Biology and Philosophy (2007) 22:115–130
DOI 10.1007/s10539-005-9010-x

Fitness and Propensity's Annulment?

MARSHALL ABRAMS
Department of Philosophy, George Washington University, Washington DC, USA (e-mail: mabrams@duke.edu)

Received 29 October 2004; accepted in revised form 12 October 2005

Key words: Determinism, Fitness, Probability, Propensity

Abstract. Recent debate on the nature of probabilities in evolutionary biology has focused largely on the propensity interpretation of fitness (PIF), which defines fitness in terms of a conception of probability known as "propensity". However, proponents of this conception of fitness have misconceived the role of probability in the constitution of fitness. First, discussions of probability and fitness have almost always focused on *organism effect probability*, the probability that an organism and its environment cause effects. I argue that much of the probability relevant to fitness must be *organism circumstance probability*, the probability that an organism encounters particular, detailed circumstances within an environment, circumstances which are not the organism's effects. Second, I argue in favor of the view that organism effect propensities either don't exist or are not part of the basis of fitness, because they usually have values close to 0 or 1. More generally, I try to show that it is possible to develop a clearer conception of the role of probability in biological processes than earlier discussions have allowed.

Introduction

There have been active debates in recent years about the nature of probabilities in evolutionary biology. These focus mainly on the nature of the probabilities thought to be the basis of fitness. The most common view seems to be that fitness is based on one particular kind of objective probability known as "propensity". However, proponents of this "propensity interpretation of fitness (PIF) have misconceived the role of probability as a foundation for fitness" because they have not looked closely enough at the concept of propensity and at how propensities might arise in biological contexts.

Advocates of the PIF have nearly always assumed that the probability which is the basis of fitness is *organism effect probability*: the probability that an organism and its environment cause effects. This should be contrasted with *organism circumstance probability*: the probability that an organism encounters particular, detailed circumstances within an environment, circumstances which are not the organism's effects. (These terms will be defined more carefully in Section 4.)

This paper has three related goals. The first is to argue that much of the probability relevant to fitness is organism circumstance probability. The second is to argue in favor of the view that propensities either don't exist or are not part of the basis of fitness, because they nearly always take the values 0 or

1. This conclusion will help to support the view that organism circumstance probability is the sole basis of fitness, but not all of my arguments that organism circumstance probability is important will turn on it.

The third goal of the paper is more general. I believe that some of the debates about the PIF have been hindered by a lack of clarity about propensity. I believe that a some-what clearer conception of biological processes and the ways in which probability enters into them is possible, and I take steps toward such a conception here. To this end, I repeatedly consider probability and specifically propensity in non-biological contexts; this helps to make certain ideas clear before turning to their biological applications.

I'll proceed as follows. After describing the propensity interpretation of fitness I'll use minimal assumptions about propensity to elucidate the role that it plays in biological processes. I'll argue that organism effect propensity doesn't play a significant role in governing outcomes during the life of an organism. I will however, explore the one way that I think organism effect propensity *might* play a role in determining fitness. My arguments suggest that if fitness is to be defined in terms of probability, this probability must be, at least in large part, organism circumstance probability. After drawing this conclusion, I explore relationships between organism circumstance probability, organism effect probability, and propensity.

The propensity interpretation of fitness

The concept of fitness plays a crucial role in evolutionary biology. It appears both in informal conceptions of biological processes and in specific biological models. That a genotype or phenotype A has greater fitness than another B in the same population is often taken to be a necessary condition for natural selection for A over B to take place, and fitness differences are an essential part of explanations of adaptation in terms of natural selection.

However, there has been and continues to be controversy over the meaning of the word "fitness". Sometimes fitness is defined in terms of the actual number of offspring that an organism has, or in terms of the average number of offspring of certain individuals in a population. It's often been argued that a definition of fitness in terms of the actual offspring of an individual or group cannot be the fundamental sense of the term, though. If it were, fitness and natural selection would not be able to explain evolution. To say that organisms in a population have certain numbers of offspring with given genotypes mathematically implies the subsequent genotype frequencies in the population, while providing no explanation of why the frequencies turned out that way.

A little more than 25 years ago, Brandon (1978) and Mills and Beatty (1979) announced the marriage of fitness and a member of the propensity family: propensity for outcomes during an organism's life. The authors suggested an elegant solution to the problem of fitness, arguing that fitness should be rooted

in objective probabilities that are not mere actual frequencies. They defined fitness as proportional to the expected number of offspring, the probability-weighted average number of offspring:

