1):247–269. doi: 10.1093/bjps/axu039.

———. 2016b. *A Critical Review of the Statisticalist Debate*. **Biology & Philosophy**, 31(4):459–482. doi: 10.1007/s10539-016-9528-0.

———. 2019. *The Role of Mathematics in Evolutionary Theory*. Cambridge University Press, Cambridge.

Pérez-González, S. and V.J. Luque. 2019. *Evolutionary Causes as Mechanisms: A Critical Analysis*. **History and Philosophy of the Life Sciences**, 41(2):13. doi: 10.1007/s40656-019-0250-5.

Pigliucci, M. and J.M. Kaplan. 2006. *Making Sense of Evolution: The Conceptual Foundations of Evolutionary Theory*. University of Chicago Press, Chicago.

Plutynski, A. 2007. *Drift: A Historical and Conceptual Overview*. **Biological Theory**, 2(2):156–167. doi: 10.1162/biot.2007.2.2.156.

Ramsey, G. 2016. *The Causal Structure of Evolutionary Theory*. **Australasian Journal of Philosophy**, 94(3):421–434. doi: 10.1080/00048402.2015.1111398.

Reisman, K. and P. Forber. 2005. *Manipulation and the Causes of Evolution*. **Philosophy of Science**, 72:1113–1123. doi: 10.1086/508120.

Salmon, W.C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton University Press, Princeton, NJ.

———. 1994. *Causality without Counterfactuals*. **Philosophy of Science**, 61(2): 297–312. doi: 10.1086/289801.

Shapiro, L. and E. Sober. 2007. *Epiphenomenalism – the Dos and the Don'ts*. In Wolters, G. and P. Machamer, editors, **Thinking about Causes: From Greek Philosophy to Modern Physics**, pages 235–264. University of Pittsburgh Press, Pittsburgh, PA.

Sober, E. 1984. *The Nature of Selection*. The MIT Press, Cambridge, MA.

Spirtes, P. and R. Scheines. 2004. *Causal Inference of Ambiguous Manipulations*. **Philosophy of Science**, 71(5):833–845. doi: 10.1086/425058.

Spirtes, P., C. Glymour, and R. Scheines. 2001. *Causation, Prediction, and Search*. The MIT Press, Cambridge, MA.

Strevens, M. 2011. *Probability out of Determinism*. In Beisbart, C. and S. Hartmann, editors, **Probabilities in Physics**, pages 339–364. Oxford University Press, Oxford.

———. 2013. *Tychomancy: Inferring Probability from Causal Structure*. Harvard University Press, Cambridge, MA.

———. 2016. *The Reference Class Problem in Evolutionary Biology: Distinguishing Selection from Drift*. In Ramsey, G. and C.H. Pence, editors, **Chance in Evolution**, pages 145–175. University of Chicago Press, Chicago.

Walsh, D.M. 2004. *Bookkeeping or Metaphysics? The Units of Selection Debate*. **Synthese**, 138(3):337–361. doi: 10.1023/B:SYNT.0000016426.73707.92.

———. 2007. *The Pomp of Superfluous Causes: The Interpretation of Evolutionary Theory*. **Philosophy of Science**, 74(3):281–303. doi: 10.1086/520777.

———. 2010. *Not a Sure Thing: Fitness, Probability, and Causation*. **Philosophy of Science**, 77(2):147–171. doi: 10.1086/651320.

———. 2013. *Descriptions and Models: Some Responses to Abrams*. **Studies in History and Philosophy of Biological and Biomedical Sciences**, 44:302–308. doi: 10.1016/j.shpsc.2013.06.004.

Walsh, D.M., T. Lewens, and A. Ariew. 2002. *The Trials of Life: Natural Selection and Random Drift*. **Philosophy of Science**, 69(3):429–446. doi: 10.1086/342454.

Walsh, D.M., A. Ariew, and M. Matthen. 2017. *Four Pillars of Statisticalism*. **Philosophy, Theory, and Practice in Biology**, 9:1. doi: 10.3998/ptb.6959004.0009.001.

Waters, C.K. 2011. *Okasha's Unintended Argument for Toolbox Theorizing*. **Philosophy and Phenomenological Research**, 82(1):232–240. doi: 10.1111/j.1933-1592.2010.00472.x.

Wilson, J. 2007. *Newtonian Forces*. **British Journal for the Philosophy of Science**, 58:173–205. doi: 10.1093/bjps/axm004.

Wilson, J.M. 2002. *Causal Powers, Forces, and Superdupervenience*. **Grazer Philosophische Studien**, 63:53–77.

Woodward, J. and C. Hitchcock. 2003. *Explanatory Generalizations, Part I: A Counterfactual Account*. **Noûs**, 37(1):1–24. doi: 10.1111/1468-0068.00426.