Adaptation, Adaptation to, and Interactive Causes

Bruce Glymour

Abstract This paper develops alternative ways of understanding 'adaptation to' specific environmental conditions, with particular attention to the explanatory power offered by differing conceptions, the concomitant epistemic demands they make of explanations, and the models such explanations employ. It is shown that explanations of adaptation to particular environmental conditions can satisfy important intuitions only if the environmental conditions to which phenotypes are adapted are interactive causes of fitness. However, taking this constraint to be both necessary and sufficient for 'adaptation to' imposes epistemic burdens on our explanatory practice, and risks violating yet other intuitions. The paper briefly explores the consequences of the constraint for the idea that selection requires shared environments, the idea that selection requires a homogeneous environment, the idea that phenotypes may be extended, and the idea that niches may be constructed.

1 Introduction

Adaptation is a relational concept: a trait cannot be an adaptation without being an adaptation to some environment. Hence, to identify a trait as an adaptation is to imply the existence of some relation between environment and trait. The theory of evolution by natural selection identifies that relation: a trait is an adaptation to an environment only if the environment selected for the trait as against alternatives. This leaves it an open question whether such selection relates any given adaptation to specific features of an environment, or instead relates the adaptation to the environment as an undifferentiated whole. The issue is of some moment, if only because biological practice invites both readings of 'adaptation to.'

B. Glymour (✉)
Department of Philosophy, Kansas State University, Manhattan, KS 66506, USA
e-mail: glymour@ksu.edu

G. Barker et al. (eds.), *Entangled Life*, History, Philosophy and Theory of the Life Sciences 4, DOI 10.1007/978-94-007-7067-6_6,

On the one hand, with some frequency biologists venture hypotheses about the particular environmental conditions to which traits are adapted, as e.g. the claim that sex is an adaptation to parasitism (Levin 1975; Maynard Smith 1976; Jaenike 1978). If adaptations are not, at least sometimes, adaptations to particular environmental conditions, then such hypotheses and tests of them make no sense. On the other hand, one does not typically specify environmental conditions when estimating fitnesses—though norms of reaction for fitness, i.e. fitness functions, can be and sometimes are estimated from data, there are many perfectly standard population genetic models employing fitnesses that implicitly condition on the whole of the environment, whatever it may be. Or again, some putative selection processes, e.g. pure r-selection, seem to depend on no specific environmental feature (see e.g. Lennox and Wilson 1994). To the extent that r-selection can drive adaptations, e.g. in life history strategies, those adaptations are arguably evolved responses to the environment as a whole, rather than to any specific environmental condition.

Even if we accept the idea that adaptation is adaptation to one or another set of particular environmental conditions, there remains a further question about just what causal or nomic relations must hold between phenotype and environmental condition if it is to be sensible to speak of the phenotype as an adaptation to that condition. There are a number of alternative possible requirements that might be imposed, and the choice among them will have consequences for both our epistemic and explanatory practices. Hence, it is of some importance to ask in what sense, and to what extent, a particular environmental condition must cause (generate, explain?) selection on a phenotype if the phenotype is count as an adaptation to that environmental condition. In this paper I consider some fairly intuitive constraints on the explanatory role of appeals to 'adaptation to,' and explore the extent to which those constraints require 'adaptation to' to imply interactions between traits and environmental conditions.

I begin by adumbrating two arguments, given in detail elsewhere (Glymour 2011), and then draw out some implications of the respective conclusions. The first argument aims to show that it makes sense to speak of an adaptation *to* a particular environmental condition only if that condition interactively causes survival or reproductive success. The second argument aims to show that, in consequence, it is possible to identify the environmental conditions to which an adaptation is adapted only if one measures and models the causal influence of the environmental condition on survival or reproductive success. If the arguments are correct, they have a number of implications. Among them are constraints on an understanding of 'environment' suitable for representing relevant organism-environment interactions; some unavoidable choices about which environmental conditions are to be taken as essentially explanatorily relevant; and finally some limitations on standard methods for measuring the strength of selection. I begin with some preliminaries and then rehearse the central arguments. In the subsequent sections I explore the above mentioned implications.

2 Preliminaries

I will, for ease, confine my attention to phenotypic adaptations carried by individual organisms. There is nothing in what is to follow that prevents extensions, with some modification, to either the genic or the population/species level. But extensions of either sort do introduce complexities that require more space to deal with than is here available. In consequence, I will employ models of natural selection in which selection acts on individuals. Nothing at all hinges on this second choice of modeling level—the same results for phenotypic adaptations can be got by modeling selection at other levels, but the treatment would thereby be made unnecessarily complex.

The arguments to follow turn on the range of possible causal structures governing survival and reproduction in biological populations. The essential causal relations are between phenotypic variables, environmental variables and individual survival and reproductive success. I will use W, which I will call fitness, as the relevant effect variable throughout. In the examples to follow, W is calculated as actual or expected reproductive success, but no metaphysical commitment is intended thereby. The reader may take those calculations to be estimates of fitness in whatever sense she prefers to understand it, so long as fitness so understood depends on the joint probability density over survival and/or reproductive success.

I assume that phenotypic variables are unproblematic. No such assumption is possible with respect to environmental variables. More will have to be said later about these, but for the moment the following will be enough. There are two ways to measure an environment, either by its net effect on fitness (in whatever sense) or by the presence/absence or magnitude of some particular property. I will call variables of the first kind *measures of environmental quality* and say that they offer a qualitative representation of the environment (though the variables may well be real-valued, what these values represent is something about the quality of the environment from the perspective of the organism). I will call variables of the second kind *environmental variables*, and say that they offer an explicit rather than qualitative representation of the environment.

I adopt the now standard language of graphical causal modeling, according to which causation is an asymmetric dependence relation between variables (Pearl 2000; Spirtes et al. 2000). A variable P is said to be a direct cause of a variable W relative to some set of variables ($\mathbf{V}$, P, $W \in \mathbf{V}$) and background conditions $\mathbf{B}$ when there is some pair of interventions on P, holding all other variables except W in $\mathbf{V}$ constant, such that the probability distribution or density over W differs across the interventions. Such direct causal relations will be represented in graphs as arrows directed from the cause to the effect. Interactive or context-dependent causation is a special case of causation. P and E are interactive causes of W, relative to $\mathbf{V}$ and $\mathbf{B}$, if and only if P and E are both direct causes of W and for some moment of the distribution or density over W, there is some pair of interventions on P (or E) and some pair of values for E (or P) such that the difference in the value of the

moment of the distribution or density over W between the two interventions on P (or E) given the first value of E (or P) is unequal to the difference in the value of the moment between the two interventions on P (or E) given the second value of E (or P). Informally, the effect on W of a change in one cause depends on the value of the second cause. For example, pressure and volume are interactive causes of temperature because the effect on temperature of a unit change in pressure depends on the volume.

The Japanese camellia (*Camellia japonica*) and its predator the camellia weevil (*Curculio camellia*) provide a biological example. In order to oviposit on the camellia seeds, the weevil bores a hole in the camellia fruit. This selects for a thicker pericarp. However, the strength of this selection depends on the length of the weevil's rostrum, and both traits vary within and among local populations. The equation relating the probability of boring success (PBS), and hence the probability that a seed is predated, to these traits is given by $\mathrm{PBS} = 1/\left[1 + e^{(0.819p + 0.471t - 4.18)}\right]$, where p is the pericarp thickness and t is the rostrum length (cf. Toju and Sota 2006). The contribution of pericarp thickness to the fitness of a given camellia plant thus depends on a locally varying environmental condition—the rostrum length characteristic of the local weevil population.