$$\sum_{k=0}^{\infty} k \cdot \mathcal{P}(O = k). \tag{1}$$

This is the sum, for all possible offspring counts k, of the products of k and the probability $\mathcal{P}(O = k)$ of an organism having k offspring. (The variable k takes every nonnegative integer value starting from 0, but the probabilities are zero for large k since litter sizes and lifetimes are bounded.) Brandon and Mills and Beatty suggested that the objective probabilities $\mathcal{P}(O = k)$ be conceived of as 'propensities' – objective probabilistic dispositions. Standard conceptions of dispositions tie them to physical facts about an object or situation which determine that a certain effect will definitely be produced in certain – perhaps counterfactual – situations. The existence of a propensity, however, does not determine that the effect in question must definitely occur in the specified circumstances. Instead the effect is associated with a real number between 0 and 1, indicating the strength of the tendency to produce the effect. One advantage of this way of understanding fitness is that it seems to make sense of fitness's apparent explanatory power. Just as salt's disposition to dissolve in water (based in salt's atomic structure) helps to explain why a lump of salt dissolved, so differences between the dispositions of phenotypes to produce various numbers of offspring help to explain why some genes come to be more common in subsequent generations. Following Mills and Beatty's suggestion, the new definition of fitness in terms of propensity (and some variations on it) have become known as the PIF. Note that the propensities involved are a variety of what I called "organism effect probabilities" above; they are propensities for an organism's characteristics to cause effects in conjunction with aspects of the environment.

Although it's not clear that fitness can always be defined as expected number of offspring (Beatty and Finsen 1989; Brandon 1990; Sober 2001), it will simplify matters greatly to talk as if fitness should definitely be computed as expected number of offspring. Also note that I will assume a certain kind of event ontology which is convenient though not essential: Probabilities are assigned to outcomes of trials. A trial is an 'event token' or 'token event', i.e. a particular physical occurrence such as the toss of the die at t. Each trial event token instantiates various 'event types', such as being a toss of a particular kind of die (cf. Davidson 1967; Bennett 1988). An outcome is an event type – for example, the property that a toss of a die has when the die lands with '5' uppermost. I'll represent event tokens with lowercase variables and event types with uppercase variables.

Single-case propensities in biological processes

In this paper I focus on the "single-case" sense of "propensity" (Mellor 1971; Giere 1973; Fetzer 1981; Miller 1994), which seems to be the sense of the term usually intended by advocates of the PIF. Single-case propensities are primitive, irreducible probabilistic dispositions – causal tendencies – inhering in token trials. A single-case propensity is analogous to a law of nature which governs a token event causing another event. However, a single-case propensity does not in general ensure that one event produces another; the propensity is usually a weaker sort of connection between cause and effect.

There are arguments that there are no such things as single-case propensities (e.g. Eagle 2004), and I will discuss one such argument below. However, in much of the paper I'll simply grant to the PIF-advocate that there are single-case propensities. I'll present an argument (with clarificatory detours) from the claim that single-case propensities are causal and are relative to token *trials*, to the view that propensities of complex *outcome* event types occurring during the life of the organism have values very near 0 or 1. The PIF then implies that fitnesses would nearly always be equal to actual numbers of offspring, since whatever number of offspring has a probability of 1 would nearly always be actual. (Remember that we are taking the PIF to define fitness as probability-weighted expected number of offspring.) While fitness would still be causal – since propensity remains dispositional – fitnesses would not have the right values. In particular, genetically and phenotypically identical organisms could easily have different fitnesses on this account. Thus single-case propensity would not seem to be the basis of fitness. I'll go on to examine a way that single-case organism effect propensities might still help to ground fitness given my argument.

Biological propensity has a reference class problem

Single-case propensities attach to *token* trials, but dispositional and other causal properties are often thought to have to do with *types* associated with the causes – in this case trial types. It is thus somewhat natural to think that single-case propensities have something to do with types or properties of trials. This dependence of probability on characteristics of a trial is often captured by viewing probability as conditional probability, i.e. as the probability of outcome A given trial type B, or $\mathcal{P}(A|B)$. Beginning from reasonable assumptions about relationships between trial types and propensities, I'll give a sort of reductio argument that it's most reasonable to think that either there are no single-case propensities, or else single-case propensity is relative to – conditional on – a *particular* complex trial type – what I call the "lowest-level" trial type. I'll begin by focusing on a non-biological case which will simplify the presentation of the argument.

First consider a trial event type: toss of this particular die. Some people have an intuition that tosses of this die have a certain propensity to show a '4', which depends on whether the die is fair or is instead loaded in various possible ways. We can consider trial event types of various degrees of specificity, where each type is instantiated by the same trial event token. A more specific trial event type might be: toss of this die along a particular trajectory beginning from a certain height with '2' initially uppermost. Intuitively, a propensity relative to this new trial type need not have the same value as the propensity relative to the first trial type, since the second type places stricter constraints on the motion of the die. By making a trial type very specific, it seems that '4' could have a propensity far from what one would normally expect. The fact that one propensity for '4' is 1/6 and the other is, say, .001, etc., is not yet problematic, since they are relative to different trial types. (The lack of conflict is explicit in notation when we use conditional probability notation for propensities.) Note that among the trial event types which a token toss instantiates is that trial event type which fixes all of the quantum mechanical details of the toss. Call this the toss's "lowest-level" trial event type, and any trial event type which fails to fix some fact a "higher-level" trial event type.

