

Productivity, relevance and natural selection

Stuart Glennan

Received: 23 May 2008 / Accepted: 5 September 2008 / Published online: 30 October 2008
© Springer Science+Business Media B.V. 2008

Abstract Recent papers by a number of philosophers have been concerned with the question of whether natural selection is a causal process, and if it is, whether the causes of selection are properties of individuals or properties of populations. I shall argue that much confusion in this debate arises because of a failure to distinguish between causal productivity and causal relevance. Causal productivity is a relation that holds between events connected via continuous causal processes, while causal relevance is a relationship that can hold between a variety of different kinds of facts and the events that counterfactually depend upon them. I shall argue that the productive character of natural selection derives from the aggregation of individual processes in which organisms live, reproduce and die. At the same time, a causal explanation of the distribution of traits will necessarily appeal both to causally relevant properties of individuals and to causally relevant properties that exist only at the level of the population.

Keywords Causation · Population causation · Statisticalism · Productivity · Relevance · Natural selection

Introduction

It is a commonplace among biologists that natural selection is the principle cause of evolutionary change, but a series of recent articles by philosophers of biology have raised doubts about whether or in what sense this claim is true. Roberta Millstein (2006) has suggested that we can understand these questions by seeing them as addressing two questions:

S. Glennan (✉)
Department of Philosophy & Religion, Butler University, 4600 Sunset Avenue,
46208 Indianapolis, IN, USA
e-mail: sglennan@butler.edu

- Is natural selection a causal process or is it a purely statistical phenomenon?
- Is natural selection at the population level or at the level of individuals?

There are, as Millstein suggests, three positions that philosophers of biology have defended with respect to these questions.

1. Natural selection is a causal process at the level of individuals.
2. Natural selection is a purely statistical phenomenon at the population level.
3. Natural selection is a causal process at the population level.

For purposes of this paper, I shall refer to these positions as the *individualistic* interpretation, the *statistical* interpretation and the *dynamical* interpretations of natural selection.¹

While advocates of the individualistic and statistical interpretation have their disagreements, they make common cause against the dynamical interpretation's defense of population level causation. Statisticalists are keen to demonstrate the statistical character of evolutionary theory and explanations while individualists hope to show that the causal story behind selection (and in particular the distinction between drift and selection) must be understood by looking at relationships between individuals and their environment. Statisticalists, like Walsh, Ariew, Lewins and Matthen, would not deny the individualists' Bouchard and Rosenberg's contention that the "sampling events" involved in drift and selection involve causal processes, nor would Bouchard and Rosenberg deny Walsh et al's contention that the theory of natural selection is a statistical theory. Their views represent complementary attacks on population level causation. One argues that the real causal story goes on at the individual level while the other argues that the sort of explanation that occurs at the population level is a statistical explanation.²

The problem that Millstein as well as Shapiro and Sober see in these positions is that they appear to make natural selection epiphenomenal. These authors believe population thinking is crucial to understanding the process of natural selection. It is variation that matters, and variation is a population level property. Any view that suggests that population level properties are causally irrelevant, they reason, must be wrongheaded.

There is much to be said for the causal intuitions on both sides of this debate. To save those intuitions we should look for a middle ground. I think one can be found if we become more sophisticated in our understanding of just what causes are. In

¹ While these three interpretations are widely recognized, there are different ways in which philosophers have named and classified the varying interpretations of the theory of natural selection (and of evolutionary theory more generally). What I'm calling the individualistic interpretation has chiefly been argued for by Bouchard and Rosenberg. (Bouchard and Rosenberg 2004; Rosenberg and Bouchard 2005). The statistical interpretation has been championed in a number of articles by Ariew, Lewins, Matthen, and Walsh (Matthen and Ariew 2002; Walsh 2000; Walsh et al. 2002; Walsh 2007). The *locus classicus* of the dynamical interpretation is Sober 1984, and that interpretation has been recently defended by Millstein, Shapiro and Sober (Millstein 2006; Shapiro and Sober 2007) The term 'dynamical interpretation', which I borrow from (Walsh 2007), alludes to Sober's analogy between evolutionary forces and the forces of classical mechanics.

² Admittedly the individualists and statisticalists spend a lot of time disagreeing about what fitness is and what it means to say that natural selection is a statistical process. These differences are not, however, germane to the main issue of this paper.

particular, I suggest that our characterization of the causal structure of natural selection, and indeed of any other complex natural phenomenon, invokes two related but distinct kinds of causes. On the one hand, we identify the objects, processes and events that *produce* the phenomenon; and, on the other hand, we identify the various facts—which may include events, properties of those events, or background conditions—which are *causally relevant* to the production of the phenomena. Given the distinction between causal productivity and causal relevance, I will argue:

1. The causal processes that *produce* natural selection are at the level of individuals.
2. Natural selection depends upon population level properties that are causally relevant to evolutionary outcomes, but these properties do not *produce* those outcomes.

