The Unity of Fitness

Marshall Abrams[†][‡]

It has been argued that biological fitness cannot be defined as expected number of offspring in all contexts. Some authors argue that fitness therefore merely satisfies a common schema or that no unified mathematical characterization of fitness is possible. I argue that comparative fitness must be relativized to an evolutionary effect; thus relativized, fitness can be given a unitary mathematical characterization in terms of probabilities of producing offspring and other effects. Such fitnesses will sometimes be defined in terms of probabilities of effects occurring over the long term, but these probabilities nevertheless concern effects occurring over the short term.

1. Introduction. According to the original version of the propensity interpretation of fitness (PIF; Brandon 1978; Mills and Beatty 1979), biological fitness is a mathematical function of probabilities and numerical values associated with reproductive outcomes, in particular, expected number of offspring. This approach seems to take fitness to be a real aspect of the process of natural selection, an aspect that is approximated by various fitness and selection coefficient terms in models of selection, drift, and so on. In response to work by Gillespie (1973, 1974, 1975, 1977), some authors have argued that fitness should sometimes be defined in terms of a more complex function (Beatty and Finsen 1989; Brandon 1990; Sober 2001). Brandon (1990) argued that fitness, therefore, merely satisfies a common schema instantiated by different mathematical functions. More extreme conclusions have been drawn from arguments that fitness must sometimes be characterized by an even wider variety of mathematical functions because of conspecifics' mutual influence on reproductive success (Ariew and Lewontin 2004; Krimbas 2004). Despite the

heterogeneity of mathematical functions needed to *model* fitness, I argue as follows.

As a causal and explanatory factor, fitness must be relativized to an effect or explanandum to be caused/explained, but as long as the effect is carefully specified, there is a single, often simple, mathematical property that constitutes the fact that one type is fitter than another. In particular, for questions about what caused one type or another to be better represented at the end of an interval of time, a type *A* is fitter than another type *B* if and only if *A* has a greater probability of having proportionally increased its frequency more than *B* at the end of that interval.

Although in such cases fitness must be defined in terms of probabilities of reproductive effects over several generations, I argue that it nevertheless has to do with influence in each generation. Since probabilities of long-term effects can be derived from probabilities of short-term effects, the former are simply mathematical properties of causes acting in the short term. These short-term mathematical properties will often be difficult or impossible to discover with any precision; that is why they must be modeled by a variety of relatively simple approximations.

PIF advocates usually claim that the probabilities of which fitness is a function are propensities, a (proposed) kind of indeterministic disposition. One can criticize this aspect of the PIF by criticizing the concept of propensity generally or by arguing that if propensities exist, their behavior does not allow them to play the role that the PIF requires (Abrams 2007).[1] Such "function of what?" problems might be solved by arguing that some other kind of probability plays the appropriate role. I ignore this kind of problem here, simply assuming for now that appropriate probabilities— whether propensities or something else—do exist.

My focus will instead be on what I call the "what function?" problem. This is the question of what mathematical function of probabilities and other factors defines fitness. Here, my focus is on arguments that there is no one function common to all contexts in terms of which a central notion of fitness should be defined. Sections 2–4 concern what-function problems raised by Gillespie's work (Gillespie 1973, 1974, 1975, 1977; Beatty and Finsen 1989; Brandon 1990; Sober 2001). Section 5 addresses what-function challenges raised by Ariew and Lewontin (2004) and Krimbas (2004).[2]

I assume that in the final analysis, the kind of fitness relevant to natural

---

1. Eagle (2004) surveys criticisms of propensity concepts.

2. My responses below to the what-function challenges raised by Krimbas and Ariew and Lewontin actually involve part of an answer to a function-of-what question, especially the question of what sorts of outcomes the probabilities constitutive of fitness concern.

selection is fitness of *types*, that is, properties of organisms, since it is types that are heritable and selected for. Fitness is often attributed to token organisms, but marginal fitnesses of various heritable types can then be derived by averaging (e.g., Mills and Beatty 1979; Sober 1984, 2001; Ewens 2004; Rice 2004).