When interactive causal dependencies are mathematically modeled, i.e. when W is written as some function of its causes, interactive causes will appear together in at least one term on the right hand side of the equation (often, though as above not always, this term is a product of the causal variables). In such cases the contributions of the two (or more) causes are not separable; if P and E are *not* interactive causes of W, then it will be possible to write W as a function of P and E (and perhaps other variables) in such fashion that the terms containing E do not contain P, and vice-versa. In this case the contributions of the two causes are separable. It follows that interactive causal connections are symmetric in the following sense: if P is an interactive cause of W with E, then E is an interactive cause of W with P. It will sometimes be useful to attend to only one cause of such a pair. When necessary I will therefore write that P (or E) is the interactive cause and E (or P) the context; the difference is entirely pragmatic.

Technical preliminaries done, a philosophical preliminary is in order. In what follows I advance a (partial) conception of what it is for a phenotypic trait to be an adaptation to some but not other environmental conditions. As such, I'm engaging in a species of conceptual analysis. But I wish to be as clear as possible about just what species of conceptual analysis I intend. The explanatory power of language depends in part on how we use language to represent the world. Scientific terms, in particular, inherit their explanatory power from the fact that in using them we denote, more or less systematically, real physical, causal, nomic, or statistical features of the world. There are any number of features of causal and statistical structure that might be counted as explanatorily relevant to evolutionary outcomes, depending on which features of which outcomes one takes to be in need of explanation and on what intuitions one has about the kind of information a satisfactory explanation ought to offer. Thus, in my view, there is no fact about which phenotypes are and are not adaptations to particular environmental conditions, prior to a choice about what

we will mean by 'adaptation to.' And that choice is in large measure open—there are many explanatorily relevant features of causal or statistical structure that we could choose to denote by 'adaptation to.' In consequence, I do not aim in what follows to specify any fact about what we do mean by 'adaptation to,' and still less to specify either what we ought to mean by that locution or what adaptation to *is*, metaphysically speaking.

The aim is rather this. Depending on which phenomena we choose to denote by 'adaptation to,' different kinds of information will be required to explain why any given phenotype is an adaptation to any particular environmental condition; similarly, an appeal to the fact that a phenotype is an adaptation to some particular environmental condition will itself carry some explanatory power, but what that power is will depend on which features of causal or statistical structure our usage of the term 'adaptation to' systematically respects. Further, the epistemic demands imposed by explanations of the adapted nature of a phenotype in turn depend on these facts about usage. And, roughly, the more intricate the physical distinctions we choose to denote with the term 'adaptation to,' the more explanatory power this usage has, but also the greater the epistemic demands on correct usage. The aim of the conceptual analysis to follow, then, is to clarify the range of choices available.

Specifically, I will argue that the choice to use 'adaptation to' in ways that track a natural but quite minimal structural distinction deprives the locution of any explanatory power beyond that already inherent in our usage of 'adaptation,' while a choice to track other features of causal and statistical structure confers a particular explanatory power on the locution, but at an epistemic cost. I will further point to yet more intricate structural features of interactions between organisms and their environments that one might wish to respect, in that doing so would endow the 'adaptation to' locution with yet more explanatory power. But I will take no stand on whether such further constraints on 'adaptation to' offer a particularly efficient regimentation of our language. It will be enough, here, to point to the choices that are open to us.

3 "Adaptation to" and Interactive Causation

I assume that to say of a phenotype that it is an adaptation to some particular environmental circumstance carries more explanatory power than to say of the phenotype simply that it is an adaptation. And I further assume that this extra explanatory power depends on the contrast between the environmental conditions to which the adaptation is an adaptation and those conditions to which it is not an adaptation. That is, the idea of 'adaptation to' is explanatorily useful only if for some adaptations there are features of the environment to which the adaptation is an adaptation, and others to which it is not an adaptation. For example, if hypsodonty (having high-crowned teeth) is an adaptation to the siliceous phytoliths (hard mineral particles contained in plant tissues) of grass, it had best not also be an adaptation to every other feature of grassland environments. For were it, then that

hypsodonty is an adaptation to the siliceous phytoliths of grass would imply no more than that (a) hypsodonty is an adaptation and (b) it evolved in populations inhabiting grassland environments; but exactly this would similarly be implied by (and imply) the fact that hypsodonty is an adaptation to the dusty conditions of grassland habitats, and similarly to any and every other feature of such environments.

If adaptation is to be to some but not all features of an environment, we must have some principled way of distinguishing those features of the environment to which species adapt from other features of the environment. Henceforth I will call these conditions, the conditions to which an adaptation is an adaptation, the *adapting conditions*. For preference, any principle we employ to individuate adapting conditions from other features of an environment should respect certain constraints suggested by the explanations we give of adaptations and the kind of explanatory and inferential use we make of them. Among these are three intuitions that are both deep and fairly central to explanatory practices in biology.[1] First, environmental conditions are supposed to explain (in part) the fixation of those phenotypes that are adaptations to those conditions—the adapting conditions should play some central role in a full-bodied explanation of phenotypic adaptations to them.[2] I will call this the intuition ACEA (adapting conditions explain adaptations).

Second, such explanations are (at least potentially) doubly contrastive: they explain the fixation of adaptive trait values *as opposed to* alternative trait values, and they do so by appeal to one *rather than another* set of environmental circumstances. Just so, if hypsodonty is an adaptation to siliceous phytoliths in grass, we explain why horses and cows evolved high- rather than low-crowned teeth by appeal to the fact that grass has siliceous phytoliths rather than to the fact that herbivores in grasslands consume large quantities of dust when grazing. More narrowly, an appeal to environmental conditions $E = \mathrm{e}$ to explain the fixation of phenotype $P = \mathrm{p}$ is warranted only if there are alternative circumstances $E = \mathrm{e}'$ and phenotype $P = \mathrm{p}'$, such that had $E = \mathrm{e}'$ obtained, the phenotype p' would (probably) have been maintained at a non-zero frequency in the population. I will call this intuition 'DC' (explanatory appeals to adaptations to particular environmental conditions are doubly contrastive).[3] The pair of contrasts will play a crucial role in the arguments to follow, and so for ease of reference I will call the alternative phenotypes permitting the first contrast *contrasting phenotypes*, and the alternative environmental conditions underwriting the second contrast *contrasting environmental conditions*.

---

[1] Space prevents a careful development of these intuitions from primary sources. But readers who do not find them obvious might usefully consider discussions of particular adaptations, such as Heywood (2010), McFadden (1992), Hunt (1994), and Wheeler (1991).

[2] I omit consideration of traits, genetic or phenotypic, which are in some important sense adaptive, but such that selection cannot drive the trait frequency to that expected from the mutation rate characteristic of the relevant genetic loci. Examples here include the sickle-cell allele. The issues here are important, but beyond the scope of this essay.

[3] Those puzzled by this intuition may consult van Fraassen (1980) for discussion of the first contrast (between alternative outcomes), and Glymour (1998, 2007) for discussions of the second contrast (between alternative causes or processes).

Thirdly, I suppose that the relevant explanations are selective—i.e., that the environmental conditions appear in such explanations as part of the description of a selection process rather than a drift process. I will call this intuition AEASE (adaptive explanations are selection explanations).

