2 The Two Faces of Fitness

Elliott Sober

The concept of fitness began its career in biology long before evolutionary theory was mathematized. Fitness was used to describe an organism's vigor, or the degree to which organisms "fit" into their environments. An organism's success in avoiding predators and in building a nest obviously contributes to its fitness and to the fitness of its off-spring, but the peacock's gaudy tail seemed to be in an entirely different line of work. Fitness, as a term in ordinary language (as in "physical fitness") and in its original bio-logical meaning, applied to the survival of an organism and its offspring, not to sheer reproductive output (Cronin 1991, Paul 1992). Darwin's separation of natural from sexual selection may sound odd from a modern perspective, but it made sense from this earlier point of view.

Biologists came to see that this limit on the concept of fitness is theoretically unjus-tified. Fitness is relevant to evolution because of the process of natural selection. Selec-tion has an impact on the traits that determine how likely it is for an organism to survive from the egg stage to adulthood, but it equally has an impact on the traits that determine how successful an adult organism is likely to be in having offspring. Success concerns not just the robustness of offspring but their number. As a result, we now regard viability and fertility as two components of fitness. If p is the probability that an organism at the egg stage will reach adulthood, and e is the expected number of offspring that the adult organism will have, then the organism's *overall fitness* is the product pe, which is itself a mathematical expectation. Thus, a trait that enhances an organism's viability but renders it sterile has an overall fitness of zero. And a trait that slightly reduces viability, while dramatically augmenting fertility, may be very fit overall.

The expansion of the concept of fitness to encompass both viability and fertility resulted from the interaction of two roles that the concept of fitness plays in evolution-ary theory. It describes the relationship of an organism to its environment. It also has a

mathematical representation that allows predictions and explanations to be formulated. *Fitness is both an ecological descriptor and a mathematical predictor.* The descriptive ecological content of the concept was widened to bring it into correspondence with the role that fitness increasingly played as a mathematical parameter in the theory of natural selection.

In this chapter I want to discuss several challenges that have arisen in connection with idea that fitness should be defined as expected number of offspring. Most of them are discussed in an interesting article by Beatty and Finsen (1989). Ten years earlier, they had championed a view they dubbed ''the propensity interpretation of fitness'' (Mills and Beatty 1979; see also Brandon 1978). In the more recent article, they ''turn critics.'' Should fitness be defined in terms of a one-generation time frame—why focus on expected number of *offspring* rather than *grand*offspring, or more distant descendants still? And is the concept of mathematical expectation the right one to use? The details of my answers to these questions differ in some respects from those suggested by Beatty and Finsen, but my bottom line will be the same—expected number of offspring is not always the right way to define fitness.

In what follows, I will talk about an *organism*'s fitness even though evolutionary theory shows scant interest in individual organisms but prefers to talk about the fitness values of *traits* (Sober 1984). Charlie the Tuna is not a particularly interesting object of study, but tuna dorsal fins are. Still, for the theory of natural selection to apply to the concrete lives of individual organisms, it is essential that the fitness values assigned to traits have implications concerning the reproductive prospects of the individuals that have those traits. How are trait fitnesses and individual fitnesses connected? Because individuals that share one trait may differ with respect to others, it would be unreasonable to demand that individuals that share a trait have identical fitness values. Rather, the customary connection is that the fitness value of a trait is the average of the fitness values of the individuals that have the trait. For this reason, my talk in what follows about the fitness of organisms will be a harmless stylistic convenience.