2. Troubles with Expectationalism. As mentioned above, the original PIF defined fitness as the expectation or arithmetic mean of the number of offspring:

$$\text{Fitness of } A = \sum_{i=0}^{\infty} i \Pr(O_A = i),$$

where $O_A = i$ says that an organism of type A has i offspring.[3] However, PIF advocates (Beatty and Finsen 1989; Brandon 1990; Sober 2001) and others have questioned the original PIF because of Gillespie's arguments (1973, 1974, 1975, 1977) that a type with a lower expected number of offspring than another can have a greater probability of long-term reproductive success. This can happen because the number of organisms of a given type in one generation is in part a function of the number of organisms previously producing offspring of that type. However, the latter number varies stochastically, producing consequent effects in the current generation. This means that the overall shape of the probability distribution over numbers of offspring, not just its mean, is relevant to long-term evolution. In very simple cases, the most probable outcome of competition between two types A and B over a large number of generations n is that A will increase its frequency in the population if and only if its geometric mean number of offspring is greater than B's:

$$\prod_{i=0}^{\infty} i^{\Pr(O_A=i)} > \prod_{i=0}^{\infty} i^{\Pr(O_B=i)}$$

(cf. Gillespie 1973, 194–195; Godfrey-Smith 1996, Chapter 7).

In a short discussion (Gillespie 1977) of his earlier work (Gillespie 1973, 1974, 1975), Gillespie pointed out that in some simple cases, we can approximate a predictive measure of fitness using certain functions of expectation and variance. Philosophers have, as a result, sometimes focused on variance as *the* mathematical property other than expectation that is relevant to fitness (Brandon 1990; Sober 2001; Walsh 2007). How-

3. Note that "expect" and "expectation" have no psychological or predictive connotation in themselves. Thus, one should not necessarily expect the "expected value" to occur, e.g., for a bimodal distribution. (Consider drawing a ball from an urn with 1,000 balls labeled "1," one labeled "2," and 1,000 more balls labeled "3.")

TABLE 1. DISTRIBUTIONS FOR *A* AND *B*.

Type	Number Offspring (Probability)					Arithmetic Mean	Variance	Skew	Kurtosis	Geometric Mean
	1	2	3	4	5					
A		.2	.6	.2		3.0	.4	0	.4	2.930
B	.05		.9		.05	3.0	.4	0	1.6	2.913

ever, Gillespie's claims about variance concerned only certain cases, and in any event Gillespie's analyses often involved approximating assumptions—in which case no claim about what precise mathematical function it is that counts as fitness would be implied.

Moreover, careful reading of some of Gillespie's papers shows that even when he proposes that fitness can be taken to be a function of expectation and variance, complex functions of higher statistical moments would still be needed to provide a precise definition of fitness.[4] Beatty and Finsen (1989) gave an example in which the third moment (skew) of two offspring distributions made a difference to fitness. It is not difficult to construct distributions over offspring numbers illustrating how higher moments might matter. For example, the distributions for *A* and *B* in table 1 have the same arithmetic mean, variance, and skew, but *B*'s fourth moment (kurtosis) is greater than *A*'s, and its geometric mean is smaller. The fact that *B*'s geometric mean is slightly smaller shows that in experiments over several generations, *A* would usually outreproduce *B*.

Beatty and Finsen (1989) and Brandon (1990) considered defining fitness in terms of long-term measures of success such as geometric mean number of offspring, Cooper's (1984) expected time to extinction (ETE), and Thoday's (1953) probability of persistence 10^8 years.[5] Brandon, however, argued that fitness cannot be defined by a long-term measure because fitness differences are supposed to reflect causes of evolution, and causes produce their effect over (many) short-term periods: "Selection has no foresight; it has no means to discriminate among organisms based on their long-term probability of having surviving offspring" (1990, 25). No doubt because mean and variance seem like respectable short-term properties,

4. The nth central moment of a random variable X has the form $E[(X - EX)^n]$, where E is expectation. Note that it is mentioned in the discussion at the bottom of page 1012 of Gillespie 1977 and it is implicit in the derivation of equation (7) on page 1013 that higher statistical moments can be relevant to relative reproductive success. Gillespie argues that these moments typically make a small contribution to fitness. He is surely correct in this, but that does not mean that higher moments never matter in practice, and they certainly matter in principle.

5. Geometric mean seems like a long-term measure if viewed as multiplying numbers of offspring in distinct generations (weighted by the probability of each kind of generation).

and because Gillespie's summaries used them, Brandon argued that fitness should be defined by various functions of mean and variance in different contexts. In response, Sober (2001, 313) remarked that "long-term probabilities imply foresight no more than short-term probabilities do," but he gave little explanation of this remark. Sober also gave examples that suggested that both short-term and long-term fitness measures might be useful. I agree with Sober's points, but I think that they can be further clarified and can be given a more systematic foundation, as I try to do below.