My paper will be in two parts. In the first part, I will give a general account of these two kinds of causes and of the relationship between them. In the second part I will see how these concepts apply to the process of evolution by natural selection.

Two kinds of causes³

Let's begin with the idea of productivity. Here are some examples of claims of causal productivity:

- The bowling ball knocked over the pin.
- The explosion made Edward deaf.
- The firing of neuron A caused the firing of neuron B.

Causal productivity is a relation between events. I don't want to get bogged down here in debates over the nature of events, but I'll take for granted what I hope is the fairly uncontroversial view that events typically involve an object doing something: a kettle boils, a bomb explodes, a hen lays an egg. While we sometimes speak of objects producing events as in "the bowling ball knocked over the pin", such talk is elliptical for events involving the object in question. In this case it was the bowling ball's striking the pin that caused the pin to be knocked over.

Production is the sort of causation that has dominated the discussion of mechanistic accounts of causation. In Glennan 1996, I argued that causally connected events require intervening mechanisms involving interacting objects. Machamer et al. (2000) introduced the concept of an activity as part of an attempt to capture what they called the "productive continuity of mechanisms". Activities of the parts of mechanisms as well as interactions between the parts of mechanisms are events which collectively *produce* a mechanism's behavior. Production is also the sort of causation under analysis in the theories of causal processes advanced by Wesley Salmon and Phil Dowe (Dowe 2000; Salmon 1984). Causal processes are

³ I have developed my account of these two concepts of cause in a recent paper (Glennan forthcoming), and the material in this section summarizes some material from that paper.

understood as continuous paths of objects through space-time that can interact when they intersect, producing changes in the properties of the objects that constitute those processes.

An important feature of productive causation is its locality. If one event produces another event, the events must be either spatiotemporally contiguous, or they must be connected by contiguous intermediate events and processes. When the bowling bowl knocks down the pin, it comes into direct contact with the pin. When the explosion causes Edward's deafness, it does not do so directly, but via a causal intermediate—the propagating blast wave. Relatedly, productivity is transitive. If one event produces a second, and the second a third, then the first produces a third. A good example of this is sequences of neuron firings, if neuron A's firing produces neuron B's firing and neuron B's firing produces neuron C's firing, we say also that neuron A's firing produces neuron C's firing.

While we often speak of one event being *the* cause of a second, there are many productive causes of any event—both because multiple causes may be necessary to produce an effect and because any given cause of an effect will have ancestors or successors in their causal chains which are themselves causes of the effect. Pragmatic considerations determine which causes we care about.

Let us turn to the second sort of causal relation—that of causal relevance or dependence. Causal relevance is a counterfactual relation of dependence between a fact f and an event e . A fact f is causally relevant to e iff either

- (1) e would not have occurred had it not been the case that f
- (2) e would have been more or less likely had it not been the case that f or, in cases where e and f can be represented by quantitative variables,
- (3) the value of e would have been different if the value of f had been different.

Causal relevance, as I have defined it, is closely connected with the sort of causal relation discussed in various counterfactual approaches to causation (e.g., Collins et al. 2004; Lewis 1973; Woodward 2003). The most notable modification is that while most counterfactual theories take the causal relata to be events, I am appealing to a relation between a fact and an event.⁴

All kinds of facts can be causally relevant to an effect. Let's consider some of them:

Case 1: f is the fact that an event occurred

The fact for instance that the bowling ball struck the pin was causally relevant to the pin falling down. More generally, events that produce an effect are also causally relevant to that event. But the following case shows the converse doesn't hold.

⁴ I do not think it is necessary for me here to commit to any particular theories about the nature either of facts or events. The main reason for distinguishing between them is that there are a lot more facts than there are events—as the examples I discuss will show. For more on the nature of events, facts and other possible candidates for causal relata see Schaffer 2003.

Case 2: f is the fact that an event did not occur

For instance, the fact that the Mom didn't turn off the hose was causally relevant to her basement flooding. This is a case of causation by omission—a sort of causation that provides a problem for causal productivity accounts. That problem is that Mom's failure to turn off the hose is not an event. It cannot be located in space and time so it cannot have productive causal influence. Nonetheless, events may depend counterfactually on omissions.

Case 3: f is the fact that an object, event, or background condition has certain properties or relations

Let's consider three examples:

- The fact that the key has a certain shape is causally relevant to whether it will open the door.
- The fact that the moth has coloration similar to the leaves on which it sits is causally relevant to its likelihood of survival.
- The fact that the wind is over 30 mph increases the likelihood that a serious fire will occur.

Here the facts are not simply facts about the occurrence or non-occurrence of an event, but facts about the properties of objects or background conditions that will or may play a causal role in the production of an effect. The shape of the key is a property of a key, and whether the lock turns when the key is put into it depends upon the shape of the key.

The examples considered so far are token causal relevance claims—claims about a particular fact being relevant to a particular event. Many relevance claims, however, are type level. For instance:

- Smoking causes cancer
- Absence of exercise increases the likelihood of heart disease.
- Smoking unfiltered cigarettes increases the amount of tar in the lungs.