Given these explanatory intuitions, how are we to differentiate between environmental conditions that are and those that are not adapting conditions? One obvious individuating principle suggests itself. Some features of the environment cause survival and reproductive success, while others do not. If E_1 and E_2 are variables that measure the presence or magnitude of some environmental feature, where E_2 causes fitness but E_1 does not, then arguably no evolved phenotypic trait can be an evolutionary response to the presence or magnitude of E_1 in the environment, and hence the phenotypic trait cannot be an adaptation to the value of E_1 characteristic of the environment.[4] We might then identify the features of the environment to which an adaptation is an adaptation with those features of the environment that cause fitness; call this the *Causal Condition*.

Unfortunately, the Causal Condition will, in certain circumstances, identify particular environmental features as adapting conditions despite the fact that an explanatory appeal to those features would run afoul of the three intuitions mentioned above. Say that the distribution of an environmental variable is *homogeneous* if for any value of that variable, the proportion of one phenotype characterized by that value is equal to the proportion of any other phenotype characterized by that value. Thus if P is a phenotypic trait variable with values 1 and 2, while E is a discrete valued environmental variable, the distribution of E in a population is homogeneous when, for any value e of E, if 1/nth of the $P = 1$ phenotype is in $E = \mathrm{e}$, then 1/nth of the $P = 2$ phenotype is in $\mathrm{E} = \mathrm{e}$. So, for example, if we quantify local populations of the camellia weevil as having short ($E = \mathrm{s}$) or long ($E = \mathrm{l}$) rostra and camellia plants as having thin ($P = 1$) or thick ($P = 2$) pericarps, the distribution of E for the metapopulation of plants is homogenous when the proportion of thin pericarp plants beset by short rostrum weevils is the same as the proportion of thick pericarp plants beset by short rostrum weevils, and the proportion of thin pericarp plants beset by long rostrum weevils is the same as the proportion of thick pericarp plants beset by long rostrum weevils. Say that an environment is homogeneous if all environmental causes of W are homogeneously distributed. If E is a cause of W, but it is not the case that this connection is interactive with P as a context, then one of two things will be true. If the actual and contrasting environments are homogeneous, an explanatory appeal to E as an adapting condition for whichever value of P is

[4]Recollect that on the conception of causation here employed, causal relations hold between variables, and to say that E causes W is to say that by changing E one can change (the probability density over) W; hence there will be values of E that increase the value of W, and other values of E that decrease the value of W. Loosely, the causes of an outcome include both producers and preventers of that outcome.

fixed will violate DC (i.e. will not be doubly contrastive). Conversely, if the relevant contrasting environment is not homogeneous such an explanatory appeal will violate AEASE (i.e. will not be a selection explanation).

To see this, consider a population of ants invading a valley. Some ants suffer when soil moisture content is too high, while others are relatively resistant; call this trait P, with values 1 and 2 respectively for the resistant and non-resistant types. Ants of both types prey on seeds, and the non-resistant type is slightly more efficient at finding and processing seeds. Further, both types are equally sensitive to the lowest soil temperature in winter. Denote the winter-minimum soil temperature by E_1, with binned values from 0 (below freezing) to 5 (above 25 °C), and the local soil moisture content by E_2, again with values 0 (very dry) to 5 (very wet). Colony fecundity (i.e. the number of daughter queens sent out in a given year) is given by the equation $W = 20 + 3P + 2E_1 - 2E_2P$. Initially, the valley is unoccupied, with far more potential colony sites than offspring colonies, so all daughter colonies survive. Once the valley fills, fecundity is still determined by the equation for W, and old colonies are replaced by offspring colonies at random from among all offspring colonies, with a probability that is independent of the types of both old and offspring colonies, so as to hold the population size constant at K, whatever it may be, for the valley.

Suppose ants of both types initially invade the value under fairly good conditions—E_1 values are at 4 for every ant colony, and E_2 values are at 1 for every ant colony (call this Environment 1). But as time passes, these values fluctuate. Consider first uniform changes in E_1, with E_2 constant. Intuitively, uniform changes in E_1 can change the rate of evolution, because a change in E_1 will influence the reproductive success of both types. But since this influence will affect both types equally, it can change the magnitude of selection coefficients, but it cannot change which of the two types is fitter. Numbers may help.

Using fecundity as our measure of fitness, the fitness of any given colony is given by the equation $W = 20 + 3P + 2E_1 - 2E_2P$. Thus, for our ants in the initial environment, the resistant strain will have a fitness of $20 + 3(1) + 2(4) - 2(1 \cdot 1) = 29$ while the non-resistant strain has a fitness of $20 + 3(2) + 2(4) - 2(2 \cdot 1) = 30$. If we relativize to the fittest type, then initially the non-resistant strain has relative fitness 1 while the resistant strain has relative fitness $29/30 = 0.967$. If the environment is invariant, i.e. remains fixed at Environment 1, the non-resistant type will, slowly, go to fixation. Imagine now that winter minima decrease, so that $E_1 = 1$ for every colony in every generation (call this Environment 2). Then the absolute fitnesses become $20 + 3(1) + 2(1) - 2(1 \cdot 1) = 23$ and $20 + 3(2) + 2(1) + 2(1 \cdot 2) = 24$ for the resistant and non-resistant types respectively. This leads to a decrease in the relative fitness for the resistant type, to 0.958: selection is slightly stronger now, and the pace of evolution has quickened. But notice that we have not changed which of the two types is fitter. And in fact, there is no change in the value of E_1 that could produce such a reversal of fitness, for the very reason that E_1 causes W independently of, i.e. without interacting with, P.

Say that two environments differ uniformly (or that a change from one to the other is uniform) if, for each cause of fitness E and for all individuals i, j in the population, the difference between the E values for i in the two environments equals

the difference between the E values for j in the two environments (in consequence, uniform changes on a homogeneous environment result in another homogeneous environment). Appeals to the causal role of the variable E_1 in producing fitness in our ant population cannot contrastively explain why one type is fitter than another (and hence why that type evolves), so long as the contrasting environment (i.e. alternative distribution of E_1 over types) differs only uniformly from the actual distribution of E_1. Contrasting homogenous environments differ uniformly, and so relative to contrasting homogeneous environments, appeals to E_1 cannot explain why one rather than another phenotype evolves. This is a general feature of non-interactive causes: non-interactive environmental causes of fitness can explain why one rather than another contrasting phenotype evolves only by appeal to an actual or contrasting non-homogeneous environment, i.e. a situation in which one type differentially inhabits the better local habitats.

To see how an explanatory appeal to a non-uniform change in the environment works, suppose our study population moves from Environment 1 ($E_1 = 4$, $E_2 = 1$) to the following non-homogeneous environment (Environment 3): all $P = 1$ individuals are in $E_1 = 5$, $E_2 = 1$ habitats (so their realized fitnesses will be $20 + 3(1) + 2(5) - 2(1 \cdot 1) = 31$), while all $P = 2$ individuals are in $E_1 = 0$, $E_2 = 1$ habitats (so their realized fitnesses will be $20 + 3(2) + 2(0) - 2(1 \cdot 2) = 22$). Now the resistant type has the higher fitness, and (supposing this distribution of types to local habitats is constant), resistance evolves. This result, the evolution of $P = 1$ (resistance) rather than $P = 2$ (non-resistance) can be explained contrastively by appeal to the new rather than old distribution of E_1. But if any such explanation of the fixation of $P = 1$ also treats $P = 1$ as an adaptation to the environment, as characterized by the distribution of E_1, the explanation will be fallacious, and will violate AEASE.