3. Long-Term Short-Term Fitness. Notice that the worry that prompted Beatty and Finsen and Brandon to revise the original expectation-based definition of fitness was that expected number of offspring sometimes did not correspond to probable *long-term* success. Long-term success is clearly what matters in some contexts. Moreover, Gillespie's results, on which these authors had focused, were in fact measures of fitness in terms of probable long-term success—*approximated* with simple statistical functions such as expectation and variance. Also notice that an apparently long-term measure like geometric mean in fact just captures a mathematical fact about a short-term probability distribution over numbers of offspring. Even ETE or probability of persistence 10^8 years can be considered a mathematical property of short-term probability distributions over numbers of offspring: *probabilities of changes in frequencies over the long term are implied by short-term reproductive probabilities* and other short-term probabilities. Let me explain.

To fix ideas, here is a relatively concrete example. Suppose we have a population of three As and two Bs and that reproduction is seasonal and asexual. For each number of offspring, there is a probability of producing that number of offspring by an A, similarly for a B. There are several possible states of the population in the next generation, for example, six As and four Bs, zero of A and one of B, or two of B, or three, or four, extinction of the population, and so on. Each state can result from the As and Bs producing numbers of offspring which add up to the numbers of As and Bs corresponding to that state. Each state's probability is thus implied by probabilities of numbers of offspring. The same is true for more complex examples.

More generally, at each generation there is a set of transition probabilities $Pr_1(j|i)$, probabilities the population will go into state j one generation later given state i in the current generation. These one-generation transition probabilities imply multiple-generation transition probabilities.[6]

6. For now, I assume that one-generation transition probabilities are the same in each

For example, the probability of going from i to k two generations after the current generation is the sum, over all intermediate states j, of probabilities of going to j and then into k:

$$\Pr_2(k|i) = \sum_j \Pr_1(k|j)\Pr_1(j|i).$$

Notice what this means: short-term probabilities of numbers of offspring imply probabilities of frequencies many generations later (see, e.g., Bharucha-Reid 1960; Grimmett and Stirzacker 1992; Ewens 2004). Since probabilities of long-term success are implied by short-term reproductive probabilities—the long-term probabilities are, as it were, *contained in* the short-term probabilities—we can define fitness in terms of long-term probabilities, *considering them as specifying properties of short-term probability distributions*. My illustration used discrete generations, but there are continuous-time models of similar phenomena, and short-term probabilities would also plausibly imply long-term probabilities in the more complex mathematical processes actually instantiated in nature. Thus, the problem that prompted Brandon to reject long-term measures of fitness can be avoided. Although long-term probabilities do not initially seem like they could be involved in causing evolution, properly understood, they are properties of short-term probabilities.

Moreover, to the extent that fitness should reflect probabilities of long-term success, this perspective allows a general characterization of fitness, regardless of how context determines the relevant long-term probabilities: *fitness consists of whatever mathematical properties of short-term probabilities imply probable long-term success*. Thus, fitness can be defined in terms of long-term probabilities without losing short-term efficaciousness. In a slogan, fitness is *defined globally but acts locally*. That is, fitness is defined in terms of probable long-term ("global") effects but accomplishes them via short-term ("local") effects—whose probabilities have the properties that constitute fitness. We can call this "long-term/short-term" (LT/ST) fitness.

But how long is long enough? Surely not always Thoday's 10^8 years. What if the entire population goes extinct before that time? All fitnesses would have been zero, and thus on such an account natural selection would not have taken place.[7]

---

generation because the environment is not changing, there is no frequency dependence, etc. Section 5 relaxes this assumption.

7. For related reasons, ETE is not appropriate in general. The ETE of a type A is the probability-weighted average of the number of years until there are no organisms of type A left in the population. However, A's ETE might be larger than B's, simply because there is a small chance that A will persist for trillions of years, even though within any interval of interest to us—say, 1 million years—B is likely to last longer.