As in token relevance claims, the facts cited as relevant may be of many kinds—facts about the occurrence or non-occurrence of types of events (smoking or not exercising), or about properties of events (like the number or kind of cigarettes smoked). While type causal relevance claims are important, I shall not discuss them further in this paper, because the sorts of causal claims at issue in discussions of the causal character of natural selection are token claims about the causes of evolutionary change within particular populations.

Full understanding of the causal basis of an event requires knowledge not only of the antecedent events and processes that produce the event, but also of relevant facts about the productive events, processes and background conditions. This point is nicely made by Chris Hitchcock in a well-known critique of Salmon's theory of causal processes (Hitchcock 1995). A chalked cue stick strikes a cue ball and the cue ball begins to move. The intersection of the stick and the ball is an interaction that produces the ball's movement—but what about the cue stick and its striking

accounts for the movement of the ball? Was it the fact that it was chalked or was it the fact that it had a certain linear momentum? It is the momentum of course, for the movement of the ball depends counterfactually upon momentum on the stick, but not upon its being chalked. Had the stick not been chalked, the ball would still have moved in the same manner, but had the stick moved with different momentum, the ball would have moved in a different manner.

In cases where we have knowledge both of (some of) the productive events and of the facts, we may formulate causal claims in this canonical form:

c causes *e* in virtue of *f*

where *c* and *e* are events and *f* is a fact expressing relevant properties of *c*, *e* or perhaps of background conditions. In this particular case we may say that (*c*) the cue stick's striking the ball causes (*e*) the ball's starting to move in virtue of the fact (*f*) that the cue stick had a certain momentum, the ball had a certain mass, and so forth.

But while complete knowledge of the causes of an event rely both on knowledge of the productive events and processes and of the relevant facts, it's possible to have one kind of knowledge without the other. For instance, if a man walks into the room and a child begins to cry, we may reasonably infer that the event of the man walking in produced the crying, but we don't know why this was the case. What facts about the man, the child, the manner of the man's entrance, and so forth were relevant to producing this reaction in the child? Perhaps the man has a beard and the child is afraid of men with beards because of a history of abuse by a bearded man. Perhaps the man is the child's father who has come home unexpectedly from the war, and the tears are tears of joy. Here we clearly have one kind of causal knowledge but not the other.

Alternatively, we may be able to establish the causal relevance of certain facts to an event or kind of event without knowing anything about the processes that produce that event. The use of controlled experiments and other sorts of manipulations of potentially causally relevant variables may be used to establish the relevance of these variables to classes of events, even though one does not understand the mechanisms that produce these events. Thus for instance we may establish that race or socio-economic status are causally relevant to performance on standardized tests without understanding the ways in which these variables are productively related to these effects.

The view that there are two kinds of causes is not original to me, and the distinction I draw bears important similarities both to the Ned Hall's (2004) distinction between production and dependence and to Jackson and Pettit's (1990) distinction between efficacy and relevance. More generally, the position is in keeping with an increasingly frequent move toward what has come to be called causal pluralism.⁵ What is peculiar to my view is the particular account of the relationship between productivity and relevance. Canonically, one event causes another in virtue of a certain set of causally relevant facts. And what makes a

⁵ See Hitchcock (2007) and Godfrey-Smith (forthcoming) for overviews of the varieties of causal pluralism. Campagner and Galavotti (2007) offer an account of causal pluralism focusing specifically on mechanistic and manipulability approaches to causation.

particular set of facts causally relevant to the production of an event is that changing these facts would change the behavior of the mechanisms that produce that event. Unlike in some accounts where causal relevance is secondary or “merely” explanatory, relations of productivity and relevance are genuine and intertwined aspects of any true account of the causal structure of the world.⁶

Productivity and relevance in natural selection

The distinction between productivity and relevance can help unknot the dispute between the different interpretations of natural selection discussed at the outset of this paper. My basic strategy will be to argue that natural selection is produced by causal processes operating at the individual level, but that there are many population-level properties that are causally relevant to the dynamics of evolutionary processes. To begin making this case, consider the causal structure of a phenomenon that provides *prima facie* evidence of population level causation—namely frequency dependent selection. Futuyma, in his textbook account of frequency dependent selection cites as an example the case of the corixid *Sigara distincta*—a variety of water bug that is subject to predation by fish (Futuyma 1986). These corixids come in different colors, and the likelihood of them being eaten by fish depends upon their color. Futuyma cites an empirical study of a population showing that the relative fitness of different color forms changes depending upon the frequency of the color form within the population. The hypothesized explanation of the frequency dependence is that the fish predators form a stereotypic searching image associated with the dominant color, making the rarer color forms fitter.