To begin, let us remind ourselves of what the idea of a mathematical expectation means. An organism's expected number of offspring is not necessarily the number of offspring one expects the organism to have. For example, suppose an organism has the following probabilities of having different numbers of offspring:

number (i) of offspring	0	1	2	3
probability (p_i) of having exactly i offspring	0.5	0.25	0.125	0.125

The expected number of offspring is $\sum ip_i = 0(0.5) + 1(0.25) + 2(0.125) + 3(0.125) = 0.875$, but we do not expect the organism to have precisely 7/8ths of an offspring. Rather, ''expectation'' means *mathematical* expectation, a technical term; the expected value is, roughly, the (arithmetic) average number that the individual would have if it

got to live its life again and again in identical circumstances. This is less weird than it sounds; a fair coin has 3.5 as the expected number of times it will land heads if it is tossed 7 times.

In this example, the expected number of offspring will not exactly predict an individual's reproductive output, but it will probably come pretty close. However, there are cases in which the expected value provides a very misleading picture as to what one should expect. Lewontin and Cohen (1969) develop this idea in connection with models of population growth. Suppose, to use one of their examples, that each year a population has a probability of 0.9 of having a growth rate of 1.1 and a probability of 0.1 of having a growth rate of 0.3. The expected (arithmetic mean) growth rate per year is $(0.9)(1.1) + (0.1)(0.3) = 1.02$; thus, the expected size of the population increases by 2% per year. At the end of a long stretch of time, the population's expected size will be much larger than its initial size. However, the fact of the matter is that the population is virtually certain to go extinct in the long run. This can be seen by computing the geometric mean growth rate. The geometric mean of n numbers is the nth root of their product; because $[(1.1)^9 (0.3)]^{1/10}$ is less than unity, we expect the population to go extinct. To see what is going on here, imagine a very large number of populations that each obey the specified pattern of growth. If we follow this ensemble for, say, 1000 years, what we will find is that almost all of the populations will go extinct, but a very small number will become huge; averaging over these end results, we will obtain the result that, on average, populations grow by 2% a year. Lewontin and Cohen point out that this anomaly is characteristic of multiplicative processes.

A simpler and more extreme example that illustrates the same point is a population that begins with a census size of 10 individuals and each year has a 0.5 chance of tripling in size and a 0.5 chance of going extinct. After 3 years, the probability is 7/8 that the population has gone extinct, but there is a probability of 1/8 that the population has achieved a census size of $(3)(3)(3)10 = 270$. The expected size of the population is $(7/8)(0) + (1/8)(270) = 33.75$. This expected size can be computed by taking the expected yearly growth rate of $(0.5)(3) + (0.5)0 = 1.5$ and raising it to the third power; $(1.5)(1.5)(1.5)10 = 33.75$. In expectation, the population increases by 50% per year, but you should expect the population to go extinct.

Probabilists will see in this phenomenon an analogue of the St. Petersburg paradox (Jeffrey 1983). Suppose you are offered a wager in which you toss a coin repeatedly until tails appears, at which point the game is over. You will receive 2^n dollars, where n is the number of tosses it takes for tails to appear. If the coin is fair, the expected payoff of the wager is

$$(1/2)\$2 + (1/4)\$4 + (1/8)\$8 + \cdots$$

The expected value of this wager is infinite, but very few people would spend more than, say, \$10 to buy into it. If rationality means maximizing expected utility, then people seem to be irrational—they allegedly should be prepared to pay a zillion dollars

for such a golden opportunity. Regardless of whether this normative point is correct, I suspect that people may be focusing on what will *probably* happen, not on what the average payoff is over all possible outcomes, no matter how improbable. Notice that the probability is only 1/8 that the game will last more than three rounds. What we expect to be paid in this game deviates enormously from the expected payoff.

For both ecologists and gamblers, the same advice is relevant: Caveat emptor! If you want to make predictions about the outcome of a probabilistic process, think carefully before you settle on expected value as the quantity you will compute.

2.1   The Long-Term and the Short-Term

The definition of fitness as expected number of offspring has a one-generation time scale. Why think of fitness in this way rather than as having a longer time horizon? Consider figure 2.1 adapted from Beatty and Finsen (1989). Trait A produces more offspring than trait B (in expectation) before time t^*; however, after t^*, A produces fewer offspring than B, and in fact A eventually produces zero offspring. The puzzle is that A seems to be fitter than B in the short term, whereas B seems to be fitter than A in the long term. Which of these descriptions is correct?