4. Interval-Relative LT/ST Fitness. Natural selection is supposed to cause (or at least explain) evolution, and fitness differences are supposed to be an essential part of that cause. However, evolution over different intervals of time should be seen as different evolutionary explananda, or more precisely as different effects that might be caused by natural selection. That is, we should recognize that in different situations, we are interested in questions about different effects—about evolution over different intervals of time. But different causes may have different effects, and different explananda can have different explanantia. Thus, there is no reason to expect that fitness could be defined without reference to an interval of time, and from the current perspective it makes no sense to do so. We should talk of fitness relative to an interval of time I beginning from a specified point in time t when a population is in a specified state. I use the terms "I-fitness," "I-fitter," and so on for fitness understood as relativized to an interval I in this way.

In keeping with the arguments for LT/ST fitness, I suggest that fitness be understood this way:

> A type A is I-fitter than a type B if A has a greater probability of having increased its frequency more (as a percentage of the frequency at t) than B at the end of interval I (which begins at t).

For example, A might be likely to increase its frequency over a long period of time I_1 but then exhaust a resource it needs and thus go extinct before the end of I_2. In such a case, A might be I_1-fitter than B but I_2-less-fit than B. (Note that this definition of "fitter than" is consistent with an I-fitter type losing in a race against an I-less-fit type over the interval I; fitness differences in the current sense summarize facts about probabilistic outcomes and should not be taken to guarantee any particular outcome.)

5. Short-Term Probabilities of What? I argued above that we can understand fitness in terms of mathematical properties of probability distributions over numbers of offspring, even though those properties have implications concerning many generations of reproduction. Ariew and Lewontin (2004) and Krimbas (2004) gave a set of arguments that challenge even this possibility. These authors argued that different short-term fitness measures are needed for various cases: sexual reproduction, niche construction, overlapping generations, and so on. Among other things, some of the Krimbas/Ariew/Lewontin arguments seem to show that in some contexts fitness cannot be defined by a probability distribution over offspring, or even over grand-offspring, but rather distributions over many generations may be needed to define fitness. It thus appears that probabilities of long-term success cannot always be reduced to probabilities concerning a single generation; fitness does not always "act locally."

By way of response, I want to consider some of the toughest cases raised by Krimbas and Ariew and Lewontin. I argue that taking the relevant probabilities to be of events in organisms' lives rather than of numbers of descendants plausibly allows these cases to be handled by a sense of interval-relative fitness that is defined by short-term probabilities.

Case 1: In a species with sexual reproduction, probabilities of numbers of grand-offspring will depend on offsprings' mates as well as parents' number of offspring. Probabilities over numbers of offspring thus do not imply long-term probabilities; probabilities over number of grand-offspring seem needed. (This is so whether or not mating is random.) However, probabilities of an offspring mating with organisms of various types depend on each parent's (short-term) probabilities of producing numbers of organisms *of various types*, along with the distribution of parent types in the population. For it is these factors that determine probabilities of different frequencies of types in the generation in which offspring mate. Moreover, where mating is nonrandom, mating preferences determine (short-term) probabilities of mating with various types. Thus, probabilities of numbers of grand-offspring are implied by short-term probabilities, although over a richer space of outcomes (e.g., offspring types must be distinguished) than considered earlier. Arguments like those given in Section 3 can then show that these short-term probabilities imply long-term probabilities. Mathematical properties of the short-term probabilities are defined by these long-term probabilities and thus can constitute LT/ST interval-relative fitness in terms of these properties. (Other cases such as frequency dependence due to epistatic interactions between loci can be handled in a similar manner.)

Case 2: Traits affecting parental investment, parent-provided developmental context, and niche construction can indirectly affect probabilities of numbers of grand-offspring or later descendants without affecting probabilities of numbers of offspring. Such cases thus seem to require that the shortest-term probabilities are those over numbers of grand-offspring or later descendants—potentially dozens of generations later. However, a parent's type plausibly gives rise to a probability distribution over many possible events in its life (given the makeup of its population and its environment). These probabilities include inter alia probabilities of numbers of offspring of various types, in various locations, and so on, as well as probabilities of offspring-aid events, niche construction activities, and so on. This variety of short-term probability (over a rich outcome space) thus implies probabilities of numbers of offspring in future generations via effects produced by a current organism. Again it can be argued that properties of such short-term distributions imply probabilities of long-term success and can thus constitute LT/ST fitness relative to a given interval.