Figure 1 represents schematically the processes involved in a simplified version of the corixid case. We imagine a population of dark and light water bugs. During the course of their lives the bugs have encounters with predator fish, in which their survival depends on their not being seen. In the diagram the solid lines represent the life history of each bug, and the fish above the lines represent encounters with predators. In our hypothetical case, 1/3 of the encounters with dark bugs lead to the bug’s being eaten, while in the rarer light form, the bug stands only 1/6 chance of being eaten. These better odds mean that a higher proportion of the rarer bugs survive to reproduce, thus making that variety not so rare. To complete the picture, we’d need to see that in subsequent generations as the light form comes to predominate, stereotypic imaging on the fish’s part will lead to a decrease of fitness of the light form.

While this causal story suggests how and why the frequency of a color form is relevant to that form’s fitness, and more generally to explaining changes in the distribution of forms within the population, we cannot strictly say that increased frequency of a form within a population *produces* decreased fitness of that form. The reason is that production is a relationship between objects and events, while the

⁶ See Glennan (forthcoming) for a more detailed explanation of the relationship between productivity, relevance and mechanisms.

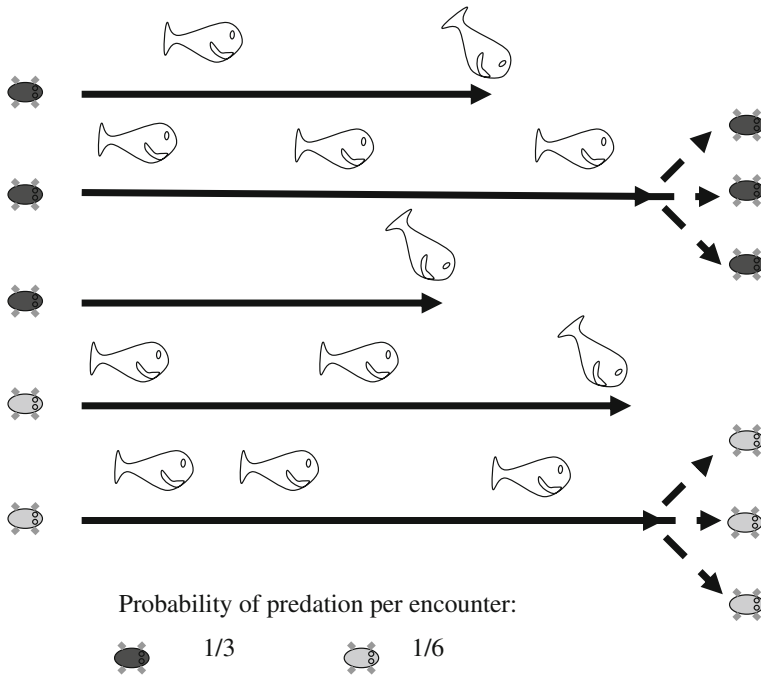


Fig. 1 Model of frequency-dependent Causation

population is not (in this case at least) an individual object and the increase of frequency or decrease of fitness are not individual events. The only entities here are the fish and the bugs, the only activities are the activities of individual fish and bugs, and the only interactions are when the fish eat the bugs and when the bugs make baby bugs. It is at the level of the activities and interactions of individual bugs that we find the mechanisms that produce new bugs.

My claim that populations aren't objects is connected with the mechanistic view of productivity suggested above. Machamer, Darden and Craver emphasize the productive role of objects or, as they prefer to call them, entities⁷:

Mechanisms are composed of both entities ... and activities. Activities are the producers of change. Entities are the things that engage in activities. Entities often must be appropriately located and oriented and the activities in which they engage must have a temporal order, rate, and duration. (Machamer et al. 2000, p. 3).

⁷ In Glennan 1996, I emphasized the fact that mechanisms were composed as parts, which I required to be objects. The major point of the object requirement was related to a point Woodward (2002) has emphasized about mechanisms, namely that parts be in principle independently manipulable. Machamer et al. have used the term 'entity' while Bechtel (Bechtel and Abrahamsen 2005) has preferred the term 'component' though these terminological differences are not indicative of substantive disagreements. I shall use the words 'object' and 'entity' interchangeably.

Entities need to be localized in space and time; they need to engage in particular activities at particular times and places. The population in the water bug case does not have these properties. The population as a whole is spread out and does not engage in collective activities. The individual members of the population will engage in activities that are essential to producing selective changes. They will float, swim, evade predators, gobble up bits of plants, and mate. But these are activities of single bugs, not of the population as a whole.

The fact that a population has parts is not what entails that it is not an entity. Single bugs have parts too. What makes a collection of parts into a single entity is that these parts have a stable structure, that the stable structure engages in activities as a unified entity, and that these collected parts share a common fate. When a fish kills a water bug, it kills the whole water bug—it can't kill its legs but not its body for instance. On the other hand, when a fish kills a water bug, it doesn't kill the whole population of water bugs. The life of one water bug is more or less independent of another.

