

Natural Selection and Decision-Making: Some Fundamental Principles

Eric Alden Smith and Bruce Winterhalder

2

2.1. INTRODUCTION

This chapter provides an introduction to the main theoretical principles employed throughout this volume. It is organized into three sections. In section 2.2 we present an overview of contemporary evolutionary theory, emphasizing the explanatory logic of natural selection. We describe the complexities imposed by the hierarchical organization of living systems and the techniques currently used to analyze the evolution of social interactions. Section 2.3 surveys selected theoretical and methodological issues in the social sciences. We emphasize methodological individualism and rational choice, topics that provide natural avenues for linking social theory and evolutionary ecology. In section 2.4, we summarize the general principles of simple optimization analysis, a framework commonly employed in both evolutionary ecology and the social sciences. We pay special attention to problems that can arise in applying optimization methods in an evolutionary context.

2.2. NATURAL SELECTION AND EVOLUTIONARY ECOLOGY

Any survey of the theory of natural selection limited to a few pages must be highly abbreviated. Our goal is to present those principles that are most relevant to the questions and findings raised in the later chapters of this volume, in a way that is accessible to readers unfamiliar with this body of theory. We begin (2.2.1) with a discussion of natural selection and its key components, as well as other evolutionary forces and constraints. Section 2.2.2 discusses problems surrounding levels of selection and adaptation:

What kinds of traits are favored by natural selection, and at what levels (genes, individuals, groups, etc.)? In section 2.2.3, we introduce the theory of evolutionarily stable strategies, the primary framework used by evolutionary ecologists in analysis of social interactions.

2.2.1. Selectionist Analysis and Its Limitations

What Is Natural Selection? In outline, natural selection is a simple process. Three conditions are required:

1. There must be phenotypic variation (differences between individuals).
2. Some of this variation must be heritable (transmitted to offspring).
3. Variants must differ in their ability to survive and reproduce (there must be fitness differences).

The term *phenotype* refers to characteristics of an organism other than DNA—its morphology, physiology, and behavior.

Some *phenotypic variation* is due to differences in genotype, some is due to environment (which influences phenotypes during ontogeny and also elicits short-term behavioral responses), and in the human case some is due to culturally acquired information (which may either be considered an aspect of environment, or be defined as heritable information analogous to genotype—see Chapter 3). The portion of phenotypic variation that is not attributable to differences between the environments of individuals is said to be *heritable*. Phenotypic characters are heritable if faithfully transmitted to offspring, even by nongenetic means. For selection to act, there must be heritability, but it need not be 100% (although the lower it is, the slower or less effective selection will be). *Fitness differences* must also exist if natural selection is to occur. It is possible to have heritable variation without fitness differences—differences in fingerprint patterns, or equally effective enzyme variants, for example—but the types of variation studied by evolutionary ecology (foraging patterns, mating systems, birth spacing, etc.) are unlikely to be selectively neutral.

Given heritable variation with fitness consequences, variants that reproduce at a higher average rate will tend to become relatively more numerous; in a finite environment, this process will eventually lead to the replacement of less “fit” variants by those with higher rates of replication. Of course, there are many subtleties and complexities hidden behind this simple statement (e.g., Darwin 1859; Williams 1966; and Krebs and Davies 1991). Some of them have to do with the concept of Darwinian fitness, our next subject.

What Is Fitness? In an important paper on life history theory (see Chapter 11), the evolutionary ecologist Stephen Stearns offered a somewhat tongue-

in-cheek definition of fitness as “something everyone understands but no one can define precisely” (Stearns 1976:4). This ambiguity has led some critics to judge evolutionary theory as fundamentally confused or even tautological. Such charges are unjustified.

In evolutionary biology, the term *fitness* derives from Herbert Spencer’s phrase (later adopted by Darwin) “survival of the fittest.” The potential for confusion arises from both terms. Survival colloquially means staying alive; but Darwin’s theory is not about survival in this sense, except as such survival increases lifetime reproductive output. The second term also confuses: fittest colloquially means strongest, healthiest, whereas what Darwin and his successors mean is relative reproductive success. But this explication of “survival of the fittest” invites the charge of tautology. What is fit? That which survives (reproductively). What survives (reproductively)? That which is fit.