Case 3: Ariew and Lewontin (2004) point out that when generations overlap, probabilities of long-term success can depend on whether the population size is increasing or decreasing. Thus, it seems that fitness cannot be defined except by reference to what the population size is actually doing. Moreover, given that population changes are probabilistic, the way in which a population's size changes might not be determined by facts at a given time. Ariew and Lewontin seem to suggest that in such cases fitness cannot be defined by short-term probabilities but instead must be derived from actual events over a relatively long period. But this is incorrect. Probabilities of size change are implied by probabilities of events in organisms' lives such as those considered so far. The simple five-organism example I gave in Section 3 illustrates this idea. Each next-generation population state had a probability derived from probabilities of numbers of offspring for the types A and B. These population states, however, did not all involve the same population size. The probability of a particular increase in population size would then be the sum of probabilities of certain states with sizes larger than the current one; probabilities of decreases would be determined similarly. For the Ariew/Lewontin example, note that mathematics of transition probabilities for overlapping generations are well known. Again it appears that probabilities of various multigeneration evolutionary paths of the populations would be determined by strictly short-term probabilities, so LT/ST fitnesses would again exist.

It thus seems plausible that interval-relative LT/ST fitness can be defined in terms of probable outcomes at the end of an interval I, where the long-term probabilities are implied by short-term probabilities of events in organisms' lives. The short-term transition probability distributions at various times need not be the same, but probabilities that various later short-term distributions come into play will be implied by earlier short-term distributions along with initial conditions at the beginning of I. Long-term probabilities will then be implied by the various short-term probabilities and short-term-derived probabilities that various short-term distributions will be invoked.[8] (None of this need be true when environmental change is caused by factors beyond the environment in question, but that is no strike against the current view since it is generally agreed that fitness should be sensitive to such environmental changes.)

---

8. Similarly, the (long-term) probability that a team will win a tournament might depend on probabilities that it will win certain meets and thus face various opponents, in turn determining probabilities of winning subsequent meets, and so on.

6. Conclusion.

6.1. Comparative Fitness. The view that results from the preceding remarks can be summarized as follows:

- A type *A* is *I*-fitter than *B* if *A* has a greater probability of having increased its frequency more than *B* at the end of the interval *I*, where the relevant long-term probabilities are implied by short-term one-generation probabilities over certain events in organisms' lives.
- The relevant interval for *I*-fitness is the interval over which a chosen evolutionary effect takes place.
- That fitness is constituted by short-term probabilities means that it concerns events that act over the short term.
- Fitness is measurable in principle but estimated in practice. Models and sampling techniques for *aspects* of the preceding notion of fitness exist.

This provides a unified characterization of what it is for one type to be fitter than another. For a given population, there is, of course, a different fitness property for each interval *I*, but it should be clear now that this is what we should expect for different evolutionary effects/explananda. The current view differs from the earlier views discussed above (Brandon 1990; Sober 2001; Ariew and Lewontin 2004; Krimbas 2004), which allowed fitness properties to be disparate and without systematic relation to one another. Note that if the relevant short-term probabilities are causal (as they would be if they were propensities, for example), the current approach can help support a view of natural selection as a single kind of cause of evolution.

I want to suggest some possible further developments of this view. First, notice that what makes the overlapping-generations case different from those considered earlier is that it requires a probability distribution over times at which a parent produces offspring. More generally, sequencing and timing of events in relation to changes of states of the environment can matter to probabilities of future proliferation. This suggests that the basic short-term probabilities that constitute fitness should be probabilities of "organism-environment histories," sequences of possible events (both organismic and environmental) in the life of an organism of a given type (Abrams 2009a). Second, I suggest that an environment should be viewed as defining a probability distribution over conditions that members of a given population might experience during the interval of time *I* over which fitness is to be calculated (Abrams 2009b). This is to treat an environment as the environment for an entire population, rather than for a specific organism or subpopulation.

6.2. Population Dynamics. The definition of "fitter than" given above does not resolve all what-function questions about fitness. Although the definition can be used to define a sense of comparative fitness over very short intervals *I*, understanding short-term dynamics usually calls for a fitness degree property. Measures of relative strength of fitnesses are also relevant to some questions about long-term evolution; we may, for example, want to know not only which type will probably go to fixation but also how fast it is likely to do so.