Intuitively, the problem is that resistance has nothing to do with the success of the resistant type; that success rather derives from the fact that the resistant types more commonly experience higher winter minimum temperatures. More precisely, either the new, non-homogenous distribution of E_1 arises by chance or as a result of some other, behavioral, phenotype perfectly correlated with P. If the former, we have a case not of selection, but of drift (see Brandon 1990), in violation of AEASE. If the latter, then there is selection, but in favor of the behavioral phenotype that leads $P = 1$ individuals to favorable habitats and $P = 2$ phenotypes to unfavorable habitats; $P = 1$ has been sorted rather than selected. Again, AEASE has been violated, and in neither case is $P = 1$ an adaptation to the environment.[5]

Thus the Causal Condition fails because the structural, causal features it employs to sort adapting from non-adapting conditions are insufficiently explanatorily powerful. In particular, the Causal Condition can be satisfied by homogeneous

[5]Note that it matters here not at all whether the behavioral trait in question produces the non-homogenous environment by habitat selection, or by niche construction—in either case, it is not P, but the phenotypic cause of E_1 that is the immediate focus of selection, and hence the immediate locus of adaptation.

environmental conditions that cannot contrastively explain the evolution of one rather than another phenotype, except by appeal to contrasting non-homogeneous environments. Doubly contrastive selection explanations are impossible in such cases because the same phenotype would evolve under any uniform change in the environmental conditions. The same failure does not beset an alternative condition, which I will call the Interactive Causal Condition. According to this condition a phenotype $P = \mathrm{p}$ can be an adaptation to an environmental condition $E = \mathrm{e}$ only if E is an interactive cause of fitness with P.

To see how the Interactive Causal Condition avoids the problems besetting the Causal Condition, consider again our ants. Suppose, as before, our ants begin in Environment 1, with relative fitnesses of 0.967 for the resistant type and 1 for the non-resistant type. Now suppose the environment changes uniformly so that $E_2 = 4$, while $E_1 = 4$ remains constant (call this Environment 4). The absolute fitness of the resistant type then becomes $20 + 3(1) + 2(4) - 2(4 \cdot 1) = 23$ while that of the non-resistant type is $20 + 3(2) + 2(4) - 2(4 \cdot 2) = 18$. The relative fitnesses are now 1 for the resistant type and 0.783 for the non-resistant type: selection strongly favors resistance, and hence an appeal to the new (homogeneous) environment ($E_2 = 4$) explains the evolution of resistance. While uniform changes in E_1, can influence the rate at which evolution occurs but not its eventual outcome, uniform changes in E_2 can influence evolutionary outcomes, because they can change which type is fitter. Thus, in this scenario, we may say that resistance, $P = 1$, is an *adaptation to* the (homogeneous) distribution of E_2 in this sense: had that distribution been different (though still homogeneous), $P = 2$ would not have evolved to fixation. Because the contrasting environment (e.g. Environment 1) is homogenous, the different evolutionary outcomes our population would experience in the two scenarios (Environment 4 versus Environment 1) are a consequence of selection on P, rather than an artifact of chance or selection on some correlated trait.

4 Explanation, Inference, and Representation

Though more demanding than the Causal Condition, the Interactive Causal Condition remains a fairly minimal constraint on usage, and thus underwrites only limited explanatory power. Insofar as we use 'adaptation to' in ways that respect the Interactive Causal Condition, by 'the phenotype $P = 1$ is an adaptation to environmental condition $E_2 = 4$' we imply that E_2 is an interactive cause, with P, of fitness in the relevant population. It follows from this that a certain kind of counterfactual is true, namely that there is some (possible) homogenous environment $\boldsymbol{E}$ in which $P = 1$ evolves to fixation, and some other environment $\boldsymbol{E}'$ in which it does not, where $\boldsymbol{E}$ and $\mathbf{E}'$ differ only in a uniform change in the distribution of E_2. This makes possible doubly contrastive explanations of adaptive phenotypes, for example: P is fixed at 1 rather than at 2 because $E_1 = 4$ rather than 1. But this power imposes an epistemic cost. If one is to diagnose from observational data the fact that E_2 is an interactive

cause, with P, of W, two things must be true. First, the observations and models fitted to them must include measures of E_2, and second, E_2 must vary over individuals in the sample. I consider the two points in turn.

Population biologists employ two different ways of representing environments in mathematical models. The first, in effect, conditions on features of the environment which are thereby presupposed to be common to all members of the modeled population; the second explicitly introduces variables whose values denote the presence, absence, or magnitude of specific features of the environment. For example, wildlife biologists often employ summary measures of environmental quality, related to the expected rate of reproduction for a focal species occupying the environment (see Johnson 2007 for a review). Similarly, logistic growth models employ the parameters r and K, both of which are, in this sense, measures of habitat quality. In the same way, in simple population genetic models employing a single fitness or selection coefficient for each genotypic class, the fitnesses or selection coefficients are in effect a measure of habitat quality, from the perspective of each genotype. More complicated models, e.g. those employing contextualized fitnesses (sensu Kerr and Godfrey-Smith 2002) or niches (Levene 1953), specify fitnesses for genotypes in more narrowly circumscribed environments. However, such measures of environmental quality, whether or not they are niche or genotype specific, do not specify or measure those features of the environment that are causally responsible for the differences in fitness, intrinsic rate of increase, or carrying capacity. When such models are fitted to data they can, sometimes, be used to diagnose the presence of an interactive environmental cause of fitness (this for example is one thing Brandon's phytometer studies do; see Brandon and Antonovics 1996). But neither the models nor the component measures of environmental quality can by themselves be used to identify which features of the environment are in fact interactive causes of fitness. Thus, while such models allow one to identify a phenotype as an adaptation, they will not permit one to identify the adapting conditions to which the phenotype is, in fact, an adaptation.

This limitation arises in the following way. Suppose we gather data which include measures on individuals of components of fitness (fecundity, survival, or what have you), measures of individual phenotype (size, height, coloration, or what have you) and location (position on a transect or grid, say), but not specific values for specific environmental conditions obtaining at that location. We can, for each location, calculate type-specific mean values for our measures of fitness and note differences in them. And we can look, in particular, for pairs of locations in which the ordinal relation between type fitnesses is reversed (i.e. one type is fitter in one location while another is fitter in a different location). One way to account for such reversals is to appeal to some changing environmental condition which is, with phenotype, an interactive cause of fitness. But, necessarily, we will have no evidence about just what this condition is, since very many environmental conditions will differ between locations and we will have measured none of them on individual organisms. And even the inference that such an interactive environmental cause varies in value over locations is suspect, for there is another way to explain such reversals in the ordinal

relations between type fitnesses. If there are two or more non-interactive causes of fitness, and they have non-homogeneous distributions over the types in each location, so that one type experiences better conditions in one location, the other better conditions in the second location, the fitness relations may be reversed without the presence of an interactive cause. To rule out this sort of case it is generally necessary to actually measure the relevant environmental conditions, either in the wild or in experimental contexts in which the conditions of interest are controlled (or manipulated).

It is now common practice to introduce explicit measures of particular environmental conditions into one's models in evolutionary ecology; it is becoming common in population demography and population management as well (see Caswell 2001; Guissan and Thuiller 2005). Given joint measures on environmental conditions and components of fitness, it is possible to test from observational (rather than experimental) data hypotheses about the causal influence of specific environmental variables on fitness, and if such a causal connection is found, to further test the hypothesis that the dependence is interactive with one or another phenotypic feature. The first epistemic price of such tests is the requirement that environmental conditions actually be measured and represented explicitly in one's model. Thus, the judgment that a phenotype is an adaptation to some specific set of environmental conditions requires that those conditions be explicitly represented as the value of a measurable (or at any rate estimable) variable, and can be warranted only by data that include measures of those variables.

