

## Philosophical foundations of microevolution

Review of: *Evolution and the Machinery of Chance*, by Marshall Abrams, 2023. Chicago University Press, Chicago. 282 pp. ISBN 9780226826639

The penultimate chapter of *The Causes of Evolution*, J.B.S. Haldane's popular 1932 book, is entitled "What is Fitness? (Haldane, 1932 [1990])." The chapter contains much interesting material but, oddly, does not attempt to answer its title question. Haldane thus initiated a venerable tradition in evolutionary biology of employing the fitness concept but without saying exactly what it means. The tradition persisted for quite some time. Over 40 years later, Stearns (1976) described fitness as "something everyone understands but no one can define precisely," a remark quoted by Marshall Abrams in his recent book. Since Stearns wrote these words, many valuable clarifications of the fitness concept have appeared in the evolutionary literature (Brommer, 2000; Day & Otto, 2001; Hanssen, 2018; Metz et al., 1992; Rice, 2005; Roff, 2008). Fitness can usually be defined unambiguously in the context of a given theoretical model, but a fully general, precise definition seems not to exist. Given the centrality of fitness to evolutionary biology this may seem paradoxical, but it is not really so surprising. For like other scientific concepts, fitness is defined by its theoretical role (roughly, determining the equilibrium gene frequencies and/or evolutionary dynamics in a population subject to natural selection); and the quantity that plays that role will typically depend on model assumptions (Okasha, forthcoming).

The fitness concept has also been extensively discussed by philosophers of biology, often as part of a broader analysis of the fundamentals of Darwinian evolution. Though this philosophical literature contains interesting ideas, I have long been troubled by its somewhat self-contained nature, its reliance on simple evolutionary models and/or toy examples, and its disconnect from the relevant biological literature. Indicative of this disconnect is that many important conceptual questions about fitness barely get a mention in this philosophical discussion. Examples include: whether invasion fitness is a fully general fitness definition under frequency-dependent selection; whether  $r$  or  $R$  is the "right" measure of fitness in life-history theory and why; when geometric mean fitness arguments work; what assumptions are needed for inclusive fitness arguments to hold; how reproductive value should be defined; how to understand fitness in age-structured populations; and more. Instead, the philosophical literature on fitness has focused on its own set of questions, many of which have to do with causality and chance. This is not necessarily a problem—for different disciplines can legitimately have their own concerns—but greater integration seems desirable.

The initial impetus behind the philosophical work on fitness was the worry that certain definitions of fitness render the theory of evolution tautological. Thus, Mills and Beatty (1979) argued that if fitness is defined as actual number of offspring, it becomes true by definition that the fittest types will out-reproduce the less fit, a result they regarded as undesirable. This led them to advocate what they called the "propensity interpretation of fitness," which says that the fitness of an individual is a disposition (tendency) to leave offspring. The strength of this disposition is naturally measured by the individual's expected, rather than actual, number of offspring, where the expectation is taken with respect to the probability distribution over possible numbers of offspring that the individual could leave, given its traits and environment. (In effect, then, the idea was simply that any individual's actual number of offspring is the realization of a random variable, the expectation of which serves to define, or at least measure, individual fitness.) Much subsequent philosophical work on fitness can be traced to Mills and Beatty's paper, though other issues are discussed too.

Abrams' book is a contribution to the philosophical literature on fitness and natural selection. Given that Abrams is highly knowledgeable about biology, I had hoped that his book would help to bridge the disconnect between the biological and philosophical literatures alluded to above, but in fact his discussion is mostly oriented towards other philosophers. Now Abrams is well aware of the biological work on fitness measures and definitions. Indeed in the first chapter, he writes that "there are quite a few names for the diverse variety of different fitness concepts in use in evolutionary biology, including 'selection coefficient', 'lifetime reproductive success', 'adaptive value', 'reproductive value', 'net reproductive rate', 'intrinsic growth rate', 'density-independent growth rate', 'Malthusian parameter', 'viability', 'number of offspring', 'selection differential', 'selection gradient', 'invasion fitness', and others" (p. 21). Abrams goes on to point out, correctly, that these terms are "not interchangeable," but says that the differences between them do not matter to his project and that he will use the term "fitness" in a way that generalizes over them all (p. 28). I can see why Abrams says this and he may be right, but I confess that this passage gave me qualms. The disquiet it provoked in me reflects my conviction that some philosophical work on fitness asks questions that are insufficiently sharp. Can there really be well-posed questions about fitness, capable of being discussed with analytical

precision (verbal or mathematical), to which it makes no difference whether “fitness” means viability, reproductive value, or Malthusian growth rate? On the face of it, it is not obvious that this is so.

