Fitness: Philosophical Problems

Grant Ramsey, *University of Notre Dame, Notre Dame, Indiana, USA*

Charles H Pence, *University of Notre Dame, Notre Dame, Indiana, USA*

Fitness plays many roles throughout evolutionary theory, from a measure of populations in the wild to a central element in abstract theoretical presentations of natural selection. It has thus been the subject of an extensive philosophical literature, which has primarily centred on the way to understand the relationship between fitness values and reproductive outcomes. If fitness is a probabilistic or statistical quantity, how is it to be defined in general theoretical contexts? How can it be measured? Can a single conceptual model for fitness be offered that applies to all biological cases, or must fitness measures be case-specific? Philosophers have explored these questions over the last several decades, largely in the context of an influential definition of fitness proposed in the late 1970s: the propensity interpretation. This interpretation as first described undeniably suffers from significant difficulties, and debate regarding the tenability of amendments and alternatives to it remains unsettled.

Introduction

Few concepts have elicited such a long and heated debate in the philosophy of biology as that of fitness. Although fitness was not a central theoretical term in Darwin's 1859 original articulation of his theory of evolution, it quickly rose to importance. Evolution by natural selection is now standardly presented (e.g. Lewontin, 1970) as requiring three conditions: variation, heritability and fitness differences. Empirically, evolutionary studies are thus to a large degree studies of fitness differences and their consequences. Moreover fitness, symbolised as W, is a key variable in the mathematical formulations of evolutionary theory. **See also**: Fitness

The fact that biologists model and measure fitness seems to imply that fitness is a (measurable) property and that some biological entities are bearers of this property. However, this leads to several questions. What sort of property is fitness? Is it, say, a causal property, allowing one to justifiably say that some biological entities of a particular kind outcompeted the others because (where this because is understood causally) they were fitter? And what sort of biological entities can legitimately be considered bearers of the property of fitness? Can genes, organisms, populations or even species have fitness values? The last of these questions is strongly linked with the 'levels/units of selection problem' and it is therefore not the focus of this article. In particular, the question of whether to define fitness as a property of token organisms or of types of organisms has been the subject of some debate. Setting this debate aside, this article will centre on the nature of the property of individual fitness. Similarly, issues of individuation will not be considered. If one is to compare the fitness of two individuals, then one must have a way of identifying them as distinct individuals. Although individuality is fairly obvious for some taxa (e.g. alligators or apple trees) it is far from obvious in others (e.g. aphids or aspens). The individuation of organisms will be taken for granted in what follows, but it should be emphasized that this is an important issue related to that of fitness. See Chapter 4 of Godfrey-Smith (2009) for an excellent introduction to issues in reproduction and individuality. **See also**: Group Selection; Philosophy of Selection: Units and Levels; Selection: Units and Levels

Setting aside the arguments over individuation and the levels of selection and assuming that organisms are one of the kinds of entities that can bear fitness, one can begin to ask what sort of property organismic fitness might be. Fitness is often broken down into (at least) two components, viability and fecundity. The former concerns survival to maturity, whereas the latter concerns the production of offspring once mature. Throughout the history of evolutionary theory, fitness has occasionally been considered to be the organism's total offspring production, or a combination of these factors. The fitness of any organism, on this view, is just the number of offspring that it actually has – and, if it is reproductively successful, its fitness will 'ratchet up' with each offspring it bears.

This understanding of fitness as a measure of actual reproductive success (henceforth 'realised fitness') has the merits of being simple and readily measurable. However, it has some significant problems. Fitness is often used to explain evolutionary outcomes – organisms of one type are said to out-reproduce those of another type because the former are fitter than the latter. However, if fitness is just realised fitness, then the claim that 'the fitter organisms out-reproduce the less-fit' is equivalent to 'the organisms that reproduce more out-reproduce those which reproduce less,' a simple tautology. This argument, an old one in evolutionary theory (see, e.g. Butler, 1879, pp. 351–355), was taken up occasionally by philosophers (e.g. Popper, 1974), and is still utilised by creationists (Pennock, 1999, p. 101). Understanding fitness as realized fitness, then, results in an arguably fatal flaw.

In addition to the tautology problem, another problem with equating fitness with actual reproductive success is that it precludes distinguishing evolutionary responses due to fitness differences from those due to genetic drift. It is generally understood that populations exhibit drift, and that the magnitude of drift increases as the population size decreases. However, if fitness is equated with evolutionary outcomes, it can no longer be contrasted with drift, either as a distinct kind of outcome or as a distinct kind of cause. It is for these reasons that philosophers have sought conceptions of fitness that are not equated with evolutionary outcomes. The philosophical debate has centred on a position known as the propensity interpretation of fitness (PIF), which will be the focus of what follows.