We can also ask what is the propensity of getting a '4' on a particular *token* toss. This is what really matters for single-case propensity. We might say that the propensity of getting a '4' on the toss is equal to the propensity associated with the trial event type that the toss instantiates, but the toss instantiates many types, each of which might generate a different propensity. For a token trial event to generate two different propensities for one outcome would be for it to have incompatible causal powers – something like having deterministic dispositions to cause incompatible results in the same context. In biological cases, where the trial is a birth or conception of an organism in an environment, allowing single-case propensities to be relative to distinct trial types would mean that the propensity for a particular organism to have k offspring would have various incompatible strengths. There might then be no fact of the matter about whether one organism were fitter than another; a particular organism might have one propensity of having many offspring relative to one trial type, but a different propensity relative to a different trial type. (The moth example below provides an extended illustration.)

A reasonable response to this kind of conflict between single-case propensity values is to deny that there are such things as single-case propensities (Gillies 2000). If there are no single-case propensities, there is no problem of incompatible causal powers. Then, of course, fitness can't be defined in terms of single-case propensity. In that case the options are (a) to give up the PIF approach completely, (b) to try to define fitness in terms of some other kind of organism effect probability, or (c) define fitness in terms of organism circumstance probability, discussed below. I don't feel that option (b) is promising, but see (Weber 2001; Millstein 2003) for discussion of the idea. In any event, in most of the rest of the paper I want to explore the idea that there are single-case propensities. Then we must claim that there is some trial type which

has a privileged status in determining outcomes relevant to fitness. But which one?

In the following sections I'll argue that because single-case propensities are (a) relative to token trials and (b) causal, there's a clear sense in which single-case propensities depend on the lowest-level trial type, i.e. one specifiable in terms of quantum mechanics.

Single-case propensity is quantum mechanical

If single-case propensity is plausible anywhere, it's plausible as an interpretation of probabilities in quantum mechanics. After all, quantum mechanical probabilities are reasonably thought to be present on token trials. In principle, quantum mechanics specifies very clearly what the values of probabilities are, and exactly what it is about a situation which determines them. And quantum mechanics trades in probabilities which are thought not to be reducible or explainable in terms of anything other than quantum mechanics. In addition, when Popper (1957, 1959) introduced the concept of propensity, he did so in part to account for the probabilities in quantum mechanics.

Now, consider token event c and its possible effect type E. Suppose c has quantum mechanical event type Q and a higher-level event type H. Suppose that H determines c's propensity to give rise to E, and that Q does not. This means that every token event with H has the same propensity to produce E; any quantum mechanical differences between tokens of H cannot affect such a propensity. Now, if we say that propensities are *always* determined by higher-level types, quantum mechanical type never makes a difference to propensities. That is, quantum mechanical type never makes a difference to a token event's tendency to produce effects of various types. But that's absurd (as long as we're willing to countenance propensities and allow that one event can cause another). If quantum mechanical types can't ever make a difference to tendencies to produce effects, then a huge part of quantum mechanics would be untestable; it might count as mathematics, but not physics.

However, if we claim that quantum mechanical probabilities are propensities in *some* cases, it would be a very odd metaphysics which claimed that that not *all* quantum mechanical probabilities are propensities. Interpreting a fundamental aspect of a single scientific theory in a non-uniform way requires special motivation. I'll assume, then, that if there are propensities, all quantum mechanical probabilities are propensities.[1]

[1] One difficulty for quantum mechanical probabilities as propensities is that quantum mechanical probabilities don't always combine in ways that accord with standard probability axioms (see e.g. Eagle 2004, §3.4).) Maybe this means that some propensities don't accord with standard axioms (a claim sometimes made for other reasons (Fetzer 1981)). Maybe it means that quantum mechanical 'probabilities' are not propensities. Then a variation on the arguments in the text might show that there are no propensities.

Still, that quantum mechanical types always determine propensities doesn't quite show that fitness can't be based on propensity. Again suppose that token event c has quantum mechanical event type Q and higher-level event type H. Suppose that every possible token of H generates the same propensities to produce effect type E. In this case, we might as well say – at least as a shorthand – that both H and Q determine c's propensity to produce E.[2] Now if such an H were the type defined by a phenotype or genotype along with the environment, and the E's corresponded to numbers of offspring, then the PIF might be vindicated. Propensities for numbers of offspring would then be determined by biological types along with the environment. The question now is whether event types defined by environments and biological types are of this kind; that's the topic of the next section.