Abrams’ own ideas about fitness, in particular about how fitness and chance relate, are situated within a broader picture of microevolution that he develops. At the heart of this picture is what he calls a “population–environment system”; roughly, this means a biological population coupled with a set of environment variables. Evolution occurs when a population–environment system changes its state over time; where “state” includes the frequency of genes and traits in the population. Abrams next argues that a population–environment system is what he calls a “chance set-up”; in effect, this means that at any time, given the system’s state, there is an objective probability, or chance, of its moving to another state at a later time. Thus, from a given starting point, the actual trajectory of a population–environment system through state space is one among a set of possible trajectories, each of which has a given probability of occurring. Natural selection is one of the factors that can influence the trajectory of a given population–environment system.

Abrams’ notion of a population–environment system seems to me useful, and fits with standard ways of thinking about stochastic evolutionary dynamics. But unlike most biological discussions of stochasticity, which simply take the notion of chance or probability as a given, Abrams devotes considerable attention to trying to analyze it. In this he draws on a long tradition, that straddles philosophy, science and statistics, of attending to foundational questions about probability. These questions arise because although probability has a clear meaning mathematically, its empirical meaning is far less clear. For example, one perennial question is whether we can talk sensibly about the probability of a “single case” outcome, for example that King Charles of England will live to age 100, as opposed to a type of outcome, for example that a 74 years old English male will live to age 100? A related question is whether the existence of nonextremal objective probabilities (i.e., not 0 or 1) for single-case outcomes requires that the universe be governed by fundamentally indeterministic laws? A further question asks what the relation is between the probability of an outcome–type and the relative frequency of that outcome–type in a long hypothetical sequence of trials—can these two things be equated or not?

Now most uses of probability in scientific practice, including in evolutionary modeling, do not require that we take a stand on these thorny questions. But the questions are real ones nonetheless, and the subject of a rich philosophical discussion that goes back decades. Abrams is well-informed about this discussion and indeed has contributed ideas to it himself. One overall aim of his book is to connect up the foundations of probability with the issue of stochasticity in evolution. Though I don’t agree with all of the moves that Abrams makes, I am sympathetic to this overall aim and regard this as one of the strongest features of his book.

However, it is to the concept of fitness that I wish to return. Abrams offers a fourfold taxonomy of fitness concepts in biology. He starts by distinguishing what he calls “token fitness” from “type fitness.” (“Token” is a word that philosophers use to mean a particular thing, as opposed to a type [or sort] of thing.) Abrams then subdivides “token fitness” into two. “Measurable token–organism fitness,” he tells us, refers to some property, relevant to evolutionary success, that can

be measured on a token organism, such as its lifetime reproductive success. “Causal token–organism fitness,” by contrast “attempts to capture the idea that a particular individual in its particular circumstances has one or more tendencies to realize properties relevant to evolutionary success.” Abrams tells us that the propensity interpretation of fitness, described above, is a fitness concept of this sort.

Type fitness also gets subdivided into two. “Statistical organism–type fitness,” Abrams tells us, “defines fitness as a property of an inheritable type, in such a way that fitness becomes a mathematical function of measurable token fitnesses” (p. 79). He continues: “for example, if we define the fitness of a trait as an average of the number of offspring that actual individuals with that trait have.....then we are treating the trait’s fitness as a statistical type fitness.” By contrast, “causal organism–type fitness...treats the fitness of a trait as an underlying property that could affect the trait’s evolutionary success.” (p. 79) Abrams goes on to argue: (a) that causal type fitness is the sort of fitness relevant to evolution, since natural selection operates on variation in fitness of this sort; (b) that biologists often seek to estimate causal type fitness from statistical type fitness; (c) that causal token fitness plays no role in evolutionary biology, *pace* what other philosophers have thought. Finally, Abrams recognizes a fifth fitness concept—“purely mathematical fitness”—but says little about it.