The issue of whether fitness should be defined as a short-term or a long-term quantity will be familiar to biologists from the work of Thoday (1953, 1958), who argued that fitness should be defined as the probability of leaving descendants in the very long run; he suggests 10^8 years as an appropriate time scale. Thoday (1958, p. 317) says that a long-term measure is needed to obtain a definition of evolutionary progress. This reason for requiring a long-term concept will not appeal to those who think that progress is not a scientific concept at all (see, for example, discussion in Nitecki 1988
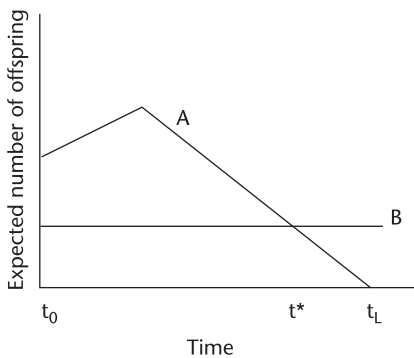

Figure 2.1
Trait A is fitter than Trait B initially, but later on the reverse is true. This means that B has a higher long-term fitness than A.

and Sober 1994). Thoday's argument also has the drawback that it repeatedly adverts to the good of the species without recognizing that this may conflict with what is good for individual organisms.

Setting aside Thoday's reason for wanting a log-term concept of fitness, does this concept make sense? Brandon (1990, pp. 24–5) criticizes Thoday's approach and the similar approach of Cooper (1984) on the grounds that selection "proceeds through generational time" and "has no foresight." I think both these criticisms miss the mark. Long-term probabilities imply foresight no more than short-term probabilities do. And the fact that selection occurs one generation at a time does not mean that it is wrong to define a quantity that describes a trait's long-term expected fate. Brandon also faults Thoday's proposal for failing to be operational. How are we to estimate the probability that a present organism or species will have descendants in the distant future? The point is well taken when the inference is *prospective*; in this case, the short-term is more knowable than the long-term future. However, when we make *retrospective* inferences, the situation reverses. An inferred phylogeny may reveal that a derived character displaced an ancestral character in one or more lineages. This information may provide evidence for the claim that the derived trait had the higher long-term fitness. In contrast, the one-generation fitnesses that obtained 60 million years ago may be quite beyond our ken.

Rather than rejecting a long-term concept of fitness and defending a short-term measure, I suggest that there is frequently no need to choose. In the accompanying figure, the y values for A and B at a given time tell us which trait had the higher short-term fitness at that time. The long-term fitness of a trait—its fitness, say, from t_0 to t^* or from t_0 to t_L—is a statistic that summarizes the relevant short-term values. There is no paradox in the fact that A has the higher short-term fitness whereas B has the higher long-term fitness. The same pattern can be found in two babies. The first has the higher probability of reaching age 20, whereas the second has the higher probability of surviving to age 60. The probability of a baby's reaching age 60 is a product—Pr (surviving to age 20 | you are a baby) Pr (surviving to age 60 | you have survived to age 20) $= (s_1)(s_2)$. The first baby may have a higher value on s_1 than the second, whereas the second has a higher value on s_2 than the first; overall, the first baby's product may be lower than that of the second. Long-term fitness is a coherent concept that may be useful in the context of certain problems; however, its coherence and desirability do not undermine the concept of short-term fitness.