Biological propensities exhibit sensitive dependence

Whether or not it's correct that biological propensities are usually near 0 or 1, it's most plausible that these propensities are sensitive to small variations in circumstances. For example, note that the number of offspring that a flying insect has can be significantly affected by a small shift in the wind which can bring its scent to another organism, or orient it in a new direction, or hide it with a leaf as a predator flies by. Small variations in nutrients or poisons in various foods might make a big difference, too. Similar points seem to apply to other organisms – plants as well as animals. The variations in environmental conditions which can make a significant difference in propensities for an organism to have various numbers of offspring would seem to be subtle and complex. It's unreasonable to think that all such variations generate identical single-case propensity distributions, or even distributions that preserve the fitness order of competing types (cf. Weber 2001, pp. S221f). Thus, that an organism has a particular biological type and is in a particular environment doesn't seem to be enough to determine propensities relevant to biological success.

This means that if biological propensities are determined by higher-level types as well as quantum mechanical types, these higher-level types are not the types that fitness is supposed to be attached to, i.e. ones constituted by the combinations of a genotype or phenotype and an environment. Though it might turn out that biological propensities are relative to fairly specific, low-level types which are nevertheless more general than quantum mechanical types, these not-very-high-level types will not be the sorts of things usually taken to explain evolution. Given that, and given that propensities would always be relative to quantum mechanical types as well, it won't hurt to take propensities to be relative to quantum mechanical types in the rest of the paper.

[2] Putnam's (1975) example of a square peg which won't pass through a round hole might be a case in which higher-level type determines propensities; perhaps the quantum mechanical propensity of passing through the hole is the same for every token of the higher-level type (rigid, square, etc.).

Probability and higher-level explanations

Later sections of this paper explore the idea that fitness is based on some kind of probability other than organism effect propensity. Before going further, I want to prepare the ground for this later discussion by clarifying what kind of relationship between low-level and high-level probabilities I think is reasonable in biology. Jackson and Pettit's (1992) example of a flask which cracks while holding boiling water provides a clear illustration from outside of biology.

When water boiling in a flask causes it to crack, it might seem that what's causally relevant to the occurrence of cracking is not the particular set of motions of water molecules, but simply that the water is boiling, i.e. that the average velocity and energy of molecules are above certain values. Jackson and Pettit maintain, however, that what *caused* the cracking in the actual world is a particular molecular motion; it's not the property of boiling that caused the actual cracking to occur. They argue that it may nevertheless be informative and explanatory to cite the mere fact that the water was boiling. Why? Because given that the water was boiling, some molecule or other was "more or less bound to" cause the flask to crack (p. 11). Here "more or less bound to" plausibly means "gives a high probability to". Thus an explanation in terms of the fact that the water was boiling implicitly involves an assumption that there is a probability distribution over sets of molecular motions, a distribution which gives the set of molecular motions which would cause the flask to crack a very high probability.

Assume that the numeric value of the probability of cracking conditional on water boiling is at least a little bit different from the value of the propensity of cracking, conditional on the actual molecular motions. (Jackson and Pettit seem to assume that the molecular interactions are deterministic, in which case the latter number is 1). Then by an incompatible causal powers argument like the one given above, the probability of cracking conditional on boiling cannot be single-case propensity. This, we'll see, parallels my view that at least some of the probabilities which allow fitness to explain evolution are not organism effect propensities.[3]

Biological propensities are extremal

It's reasonable to think that propensities relative to the quantum mechanical type of a die toss would be very close to 0 or 1 (but see below). That

[3] You may wonder whether probabilities of boiling water states are as problematic as the probabilities which underly fitness. What about probabilities in social sciences – Jackson and Pettit's real interest? Yes, all those probabilities are problematic, too. However, probabilities of boiling water states are not interesting in the same way that biological probabilities are, since the relevant thermodynamical probabilities are usually close to 1. Probabilities in social sciences are quite interesting and problematic, though. Philosophers of probability and philosophical statisticians have focused on them to some extent.

indeterminism relative to lowest-level *trial* event types usually adds up to near-determinism concerning *outcome* event types above the atomic level is suggested by countless causal interactions we observe every day. There is as much quantum mechanical indeterminism in the processes that govern the tossing of a ball, *prima facie*, as there are in the processes that govern the tossing of a die. Yet the behavior of the ball approximates that specified by simple Newtonian laws. The reason is that although outcomes concerning a single particle have nontrivial propensities, these propensities generally combine mathematically to produce propensities near 0 or 1 for outcomes involving many particles (cf. Levi 1983; Rosenberg 1994, Ch. 4; Glennan 1997; Graves et al. 1999). There are exceptions to this pattern of "effective determinism", for example involving Geiger counters. However, these exceptions are rare, and they are usually the result of very careful design. The many processes making up an organism and its environment are implemented by large numbers of atoms, and hence it would be reasonable to think that they are as nearly-deterministic as balls and automobiles. Thus it is also reasonable to think that the propensity for the organism to have, say, 4 offspring, relative to a token trial, is either 0 or 1. Then, clearly, single-case propensities would not provide a suitable basis for a theory of fitness.