How if at all does Abrams’ taxonomy of fitness map onto more established uses of the term? His distinction between token and type fitness corresponds fairly closely to the well-known distinction between the fitness of an individual and of a genotype (or trait), as Abrams himself notes. Moreover, since genotypic fitness is standardly defined as the average fitness of individuals with the genotype, this is “statistical type fitness” in Abrams’ terms. What about causal token fitness? So far as I can see, this sort of fitness is at work in standard evolutionary models of demographic stochasticity. In such models, the fitness (actual number of offspring) of any individual is treated as a random variable, which is assumed to be independent and identically distributed (i.i.d.) across individuals of the same genotype. The point of this is to allow that individuals of the same genotype might leave different numbers of offspring, that is have different *realized* fitnesses, by chance. This corresponds closely to Abrams’ description of causal token fitness, namely “that a particular individual in its particular circumstances has one or more tendencies to realize properties relevant to evolutionary success” (p. 78).

What about causal type fitness—the sort of fitness that Abrams sees as being relevant to natural selection? Abrams tells us that biologists estimate this via statistical type fitness, that is, average fitness of individuals of the type. It appears then, that causal type fitness is at work in those evolutionary models that incorporate environmental rather than demographic stochasticity. In such models, the environmental state in each generation is drawn from a probability distribution, typically assumed independent across generations, and the realized state affects the fitness of all organisms in the population in a given generation. So the fitness of any given genotype, or what Abrams calls statistical type fitness, becomes a random variable whose realized value, in any generation, depends on which environmental state occurs. This makes sense of Abrams’ notion that causal type fitness can be estimated from statistical type fitness, and that a type’s causal fitness is an underlying property that is relevant to its evolutionary success.

Abrams says nothing about the distinction between demographic and environmental stochasticity. This is puzzling in itself, given that much of his book is about stochastic factors in evolution. Surely the question of whether this stochastic variation is or is not i.i.d. across individuals must be relevant, given that this makes a fundamental difference to evolutionary outcomes? How can this not matter to Abrams' discussion? Moreover, I fail to see how Abrams' claim—to which he devotes a whole chapter—that causal token organism fitness is irrelevant to understanding evolution can be squared with the fact that many evolutionary models incorporate a component of demographic stochasticity, and need to do so if they are to be faithful to their real-world targets. For to repeat, such models precisely involve treating an individual's actual reproductive output as a realization of an i.i.d. random variable, whose probability distribution depends on the individual's genotype and the environment. This seems to involve modeling exactly the thing that Abrams argues to be irrelevant, namely an individual organism's tendency (disposition) to leave different numbers of offspring given its circumstances.

Perhaps Abrams arrives at his view because he is implicitly confining attention to evolutionary models that ignore demographic stochasticity, for example because they assume a very large or effectively infinite population? However, this seems unlikely, since Abrams says quite a lot about genetic drift, and thus his discussion cannot be premised on the assumption of large population size. In any case, given the widespread agreement that demographic stochasticity is important in evolution, one could hardly justify basing a philosophical account of fitness on assuming it to be absent, which so far as I can see is what Abrams has in effect done.

Since it is individual organisms, not types, that survive and reproduce, it seems clear that the fitness of a type must ultimately depend on the fitness of individual (or token) organisms. Abrams appears to accept this point. Nonetheless, he argues that causal token fitness is irrelevant to evolution, and cannot be used as the basis on which to define the causal type fitness that he thinks does matter. These arguments have the form “if fitness is understood in such-and-such a way, it cannot do the work that it is meant to.” I am sympathetic to this *style* of argument, for it fits naturally with my own view that fitness is to be defined by its theoretical role in evolutionary models, but I was not persuaded by Abrams' reasoning. In part, this is because some of his arguments do not actually pertain to the relation between causal token fitness and causal type fitness as he defines them, but rather to the relation between individual fitness and genotype fitness in the ordinary sense. For example, one of his contentions is that the phenomenon of genetic hitchhiking, in which selection at one locus causes allele frequency change at a second neutral locus due to linkage disequilibrium, tells against the idea that a type's fitness may be defined as the “actual average fitness” of individuals in the type. Abrams says that the same applies to correlated phenotypic traits, that is where direct selection on one trait induces an evolutionary response on another trait that is itself selectively neutral.