One cannot say categorically that populations either are or are not individual entities. The question of whether they are individuals only makes sense in the context of analyzing a particular causal process (cf. Glennan 1996). For instance, in migration processes, populations do act as individuals. With respect to selection processes, the question of whether or not populations or sub-populations should be treated as individual entities depends upon whether or not group selection is at work. To the extent that certain populations form communities or groups with stable structures and who share common fates, we might justifiably treat them as individuals that engage in unified and productive activities. An ant colony or a baboon troop may be an individual, but my supposition is that in this case the bugs in the pond are not.⁸

Denis Walsh has criticized Sober's view that natural selection is a force on the basis of a similar view of causal processes:

Natural selection, it seems, is merely the consequence of an assemblage of causal processes taking place at the individual level. There is no need to invoke a distinct force operating over populations in order to explain the changes in gene frequency thought to be explained by natural selection (Walsh 2000)

I concur with Walsh and others' criticisms of the force analogy. The problem with thinking of natural selection as a force is that a force, if we are to give it anything like its Newtonian meaning, describes the way in which one individual object interacts with another. Sober (1984) argues for the force analogy in evolution by suggesting that, in order to understand the causal structure of evolutionary processes, we must separate out the component forces (selection, drift, migration, etc.) that produce the resultant force. The component forces are what describe the

⁸ While paradigm cases are fairly clear, it is difficult or perhaps impossible to give strict criteria for deciding when a population is sufficiently stable in its structure and organization that it can be considered an individual undergoing selection. It seems to me that Sober and Wilson's idea of common fate (Sober 1984; Sober and Wilson 1998) is important. For further discussion of this problem in the context of the units of selection problem see Glennan 2002.

true causes. But in the same way, the force of selection operating on a population is really just the sum of component forces operating on individuals within the population. Those component forces describe the real causes.

Shapiro and Sober (2007) are sensitive to this criticism and have responded to it by arguing for a more expansive view of what can count as a causal evolutionary process:

Sometimes x and y experience the same token selection process because they causally interact. At other times they participate in the same selection process because they are affected by a common token cause. ... Darwin drew this distinction in a famous passage from the *Origin*: “I should premise that I use the term struggle for existence in a large and metaphorical sense.... Two canine animals in time of dearth may truly be said to struggle with each other which shall get food and live. But a plant on the edge of a desert is said to struggle for life against the drought” (Darwin 1859, p. 62). When two plants in the same population differ in fitness, Walsh sees one process affecting the one and a different process affecting the other. There is no single process affecting them both. Our reply is that we find it entirely natural to point to the drought (Shapiro and Sober 2007, pp. 252–253).

Shapiro and Sober’s reply is mistaken on two counts. First, when two events have a common cause that does not mean that both are part of the same causal process. If we have a common cause situation, where c produces both e_1 and e_2 , the sequence of intermediates leading from c to e_1 will be one process, and the sequence leading from c to e_2 will be a distinct process. The process $c-e_1$ and the process $c-e_2$ both include event c , but it doesn’t follow that e_1 and e_2 belong to a single process. If one thought they did, then by parity of reasoning one would have to conclude that, because there is a continuous causal process linking every human being on the planet to a single ancestor, every human being on the planet is part of a single causal process. When the two canines fight, they interact with each other, and thus they and their fight are part of a shared process, but that is not the case with the two plants struggling with drought. Their struggles are causally independent of each other in a way that the canines’ struggle is not. Quite apart from the common cause problem, there are difficulties with thinking that a drought can be a cause at all. Arguably, droughts cannot be productive causes because droughts are not events. They are just extended non-occurrences of rain. Thus, if droughts cause anything they do so by omission, and, as I argued above, omissions can be causally relevant but not productive.

The crux of my objection to Shapiro and Sober’s account of causal processes is that the processes that produce change must consist of individual entities interacting with each other. This view is consistent with the views of advocates of the individualistic and statisticalist interpretations of evolutionary processes. But there is a worry that any view that attributes real causal powers only to individuals rather than collections of individuals will make all higher level entities and properties epiphenomenal. If populations, being collections of organisms, can’t genuinely produce effects, then it would seem that the parts of organisms, organs and tissues, and the parts of their parts, and so on, would all be incapable of producing effects.

This line of reasoning would seem to suggest that the only causally productive entities would be the most basic physical entities.

In response to this, it is important to recognize that the reason population level properties don't produce change does not have to do with the fact that populations are composed of individuals. It has to do instead with the fact that the population is not a part of the mechanism that produces changes in genotype and phenotype frequencies. The population as a whole does not interact with other entities as a whole in order to change its genotype and phenotype frequencies. In contrast, an organism engaging in interactions with its environment that ultimately lead to its reproducing (or failing to reproduce) is a part of the mechanism that produces changes in genotype and phenotype frequencies. That organism is in turn composed of parts that produce the organism's behavior, but the fact that these organism parts engage in productive activities does not make the activities of the organism as a whole any less productive.⁹

My argument has so far aimed at showing that populations and population level properties are not typically causally productive. This is because populations are not typically entities that enter into productive causal relations, but are rather statistical aggregations. In saying this I agree with Walsh and the other statisticalists. But, like Millstein, Shapiro and Sober, I am convinced that population level properties are not epiphenomenal. To save this intuition, we must show that the individualistic view of causal processes is consistent with the claim that population level properties are causally relevant.