The Propensity Interpretation of Fitness

The PIF was introduced in the late 1970s by Brandon (1978) and Mills and Beatty (1979). They argued that instead of considering the fitness of organisms to be identified with the actual number of offspring produced, fitness should instead be equated with the probabilistic propensity to produce offspring – a distribution of probability values describing how likely it is that an organism will produce no offspring, one offspring, etc. The PIF, then, appears to solve the problems just discussed. First, if fitness is a propensity, then it is a kind of dispositional property. Consider solubility, a common example of a dispositional property. The fact that an object (a grain of salt, say) is soluble does not mean that it actually will dissolve, but that it would, were it placed in the appropriate circumstances. Similarly, the fitness of an organism does not determine a particular reproductive outcome (unless one particular value in the distribution has a probability of 1), but describes how that organism is disposed to reproduce in different kinds of circumstances. And the degree to which outcomes can be explained by probabilistic causes (itself a contentious philosophical question; see also Mayr (1961)), the

organism's actual reproductive success is causally explained by its fitness.

If fitness is a dispositional property, then, unlike in the case of realised fitness, drift can be distinguished from selection. There are debates about how, precisely, this distinction should be drawn. Some, for example, argue that selection and drift represent distinct causes (Hodge, 1987; Millstein, 2006), whereas others hold that the distinction is best understood in terms of outcomes (Matthen and Ariew, 2002; Walsh, 2007). This issue aside, it is clear that if fitness is understood in terms of realised outcomes only, then the theoretical resources for distinguishing selection and drift do not exist.

As an example of how drift might be distinguished from selection, consider that the PIF provides a probability distribution over possible values of offspring production. An organism may have a probability of 0.2 of having 0 offspring, a probability of 0.1 of having 1 offspring, etc. Now consider how we would describe the reproductive output of all the individuals with a particular trait (having brown, as opposed to grey fur). This will again be a probability distribution, derivable from the individual-level distributions. If we take the arithmetic mean of the distribution, we are provided with the expected number of copies of the trait that will appear in the next generation. Assuming the trait is perfectly heritable, and that there is no migration or mutation (and other usual caveats), fitness values thus provide us with expectation values for the proportion of individuals possessing the traits found in the next generation (the ratio of individuals with brown fur, say). Depending on how drift is understood (a topic too far afield for us here), deviations from those expectation values either count as genetic drift or provide evidence for drift (though the nature and weight of this evidence will vary depending on the precise definition of genetic drift in use). Moreover, as there will be fewer individuals to be sampled from in small populations, there will be a larger deviation from this expectation value in these populations. (For the same reason, the deviation from the expectation value of a fair coin will tend to decrease with the number of flips. You would predict a large deviation from 50–50 for a sequence of three flips, but a small deviation for 300 flips.) The PIF thus allows for the conceptual distinction between drift and selection, and also reflects the observation that drift tends to have a higher magnitude in small populations. **See also**: Drift: Introduction; Drift: Theoretical Aspects

Challenges to the PIF

Despite the clear benefits of the PIF over realised fitness, the PIF as originally described suffers from serious problems. Considered here are some of the key challenges to the PIF. The possible responses to these challenges are discussed in the next two sections.

Challenge 1: The PIF makes fitness unknowable

In practice, the fitness of individual organisms will be difficult to ascertain under the PIF framework, and fitness values cannot be directly determined by observing a small number of cases of actual offspring production. Given that the propensity proposed by the PIF manifests as a probability distribution over all possible numbers of offspring, a large number of similar (if not clonal) organisms in similar environments will need to be observed in order to have any confidence in our estimate of an individual's fitness. In almost all biological cases, this will be exceedingly difficult.

Challenge 2: The arithmetic mean is not always a good way of modelling the propensity

Although describing the probability distribution associated with the reproductive success of an organism is useful, there are also many circumstances in which fitness must be considered as a single numerical value, to enable comparisons between the fitnesses of different organisms. As mentioned above, the traditional way to formalise this in the PIF is to let the numerical value of fitness be equal to the arithmetic mean or expectation value of the probability distribution. It has long been known, however, that the arithmetic mean is not always the best predictor of future population success (see, e.g. Lewontin and Cohen, 1969; Gillespie, 1974), and this problem was even noted by some of the creators of the PIF (Beatty and Finsen, 1989). In some of these cases, a geometric mean (or some other measure) may offer a better prediction of a trait's future frequency. The defender of the PIF must, therefore, either provide a method for reducing this probability distribution to a single numerical value – one that takes into account the fact that the arithmetic mean will not always be the best choice – or they must argue why such a reduction is not necessary for the cogency of the PIF.