Why does Abrams think that genetic hitchhiking/phenotypic correlation undermines the standard definitional link between type fitness and individual (token) fitness? He reasons as follows. Consider two alleles  $B$  and  $b$  at the neutral locus, and two alleles  $A$  and  $a$  at the locus under selection. Suppose that due to linkage disequilibrium, the  $B$  allele is statistically associated with the favorable  $A$  allele. Therefore, the

average fitness of individuals carrying the  $B$  allele will exceed that of those carrying the  $b$  allele. So if type fitness equals average individual fitness, we are forced to say that the  $B$  type is fitter than the  $b$  type; but this makes no sense since by assumption they are selectively neutral! The definition of fitness in question therefore “treats as nonsense an enormous body of work in modeling and empirical research in evolutionary biology,” Abrams declares (p. 89).

I am puzzled by this argument, which is strangely divorced from the textbook population-genetic treatment of genetic hitchhiking. This treatment uses two-locus theory, so focuses on the relative frequencies of the different gamete types (haplotypes) in a large population. (The gametic frequencies determine the allele frequencies but not vice-versa.) In Abrams' example, there are four gamete types:  $AB$ ,  $Ab$ ,  $aB$ , and  $ab$ . Fitness values are then assigned to the diploid genotypes formed from the union of two gametes. Genotypic fitness has its usual meaning, namely average fitness of individuals with the genotype. Selective neutrality of the  $B$  locus simply means that genotype fitness depends only on the alleles present at the  $A$  locus, that is  $w(AB/AB) = w(AB/Ab) = w(aB/AB) = w(aB/Ab)$ , and so on. One can then calculate the *marginal fitness* of each gamete type, by averaging over all the genotypes in which the gamete type can be found, weighted by their frequency. Note that the marginal fitnesses of  $AB$  and  $Ab$  (for example) are not necessarily equal, despite the selective neutrality of the  $B$  locus. The marginal gametic fitnesses, together with the magnitude of linkage disequilibrium and the rate of recombination, then determine the next generation's gamete frequencies. Note also that if we wanted, we could compute marginal allelic fitnesses in a two-locus model by the same averaging procedure; clearly, the marginal fitness of  $B$  and  $b$  would not necessarily be equal, owing to linkage disequilibrium.

I belabor all of this to emphasize that the fitness concepts of standard population genetics—individual fitness, genotype fitness, marginal gametic fitness and marginal allelic fitness—are perfectly adequate for giving a consistent description of genetic hitchhiking. There is nothing conceptually problematic here, nothing that stands in need of delicate philosophical analysis, and nothing that calls for a revision or redefinition of fitness. It is all rather straightforward. Abrams appears to think otherwise, but I think that this can only reflect the mismatch between his own taxonomy of fitness concepts and the ones that population geneticists ordinarily employ.

There is a general methodological moral here. Evolutionary biology has long attracted the interest of philosophers, including myself. Much interesting philosophical work has been done on evolutionary topics including function and goal-directedness, the logic of evolutionary explanation, the nature of species, hierarchical organization, and more. Many evolutionists have welcomed the input of philosophers on these topics. Crucial to the success of such philosophical enquiries is to focus squarely on issues that are “fair game” for philosophers, which requires, at a minimum, that the issues not be purely scientific matters (empirical or theoretical). I think that some parts of the philosophical literature on fitness (though not all) fail to respect this. How genotype fitness relates to individual fitness, and how genotype fitness must be defined in the context of environmental stochasticity, are questions that, in the context of well-specified evolutionary models at least, have perfectly clear, determinate answers, give the theoretical role that fitness is meant to play.

One topic that certainly calls for philosophical analysis is causality. Like most sciences, evolutionary biology is in the business of giving causal explanations, but many have thought that evolutionary causality is a bit different from causality in other domains (Mayr, 1961.) Abrams has plenty to say about causality, much of which seems to me correct. He insists that natural selection should be regarded as a cause of evolutionary change and articulates this idea in an original way via a notion that he calls “causal probability.” Other philosophers of probability will have to be the judge of this notion, but Abrams is surely right to emphasize that natural selection is a cause of evolution. This latter idea comports with standard biological understanding, and indeed one might wonder whether anyone could dispute it. In fact, the answer is yes. An influential minority of recent philosophers of biology, who go by the rather odd name of “statisticalists,” have repeatedly argued that natural selection is *not* in fact a cause of evolutionary change, appearances to the contrary notwithstanding. Their grounds for this heterodox opinion are not easy to fathom, and in my view have not been spelled out clearly enough to permit serious discussion. Abrams apparently thinks otherwise, since he does treat the “statisticalist” view seriously, carefully defending the rival “causalist” position. While I would hesitate to recommend the “statisticalist versus causalist” debate to noninitiates, I do think that Abrams comes down on the right side of it.