One sensible solution to this quandary is that proposed by Mills and Beatty (1984). They define the fitness of an organism or a type (a “variant” or “trait” in our terminology) as “its *propensity* to survive and reproduce in a particularly specified environment and population” (p. 42). Fitness then refers to the *expected* number of descendants of a type (relative to other variants in the population), whereas the *actual* number of descendants generated by a given individual is determined by a number of factors, some of which might be unrelated to adaptive design (see also Williams 1966a: 102ff.; Brandon 1990:14–24). For example, we would not want to conclude that water-conserving traits have lower fitness in an arid environment just because one year’s data indicate that a disproportionate number of organisms with these traits died from lightning strikes. The propensity measure of fitness cautions us to look beyond limited data such as these, and to ask if *on average* water conservers have higher fitness than nonconservers. In effect, fitness propensity is a probabilistic measure. It focuses attention on adaptive *design* for reproduction rather than actual number of offspring per se.

The propensity interpretation of fitness helps us understand what evolutionary biologists mean by adaptations. For Darwin, adaptations were any characteristics that make an organism better fitted to survive and reproduce in its environment. What saves selection theory from tautology is that there is a highly regular relationship between an organism’s way of life, its environment, and the kinds of characteristics that will actually improve its fitness: Desert dwellers improve their adaptedness by traits that conserve water, water dwellers by hydrodynamic efficiency, nocturnal creatures by sensitive sight, hearing, or echolocation, and so on. Hence, fitness is the product of interactions between two sets of characteristics: those of the environment (including other organisms) and those of the organism or type being studied. The more we understand about an organism and its environment, the better we are able to predict what kind of traits—should they

arise—will be favored by selection. Fitness itself is tautological (as is any definition), but the explanatory framework of natural selection in which fitness is embedded is not.

What then is the significance of natural selection? Darwin's fundamental insight was this: *If* some process is at work generating new variants (with genetically transmitted variation, we now know these to be mutation and recombination), and *if* the relevant aspects of the interaction between organism and environment ("selection pressures") remain relatively stable, then natural selection will increase the adaptive fit of descendant members of the population relative to their ancestors. In addition, selection (in concert with other factors) will lead to the evolution of organisms with novel characteristics or abilities, and to new, reproductively isolated species. In this way, a blind mechanistic process creates adaptive design in nature, and provides a nonteleological and naturalistic explanation for such design.

Other Evolutionary Forces. Natural selection is only one of the elements involved in evolution. Mutation, recombination, and diploid reproduction all introduce random genotypic variation into populations. Genetic drift (sampling error due to small effective population size) and founder effects (when a small emigrant group is not representative of its original population) operate primarily to reduce genetic differences within groups, but increase it between them. There are also macroevolutionary processes (such as differential survival of species or higher taxa) that shape the diversity of living things. Most of these are random with respect to adaptive design (e.g., mass extinctions due to catastrophic events) but some may be directional (Gould 1982; Stanley 1975, 1979; Eldredge 1985; cf. Charlesworth et al. 1982).

While all of these factors play important roles in evolutionary theory, only natural selection has broad relevance to the questions addressed in this volume, which mainly concern variation and adaptive design in human behavior.

Constraints on Selection. Evolutionary ecologists make selection the centerpiece of their analytical efforts, but do so aware of various limitations on its effectiveness. Their models make implicit assumptions about the availability of well-behaved genetic variation, consistency of selection pressures, and predominance of selection over other evolutionary factors (2.4.1). They assume that various types of interactions (among physically or functionally linked portions of the genome; among various competing adaptive demands) do not impede selection for the trait(s) under analysis. They give secondary attention to questions of psychological and physiological mechanism that fall under the rubric of proximate analysis (1.2.1).