In such cases the condition to be caused, explained, or predicted is more precise than the question of which type will probably increase its frequency more, and so on, and the relevant notion of fitness should reflect that fact. We should not expect, however, that there is one scalar fitness property that can be taken to be *the* property that objectively drives all short-term dynamics. For what really drives short-term dynamics is the set of full moment-to-moment transition probability distributions for members of a population, rather than sequences of summary statistics such as expectations. The full set of transition probability distributions implies a probability distribution over all possible sequences of frequency and population size changes over an interval of time. Then, we have to decide what kind of summary property of the latter distribution is of interest to us. Scalar time-indexed fitness properties can be relevant to prediction/explanation/causation of dynamics but only relative to a precise specification of what is to be predicted/explained/and so on. Thus, for example, the common use of expected number of offspring to model short-term dynamics is in fact ideal for answering one particular question: What is the path of the average population state conditional on the previous state through the space of possible relative frequencies of types? However, that is only one of many simple questions one might ask about population dynamics.

These last remarks might make one conclude that the notion of fitness is far less unified than is suggested by the rest of the article. Such a conclusion seems premature. For my remarks suggest only that, as in the rest of the article, a causal and explanatory property going by the name of "fitness" must be relative to the effect to be caused or explained. I have explained how to provide a systematic characterization of comparative fitness relativized to intervals of time. There is at present no reason to think that the tools developed for that purpose cannot be used as the basis of a systematic characterization of fitness for a broader range of evolutionary effects.

REFERENCES

Abrams, M. (2007), "Fitness and Propensity's Annulment?", *Biology and Philosophy* 22: 115–130.

——— (2009a), "Fitness 'Kinematics': Altruism, Biological Function, and Organism-Environment Histories", *Biology and Philosophy* 24 (4): 487–504.

———. (2009b), "What Determines Fitness? The Problem of the Reference Environment", *Synthese* 166 (1): 21–40.

Ariew, A., and R. C. Lewontin (2004), "The Confusions of Fitness", *British Journal for the Philosophy of Science* 55: 347–363.

Beatty, J., and S. Finsen (1989), "Rethinking the Propensity Interpretation: A Peek Inside Pandora's Box", in M. Ruse (ed.), *What the Philosophy of Biology Is*. Dordrecht: Kluwer, 17–30.

Bharucha-Reid, A. T. ([1960] 1996), *Elements of the Theory of Markov Processes and Their Applications*. Mineola, NY: Dover.

Brandon, R. N. (1978), "Adaptation and Evolutionary Theory", *Studies in History and Philosophy of Science* 9 (3): 181–206.

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

Cooper, W. S. (1984), "Expected Time to Extinction and the Concept of Fundamental Fitness", *Journal of Theoretical Biology* 107: 603–629.

Eagle, A. (2004), "Twenty-one Arguments against Propensity Analyses of Probability", *Erkenntnis* 60: 371–416.

Ewens, W. J. (2004), *Mathematical Population Genetics*, vol. 1, *Theoretical Introduction*. 2nd ed. New York: Springer.

Gillespie, J. H. (1973), "Polymorphism in Random Environments", *Theoretical Population Biology* 4 (2): 193–195.

——— (1974), "Natural Selection for Within-Generation Variance in Offspring Number", *Genetics* 76: 601–606.

——— (1975), "Natural Selection for Within-Generation Variance in Offspring Number", pt. 2, "Discrete Haploid Models", *Genetics* 81: 403–413.

——— (1977), "Natural Selection for Variances in Offspring Numbers: A New Evolutionary Principle", *American Naturalist* 111: 1010–1014.

Godfrey-Smith, P. (1996), *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.

Grimmett, G. R., and D. R. Stirzacker (1992), *Probability and Random Processes*. 2nd ed. Oxford: Oxford University Press.

Krimbas, C. B. (2004), "On Fitness", *Biology and Philosophy* 19 (2): 185–203.

Mills, S., and J. Beatty (1979), "The Propensity Interpretation of Fitness", *Philosophy of Science* 46 (2): 263–286.

Rice, S. H. (2004), *Evolutionary Theory: Mathematical and Conceptual Foundations*. Sunderland, MA: Sinauer.

Sober, E. (1984), *The Nature of Selection*. Cambridge, MA: MIT Press.

——— (2001), "The Two Faces of Fitness", in R. S. Singh, C. B. Krimbas, D. B. Paul, and J. Beatty (eds.), *Thinking about Evolution*. Cambridge: Cambridge University Press, 309–321.

Thoday, J. M. (1953), "Components of Fitness", *Symposium for the Society for Experimental Biology* 7: 96–113.

Walsh, D. M. (2007), "The Pomp of Superfluous Causes: The Interpretation of Evolutionary Theory", *Philosophy of Science* 74: 281–303.