What is more, that a given environmental variable is a cause of fitness can be determined from observational data only if those conditions vary over sample membership. While it is true that correlation is not causation, it is also true that the one statistical signature characteristic of the absence of a causal dependency is statistical independency (e.g. the absence of a correlation). Associations between variables can be detected in a sample only if both variables vary over the sample membership—without variance there can be no covariance. Hence, the hypothesis that an environmental variable causally influences fitness cannot be tested against data in which the environmental variable is constant. This is the second epistemic price of tests for interactive causal dependencies between environment, phenotype, and fitness. It is not trivial, for there may be minimal variation over individuals or sub-populations of a species with respect to large scale environmental features; this in turn may require that data be gathered over long time periods so that the requisite variation will appear as temporal, inter-generational variation rather than intra-generational variance in environmental conditions. There is a conceptual price here as well. It is sometimes thought that selection requires a common environment.[6] Whether or not this is so, the kind of selection that drives adaptation to particular environmental conditions will be undiscoverable if all organisms in the population are subject to identical environmental conditions. We will recur to this point below.

[6]This is an implicit consequence of any view that pairs dispositional fitnesses with the standard view that selection requires heritable differences in fitness. It is sometimes made explicit, as e.g. in (Brandon 1990).

5 Limits and Extensions

I said above that the Interactive Causal Condition imbues appeals to 'adaptation to' with a limited explanatory power, but did not elaborate on the nature of the limits. I do so now. If we confine our usage of 'adaptation to' so as to respect the Interactive Causal Condition, it will follow from the claim that $P = 1$ is an adaptation to $E_2 = 4$ that there is some (possible) homogenous environment $\boldsymbol{E}$ in which $P = 1$ evolves to fixation, and some other environment $\boldsymbol{E}'$ in which it does not, where $\boldsymbol{E}$ and $\boldsymbol{E}'$ differ only in a uniform change in the distribution of E_2, and such counterfactuals make possible the doubly contrastive explanations of the form 'the phenotype is fixed at this value rather than that because the environment was characterized by this condition rather than that.' But in fact very few populations inhabit homogeneous environments, and for any phenotype of interest it is likely that there are some adapting conditions that do not have a homogeneous distribution, over time and space, in the adapting population.[7] This raises some puzzles about which potential contrasting environments are explanatorily relevant. To explore the implications of such non-homogeneous distributions, we need to expand our conception of the environment in a number of ways.

First, we need two distinct conceptions of the environment occupied by an individual organism. Let $\boldsymbol{E}$ be a vector $< E_1, E_2, \ \ldots \ E_n >$ of environmental variables. I will say that a set of values for each of the E_i, as measured on an individual organism j (thus, $\boldsymbol{E}(j) = \mathbf{e} = <E_1(j) = e_1, E_2(j) = e_2, \ \ldots \ E_n(j) = e_n>$), comprises the *narrow individual environment* occupied by individual j. Let $\mathbf{P}_j(\boldsymbol{E})$ be a probability density over $\boldsymbol{E}$ for j, characterizing for each possible narrow individual environment the chance that individual j comes to occupy that environment. I will call $\mathbf{P}_j(\boldsymbol{E})$ a *wide individual environment*. Second, we need two distinct conceptions of the environment occupied by a population. Let $\boldsymbol{\mathcal{F}}_p(\boldsymbol{E})$ be a frequency distribution of individuals in population p over narrow individual environments; I will call $\boldsymbol{\mathcal{F}}_p(\boldsymbol{E})$ the narrow population environment. Finally, let $\mathbf{D}_p(\boldsymbol{E})$ be a probability density over all possible frequency distributions $\boldsymbol{\mathcal{F}}_p(\boldsymbol{E})$. I will call $\mathbf{D}_p(\boldsymbol{E})$ a *wide population environment*. It will be helpful in what follows to relativize population environments to phenotypically specified classes in a population. To that end I will write $\boldsymbol{\mathcal{F}}_c(\boldsymbol{E})$ to represent the frequency distribution over narrow individual environments of individuals in the cth class of the population p, and $\mathbf{D}_c(\boldsymbol{E})$ to represent the probability density over such frequency distributions.

Let us reconstruct the explanatory import of the 'adaptation to' locution, given the Interactive Causal Condition, but now employing the above conceptions of 'environment.' Suppose that every individual in the population occupies an identical wide individual environment, with $\boldsymbol{E}(i)$ probabilistically independent of $\boldsymbol{E}(j)$ for any two individuals i and j in the population. It follows that for any pair of similarly sized classes c and c′ in the population, $\mathbf{D}_c(\boldsymbol{E}) = \mathbf{D}_{c'}(\boldsymbol{E})$, even though

[7] Marshall Abrams has in conversation pressed various critical points regarding actual non-homogeneous distributions of adapting conditions. Though what I say will doubtless leave him unsatisfied, his worries influenced some of what follows and I thank Marshall for pressing them.

individuals may differ in their narrow environments, and hence even though the frequency distributions over the classes c and c′ might differ (i.e. $\mathbf{F}_c(\boldsymbol{E}) \neq \mathbf{F}_{c'}(\boldsymbol{E})$). When $\mathbf{D}_c(\boldsymbol{E}) = \mathbf{D}_{c'}(\boldsymbol{E})$, I will say that the wide population environment $\mathbf{D}_p(\boldsymbol{E})$ is homogeneous. Say that $\mathbf{D}_p(\boldsymbol{E})$ and $\mathbf{D}'_p(\boldsymbol{E})$ differ uniformly if and only if both are homogeneous. For cases in which the actual wide population environment is homogeneous, the actual wide population environment and some subset (not necessarily proper) of the remaining homogeneous environments will comprise a set of reasonable contrasting environments. What is required is that the actual and contrasting environment(s) differ in the expected (homogeneous) distribution of one or more interactive environmental causes. Given such homogeneous wide population environments, the explanation of the fixation of P at 1 rather than 2 can appeal to the fact that the adapting conditions have expected values determined by $\mathbf{D}_p(\boldsymbol{E})$ rather than different expected values, determined by a different homogeneous wide population environment $\mathbf{D}'_p(\boldsymbol{E})$, where, had $\mathbf{D}'_p(\boldsymbol{E})$ been the actual wide population environment, we would not have expected P to fix at 1.

The forgoing recapitulation presupposes that the adapting population occupies a homogeneous wide population environment, and that supposition will, for various reasons, often not be satisfied. Such cases raise a number of puzzles; before turning to them, it is worth pausing to expand on the conceptual puzzle noted at the end of Sect. 4. The idea that selection requires a common environment is certainly implicit in standard readings of relatively simple population genetic models (the charge of illicit averaging, for example, depends on it). Explanations of adaptations to particular adapting conditions, as above, also depend on a shared environment, when the environment is understood as a wide population environment—i.e. a probability density over the proportion of each geno- or phenotype characterized by specific values for environmental causes of fitness. But such explanations are available even when the actual proportion of types in particular conditions varies quite radically from type to type. Hence, the relevant classes may differ in their narrow population environments, and individual organisms may differ in both their wide and narrow individual environments (and in fact, as noted at the end of Sect. 4, the last such difference is essential for the possibility of detecting the kind of selection that drives adaptation to particular conditions). Hence, to make sense of adaptation to specific environmental conditions, we must relinquish the idea that selection requires shared environmental conditions. This will be relevant below.