2.2   When a One-Generation Time Frame Is Inadequate

The concept of short-term fitness discussed so far has a one-generation time frame—an organism at the egg stage has a probability p of reaching reproductive age and, once it is an adult, it has e as its expected number of offspring—the product pe is its overall

fitness. However, a one-generation time frame will not always be satisfactory for the concept of short-term fitness. Fisher's (1930) model of sex ratio shows why (Sober 1984). If, in expectation, one female has 5 sons and 5 daughters whereas another produces 10 daughters and 0 sons, how can their different sex-ratio strategies make a difference in their fitnesses? Fisher saw that the answer is invisible if we think one generation ahead but falls into place if we consider two. The sex ratio exhibited by a female's progeny influences how many *grand*offspring she will have.

Other examples may be constructed of the same type. Parental care is a familiar biological phenomenon, but let us consider its extension—care of grandoffspring. If A individuals care for their grandoffspring, but B individuals do not, it may turn out that A individuals are fitter. However, the advantage of A over B surfaces only if we consider the expected numbers of grandoffspring that survive to adulthood. This example may be more of a logical possibility than a biological reality; still, it and sex ratio illustrate the same point. In principle, there is no a priori limit on the size of the time frame over which the concept of fitness may have to be stretched. If what an organism does in its lifetime affects the life prospects of organisms in succeeding generations, the concept of fitness may have to encompass those far-reaching effects.

2.3   Stochastic Variation in Offspring Number

Let us leave the question of short-term versus long-term behind and turn now to the question of whether fitness should be defined as a mathematical expectation. This is not an adequate definition when there is stochastic variation in viability or fertility. Dempster (1955), Haldane and Jayakar (1963), and Gillespie (1973, 1974, 1977) consider stochastic variation among generations; Gillespie (1974, 1977) addresses the issue of within-generation variation. These cases turn out to have different mathematical consequences for how fitness should be defined. However, in both of them, selection favors traits that have lower variances. In what follows, I will not attempt to reproduce the arguments these authors give for drawing this conclusion. Rather, I will describe two simple examples that exhibit the relevant qualitative features.

Let us begin with the case of stochastic variation among generations. Suppose a population begins with two A individuals and two B's. A individuals always have two offspring, whereas the B individuals in a given generation all have one offspring or all have three, with equal probability. Notice that the expected (arithmetic average) offspring number is the same for both traits—2. However, we will see that the expected frequency of B declines in the next generation.

Assume that these individuals reproduce asexually and then die and that offspring always resemble their parents. Given the numbers just described, there will be four A individuals in the next generation and either two B individuals or six, with equal probability. Although the two traits begin with the same population frequency and have

the same expected number of offspring, their expected frequencies in the next generation differ:

Expected frequency of A $= (1/2)(4/6 + 4/10) = 0.535$

Expected frequency of C $= (1/2)(2/6 + 6/10) = 0.465$

The trait with the lower variance can be expected to increase in frequency. The appropriate measure for fitness in this case is the geometric mean of offspring number averaged over time; this is the same as the expected log of the number of offspring. Trait B has the lower geometric mean because $[(3)(1)]^{1/2} = 1.7 < [(2)(2)]^{1/2} = 2$. The geometric mean is approximately the arithmetic expected number minus $\sigma^2/2$.

Let us now consider the case of within-generation variance in offspring number. Gillespie (1974) describes the example of a bird whose nest has a probability of escaping predators of about 0.1. Should this bird put all its eggs in one nest or establish separate nests? If the bird lays 10 eggs in just one nest, it has a probability of 0.9 of having 0 offspring and a probability of 0.1 of having 10. Alternatively, if the bird creates 2 nests containing 5 eggs each, it has a probability of $(0.9)^2$ of having 0 offspring, a probability of $2(0.9)(0.1)$ of having 5, and a probability of $(0.1)^2$ of having 10. The expected value is the same in both cases—1.0 offspring—but the strategy of putting all eggs in one nest has the higher variance in outcomes. This example illustrates the idea of within-generation variance because two individuals in the same generation who follow the same strategy may have different numbers of offspring.