To begin, observe that natural selection occurs only when there is heritable variation of fitness. Without variation there is no natural selection, and variation is essentially a population concept. An individual cannot vary with respect to itself. It is this fact about natural selection that I believe inclines people like Millstein to insist so strenuously that natural selection is a population level causal process.

My earlier discussion of the water bug example emphasized that we should not think of a population level property as *producing* anything. But clearly, population level properties can be causally relevant. At the beginning of this paper I suggest that one way that a fact f could be relevant to an event e was if e would have been more or less likely had it not been the case that f . Applying this to the water bug case, let e be a predation event and f be the fact that the color phenotype of the bug has a certain frequency. The fact that the frequency has a certain value will indeed make it more or less likely that e occurs. Given that changes in the frequency of the genotype will depend on these individual predation events, the population level trait of frequency is thus a causally relevant fact about a population.

Moreover, as Millstein (2006) has argued, even in cases of simple selection, fitness is essentially comparative. When we evoke fitness to explain why one variant increases in the population, we are saying that the trait's propensity to reproduce is

⁹ The concern that perhaps all non-fundamental entities and properties lack causal powers has been the subject of considerable discussion, especially among philosophers of mind. Jaegwon Kim (1984, 1998) in particular has flirted with thoroughgoing epiphenomenalism. I cannot, in the space of this paper, provide a general answer to these worries, though I have tried to show how a mechanistic account of causation can do so in Glennan (forthcoming). Shapiro and Sober's (2007) analysis of epiphenomenalism offers a complementary response to Kim.

greater than the alternatives in the population. What is causally relevant to changes in the frequency of traits is not the absolute rate of reproduction of the trait but the ratio of its reproductive rate to that of its variants, which is again a population level property.¹⁰

These seem like pretty elementary observations about the causal role of population level properties, but Walsh thinks that individualism about processes shows that the causal significance of such properties is illusory. He offers the following metaphor for thinking about selection:

I suggest that the process of selection is more like the motion of a shadow than it is like your own motion. It is simply the consequence of the differential rates of distinct causal processes occurring within individuals. Of course changes in gene frequencies are themselves simply the consequences of the differences in these same processes. Natural selection and changes in gene frequencies are not related as cause and effect; they are joint effects of a common cause (Walsh 2000).

The mistake in this metaphor is that it suggests that natural selection is a distinct process that *is caused by* the individual level processes. In fact thought, the relation of the individuals to the population is one of part to whole—not of cause to effect.

To see the consequence of this error, let us compare the natural selection case with a genuine case of a common cause—the venerable example of the barometer and the storm. Atmospheric pressure drops causing both the barometer to fall and the storm to occur. The barometer dropping doesn't cause the storm. Because changes in barometer readings don't cause storms, we can manipulate the barometer in any way we want and we won't make it rain. This is very different from the case of manipulating a population level property. If we manipulate a population level property like the relative frequency of a frequency dependent trait, we will have a causal influence on selective outcomes. Because a population is a logical aggregation of rather than a causal product of the individuals of which it is composed, when one manipulates a property of the population one *ipso facto* manipulates the properties of individuals within the population. That is why we need not worry that population level properties are epiphenomenal.

Shapiro and Sober (2007) make essentially the same point in their discussion of Weismann's experiments concerning the inheritance of acquired characteristics. Weismann's experiments involved altering a phenotype (cutting off the tails of mice) while holding a genotype fixed. The fact that this intervention does not cause mice in the next generation to have shorter tails suggests that the parental phenotype is epiphenomenal with respect to the tail length of the next generation. Cutting off the tails is like messing with the needle on the barometer.

Sober and Shapiro suggest that concerns about the causal powers of higher-level properties arise because of a misleading analogy between the screening-off

¹⁰ Forber and Reisman (2007) have made points similar to those of Millstein. They have argued on the basis of manipulability approaches to causation that population level parameters (such as population size) can cause evolutionary change. Manipulability criteria do indeed demonstrate the causal relevance of these parameters, but this fact is not inconsistent with the claim that the producers of evolutionary change in such cases are individual organisms.

relationship that holds between genotypes and phenotypes and the relationship that holds between higher-level properties and their supervenience bases. A genotype is definitively not a supervenience base for a phenotype. Higher-level properties depend synchronically and non-causally upon their supervenience bases, while phenotypes depend diachronically upon their genotypes as well as environmental factors. One cannot change a higher-level property while holding its supervenience base fixed in the way that one can change a phenotype while holding the genotype fixed. A phenotype truly can be like a shadow of the genotype, because it is a property of a distinct individual (the mature organism) that does not affect the next generation's genotype. But the screening-off argument does not apply to higher-level properties of populations—like population size or genotype frequency—because the relationship between individuals and populations is constitutive, synchronic and non-causal.