Let us now recur to the case in which the actual wide population environment is not homogeneous. The Interactive Causal Condition insists that for a phenotype to be an adaptation to some environmental condition, the phenotype must be encoded as some value of a trait variable P_a and the environmental condition as the value of an environmental variable E_a where P_a and E_a are connected to fitness W by a causal structure of the kind depicted in Fig. 1.

Here f is the functional dependence of W on E_a, which is controlled by P_a, and P_a may or may not also directly influence W (represented by the dotted arrow). One important way in which homogeneity can fail involves a probabilistic association between environmental conditions and phenotype, represented by a double-ended edge in Fig. 2.

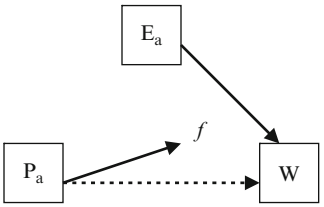

Fig. 1 A simple interactive dependency. A simple causal structure satisfying the interactive causal condition; the *dashed edge* represents a possible but non-essential secondary, direct path by which P_a may contribute to W independently of E_a

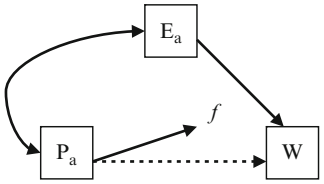

Fig. 2 An interactive dependency with an unexplained association between causes. E_a and P_a interactively cause W, and are themselves associated in virtue of an unspecified causal connection between them and/or some unmeasured common cause

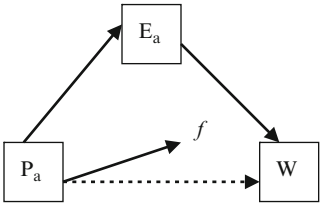

Fig. 3 Interactive phenotypic cause of an environmental condition. P_a both causes E_a and interactively controls the influence of E_a on W

This sort of association may arise in three ways: E_a may cause P_a, P_a may cause E_a, or they may share some common prior cause. Phenotypic plasticity offers an example of the first kind of case, and I discuss it in my (2011). The third kind of case involves more complexities than I have space to deal with, but any resolution of such cases requires as a background some view of the second kind of case, which I will therefore briefly address here. Graphically, the second kind of case can be represented by Fig. 3.

This structure arises whenever the phenotypic composition of a population causally influences (or in the case of frequency dependent selection, constitutes) an adapting condition. Such structures also represent one kind of 'extended phenotype' (Dawkins 1999) and one kind of 'constructed niche' (Odling-Smee et al. 2003), in

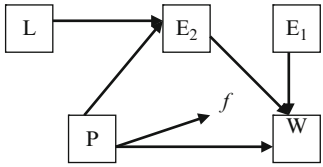

Fig. 4 Causal structure with a landscape variable. The landscape variable L causes E_2, which is both an effect of P and an interactive cause with P of W

that they occur whenever the phenotype of an individual causally influences the narrow individual environment occupied by the individual.[8]

Consider again our ant population, but now attending to variations in the valley's landscape. At least initially, the landscape includes many more potential colony sites, or locations, than there are progeny seeking such sites. For ease, suppose that all locations in the landscape are characterized by a similar lowest annual soil temperature, encoded by the value 2 for the variable E_1. We will, however, allow E_2 to vary uniformly, so that 1/6 of the locations have a soil moisture content falling in the range denoted by $E_2 = 0$, 1/6 have an E_2 value of 1, and so on. We will encode this shared landscape-level environmental property (the frequency distribution of E_2 values over locations) with the variable L; every individual in the population shares the same value for L because all inhabit the same landscape. The reproductive success (new colonies established) is governed, as in the examples above, by the equation $W = 20 + 3P + 2E_1 - 2E_2P$. But now we will not only let P directly cause fitness (as per the second term in the equation for W), and govern the contribution E_2 makes to fitness (as per the fourth term in the equation for W), but also further suppose that P causes E_2: new queens with the $P = 2$ phenotype simply choose the first location they happen across, while those with the $P = 1$ phenotype are slightly more likely to choose drier locations. In particular, let $\Pr(E_2(\mathrm{i}){=}\mathrm{n}){=}1/6{+}b$, where the preference for drier habitats is quantified by b, with $b{=}(P(\mathrm{i})\text{-}2)(n\text{-}2.5)/15$. Hence, the probability that a $P = 2$ queen occupies a location with $E_2 = \mathrm{n}$ is 1/6, for all values of n. But the probability that a $P = 1$ queen occupies a location with $E_2 = \mathrm{n}$ varies: when $\mathrm{n} = 0$, this probability is at its highest, $1/6 + 1/6 = 10/30$; for $\mathrm{n} = 1$ the probability is $1/6 + 1/10 = 8/30$; for $\mathrm{n} = 2$ the probability is $1/6 + 1/30 = 6/30$; for $\mathrm{n} = 3$ the probability is $1/6 - 1/30 = 4/30$; for $\mathrm{n} = 4$ the probability is $1/6 - 1/10 = 2/30$; and for $\mathrm{n} = 5$ the probability is $1/6 - 1/6 = 0$. Thus, resistant ($P = 1$) ants enjoy a kind of double advantage—they are less likely to find themselves in especially wet soil conditions, and better able to deal with those conditions when they do occupy them. The causal structure for the system, including the landscape variable L, is given in Fig. 4.

[8]Both ideas remain largely metaphorical, and hence conceptually quite rich. Consequently, it is not the case that Fig. 3 represents the causal structure operative in any realization of either metaphor. It is rather that any system for which the structure in Fig. 3 holds is a system in which the E_a value counts as an extended phenotype, in one sense of that term, and as a constructed niche, in one sense of that term.

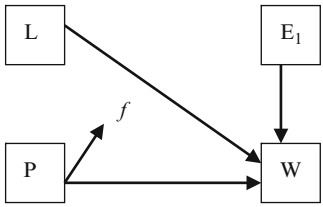

Fig. 5 Causal structure with a landscape variable, omitting mediating environmental variable. The causal structure from Fig. 4, with E_2 omitted; L is now an interactive cause with P of W

The respective fitnesses of the two types, calculated as expected per-capita reproductive success, are then 24.33 and 20 for the $P = 1$ and $P = 2$ phenotypes. Barring drift, the population fixes at $P = 1$, and clearly it does so as a result of selection on P, and for $P = 1$ in particular. And clearly, P is not an adaptation to $E_1 = 2$, for $P = 1$ would fix under any uniform change in the distribution of E_1. E_1 influences only the rate at which evolution occurs, not its eventual outcome. We might, employing the Interactive Causal Condition as not only necessary but sufficient, claim that P is an adaptation to E_2, since P and E_2 are interactive causes of W. And there is at least the following to say for that option: had the distribution of E_2 been different in specifiable ways, $P = 1$ would not have evolved (barring drift); what is more, some of those alternative distributions of E_2 are homogeneous. If, for example, $\mathbf{D}_{p=2}$ were biased in the same way $\mathbf{D}_{p=1}$ is, then $P = 2$ would evolve; similarly, were the wide population environments for both the $P = 1$ and the $P = 2$ phenotypes Environment 1, then again $P = 2$ would fix.