Does the process of natural selection vindicate the maxim that there is a disadvantage in putting all one's eggs in one basket? The answer is yes. To see why, let us examine a population that begins with two A individuals and two C's. A individuals always have two offspring, whereas each C individual has a 50% chance of having 1 offspring and a 50% chance of having 3. Here C individuals in the same generation may vary in fitness, but the expected value in one generation is the same as in any other. In the next generation, there will be four A individuals. There are four equiprobable arrangements of fitnesses for the two C individuals, and thus there are four equiprobable answers to the question of how many C individuals there will be in the next generation—two, four, four, and six. The expected number of C individuals in the next generation is four, but the expected frequencies of the two traits change:

Expected frequency of A $= (1/4)(4/6 + 4/8 + 4/8 + 4/10) = 0.52$

Expected frequency of C $= (1/4)(2/6 + 4/8 + 4/8 + 6/10) = 0.48$

Once again, the trait with the lower variance can be expected to increase in frequency.

In this example, the population grows from four individuals in the first generation to somewhere between 6 and 10 individuals in the second. Suppose we require that population size remain constant; after the four parents reproduce, random sampling

reduces the offspring generation to four individuals. When this occurs, the trait with the higher variance has the higher probability of going extinct.

Gillespie (1974, 1977) constructed a model to describe the effect of within-generation variance. A trait's variance (σ^2) influences what happens only when population size (N) is finite; in the infinite limit, variance plays no role. On the basis of this model, Gillespie says that a trait's fitness is approximately its arithmetic mean number of offspring minus the quantity σ^2/N. Notice that this correction factor will be smaller than the one required for between-generation variance if $N > 2$.

Why, in the case of within-generation variance, does the number of individuals (N) in the whole *population* appear in the expression that describes the fitness of a single *trait*, which may be one of many traits represented in the population? In our example, why does the fitness of C depend on the total number of C and A individuals? And why does the effect of selecting for lower variance decline as population size increases? The reasons can be glimpsed in the simple calculation just described. To figure out the expected frequency of C, we summed over the four possible configurations that the population has in the next generation. There is a considerable difference among these four possibilities—trait C's absolute frequency is either 2/6, 4/8, 4/8, or 6/10. In contrast, if there were 2 C parents but 100 A's, there still would be four fractions to consider, but their values would be 2/202, 4/204, 4/204, and 6/206; these differ among themselves much less than the four that pertain to the case of 2 A's and 2 C's. The same diminution occurs if we increase the number of C parents; there would then be a larger number of possible configurations of the next generation to consider, and these would differ among themselves less than the four described initially. In the limit, if the population were infinitely large, there would be no difference, on average, among the different possible future configurations.

The presence of N in the definition of fitness for the case of within-generation variance suggests that the selection process under discussion is density dependent. Indeed, Gillespie (1974, p. 602) says that the population he is describing is "density-regulated," for a fixed population size is maintained. However, we need to recognize two differences between the case he is describing and the more standard notion of density dependence that is used, for example, to describe the effects of crowding. In the case of crowding, the size of the population has a causal impact on an organism's expected number of offspring. However, the point of Gillespie's analysis of within-generation variance is to show that fitness should not be defined as expected number of offspring. In addition, the case he is describing does not require that the size of the population have any causal influence on the reproductive behavior of individuals. The two A's and two C's in my example might be four cows standing in the four corners of a large pasture; the two A's have two calves each, whereas each of the C's flips a coin to decide whether she will have one calf or three. The cows are causally isolated from each other, but the fitnesses of the two strategies reflect population size.

   In the two examples just presented, within-generation variance and between-generation variance have been understood in such a way that the former entails the latter, but not conversely. Because each C individual in each generation tosses a coin to determine whether she will have one offspring or three, it is possible for the mean offspring number produced by C parents in one generation to differ from the mean produced by the C parents in another. However, B parents in the same generation always have the same number of offspring. What this means is that B is a strategy that produces a purely between-generation variance, whereas C is a strategy that produces both within- and between-generation variance.