My own view is that the key to understanding the causal dimension of natural selection is to observe a distinction. When we say that a given evolutionary change, for example the increase in frequency of an allele in a particular population, is “due” to natural selection, two distinct causal claims are being made. Firstly, we mean that the allele increased because of the differential reproduction of individuals (as opposed to mutation, for example.) Secondly, we mean that this differential reproduction occurred because the allele had a systematic effect on individuals’ reproductive success (rather than being due to chance, for example). The distinction is usefully conceptualized in terms of the Price equation: the covariance between an individual’s reproductive output and the number of copies of the allele they carry should be positive, and this covariance should reflect a causal influence of the latter variable on the former. This distinction is not always made in the literature on causality and natural selection, which has led to confusion. (For example, it is not always clear which of these two causal claims the “statisticalists” are criticizing.) It would be interesting to know whether Abrams agrees with this assessment.

In recent philosophy of biology, a number of authors have suggested that the framework of *causal modeling* can be utilized to address issues about causation in evolutionary biology (Edelaar et al., 2023; Okasha, 2016; Otsuka, 2015, 2016). Causal modeling has its roots in the path analysis of Sewall Wright and the later “structural equations” tradition in social science. In the last three decades, it has matured into a general, mathematically precise language for expressing causal relations (free from path analysis’s linearity assumption) and an associated set of graphical techniques. Key contributions including works by the computer scientist Pearl (2000) and the philosophers Spirtes et al. (2001). Though causal modeling is not a panacea for all philosophical worries about causality, it is arguably the best framework that we have for rendering causal talk precise, so applying it to debates about evolutionary causality is a natural move. Abrams does not mention causal models in his book, but my impression is that

his overall approach would be congenial to the causal modeling framework. However, confirming this impression would require systematic analysis.

One part of Abrams’ book that I particularly enjoyed was his treatment of the concept of population. An interesting philosophical literature on this concept has emerged in recent years (e.g., Millstein, 2009, 2014). The motivation for this literature stems from the fact that, despite the centrality of populations to evolutionary analysis, little is usually said about what a population *is*, that is which organisms it includes and why. In some ways this is quite surprising. Since evolutionary change is inherently a population-level phenomenon—for it is populations not individuals that evolve—one might have thought that the term “population” must have a canonical definition in evolutionary biology. But it does not seem to, and moreover no-one seems especially troubled by this. As the philosopher Roberta Millstein has noted, there is a sharp contrast here with the notion of species, which generations of biologists *have* thought it important to try to define (whether or not they have succeeded). Millstein herself argues that the relevant notion of population, for evolutionary analysis, is a collection of conspecific organisms that causally interact with each other more frequently than with members of other populations, where these interactions can include competition, mating, offspring rearing, and more. A similar position is taken by Godfrey-Smith (2009).

Abrams agrees that populations are comprised of conspecifics but argues persuasively against this “causal interaction” requirement on what counts as a population. According to Abrams, the requirement is too restrictive, impossible to apply given our typically limited knowledge of patterns of interaction, and in conflict with biological practice. Based on detailed analysis of case studies from the evolutionary biology literature, Abrams argues that biologists need to be able to “flexibly define” the term population depending on their particular research interests. That is, researchers are free to treat *any* collection of (conspecific) organisms as a population if they so wish, and no definition should preclude this. Now one might worry that this introduces an unwelcome element of subjectivity into biology. Moreover, it may seem to threaten the idea that natural selection is a cause of evolutionary change. For natural selection occurs in populations, and causes and effects have to exist; so surely natural selection can only cause evolutionary change if the population in which the change occurs is a real entity rather than an arbitrary grouping? Abrams anticipates this worry and shows how it can be defused, in effect by rejecting the dichotomy between “real” and “arbitrary.” In his view, populations are delimited by “pragmatic” criteria that are not very precise and may legitimately vary from study to study, but they are real entities in nature nonetheless. This is a subtle line of argument that Abrams develops with skill.