These assumptions sometimes will be inaccurate. Less often will they be so wrong as to debilitate the analysis. There are two reasons for this. First,

among all the traits available for study, evolutionary ecologists tend to focus on those most likely to meet the conditions just stated—that is, traits with large effects on fitness, which do not rely on specific forms of genetic inheritance. Second, partially as a result of the simplification gained by ignoring some factors, the results of evolutionary ecology models usually can be reliably tested. The risks that arise from simplifying assumptions are more than offset by the likelihood of discovering and correcting the mistake.

2.2.2. Levels of Selection

The Group Selection Controversy. In 1962 the Scottish ethologist Wynne-Edwards provoked a major evolutionary debate with this simple proposition: Many behavioral features of social animals are signal mechanisms by which individuals become aware of population crowding and subsequently act (individually and voluntarily) to reduce their own reproduction. They do this, said Wynne-Edwards, because otherwise depletion of resources and starvation or even group extinction would follow; hence natural selection has favored reproductive self-restraint and the social signals whereby it is maintained.

A flurry of rebuttals (Maynard Smith 1964; Williams 1966a; Wiens 1966; Lewontin 1970) followed on this work. The criticism focused attention on the units upon which natural selection acts. If we arrange the possibilities in a hierarchy of size and inclusiveness—genes, genotypes, isolated aspects of phenotype, whole phenotypes (individuals), kin (groups of related individuals), groups of randomly related individuals, breeding populations, communities, ecosystems, the biosphere—then we can say that the Wynne-Edwards proposal stimulated an extensive and productive debate on the *level of selection*. Wynne-Edwards was claiming that selection acted effectively at the level of groups, to maximize group rather than individual benefit. The critics claimed that selection is ineffective at this level, and any group benefits are simply by-products of individual ones. Newly sensitized by the level of selection debate, evolutionary biologists quickly realized that some of their favorite evolutionary explanations unwittingly relied on group selection and would need to be rethought.

As will become clear shortly, the meaning of the term *group selection* is itself controversial. For present purposes, the following definitions are sufficient. *Individual selection* is the form of natural selection that Darwin described; it "selects" between individuals who differ in heritable characteristics, and favors the characteristics that cause individuals to leave the largest relative number of surviving offspring. *Group selection* acts on heritable variation between groups of individuals; it favors characteristics that result in increased survival of groups (vs. extinction) and in increased reproduction of groups (through immigrants who disperse to found new

groups or join other existing ones). "Groups" is general enough to include the various units lying between individuals and the biosphere in the list above, but has usually referred to demes (isolated breeding populations) or subsets of demes. Finally, *kin selection* acts via the effects of characteristics expressed in one individual on the reproductive success of itself and its close relatives; it favors characteristics that have highest "inclusive fitness" (see below).

The problem with Wynne-Edwards's proposition is evident if we consider the effect of natural selection on the frequency of alternative traits (and underlying genetic variants), not on the condition of individuals or groups per se. Imagine that a growing population faces a shortage of some essential common resource. Further population growth would lead to resource depletion, increased mortality, and possibly even extinction of the population. This could be averted if individuals voluntarily and unilaterally limited the number of offspring that they bear and successfully raise to maturity. We might well expect such restraint, and explain it in terms of the benefits to group survival and well-being. But is it feasible that natural selection would produce such an outcome?

In most circumstances, the answer is no. The rationale for restraint is that otherwise the population will grow and cause resource depletion; but this means that it is possible (in the short run, at least) to raise additional offspring. Under these conditions, a strategy of reproductive restraint will always be of lower relative fitness within the group than one that causes the production of slightly greater numbers of offspring that survive to reproduction. The more profligate members gain the immediate fitness advantages of additional offspring. Whenever the undesirable costs of overpopulation are shared equally by all, they gain a long-term reproductive advantage as well: Even if resource depletion brings population stasis or decline, their descendants will be represented in relatively greater numbers. Thus, even though in the long run all individuals would be healthier and the population would be larger if resources were not overexploited, selection within groups will relentlessly favor the more prolific. This is the evolutionary equivalent of the tragedy of the commons (Hardin 1968).