   In both of the examples I have described, the argument that fitness must reflect variance as well as the (arithmetic) mean number of offspring depends on the assumption that fitnesses should predict *frequencies* of traits. If, instead, one merely demanded that the fitness of a trait should allow one to compute the expected *number* of individuals that will have the trait in the future, given the number of individuals that have the trait initially, the argument would not go through. The expected number of individuals in some future generation is computed by using the arithmetic mean number of offspring. When the population begins with two B individuals or with two C individuals, the expected number of B or C individuals in the next generation is four. The value that generates this next-generation prediction is two—the arithmetic mean of one and three. Note that the variance in offspring number and the size of the whole population (N) are irrelevant to this calculation.

   That fitness is influenced by variance may seem paradoxical at first, but it makes sense in the light of a simple mathematical consideration. If traits X and Y are exclusive and exhaustive, then the number of X and Y individuals in a given generation determines the frequencies with which the two types occur at that time; however, it is not true that the *expected* number of X and Y individuals determines their *expected* frequencies. The reason is that frequency is a quotient:

frequency of X individuals =

   (number of X individuals)/(total number of individuals).

The important point is that the expected value of a quotient is not identical with the quotient of expected values:

E(frequency of X individuals) $\neq$

   E(number of X individuals)/E(total number of individuals).

This is why a general definition of fitness cannot equate fitness with expected offspring number. The fitness values of traits, along with the number of individuals initially possessing each trait, are supposed to entail the expected frequencies of the traits one or more generations in the future (if selection is the only force influencing evolutionary

change). Expected number of offspring determines the value of the quotient on the right, but the expected frequency is left open.

Notice that this point about the definition of fitness differs from the one that Lewontin and Cohen (1969) made concerning population growth. Their point was to warn against using the *expected number* of individuals as a predictor. The present idea is that if one wants to predict the *expected frequencies* of traits, something beyond the expected number of individuals having the different traits must be taken into account.

2.4 Conclusion

Evolutionists are often interested in long-term trends rather than in short-term events. However, this fact about the interests of *theorists* does not mean that the *theory* enshrines an autonomous concept called "long-term fitness." The long-term is a function of what happens in successive short terms. This metaphysical principle is alive and well in evolutionary theory. However, traits like sex ratio show that the short term sometimes has to be longer than a single generation.

The example of sex ratio aside, we may begin thinking about the fitness of a trait by considering a total probability distribution, which specifies an individual's probability of having $0, 1, 2, 3 \ldots$ offspring. The expected value is a summary statistic of this distribution. Although this statistic sometimes is sufficient to predict expected frequencies, it is not always a sufficient predictor; when there is stochastic variation in offspring number, the variance is relevant as well.

Are the mean and variance together sufficient to define the concept of fitness? Beatty and Finsen (1989) point out that the skew of the distribution is sometimes relevant. In principle, fitness may depend on all the details of the probability distribution. However, Gillespie's analysis of within-generation variance leads to a more radical conclusion. When there is stochastic variation within generations, Gillespie says that the fitness of a trait is approximately the mean offspring number minus σ^2/N. Notice that the correction factor adverts to N, the population size; this is a piece of information not contained in the probability distribution associated with the trait. It is surprising that population size exerts a general and positive effect on fitness.

The results of Dempster, Haldane and Jayakar, and Gillespie show how the mathematical development of a theoretical concept can lead to a reconceptualization of its empirical meaning. In Newtonian mechanics, an object's mass does not depend on its velocity or on the speed of light; in relativity theory, this classical concept is replaced with relativistic mass, which is the classical mass divided by $(1 - v^2/c^2)^{1/2}$. As an object's velocity approaches zero, its relativistic mass approaches the classical value. In similar fashion, the corrected definition of fitness approaches the "classical" definition as σ^2 approaches zero. People reacted to Einstein's reconceptualization of mass by saying that it is strange and unintuitive, but the enhanced predictive power of relativity

theory meant that these intuitions had to be re-educated. A definition of fitness that reflects the expected number of offspring, the variance in offspring number, and the population size yields more accurate predictions of expected population frequencies than the classical concept, and so it is preferable for the same reason.