Challenge 3: How the PIF is understood changes with differing environmental circumstances, population structure, etc.

Many researchers (Rosenberg, 1982; Sober, 2001; Matthen and Ariew, 2002; Ariew and Lewontin, 2004; Bouchard and Rosenberg, 2004) have argued that one basic theoretical premise in the PIF is flawed. The PIF assumes, they argue, that a single definition of the concept of fitness can be adequate to describe the fitness of individual organisms in all biological circumstances, and this is a false assumption. Some have argued further that fitness can only be specified by relativising to a particular set of fitness components of interest in a particular set of local environmental circumstances, or even to a particular pair of organisms of interest. As environment, population structure and other local factors change, the very definition of fitness, they argue, must change with them. To salvage the PIF, then, its defenders must show that it is sensitive in the appropriate way to all these influences, and that the theoretical worries raised by these critics concerning the possibility of a general measure of fitness are unfounded.

Challenge 4: What facts determine the probability distribution in the PIF? What are its environmental scope and time frame?

In the initial description of the PIF above, the precise facts that are to be used to fix the values of the probability distribution were left unspecified. Are these only facts about currently living organisms and their projected offspring numbers? To deal with mutations with effects on future generations (Crow and Kimura, 1956; Ahmed and Hodgkin, 2000), should the PIF take into account descendants in later generations? Building on the last challenge, which environmental factors should be considered part of fitness calculations, and which should be considered 'external' to individual fitness? The defender of the PIF will need to provide general answers to these questions in order for the probability distribution to be well founded in all biological cases.

Responses Abandoning the PIF

One obvious way to respond to this set of rather thorny challenges is by abandoning the PIF entirely, and developing a new way to understand individual fitness. One such attempt was put forth by Bouchard and Rosenberg (2004). They argue that the focus on propensities – in particular, the focus on measuring fitness using propensities – is mistaken. The fundamental notion of fitness is, rather, comparative. The best we can hope for in a general interpretation of fitness, then, defines fitness in terms of two organisms a and b, and an environment E:

> a is fitter than b in E = a's traits result in its solving the design problems set by E more fully than b's traits.
> (Bouchard and Rosenberg, 2004, p. 699)

They call this notion 'ecological fitness,' and endeavour to respond to several of the difficult philosophical issues inherent both in providing a definition of 'design problems' and measuring how well organisms might 'solve' those problems. **See also**: Adaptation and Constraint: Overview

Another highly influential attempt to develop an alternative to the PIF has come to be known as the 'statisticalist' interpretation of evolutionary theory. As framed in Matthen and Ariew (2002) (see also Walsh et al., 2002; Pigliucci and Kaplan, 2006, chap. 1), the PIF's troubles come from a conflation of two senses of 'fitness' that ought, in fact, to be kept distinct. The 'vernacular' notion of fitness is a highly general, usually descriptive concept of fitness. It appears in some characterisations of natural selection such as the

'principle of natural selection' (i.e. that 'if A has higher fitness than B in E, then A will probably outcompete B over time', see Bouchard and Rosenberg, 2004). This is to be contrasted with the 'predictive' notion of fitness as it appears in population genetics or mathematical biology. Predictive fitness is strictly specified, only valid in very precise circumstances, and quantifiable. **See also**: Population Genetics: Overview

The trouble with the propensity interpretation, Matthen and Ariew argue, is that it mistakenly takes these two notions of fitness to be connected – that is, both that the propensity itself is a characterisation of vernacular fitness, and that the expected number of offspring given that propensity is a characterisation of predictive fitness. They offer several arguments to the effect that, in many cases, predictive and vernacular fitness simply cannot be related, and that in any event, the connection offered by the PIF is the wrong one.

The statisticalist interpretation goes further than this, on two fronts. First, vernacular fitness, they claim – due to its generality and, as we saw in the challenges above, the impossibility of specifying a universally valid formula connecting it to components of fitness – is not particularly useful in coming to understand biological populations. This is yet another place where the PIF has gone astray. It is rather the predictive fitness (and, in particular, the predictive fitness of traits and not of individual organisms) that is the appropriate target for biological study (Walsh *et al.*, 2002; Walsh, 2003, 2004, 2007).