But while Sober and Shapiro's analysis shows that fears of epiphenomenalism are overblown, it only shows that population level properties can be causally relevant, not that they can *produce* change. If my analysis is correct, it is only entities and not properties that can truly be said to be productive. The producers of change are the individual entities that are the parts of the evolutionary mechanism, and, assuming that the only selective forces at work are organismic (rather than genic or group), these parts will be individual organisms.

The distinction between productivity and relevance can shed some light on one final issue that has worried philosophers—namely the causal role of fitness. The term fitness is used in a number of ways—at least two of which are what Matthen and Ariew (2002) call predictive fitness and vernacular fitness. They characterize predictive fitness as the concept from population genetics that is a “statistical measure of evolutionary change, the expected rate of increase (normalized relative to others) of a gene, a trait or an organism's representation in future generations” (56). They describe vernacular fitness more loosely as “an organism's overall comparative advantage traceable to heritable traits” (*ibid*).

Matthen and Ariew rightly argue that predictive fitness can't be a cause of selection, because it is in fact just a measure of the outcome of selection. Causes must be logically distinct from their effects, and predictive fitness does not meet this test. Vernacular fitness is a more complex issue. We can think of vernacular fitness either in terms of a pairwise relation between organisms— x is fitter than y —or a relation between traits— t_1 is fitter than t_2 —or we can think of it as a relationship between an organism and the environment, or of a trait and the environment. While biologists rightly point out that there is no selection without fitness differences, the latter ecological or engineering fitness concept is fundamental. A plant living in an arid environment may have a mechanism for collecting moisture from the air. If we consider the causal process that leads to this one plant's reproducing, this property of the plant will be causally relevant to the reproduction event, in the sense that if the plant hadn't had this property, it wouldn't have survived to reproduce. Actually, to be more precise, this property only might have been causally relevant. Whether it was causally relevant depends on the actual details of the plant's life. For if the plant was born perhaps in a season in which the rains were greater than average, perhaps it would have survived even if it didn't have the mechanism. More generally, if we are asking a question about why a particular event occurred—why this particular

plant survived and reproduced—an explanation will cite both the productive events and causally relevant properties like the water collection mechanism. We will of course need to go beyond this ecological fitness concept and look at variations of fitness if we are to explain why certain traits change frequency within a population, but this fact should not prevent us from seeing the relevance of individual traits to the production of further individuals, without which no populations, and hence, no population changes, could be produced.

Conclusion

In their oft quoted characterization of evolutionary theory, Sterelny and Kitcher claim that

in principle we could relate the biography of each organism in the population, explaining in full detail how it developed, reproduced and survived But evolutionary theory, like statistical mechanics has no use for such a fine grain of description: the aim is to make clear the central tendencies in the history of evolving populations. (Kitcher and Sterelny 1988, p. 345)

The genius of evolutionary theory is that it recognizes that changes in populations depend (in a causal sense) not on what one organism or the other does, but on how different types of organisms do on average. This explains why population level properties are causally relevant, as well as why we can provide causal explanations of population changes that do not refer to the detailed biographies of organisms. But while these biographies may be unnecessary for our explanation, it is these organisms and their life histories that produce change. Without the individual organisms, there are no populations, and without the activities and interactions of these individual organisms, there are no changes in populations. By distinguishing between productivity and relevance, we can see how accepting this obvious truth about individuals does not lead us to epiphenomenalism about properties of biological populations.

Acknowledgements An earlier version of this paper was presented at the 2007 meeting of the International Society for the History Philosophy and Social Studies of Biology in Exeter, the United Kingdom. I am grateful to participants in the conference for their feedback. I also had useful conversations about this paper with Lisa Lloyd and members of the Biological Studies Reading Group at Indiana University.

References

- Bechtel W, Abrahamsen A (2005) Explanation: a mechanist alternative. *Stud His Philos Bio Biomed Sci* 36(2):421–441
- Bouchard F, Rosenberg A (2004) Fitness, probability and the principles of natural selection. *Br J Philos Sci* 55(4):693–712. doi:[10.1093/bjps/55.4.693](https://doi.org/10.1093/bjps/55.4.693)
- Campagner R, Galavotti MC (2007) Plurality in causality. In: Wolters G, Machamer P (eds) *Thinking about causes: from Greek philosophy to modern physics*. University of Pittsburgh Press, Pittsburgh
- Collins J, Hall N, Paul LA (2004) *Causation and counterfactuals*. Bradford Book/MIT Press, Cambridge, MA