An illustration of effective determinism

It will be useful to have in mind a concrete story to flesh out the picture I've been sketching. Brandon's (2004) discussion of Beatty's (1984) moth example provides a nice starting point.

In Beatty's example, a group of moths lives in a forest in which 40% of the trees have light bark and 60% of the trees have dark bark. The trees are distributed in a way which we would naturally be inclined to call "random" within a region in which all of the moths remain. The moths are either dark or light and are subject to predation by birds which hunt by sight. Consider a dark moth and a light moth both sitting on a light tree. Should we say that the dark moth's fitness is lower than the light moth's, since while on that tree it is more easily visible to birds? Or should we say that the dark moth's fitness is greater, since there are more dark trees than light trees? The answer that Brandon's (2004) discussion suggests is that if the moths fly around in such a way that the "statistical average" (p. 23) of moths' backgrounds depends on the proportions of light and dark trees, then the fitness of a dark moth, even one sitting on a light tree, can be greater than the fitness of any light moth, including one on the same tree. It is then not the dark moth's current position on the light tree but more general facts that determine its fitness.

Although I am in sympathy with this conclusion, I do not think that single-case propensities will give it to us, given that they are conditional on token trials (cf. Graves et al., 1999). Particular moths will not necessarily spend time in front of light and dark trees in proportions that reflect the proportions of

trees in the forest. One dark moth may end up alighting on light trees more often than dark trees, for example, despite the fact that there are more dark trees. How might this come about? First, a fertilized moth egg interacts with various actual environmental inputs to develop, eventually, into a moth with a particular internal physiology, etc. Subtle differences in environmental inputs may lead to genetically identical moths being physically different in subtle or not so subtle ways. For the adult moth, internal processes interact with stimuli from the environment and air currents to produce the precise behavior of the moth from moment to moment. Stimuli from the environment are in turn determined by the internal processes of other animals and plants in interaction with abiotic elements of the environment. Each of the many processes alluded to here are made up of large numbers of atoms, and hence it is reasonable to think that nothing in this picture is significantly indeterministic. Thus the propensity for a particular moth to spend time in front of the actual sequence of backgrounds in which it is found is near 1. A similar argument leads to the view that the behavior of each particular predatory bird is also effectively deterministic, so that the fate of a given moth is effectively determined. Thus there is a propensity near 1 for a particular moth to have exactly the number of offspring that it has.

What if biological propensities are not extremal?

I gave reasons above for thinking it plausible that single-case propensities are a negligible part of the basis of fitness since they are nearly always 0 or 1. One might, however, feel that biological processes are different from baseballs and automobiles in ways that allow nontrivial single-case (organism effect) propensities for higher-level outcomes to arise. The problem is that even if significant quantum mechanical indeterminism infected biological processes, we would still need organism circumstance probabilities – in effect, probabilities for organisms to experience various organism effect propensity distributions. In order to make sense of the suggestion that the probabilities underlying fitness depend only on indeterminism in organism effects, one would have to argue that within a given environment, indeterminism always intervenes in processes in such a way that that the resulting propensities for numbers of offspring are the same, *no matter what particular circumstances an organism began in* (or that at the very least that the order of fitnesses of types was the same). But as pointed out above (§3.3), the sensitivity of biological propensities to variations in circumstances makes this implausible.[4]

[4] Lewis (1986a,b, §B) tells a plausible story about how the combination of quantum mechanical indeterminism with sensitive dependence on conditions during a coin toss might make the outcome of a token coin toss truly indeterministic. However, he makes it clear that his story leaves open the possibility that differences in initial and surrounding conditions for different coin tosses might produce different probabilities on different tosses. Thus we might need, in addition, a probability distribution over such conditions in order to construct a probability of 0.5 for heads.

Fitness depends on organism circumstance probability

The preceding sections show that the propensity interpretation of fitness is at best incomplete. What's needed to save the general approach pioneered by the PIF is an objective probability distribution over circumstances in which organisms find themselves. Without that there will be no fact of the matter about which biological types are likely to enjoy greater reproductive success. That is, in one sense of "fitness", if there is no such objective probability distribution, one type cannot be fitter than another. For the fact that an organism has a certain type and is in a given environment will not in itself make for determinate propensities for numbers of offspring. Thus, in order to understand what natural selection is we need to know what kind of probability is involved in a distribution over particular circumstances organisms find themselves in.