It is sometimes said that relativity theory would not be needed if all objects moved slowly. After all, the correction factor $(1 - v^2/c^2)^{1/2}$ makes only a trivial difference when $v \ll c$. The claim is correct when the issue is prediction, but science has goals beyond that of making accurate predictions. There is the goal of understanding nature—of grasping what reality is like. Here we want to know which laws are true, and relativity theory has value here, whether or not we need to use that theory to make reasonable predictions. A similar point may apply to the corrected definition of fitness; perhaps evolving traits rarely differ significantly in their values of σ^2; if so, the corrected definitions will not be very useful when the goal is to predict new trait frequencies. This is an empirical question whose answer depends not just on how traits differ with respect to their variances but on the population size; after all, even modest differences in fitness can be important in large populations. But quite apart from the goal of making predictions, there is the goal of understanding nature—we want to understand what fitness is. In this theoretical context, the corrected definition of fitness is interesting.

What is the upshot of this discussion for the "propensity interpretation of fitness?" This interpretation has both a nonmathematical and a mathematical component. The nonmathematical idea is that an organism's fitness is its propensity to survive and be reproductively successful. Propensities are probabilistic dispositions. An organism's fitness is like a coin's probability of landing heads when tossed. Just as a coin's probability of landing heads depends on how it is tossed, so an organism's fitness depends on the environment in which it lives. And just as a coin's probability may fail to coincide exactly with the actual frequency of heads in a run of tosses, so an organism's fitness need not coincide exactly with the actual number of offspring it produces.

These ideas about fitness are not threatened by the foregoing discussion. However, the propensity interpretation also has its mathematical side, and this is standardly expressed by saying that fitness is a mathematical expectation (see, for example, Brandon 1978, Mills and Beatty 1979, Sober 1984). As we have seen, this characterization is not adequate in general, although it is correct in special circumstances. But perhaps all we need do is modify the mathematical characterization of fitness while retaining the idea that fitness is a propensity (Brandon 1990, p. 20).

This modest modification seems unobjectionable when there is between-generation variation in fitness; after all, if an organism's expected (= arithmetic mean) number of offspring reflects a "propensity" that it has, so too does its geometric mean averaged over time. However, when there is within-generation variation, the propensity interpretation is more problematic. The problem is the role of population size (N) in the

definition. To say that a coin is fair—that $p = 1/2$, where p is the coin's probability of landing heads when tossed—is to describe a dispositional property that it has. However, suppose I define a new quantity, which is the coin's probability of landing heads minus σ^2/N, where N is the number of coins in some population that happens to contain the coin of interest. This new quantity $(p - \sigma^2/N)$ does not describe a property (just) of the coin. The coin is described by p and by σ^2, but N adverts to a property that is quite extrinsic to the coin.

Is it really tenable to say that p describes a propensity that the coin has but that $(p - \sigma^2/N)$ does not? After all, the coin's value for p reflects a fact about how the coin is tossed just as much as it reflects a fact about the coin's internal composition. Perhaps the propensity is more appropriately attributed to the entire coin-tossing device. However, $(p - \sigma^2/N)$ brings in a feature of the environment—N—that has no causal impact whatever on the coin's behavior when it is tossed. It is for this reason that we should decline to say that $(p - \sigma^2/N)$ represents a propensity of the coin.