Second, there is a further, and more important, reason that a singular conceptual picture of vernacular fitness fails. A particular fitness distribution can be the result of many different types of causes. We might, in one population, determine that some sex-ratio strategy is beneficial, and that in another population a particular parental care strategy is beneficial (Matthen and Ariew, 2002, p. 67). In both cases, we are justified in saying that the outcomes at issue increase the fitness of the respective organisms. However, to phrase this change in terms of a general notion of vernacular fitness is, the statisticalist interpretation argues, to miss the point. Although both strategies may have 'high' vernacular fitness – perhaps even both instantiate the same fitness distribution – the causes of this distribution are radically different, and the PIF fails, it is claimed, to recognise or sufficiently account for this fact.

Finally, one more target of the statisticalist line of argument brings us into contact with an aspect of the PIF that was only briefly mentioned above. In addition to the idea that fitness is derivable from the propensity to produce offspring, and that it should be measured by the expected offspring number, the 'traditional' PIF claimed that fitness plays a particular sort of causal role in natural selection. In this view, when we say that on a standard view of natural selection, fitter organisms will outperform the less fit, we mean (at least in part) that this higher fitness is a cause of this higher performance. The fitness of individuals is causally responsible for their evolutionary success.

The statisticalist interpretation also strongly rejects this claim. Rather than playing a causal role in the biological world, fitness – along with natural selection and genetic drift – are merely pragmatically useful ways to summarise events that take place in the biological world. There is nothing causally significant about these summaries, and it is only prudence and the good judgment of investigators that leads us to sometimes measure fitness in one way, sometimes in another. **See also**: Fitness and Selection

The claims of the statisticalist interpretation constitute a dramatic revision of our conceptual structure for fitness, selection and drift, and have therefore been hotly contested in the philosophical literature – see, for example, Matthen (2009), Matthen and Ariew (2009), Walsh (2010), and Lewens (2010) in support, and Reisman and Forber (2005), Brandon and Ramsey (2006), Millstein (2006), Abrams (2007), Shapiro and Sober (2007), Gildenhuys (2009) and Ramsey (2013) in opposition.

Responses in Defence of the PIF

Several other researchers, meanwhile, have offered ways in which to salvage the central insights of the PIF, attempting to sidestep the problems developed above. One modification of the PIF, originally proposed by Beatty and Finsen (1989) and elaborated by Brandon (1990), is intended to address the problem of reducing the PIF to a single numerical value. Perhaps it was a mistake, then, to think that one single way of making this reduction was suitable in all circumstances and that, as discussed in challenge 2, the arithmetic mean is not the appropriate single mathematical model of the PIF. Brandon (1990, p. 20) modified the original expected-value formula for fitness, introducing a 'correction factor' intended to compensate for the effects of influences like variance. This correction factor takes the form of a function of the environment and the variance, added to the expected number of offspring, which Brandon termed $f(E, \sigma^2)$. This makes the PIF provide not a single numerical value for individual fitness, but a 'schema' of possible equations, each of which needs to be filled out given the details of the particular population to be measured, the distribution of variation within it and so forth. Such a defence could preserve the central insights of the PIF, at the cost of losing a mathematically unified definition of fitness (see Abrams, 2009).

A related approach, taken by Pence and Ramsey (2013), is to develop a more complicated conceptual and mathematical formulation of the PIF that can manage to avoid the objections developed above. Although the PIF understands organismic fitness as a propensity of individuals to have offspring, it says nothing about how that propensity should be analysed conceptually. The model offered by Pence and Ramsey thus attempts to provide a more detailed vocabulary in which to describe this propensity, and with it a more intricate way in which to compute a numerical measure of fitness, drawn from adaptive

dynamics. The Pence and Ramsey proposal is attempting to at once address challenges 2–4.

Implicit in the Pence and Ramsey approach is another way to defend the PIF, originally mentioned by Brandon (1990) and developed in more detail by Millstein (in press). As we have stated, the PIF has multiple interrelated components. One of these, which we might call the 'non-mathematical' portion of the PIF, is the claim that the fitness of an organism can be understood in terms of its propensity to produce offspring. Another, the 'mathematical' portion of the PIF, is the claim that the best measure of this propensity is the expected number of offspring. Both the responses we have just seen have attempted to salvage this mathematical portion of the PIF. But is this really necessary? Millstein (in press) argues that it would be better served if these two questions can be separated. Discarding this mathematical approach, and arguing for the propensity interpretation as an answer to the question of 'what fitness is', she claims, can evade many of these posed difficulties while leaving the question of the comparison of fitness distributions as a problem for mathematical biology. This addresses challenges 2–4 by arguing that the PIF does not stand or fall based on the quality of the mathematical models with which it is associated.