We might save the proposition if the differential extinction of groups was frequent enough, and the opportunities for emigration from crowded groups to less crowded ones scant enough. In other words, group selection for population control (with selection favoring groups with high frequencies of restrained reproducers) would have to be stronger than individual selection for increased reproduction. But considerable analysis has cast doubt on this possibility, at least in the form envisioned by Wynne-Edwards. Individuals generally both reproduce and die at higher rates than the groups of which they are members. Furthermore, genetic variation is generally greater within groups (among its members) than between them (within a single species).

Since selection can only operate on heritable variation associated with differential fertility and mortality, the two facts just stated tell us that while group selection of the Wynne-Edwards sort might occur, it will generally be a weak evolutionary force when measured against the pervasiveness and effectiveness of selection on individuals.

Since the critical reaction to Wynne-Edwards, a second family of group selection models has been produced. These "structured-deme" models make different assumptions and yield different results. First, the structured-deme models include an intermediate level of organization between individuals and the deme or breeding population: a temporary association of interacting individuals, often termed a "trait group" (Wilson 1977, 1980). Members of trait groups interact, and traits that are "altruistic" (costly to individual fitness but beneficial to the fitness of other group members) are expressed at this point. Subsequently, members disperse, and contribute to the next generation in the population at large. Second, rather than differential *extinction* of groups as in the Wynne-Edwards formulation, in the newer models group selection is fueled by differential *propagation* of group-advantageous traits in the dispersal (or population-mixing) phase. This is actually the original approach to group selection developed by Sewall Wright (1945) in his pioneering model; it has only recently been given detailed theoretical treatment (Harpending and Rogers 1987; Rogers 1990b).

The structured-deme or trait group approach to group selection avoids many shortcomings of Wynne-Edwards's version; and it overlaps considerably with the kin selection inclusive-fitness approach pioneered by Hamilton (1964; see below). Thus, there has been renewed debate in the last decade concerning the significance of group selection as an evolutionary force. The analytical nuances of this debate are extensive and important, but cannot be detailed here (see Maynard Smith 1976, 1987; Wade 1978, 1985; Uyenoyama and Feldman 1980; Michod 1982; Wilson 1983; Grafen 1984; Brandon and Burian 1984; Sober 1984; and Nunney 1985). Nevertheless, it seems fair to make some summary statements. Evolutionary biologists are quite skeptical of unsupported claims that, because a given characteristic is or would be beneficial to the survival or well-being of one or more supraindividual units, it has or will be favored by natural selection. On the other hand, the extent to which the newer, more sophisticated models of group selection may account for various biological phenomena is more open to debate and empirical assessment. In addition, some researchers have argued that human societies show much higher levels of cooperation between unrelated individuals than is found in other species or can be derived from individual and kin selection. Prominent explanations for this include a history of particularly potent intergroup competition (Alexander 1974; Hamilton 1975) and particular features of culturally inherited variation (Boyd and Richerson 1982, 1985; see 3.4.4). Resolution of the role of group selection

thus awaits further work, although many evolutionary biologists remain convinced it plays a role only in special and relatively rare circumstances.

Kin Selection. With individual selection a trait spreads solely through its fitness-enhancing effects on the actor, measured in terms of direct descendants. This is selection as envisioned by Darwin. By contrast, kin selection (Maynard Smith 1964) expands the assessment of fitness to include the fitness-enhancing effects of an organism on biologically related individuals with which it interacts. Kin selection is often measured in terms of inclusive fitness ("inclusive" of related individuals), referring to the net effect of a particular trait or pattern of behavior on Ego's own fitness, plus its net effect on the fitness of related individuals, each devalued by their degree of relationship to Ego (Hamilton 1964).