On the other hand, the *actual* distribution of E_2 over phenotypic classes, i.e. the narrow population environment, will not be homogeneous in any generation (again, barring drift in the form of sampling error). Moreover, these biases are not accidental, which suggests first that the net effect of the bias should not be chalked up to drift, and second that we might be attending to the wrong causal variables. We might then prefer to say that $P = 1$ is not an adaptation to E_2, but rather to whatever causes the biased distribution of E_2. But—and this is perhaps even worse—P is itself just such a cause, since the P value of any individual in our population is a cause of the E_2 value for that individual. And it seems, at least, disconcerting to say of a phenotype that it is an adaptation to an environmental condition that it itself causes. Here again the notions of an extended phenotype and constructed niche seem relevant.

One might, instead, hold that $P = 1$ is an adaptation not to E_2, but to L, since L is the other cause of the biased distribution of E_2 over phenotypes. What is more, this is consistent with the Interactive Causal Condition. For if we simply ignore the variable E_2, the causal structure in the system reduces to that in Fig. 5, and with respect to that structure, L is an interactive cause, with P, of W. On the other hand, if we say only that $P = 1$ is an adaptation to L, our description of the situation omits the explanatorily relevant fact that P not only influences the value that E_2 takes, it influences as well the degree to which that value, whatever it is, in turn influences W.

Other options are possible. One might hold that $P = 1$ is an adaptation to each of L and E_2 individually—this follows if we insist that the Interactive Causal Condition is both necessary and sufficient for adaptation to a particular environmental

condition. Somewhat differently, we could hold that $P = 1$ is an adaptation to the conjunction of L and E_2. Either option entails that the selection pressure driving adaptation to a particular environmental condition can occur even when the actual distribution of the relevant environmental condition is not homogeneous, i.e. even when environments are not shared. Indeed, either view implies that selection does not require even that phenotypes share wide population environments.

Very differently, one might insist that adaptation to a particular environmental condition requires not only that the Interactive Causal Condition be satisfied, but requires as well that the putative adapting condition have a homogeneous distribution. Motivation for this view can be found in the idea that selection requires a shared environment, for that idea is preserved on this conception of 'adaptation to' by the requirement that adapting conditions have an actual homogenous distribution over phenotypes. On the other hand, any such requirement will impose exacting epistemic demands on explanatory appeals to 'adaptation to,' and will severely restrict the number of adaptations that are, in fact, adaptations to particular adapting conditions. Nonetheless, the requirement would nicely circumvent the above quandaries.

My own inclination is to regard the Interactive Causal Condition as a necessary and sufficient constraint on 'adaptation to,' i.e. to say that a phenotype is an adaptation to a particular environmental condition just in case the phenotype is an adaptation and an interactive cause with the condition of fitness. But I don't hope to defend my preferences here. In my view there is no independent fact about what P is or is not an adaptation to, with respect to which our analysis can get things right or wrong. Rather, there is a set of causally and statistically distinguishable ways in which phenotypes and environmental conditions can interact to produce fitnesses, and we can choose to use 'adaptation to' in ways that respect more or fewer of those distinctions, with concomitant implications for the explanatory power and epistemic demands of the 'adaptation to' locution. Our choices should, however, be informed, and I hope that the forgoing has limned some of the consequences of some of the available choices. I do wish, however, to close in the next section by attending to one further epistemic consequence of the Interactive Causal Condition.

6 Measuring Selection

There are several broadly distinguishable ways of measuring the strength of selection. One set of measures quantifies the strength of selection by the evolutionary change, *i.e.*, the change in type frequencies, engendered by a selection process.[9] Among such measures are selection differentials and the response to selection, for example. Unfortunately, these statistics are not well suited as measures of

[9]Though heritable variation in fitness is commonly taken to be either a necessary or a necessary and sufficient condition for selection (e.g. Lewontin 1970), on some views selection just *is* differential reproductive success (see e.g. Eldredge 1986 or Grant 1991) or differential fitness of types (e.g. Schluter 1988).

the strength of selection driving adaptation to particular environmental conditions, because when treated as measures of selection they confound the effect of the adapting conditions with the effect of other environmental features that, while influencing survival and reproductive success, are not adapting conditions.

To illustrate, consider again our ants in Environment 1 ($E_1 = 4$, $E_2 = 1$), and recollect that success is governed by the equation $W = 20 + 3P + 2E_1 - 2E_2P$. The respective fitnesses for $P = 1$ and $P = 2$ individuals are 29 and 30 respectively. Supposing we begin with equal numbers of each type, the corresponding response to selection (given by the difference between mean phenotype among parents and offspring) is then $89/59 - 30/20 = 0.008$. Consider an alternative environment in which E_2 remains 1 but $E_1 = 0$. Now the fitnesses are 21 and 22, respectively, and the selection differential is $65/43 - 30/20 = 0.012$. In the second case, evolution proceeds somewhat more quickly, and this yields a higher estimate of the strength of selection. But the net effect of E_2, the only adapting condition, is the same in the two cases.

A second way of measuring the strength of selection is to consider the differences in type-specific rates of survival and reproductive success generated by (or on some views constitutive of) the selection process. Selection coefficients and fitness differences are standard examples. But such measures suffer from exactly the same flaws. As measures of the strength of selection simpliciter they may be fine, but as measures of the strength of selection driving adaptation to particular conditions, they are confounded. This can be seen in the example above. When $E_1 = 4$ and $E_2 = 1$, the type fitnesses are 29 and 30, yielding relative fitnesses (dividing through by the maximal fitness) of $w_1 = 0.967$ and $w_2 = 1$, for a selection coefficient $S = 0.033$. Although the effect of E_2 on reproductive success remains unchanged in the alternative environment $E_1 = 0$, $E_2 = 1$, the relative fitnesses and selection differentials have changed. The fitnesses are now 21 and 22, with relative fitnesses $w_1 = 0.955$, $w_2 = 1$, and the selection coefficient is $S = 0.045$. Again, our measures of the strength of selection are responsive to changes in E_1, when to measure the strength of selection driving the fixation of $P = 1$ as an adaptation to E_2 they ought not be.

A third kind of measure of the strength of selection is closer to what we require. This third way of tracking the strength of selection identifies the strength of selection with some measure of the association between phenotype and (components of) fitness, where the measure of association is in turn interpretable as the strength of the causal influence of the phenotypic variable on fitness. The use of linear regression methods in evolutionary ecology (c.f. Roughgarden 1979, though the methods have been employed at least since the 1960s), and of selection gradients (partial regression coefficients) in population genetic treatments (c.f. Lande and Arnold 1983), are illustrations. Prospects here are more promising, but there are substantive methodological problems to be resolved. I will illustrate just one.[10]

[10]The interactive causal connection between adapting conditions and fitness is of particular concern, but space prevents any useful elaboration here.