- Darwin C (1859[1964]) *On the origin of species*. Facsimile of the 1st edition. Harvard University Press, Cambridge, MA
- Dowe P (2000) *Physical causation*. Cambridge studies in probability, induction, and decision theory. Cambridge University Press, Cambridge
- Forber P, Reisman K (2007) Can there be stochastic evolutionary causes? *Philos Sci* 74(5):616–627. doi: [10.1086/525608](https://doi.org/10.1086/525608)
- Futuyma DJ (1986) *Evolutionary biology*, 2nd edn. Sinauer, Sunderland, MA
- Glennan SS (1996) Mechanisms and the nature of causation. *Erkenntnis* 44(1):49–71. doi: [10.1007/BF00172853](https://doi.org/10.1007/BF00172853)
- Glennan SS (2002) Contextual unanimity and the units of selection problem. *Philos Sci* 69(1):118–137. doi: [10.1086/338944](https://doi.org/10.1086/338944)
- Glennan SS (forthcoming) Mechanisms, causes and the layered model of the world. *Philos Phenomenol Res*
- Godfrey-Smith P (forthcoming) Causal pluralism. In: Beebe H, Hitchcock C, Menzies P (eds) *Oxford handbook of causation*. Oxford University Press, New York
- Hall N (2004) Two concepts of causation. In: Collins J, Hall N, Paul LA (eds) *Causation and counterfactuals*. MIT Press, Cambridge, MA
- Hitchcock CR (1995) Discussion: Salmon on explanatory relevance. *Philos Sci* 62(2):304–320. doi: [10.1086/289858](https://doi.org/10.1086/289858)
- Hitchcock CR (2007) How to be a causal pluralist. In: Wolters G, Machamer P (eds) *Thinking about causes: from Greek philosophy to modern physics*. University of Pittsburgh Press, Pittsburgh
- Jackson F, Pettit P (1990) Program explanation: a general perspective. *Analysis* 50(2):107–117. doi: [10.2307/3328853](https://doi.org/10.2307/3328853)
- Kim J (1984) Epiphenomenal and supervenient causation. *Midwest Stud Philos* 9:257–270. doi: [10.1111/j.1475-4975.1984.tb00063.x](https://doi.org/10.1111/j.1475-4975.1984.tb00063.x)
- Kim J (1998) *Mind in a physical world: an essay on the mind–body problem and mental causation*. MIT Press, Cambridge, MA
- Kitcher P, Sterelny K (1988) The return of the gene. *J Philos* 85:339–361. doi: [10.2307/2026953](https://doi.org/10.2307/2026953)
- Lewis D (1973) Causation. *J Philos* 70:556–567. doi: [10.2307/2025310](https://doi.org/10.2307/2025310)
- Machamer P, Darden L, Craver CF (2000) Thinking about mechanisms. *Philos Sci* 67(1):1–25. doi: [10.1086/392759](https://doi.org/10.1086/392759)
- Matthen M, Ariew A (2002) Two ways of thinking about fitness and natural selection. *J Philos* 99(2):55–83. doi: [10.2307/3655552](https://doi.org/10.2307/3655552)
- Millstein RL (2006) Natural selection as a population-level causal process. *Br J Philos Sci* 57:627–653. doi: [10.1093/bjps/axl025](https://doi.org/10.1093/bjps/axl025)
- Rosenberg A, Bouchard F (2005) Matthen and Ariew's obituary for fitness: reports of its death have been greatly exaggerated. *Biol Philos* 20(2–3):343–353. doi: [10.1007/s10539-005-2560-0](https://doi.org/10.1007/s10539-005-2560-0)
- Salmon WC (1984) *Scientific explanation and the causal structure of the world*. Princeton University Press, Princeton, NJ
- Schaffer J (2003) The metaphysics of causation. In: Zalta EN (ed) *The Stanford encyclopedia of philosophy*, Spring 2003 edn. <http://plato.stanford.edu/archives/spr2003/entries/causation-metaphysics/>
- Shapiro LA, Sober E (2007) Epiphenomenalism—the dos and the don'ts. In: Wolters G, Machamer P (eds) *Thinking about causes: from Greek philosophy to modern physics*. University of Pittsburgh Press, Pittsburgh
- Sober E (1984) *The nature of selection: evolutionary theory in philosophical focus*. MIT Press, Cambridge, MA
- Sober E, Wilson DS (1998) *Unto others: the evolution and psychology of unselfish behavior*. Harvard University Press, Cambridge, MA
- Walsh DM (2000) Chasing shadows: natural selection and adaptation. *Stud Hist Philos Biol Biomed Sci* 31(1):135–153. doi: [10.1016/S1369-8486\(99\)00041-2](https://doi.org/10.1016/S1369-8486(99)00041-2)
- Walsh DM (2007) The pomp of superfluous causes: the interpretation of evolutionary theory. *Philos Sci* 74(3):281–303. doi: [10.1086/520777](https://doi.org/10.1086/520777)
- Walsh DM, Lewens T, Ariew A (2002) The trials of life: natural selection and random drift. *Philos Sci* 69(3):452–473. doi: [10.1086/342454](https://doi.org/10.1086/342454)
- Woodward J (2002) What is a mechanism? a counterfactual account. *Philos Sci* 69(3 Supplement):S366–S377. doi: [10.1086/341859](https://doi.org/10.1086/341859)
- Woodward J (2003) *Making things happen: a theory of causal explanation*. Oxford University Press, Oxford