Organism circumstance probabilities will be, to a first approximation, probabilities of (types of) conditions, states, events, etc. which (a) are not caused in part by the organism itself, and (b) whose presence or absence can make a difference to what effects the organism does cause. I intend "can make a difference to what effects..." to mean also "can make a difference to propensities for effects," in case organism effect propensities sometimes have values other than 0 and 1. More precisely, it will simplify matters to apply "organism circumstance probabilities" only to probabilities of conditions, states, etc. at the point in time when the organism begins, i.e. at the time of its conception. These circumstances can include the initial state of the organism at conception. If the subsequent processes relevant to reproductive success are effectively deterministic, then there is no loss in focusing on the time of conception, since circumstances at that time determine later circumstances. If processes subsequent to conception are significantly indeterministic, then although it is not determined exactly which relevant circumstances later in life might occur, still, circumstances at the time of conception determine single-case propensities for such later circumstances (i.e. single-case propensities conditional on circumstances at the time of conception). Thus whether or not is true, as argued above, that single-case organism effect propensities nearly always take values of 0 or 1, organism circumstance probabilities would be needed to make fitness determinate.[5]

[5] Note that since organism circumstance probabilities are probabilities of conditions which can indirectly make a difference to the effects which an organism produces (clause (b)), organism circumstance probabilities, like organism effect probabilities, can make a difference to whether niche construction (organism-caused environmental modification) occurs. I intend to discuss niche construction in a future paper as part of a general discussion of the ways in which probabilities concerning the effects of organisms in a population combine. Models which use probabilities concerning individual organisms' fates to derive probabilities for changes of frequencies within a population sometimes assume that the individual-level probabilities are independent. Niche construction is only one of the reasons that this assumption is often false, even for organisms in non-overlapping generations.

If organism effect propensities are limited to values of 0 and 1, then organism circumstance probabilities are *all* of the probabilities relevant to fitness. If organism effects propensities sometimes take other values, then the overall probability of an offspring count for a given genotype or phenotype is the weighted average of the organism effect propensities for that offspring count, the weights being organism circumstance probabilities for various circumstances in a given environment. If we represent the organism effect propensity for a given organism type to have k offspring in circumstance i as $\mathcal{P}_i(O = k)$, and use $\mathrm{P}(S = i)$ for the organism circumstance probability that an organism of that type is in circumstance i, then the expected number of offspring for the type is

$$\sum_{k=0}^{\infty} k \sum_{i} \mathcal{P}_i(O = k)\mathrm{P}(S = i). \tag{2}$$

Here $\sum_i \mathcal{P}_i(O = k)\mathrm{P}(S = i)$ is the objective probability that an organism will have k offspring; it plays the role that propensity alone plays in formula (1). It is this overall, composite probability which would bear the right sort of relationship to relative frequencies of offspring counts for a given genotype or phenotype. It is what would capture what is common to explanations of the relative success of fitter types over less fit types, and thus it is this probability which fitness, in the present sense, should be a function of.

Organism circumstance probabilities are analogous to various other kinds of probabilities that one encounters in evolutionary theory. These are probabilities that an organism will encounter some general kind of situation labeled a "habitat", a "patch", or an "environment" within a larger environment in which a population ranges (e.g. Roughgarden 1979, Chs. 12, 13; Brandon, 1990, Ch. 2; Gillespie 1998, §3.6). Such probabilities might be called "patch-distribution probabilities", but are usually described as, for example, migration probabilities or probabilities concerning habitat choice by an egg-laying parent. In models using such probabilities, it is often useful to define "fitness" relative to each patch. Note, though, that differences of fitness in this sense of the term cannot by themselves explain the evolution of a population; the patch distribution probabilities are essential to determining probabilities concerning relative success of competing organism types. Thus, there is in such models, at least implicitly, a more comprehensive sense of fitness for each type, defined across all of the patches in the larger environment. A less detailed model might use "fitness" only in this comprehensive sense. Whatever terminology it is that is convenient in a particular modeling context, it is this more general sense of fitness which would be capable of explaining evolution on its own (even if, say, drift and mutation played no role in evolution). In any event, note that the analogy between what organism circumstance probabilities concern and the relatively large scale patches referenced in evolutionary models is imperfect. I've argued that outcomes relevant to fitness are sensitive to small variations in circumstances. That means that circumstances are fleeting states with small

spatial extent. Thus unlike the patches in many evolutionary models, different organisms hardly ever experience the same circumstances.

Organism circumstance propensity?

The approach pioneered in the PIF faces a challenge which has not generally been recognized before now, the challenge of explaining the source of organism circumstance probabilities. While it's not the point of this paper to answer this challenge, I do want to briefly suggest questions that would face the most obvious of responses, namely a proposal that organism circumstance probabilities are single-case propensities.

First, if I'm right that biological propensities are generally very close to 0 and 1, do arguments similar to mine apply to propensities for organism circumstances? Organism circumstance propensities whose values were usually 0 or 1 would not seem to be able to help us understand fitness.

Apart from that issue, two very basic questions must be answered to make sense of the suggestion that organism circumstance probabilities are propensities: First, what trial tokens determine the propensities? Second, what are the outcomes governed by the propensities? Recall, for comparison, that for organism effect propensities a trial token was the conception of an organism in its surrounding environment, while outcomes were defined in terms of numbers of offspring. For organism circumstance propensities, a rough answer to the second question is not hard to provide: An outcome is a state of the environment and an organism at the time of a conception.