Finally, concerning challenge 1, does the difficulty of measuring the quantity described by the PIF undercut it as an interpretation of fitness? One response is to point out that there are multiple possible roles that an interpretation of fitness can play. If an interpretation plays one role well (serving as a theoretical foundation for fitness, say), it does not follow that it should also be useful for another (such as studying evolution in natural populations with limited data). On this response, the PIF is offered not as a measurable model for fitness in experimental studies (Endler, 1986), but rather as a way to ground theoretical considerations of natural selection in general – discussions of selection that are supposed to apply to every possible natural population, like those of Lewontin (1970) or Thoday (1953).

Conclusions

The debate over the best understanding of biological fitness has stood as a central problem in philosophical work on evolutionary theory for nearly 25 years, beginning with the rejection of the definition of fitness as actual contribution of offspring to the next generation. The 'traditional' version of the PIF now certainly shows its age, and has accumulated several counterexamples that are quite probably fatal. Several plausible ways forward are apparent – from moderate revisions of the PIF or novel ways to argue for its central insights, to the complete rejection of the PIF and a view of fitness as a statistical predictor set by the interests of particular investigators. This debate, as well, has ramifications for the ways in which we see natural selection, genetic drift and indeed the entire conceptual structure of

evolutionary theory. **See also**: Adaptation and Natural Selection: Overview; Drift: Introduction; Evolution: Views of; Natural Selection: Introduction; Philosophy of Selection: Units and Levels

References

Abrams M (2007) How do natural selection and random drift interact? *Philosophy of Science* **74**: 666–679.

Abrams M (2009) The unity of fitness. *Philosophy of Science* **76**: 750–761.

Ahmed S and Hodgkin J (2000) MRT-2 checkpoint protein is required for germline immortality and telomere replication in *C. elegans*. *Nature* **403**: 159–164.

Ariew A and Lewontin RC (2004) The confusions of fitness. *British Journal for the Philosophy of Science* **55**: 347–363.

Beatty JH and Finsen S (1989) Rethinking the propensity interpretation of fitness: a peek inside Pandora's box. In: Ruse M (ed.) *What the Philosophy of Biology Is: Essays for David Hull*, pp. 18–30. Dordrecht: Kluwer Academic Publishers.

Bouchard F and Rosenberg A (2004) Fitness, probability and the principles of natural selection. *British Journal for the Philosophy of Science* **55**: 693–712.

Brandon RN (1978) Adaptation and evolutionary theory. *Studies in History and Philosophy of Science* **9**: 181–206.

Brandon RN (1990) *Adaptation and Environment*. Princeton, NJ: Princeton University Press.

Brandon RN and Ramsey G (2006) What's wrong with the emergentist statistical interpretation of natural selection and random drift? In: Hull DL and Ruse M (eds) *The Cambridge Companion to the Philosophy of Biology*, pp. 66–84. Cambridge: Cambridge University Press.

Butler S (1879) *Evolution old and new; or, the theories of Buffon, Dr. Erasmus Darwin, and Lamarck, as compared with that of Mr. Charles Darwin*. London: Hardwick and Bogue.

Crow J and Kimura M (1956) Some genetic problems in natural populations. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability* **4**: 1–22.

Darwin C (1859) *On the Origin of Species*. London: John Murray.

Endler JA (1986) *Natural Selection in the Wild*. Princeton, NJ: Princeton University Press.

Gildenhuys P (2009) An explication of the causal dimension of drift. *British Journal for the Philosophy of Science* **60**: 521–555.

Gillespie JH (1974) Natural selection for within-generation variance in offspring number. *Genetics* **76**: 601–606.

Godfrey-Smith P (2009) *Darwinian populations and natural selection*. Oxford: Oxford University Press.

Hodge MJS (1987) Natural selection as a causal, empirical, and probabilistic theory. In: Krüger L, Gigerenzer G and Morgan MS (eds) *The Probabilistic Revolution*, pp. 233–270. Cambridge, MA: The MIT Press.

Lewens T (2010) The natures of selection. *British Journal for the Philosophy of Science* **61**: 313–333.

Lewontin RC (1970) The units of selection. *Annual Review of Ecology and Systematics* **1**: 1–18.