To conclude, Abrams’ book contributes a wealth of interesting ideas to the philosophical analysis of microevolution. Though I was not convinced by all of his arguments, I thoroughly recommend his book nonetheless.

Samir Okasha

Department of Philosophy, University of Bristol, Bristol, United Kingdom

Corresponding author: Department of Philosophy, University of Bristol, Cotham House, Bristol BS65DR, United Kingdom. Email: [Samir.Okasha@bristol.ac.uk](mailto:Samir.Okasha@bristol.ac.uk)

## Funding

This paper is part of a project that has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no 101018523).

## Conflict of interest:

The author declares no conflict of interest.

## Acknowledgment

The author is thankful to John McNamara, Walter Veit, Ross Pain, and Jim Mallet for comments and discussion.

## References

- Brommer, J. E. (2000). The evolution of fitness in life-history theory. *Biological Reviews of the Cambridge Philosophical Society*, 75(3), 377–404. <https://doi.org/10.1017/s000632310000551x>
- Day, T., & Otto, S.P. (2001). Fitness. In *Encyclopaedia of life science*. John Wiley & Sons, Ltd. <https://doi.org/10.1038/npg.els.0001745>
- Edelaar, P., Otsuka J., & Luque J.V. J. (2023). A generalised approach to the study and understanding of adaptive evolution. *Biological Review*, 98, 352–375.
- Godfrey-Smith, P. (2009). *Darwinian populations*. Oxford University Press.
- Haldane, J. B. S. (1932 [1990]). *The causes of evolution*. Princeton University Press.
- Hansen, T. F. (2018). Fitness in evolutionary biology. *Preprints*, 2018040271. <https://doi.org/10.20944/preprints201804.0271.v1>
- Mayr, E. (1961). Cause and effect in biology. *Science*, 134(3489), 1501–1506. <https://doi.org/10.1126/science.134.3489.1501>
- Metz, J. A., Nisbet, R. M., & Geritz, S. A. (1992). How should we define “fitness” for general ecological scenarios?. *Trends in Ecology and Evolution*, 7(6), 198–202. [https://doi.org/10.1016/0169-5347\(92\)90073-K](https://doi.org/10.1016/0169-5347(92)90073-K)
- Mills, S. K., & Beatty, J. H. (1979). The propensity interpretation of fitness. *Philosophy of Science*, 46(2), 263–286. <https://doi.org/10.1086/288865>
- Millstein, R. L. (2009). Populations as individuals. *Biological Theory*, 4(3), 267–273. <https://doi.org/10.1162/biot.2009.4.3.267>
- Millstein, R. L. (2014). How the concept of “population” resolves concepts of “environment.” *Philosophy of Science*, 81, 741–755.
- Okasha, S. (2016). The relation between kin and multilevel selection: An approach using causal graphs. *The British Journal for the Philosophy of Science*, 67(2), 435–470. <https://doi.org/10.1093/bjps/axu047>
- Okasha, S. (forthcoming) *Fitness*. Cambridge University Press.
- Otsuka, J. (2015). Using causal models to integrate proximate and ultimate causation. *Biology and Philosophy*, 30(1), 19–37. <https://doi.org/10.1007/s10539-014-9448-9>
- Otsuka, J. (2016). Causal foundations of evolutionary genetics. *The British Journal for the Philosophy of Science*, 67, 247–269.
- Pearl, J. (2000). *Causality*. Cambridge University Press.
- Rice, S. H. (2005). *Evolutionary theory*. Sinauer Associates.
- Roff, D. A. (2008). Defining fitness in evolutionary models. *Journal of Genetics*, 87(4), 339–348. <https://doi.org/10.1007/s12041-008-0056-9>
- Spirtes, P., Glymour, C., & Scheines, R. (2001). *Causation, prediction and search* (2nd ed). MIT Press.
- Stearns, S. C. (1976). Life-history tactics: A review of the ideas. *Quarterly Review of Biology*, 51(1), 3–47. <https://doi.org/10.1086/409052>