The answer to the first question is less clear; what are the relevant trial tokens? One natural suggestion would be to take a trial which generates organism circumstance propensities to have to do with actual organisms which are potential parents. The propensity for a conception in circumstance C would be a propensity for a particular parent to conceive an offspring in C. Thus for a haploid species the trial which generates organism circumstance propensities might be the conception of such a potential *parent* along with its surrounding environmental circumstances. An organism circumstance trial would then be the same sort of thing as an organism effect trial, but one generation back. Still, why is that the relevant trial? Why not take the trial to have to do with the potential parent and its environment at some intermediate point in the parent's life? Why not some even earlier trial, perhaps involving a grandparent? (Note also that for a diploid species the relevant trial might have to involve two parents, and that the two parents' conceptions could occur at different times.)

Just as organism effect propensities are affected by an organism's circumstances, organism circumstance propensities would be affected by the earlier circumstances involved in earlier trial tokens. Thus it seems that different earlier trials could generate different organism circumstance propensities. But an organism circumstance probability distribution relevant to fitness should be a distribution over all conceptions in a population for a relatively long period

of time. It's that kind of distribution which would be relevant to the combined process that is natural selection. The question then arises: Would we need a probability distribution over the earlier trial circumstances which determine organism circumstance propensities? Does this lead to a regress?

Others may have answers to these questions. My approach has been to develop an account of organism circumstance probabilities which does not require single-case propensities. This account uses an interpretation of probability which I call "mechanistic probability", based partly on ideas in (Strevens 2003). I feel that it's difficult to do justice to my approach in a few sentences, and I won't try to describe it here; see (Abrams 2000; Abrams 2005). Other ideas about fitness which seem promising include Millikan's (2000) ideas about persistence of frequencies in an environment, and ideas from statistical mechanics discussed by Weber (2001), Rosenberg (2001) and Bouchard and Rosenberg (2004).

Fitness is not a simple average of actual propensities

One last clarification: The view that fitness is determined by an organism circumstance probability distribution has an obvious but superficial similarity to an idea which goes back to the earliest formulations of the PIF. Mills and Beatty (1979) defined the fitness of a genotype as the average of the fitnesses of individual organisms with that genotype, where fitnesses of individuals are expected numbers of offspring computed using propensities. The most plausible reading of this would take the average to be over fitnesses of actual organisms with the genotype in question, each organism's fitness being weighted equally (cf. Sober, 2000, §3.6; Rosenberg and Bouchard 2004; Bouchard and Rosenberg 2004). If all lowest-level propensities were 0 or 1, such a genotype fitness would amount to no more than the average of actual numbers of offspring. This would be tantamount to a return to the original definition of fitness to which the PIF was a reaction, which is not a viable option. However, even if nontrivial indeterminism infected biological processes, the Mills and Beatty actual average approach wouldn't work.

Let's simplify the moth example by letting it be large insect eggs rather than moths which are dark or light. Again 40% of trees have light bark and 60% have dark bark. Suppose that in one or several generations, many dark eggs happen to be laid on light-colored trees, thus giving each such dark egg a higher propensity to be eaten and a lower propensity-based expected number of offspring. Our concept of fitness should allow this to happen sometimes even if there is a general sense in which dark-egg insects are fitter than their light-egg conspecifics (perhaps partly because there are more dark trees than light trees, etc.). But if this kind of idiosyncratic occurrence is possible sometimes, it should be possible even over many generations. The problem is that a definition of fitness as the actual average of expected numbers of offspring makes it impossible for there to be such idiosyncrasies in circumstances organisms

experience – idiosyncrasies which do not accord with fitness. Whatever fitness is, it should allow what is actual to depart from what fitness differences would predict. The view that I am advocating, in contrast, emphasizes the need for an explicit, principled account of organism circumstance probabilities.

Conclusion

I've argued that because of reference class problems, (single-case) organism effect propensities either do not exist, or else there are good reasons to think that they cannot ground fitness (because they are nearly always 0 or 1). In addition, I've argued that whether or not organism effect propensities are nearly always 0 or 1, organism circumstance probabilities must play a major role in constituting fitness. Those who want to give an account of fitness in terms of objective probabilities face the challenge of providing an account of organism circumstance probabilities.

I suggest now that the public record of the marriage of fitness and organism effects propensities be changed to show that the two were never properly married, or at the very least were never married *only* to each other. (Theoretical marriages obey their own rules.) Nevertheless, the PIF was and remains an important advance in thinking about fitness. If those who championed a fitness-propensity union were mistaken, they may not have been wholly wrong. There are difficulties to be overcome by an account of fitness in terms of objective probability, but if they can be overcome, it may turn out that the mistaken marriage announcements were due only to a case of mistaken identity, and that a member of the objective probability family other than organism effect propensity was part of fitness's basis.