Lewontin RC and Cohen D (1969) On population growth in a randomly varying environment. *Proceedings of the National Academy of Sciences of the USA* **62**: 1056–1060.

Matthen M (2009) Drift and 'statistically abstractive explanation'. *Philosophy of Science* **76**: 464–487.

Matthen M and Ariew A (2002) Two ways of thinking about fitness and natural selection. *Journal of Philosophy* **99**: 55–83.

Matthen M and Ariew A (2009) Selection and causation. *Philosophy of Science* **76**: 201–224.

Mayr E (1961) Cause and effect in biology. *Science* **134**: 1501–1506.

Mills SK and Beatty JH (1979) The propensity interpretation of fitness. *Philosophy of Science* **46**: 263–286.

Millstein RL (2006) Natural selection as a population-level causal process. *British Journal for the Philosophy of Science* **57**: 627–653.

Millstein RL (in press) Probability in biology: the case of fitness. In: Hájek A and Hitchcock C (eds) *The Oxford Handbook of Probability and Philosophy*. Oxford: Oxford University Press.

Pence CH and Ramsey G (2013) A new foundation for the propensity interpretation of fitness. *British Journal for the Philosophy of Science*. doi:10.1093/bjps/axs037.

Pennock RT (1999) *Tower of Babel: the evidence against the new creationism*. Cambridge, MA: MIT Press.

Pigliucci M and Kaplan JM (2006) *Making sense of evolution: the conceptual foundations of evolutionary theory*. Chicago: University of Chicago Press.

Popper K (1974) Darwinism as a metaphysical research programme. In: Schilpp PA (ed.) *The Philosophy of Karl Popper*, pp. 133–143. New York and Chicago: Open Court Press.

Ramsey G (2013) Can fitness differences be a cause of evolution? *Philosophy and Theory in Biology* **5**: e401.

Reisman K and Forber P (2005) Manipulation and the causes of evolution. *Philosophy of Science* **72**: 1113–1123.

Rosenberg A (1982) On the propensity definition of fitness. *Philosophy of Science* **49**: 268–273.

Shapiro L and Sober E (2007) Epiphenomenalism – the do's and the don'ts. In: Wolters G and Machamer P (eds) *Thinking About Causes: From Greek Philosophy to Modern Physics*, pp. 235–264. Pittsburgh, PA: University of Pittsburgh Press.

Sober E (2001) The two faces of fitness. In: Singh RS (ed.) *Thinking About Evolution: Historical, Philosophical, and Political Perspectives*, pp. 309–321. Cambridge, MA: The MIT Press.

Thoday JM (1953) Components of fitness. *Symposia of the Society for Experimental Biology* **7**: 96–113.

Walsh DM (2003) Fit and diversity: explaining adaptive evolution. *Philosophy of Science* **70**: 280–301.

Walsh DM (2004) Bookkeeping or metaphysics? The units of selection debate. *Synthese* **138**: 337–361.

Walsh DM (2007) The pomp of superfluous causes: the interpretation of evolutionary theory. *Philosophy of Science* **74**: 281–303.

Walsh DM (2010) Not a sure thing: fitness, probability, and causation. *Philosophy of Science* **77**: 147–171.

Walsh DM, Lewens T and Ariew A (2002) The trials of life: natural selection and random drift. *Philosophy of Science* **69**: 429–446.

Further Reading

Abrams M (2012) Measured, modeled, and causal conceptions of fitness. *Frontiers in Genetics* **3**: 196.

Beatty JH (1992) Fitness: theoretical contexts. In: Keller EF and Lloyd EA (eds) *Keywords in Evolutionary Biology*, pp. 115–119. Cambridge, MA and London: Harvard University Press.

De Jong G (1994) The fitness of fitness concepts and the description of natural selection. *Quarterly Review of Biology* **69**: 3–29.

Henle K (1991) Some reflections on evolutionary theory, with a classification of fitness. *Acta Biotheoretica* **39**: 91–106.

Ramsey G (2006) Block fitness. *Studies in History and Philosophy of Biological and Biomedical Sciences* **37**: 484–498.

Roff DA (2008) Defining fitness in evolutionary models. *Journal of Genetics* **87**: 339–348.

Sober E (2000) *Philosophy of Biology*, 2nd edn. Boulder, CO: Westview Press.

Sober E (2010) Evolutionary theory and the reality of macro probabilities. In: Eells E and Fetzer J (eds) *The Place of Probability in Science*, pp. 133–162. Dordrecht: Springer.