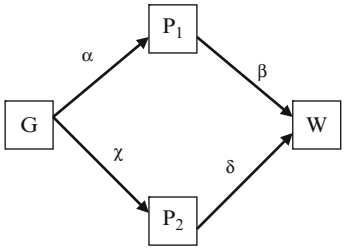

Fig. 6 Two phenotypic causes of fitness with a shared common cause. P_1 and P_2 both cause W, and share a measured common cause G; α, β, χ and δ are path coefficients measuring the conditional association between pairs of variables

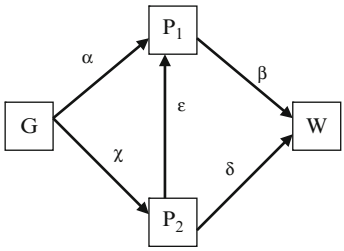

Fig. 7 Complex structure with a common cause. P_2 causes P_1, and both cause W; in consequence P_2 influences W directly, and indirectly through its influence on P_1

The idea behind selection gradients was to sort out the extent to which different phenotypes that might share a common genetic cause were individually influencing fitness, and hence evolutionary outcomes. Put another way, given that fitnesses vary among phenotypic classes as defined by values of P_2, as in Fig. 6 below, how much of that difference in fitness is explained by selection on P_2, and how much by selection on the associated trait variable P_1? Given the relatively simple structure in Fig. 6 (and assuming linear dependencies), one could simply consider the total association (the correlation or regression coefficient) between P_2 and W. But that total association will be proportional to $\chi\alpha\beta + \delta$, i.e. it will confound the effect of P_2 with an association induced by the effect of P_1 and the fact that P_1 and P_2 share a common genetic cause. In such simple structures, one can produce an unconfounded estimate of the path coefficient δ by taking the partial regression coefficient between P_2 and W, conditioning on P_1. If we identify the strength of selection on P_2 with its causal effect on fitness, and measure that effect by estimating δ, we seem to be on safe ground.

We might then try something similar with respect to environmental conditions. If we have the structure represented in Fig. 4 above, we could measure the strength of selection driving adaptation to E_2 by the association between E_2 and W conditioning on E_1 (a non-interactive cause of fitness, and therefore not an adapting condition). Problems arise however. Consider the structure represented in Fig. 7. Here, P_2 actually causes P_1, and so the total causal influence of P_2 on fitness is really best represented by $\delta + \varepsilon\beta$; here δ represents the direct influence of P_2 on W, and $\varepsilon\beta$ the indirect influence of P_2 on W through P_1. Here, the partial regression of P_2 on W controlling for P_1 will yield an estimate of δ. Thus, the selection gradient is a *biased* estimate of the causal influence of P_2 on W, and hence of the strength of selection on W.

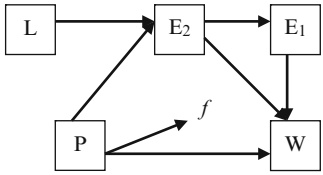

Fig. 8 Complex structure relating adapting condition to fitness. E_2 causes W directly and interactively with P, and indirectly but not interactively through its influence on E_1

These worries, and others, beset the extension of these measures to environmental conditions. If, as in Fig. 8, the adapting condition influences W directly and indirectly by way of some further environmental variable, conditioning on that variable will lead to biases in estimates of the strength of selection. There are therefore methodological problems associated with the idea that the strength of selection to particular environmental conditions should be measured by the causal influence of those conditions on fitness.

These difficulties are, nonetheless, less pressing than those confronting alternatives, and hence provide reasons for preferring a causal measure of the strength of selection driving adaptation to over the much less sensitive statistical measures using the response to selection or fitness differences.

7 Summary

The arguments above show, I believe, that if 'adaptation to' is to carry more explanatory power than 'adaptation,' something like the Interactive Causal Condition will have to be endorsed. They show further that any such endorsement of the Interactive Causal Condition carries with it a commitment to explicit measures of environmental conditions as against measures of habitat quality. Finally, I hope the subsequent sections have suggested the range of quite intricate statistical and causal relationships between environment, phenotype, and fitness to which we might wish to attend. Failure to be clear about which structures we intend to denote when speaking of 'adaptation to' will lead to explanatory incoherence. And failure to explicitly model those structures will lead to biased or confounded estimates of the strength of selection.

References

Brandon, Robert N. 1990. *Adaptation and environment*. Princeton: Princeton University Press.

Brandon, Robert N., and Janis Antonovics. 1996. The coevolution of organism and environment. In *Concepts and methods in evolutionary biology*, ed. Robert N. Brandon, 161–178. Cambridge: Cambridge University Press.

Caswell, Hal. 2001. *Matrix population models*. Sunderland: Sinauer.

Dawkins, Richard. 1999. *The extended phenotype*. Oxford: Oxford University Press.

Eldredge, Niles. 1986. Information, economics, and evolution. *Annual Review of Ecology and Systematics* 17: 351–369.

Glymour, Bruce. 1998. Contrastive, non-probabilistic statistical explanations. *Philosophy of Science* 65: 448–471.

Glymour, Bruce. 2007. In defense of explanatory deductivism. In *Causation and explanation*, ed. Joseph Campbell, Michael O'Rourke, and Harry Silverstein, 133–154. Cambridge, MA: MIT Press.

Glymour, Bruce. 2011. Modeling environments: Interactive causation and adaptation to environmental conditions. *Philosophy of Science* 78: 448–471.

Grant, Peter R. 1991. Natural selection and Darwin's finches. *Scientific American* 265: 82–87.

Guisan, Antione, and Wilfried Thuiller. 2005. Predicting species distribution: Offering more than simple habitat models. *Ecology Letters* 8: 993–1009.

Heywood, James. 2010. Explaining patterns in modern ruminant diversity: Contingency or constraint? *Biological Journal of the Linnean Society* 99: 657–672.

Hunt, Kevin D. 1994. The evolution of human bipedality: Ecology and functional morphology. *Journal of Human Evolution* 23: 183–202.

Jaenike, John. 1978. An hypothesis to account for the maintenance of sex within populations. *Evolutionary Theory* 3: 191–194.

Johnson, Matthew D. 2007. Measuring habitat quality: A review. *The Condor* 109: 489–504.

Kerr, Benjamin, and Peter Godfrey-Smith. 2002. Individualist and multi-level perspectives on selection in structured populations. *Biology and Philosophy* 17: 477–517.

Lande, Russell, and Stevan J. Arnold. 1983. The measurement of selection on correlated characters. *Evolution* 37: 1210–1226.

Lennox, James G., and Bradley E. Wilson. 1994. Natural selection and the struggle for existence. *Studies in History and Philosophy of Science* 25: 65–80.

Levene, Howard. 1953. Genetic equilibrium when more than one ecological niche is available. *American Naturalist* 87: 331–333.

Levin, Donald. 1975. Pest pressure and recombination systems in plants. *American Naturalist* 109: 437–451.

Lewontin, Richard C. 1970. The units of selection. *Annual Review of Ecology and Systematics* 1: 1–18.

Maynard Smith, John. 1976. What determines the rate of evolution? *American Naturalist* 110: 331–338.

McFadden, Bruce. 1992. *Fossil horses*. Cambridge: Cambridge University Press.

Odling-Smee, F. John, Kevin N. Laland, and Marcus W. Feldman. 2003. *Niche construction: The neglected process in evolution*. Princeton: Princeton University Press.

Pearl, Judea. 2000. *Causality*. Cambridge: Cambridge University Press.

Roughgarden, Jonathan. 1979. *Theory of population genetics and evolutionary ecology*. New York: Macmillan.

Schluter, Dolph. 1988. Estimating the form of natural selection on a quantitative trait. *Evolution* 42: 849–861.

Spirtes, Peter, Clark Glymour, and Richard Scheines. 2000. *Causation, prediction, and search*, 2nd ed. Cambridge, MA: MIT Press.

Toju, Hirokazu, and Teiji Sota. 2006. Imbalance of predator and prey armament: Geographic clines in phenotypic interface and natural selection. *American Naturalist* 167: 105–117.

van Fraassen, Bas. 1980. *The scientific image*. Oxford: Oxford University Press.

Wheeler, Peter E. 1991. The influence of bipedalism on the energy and water budgets of early hominids. *Journal of Human Evolution* 21: 117–136.