References

Abrams M. 2000. Short-Run Mechanistic Probability. Presented at the Philosophy of Science Association, Seventeenth Biennial Meeting, Vancouver BC, Canada, November 2–5, 2000, http://hypatia.ss.uci.edu/lps/psa2k/short-run.pdf.

Abrams M. 2005. The Causal Structure of Biological Fitness, MS.

Beatty J. 1984. Chance and natural selection. Philos. Sci. 51: 183–211.

Beatty J. and Finsen S. 1989. Rethinking the propensity interpretation: A peek inside Pandora's box. In: Ruse M. (ed.), What the Philosophy of Biology is? Kluwer Academic Publishers, pp. 17–30.

Bennett J. 1988. Events and their Names, Hackett.

Bouchard F. and Rosenberg A. 2004. Fitness, probability, and the principles of natural selection. Brit. J. Philos. Sci. 55(4): 693–712.

Brandon R.N. 1978. Adaptation and evolutionary theory. Stud. Hist. Philos. Sci. 9(3): 181–206.

Brandon R.N. 1990. Adaptation and Environment, Princeton University.

Brandon R.N. 2004. The difference between selection and drift: A reply to Millstein. Biol. Philos. 20(1): 153–170.

Davidson D. 1967. Causal relations. J. Philos. 64.

Eagle A. 2004. Twenty-one arguments against propensity analyses of probability. Erkenntnis 60: 371–416.

Fetzer J.H.: 1981, Scientific Knowledge: Causation, Explanation, and Corroboration, Boston Studies in the Philosophy of Science. D. Reidel.
Giere R.N. 1973. Objective single-case probabilities and the foundations of statistics. In: Suppes P., Henkin L., Joja A. and Moisil G.C. (eds), Logic, Methodology and Philosophy of Science IV. North-Holland, pp. 467–483.
Gillespie J.H. 1998. Population Genetics: A Concise Guide. The Johns Hopkins University Press.
Gillies D.A. 2000. Varieties of propensity. Brit. J. Philos. Sci. 51: 807–835.
Glennan S. 1997. Probable causes and the distinction between subjective and objective chance. Noûs 31(4): 496–519.
Graves L., Horan B.L. and Rosenberg A. March 1999. Is indeterminism the source of the statistical character of evolutionary theory? Philos. Sci. 66: 140–157.
Jackson, F. and P. Pettit: 1992, In defense of explanatory ecumenism. Econ. Philos. 8(1): 1–21.
Levi I. 1983. Review of studies in inductive logic and probability. Phil. Rev. 91: 116–121.
Lewis D. 1986a. Philosophical Papers, Vol. II. Oxford University Press.
Lewis D. 1986b. Postscripts to 'A Subjectivist's Guide to Objective Chance'. In: Lewis (1986a).
Mellor D.H. 1971. The Matter of Chance. Cambridge University Press.
Miller D.W. 1994. Critical Rationalism: A Restatement and Defence, Open Court.
Millikan R.G. 2000. What has natural information to do with intentional representation? In: On Clear and Confused Ideas, Cambridge University Press, pp. 217–237.
Mills S. and Beatty J. 1979. The propensity interpretation of fitness. Philos. Sci. 46(2): 263–286.
Millstein R.L. 2003. Interpretations of probability in evolutionary theory. Philos. Sci. 70: 1317–1328.
Popper K.R. 1957. The propensity interpretation of the calculus of probability, and the quantum theory'. In: Körner S. (ed.), Observation and Interpretation. Academic Press Inc., Butterworths Scientific Publications, pp. 65–70.
Popper K.R. 1959. The propensity interpretation of probability. Brit. J. Philos. Sci. 10: 25–42.
Putnam H. 1975. The meaning of 'meaning'. In: Mind, Language and Reality, Cambridge University, pp. 291–303.
Rosenberg A. 1994. Instrumental Biology or the Disunity of Science, University of Chicago Press.
Rosenberg A. December 2001. Discussion note: Indeterminism, probability, and randomness in evolutionary theory. Philos. Sci. 68: 536–544.
Rosenberg A. and Bouchard F. 2004. Matthen and ariew's obituary for fitness: Reports of its death have been greatly exaggerated. Biol. Philos. 20(2–3): 343–353.
Roughgarden J. 1979. Theory of Population Genetics and Evolutionary Ecology: An Introduction. Macmillan.
Sober E. 2000. Philosophy of Biology, 2nd ed., Westview Press.
Sober E. 2001. The two faces of fitness. In: Singh R.S., Krimbas C.B., Paul D.B. and Beatty J. (eds), Thinking About Evolution. Cambridge University Press, pp. 309–321.
Strevens M. 2003. Bigger Than Chaos: Understanding Complexity through Probability, Harvard University Press.
Weber M. 2001. Determinism, realism, and probability in evolutionary theory: The pitfalls, and how to avoid them. Philos. Sci. 68(3): S213–S224.