I conclude that an organism's fitness is not a propensity that it has—at least not when fitness must reflect the existence of within-generation variance in offspring number. In this context, fitness becomes a more ''holistic'' quantity; it reflects properties of the organism's relation to its environment that affect how many offspring the organism has; but fitness also reflects a property of the containing population—namely, its census size—that may have no effect on the organism's reproductive behavior. Of course, the old idea that fitness is a mathematical expectation was consistent with the possibility that this expectation might be influenced by various properties of the population; frequency-dependent and density-dependent fitnesses are nothing new. What is new is that the *definition* of fitness, not just the factors that sometimes affect an individual's expected number of offspring, includes reference to census size.

2.5 Acknowledgments

Dick is a ''natural philosopher.'' I do not mean this in the old-fashioned sense that he is a *scientist* (though of course he is that) but in the sense that he is a *natural at doing philosophy*. It was a striking experience during that year to find that Dick, a scientist,

was interested in the philosophical problems I was thinking about and that he was prepared to consider the possibility that they might be relevant to scientific questions. I came to the laboratory with the rather ''theoretical'' conviction that there should be common ground between science and philosophy, but the experience I had in the laboratory made me see that this could be true, not just in theory, but in practice.

During that year, I attended Dick's courses in biostatistics and population genetics; I gradually started to see how deeply the concept of probability figures in evolutionary biology. The present chapter, I think, is on a subject that is up Dick's alley. It is a pleasure to contribute this chapter to a volume that honors him.

My thanks to Martin Barrett, John Beatty, James Crow, Carter Denniston, Branden Fitelson, John Gillespie, David Lorvick, Steve Orzack, and to Dick as well for useful discussion of earlier drafts of this chapter.

References

Beatty, J., and Finsen, S. (1989). Rethinking the propensity interpretation—a peek inside Pandora's box. In *What the Philosophy of Biology Is*, ed. M. Ruse, pp. 17–30. Dordrecht, The Netherlands: Kluwer Publishers.

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

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

Cooper, W. (1984). Expected time to extinction and the concept of fundamental fitness. *Journal of Theoretical Biology* 107: 603–29.

Cronin, H. (1991). *The Ant and the Peacock*. Cambridge, UK: Cambridge University Press.

Dempster, E. R. (1955). Maintenance of genetic heterogeneity. *Cold Spring Harbor Symposium on Quantitative Biology* 20: 25–32.

Fisher, R. (1930). *The Genetical Theory of Natural Selection*. New York: Dover, 2d edition, 1957.

Gillespie, J. (1973). Natural selection with varying selection coefficients—a haploid model. *Genetical Research* 21: 115–20.

Gillespie, J. (1974). Natural selection for within-generation variance in offspring. *Genetics* 76: 601–6.

Gillespie, J. (1977). Natural selection for variances in offspring numbers—a new evolutionary principle. *American Naturalist* 111: 1010–14.

Haldane, J. B. S., and Jayakar, S. D. (1963). Polymorphism due to selection of varying direction. *Journal of Genetics* 58: 237–42.

Jeffrey, R. (1983). *The Logic of Decision*. Chicago: University of Chicago Press. 2d edition.

Lewontin, R., and Cohen, D. (1969). On population growth in a randomly varying environment. *Proceedings of the National Academy of Sciences* 62: 1056–60.

Mills, S., and Beatty, J. (1979). The propensity interpretation of fitness. *Philosophy of Science* 46: 263–86.

Nitecki, M. (ed.) (1988). *Evolutionary Progress*. Chicago: University of Chicago Press.

Paul, D. (1992). Fitness—historical perspectives. In *Keywords in Evolutionary Biology*, ed. E. Keller and E. Lloyd, pp. 112–14. Cambridge, MA: Harvard University Press.

Sober, E. (1984). *The Nature of Selection*. Cambridge: MIT Press. 2d edition, Chicago: University of Chicago Press, 1994.

Sober, E. (1994). Progress and direction in evolution. In *Progressive Evolution*?, ed. J. Campbell. Boston: Jones and Bartlett.

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

Thoday, J. (1958). Natural selection and biological process. In *A Century of Darwin*, ed. S. Barnett, pp. 313–33. London: Heinemann.