Kin selection has been frequently invoked as a neo-Darwinian explanation of the evolution of altruistic behavior between closely related individuals (e.g., the members of a colony of social insects, which is generally a large family group). The altruist actor by definition suffers some loss of individual fitness, but the genotype for altruism can still increase in the population if the behavior results in sufficient fitness benefits to nondescendant kin. Depending on their genealogical closeness to the actor, kin have specific probabilities of sharing the same genotype by common descent. The aggregate fitness consequences of the altruist's phenotype on the survival and reproduction of kin carrying these shared genes may more than offset the loss in individual fitness incurred by the altruist. Specifically, Hamilton's rule says that altruism will be favored by selection when $C < rB$, where C is the fitness cost paid by the altruist, B is the benefit received by the recipient, and r is the coefficient of relationship by immediate descent (e.g., 0.5 for full sibs or parent to offspring, 0.25 for grandchildren, nephews, and nieces).

Although theory in evolutionary ecology has made relatively little use of the logic of kin selection or inclusive fitness, it is worth stressing two points. First, there is a widespread belief that kin selection theory predicts that altruism will characterize most interactions between kin, and in proportion to the genealogical closeness (the coefficient r); but this notion is mistaken (Dawkins 1979; Altmann 1979; Harpending 1981). For example, it is easy to see that Hamilton's rule allows the evolution of competition between kin, even infanticide or fratricide; simply make C the fitness gain to Ego, B the loss suffered by Ego's kinsman of relatedness r , and ask whether the gain to Ego is sufficiently great to offset the fitness loss to the kinsman, devalued by r . There is no reason the answer must be no (and history is replete with cases of siblings fighting to the death over inheritance of the throne or family estate).

Second, much discussion of kin selection, by both proponents and critics (e.g., Sahlins 1976), is phrased as if r were the only factor that really

mattered. To the contrary, treatment of kin of a given degree of relatedness (e.g., offspring) varies greatly between and within animal species and human societies, and Hamilton's rule predicts that these differences must be due primarily to variation in costs and benefits (C and B). In turn, variation in C and B is determined by social, ecological, and life history circumstances. Hence any deep understanding of variation in kin interactions will require the insights of social and ecological theory; focusing on relatedness alone yields an extremely truncated version of Hamilton's model (Smith 1979b).

The Phenotypic Gambit. Beginning in the 1920s with the work of Fisher, Haldane, and Wright, the mathematical theory of population genetics has given precision to that part of natural selection theory dealing with inheritance and genetic variation in populations. The analysis of organism-environment interactions that is the focus of evolutionary ecology arose later, and has required somewhat different conceptual tools.

The phenotypic traits of greatest ecological interest (including behavioral traits) are not controlled by single genes. Their expression is complex and multicausal, and dependent on environment in ways not easily captured in exact genetic models. Because of this, evolutionary ecologists typically treat the observable phenotypes of organisms (including behavior) as adaptations, avoiding detailed assumptions about heritability. In effect, they postulate that some underlying "strategies" or "decision rules" have been shaped by selection so as to produce adaptive phenotypes. According to this approach, it is neither necessary nor feasible to demonstrate the exact heritable basis of every trait of interest. Instead, one should proceed as if the precise nature or details of the link between heritability and phenotype were unimportant.

Grafen (1984) calls this research strategy the "phenotypic gambit" (see also Maynard Smith 1978). The phenotypic gambit analyzes a trait

as if the very simplest genetic system controlled it: as if there were a haploid locus at which each distinct strategy was represented by a distinct allele, as if the payoff rule gave the number of offspring for each allele, and as if enough mutation occurred to allow each strategy the chance to invade. (Grafen 1984:63-64)

Since few if any of the traits studied by evolutionary ecologists are controlled by single loci in a haploid system, the phenotypic gambit is really based on the premise that selection will favor traits with high fitness (or evolutionary stability; see 2.2.3) irrespective of the particulars of inheritance. In fact, behavioral ecologists usually assume extreme phenotypic plasticity, a wide array of feasible strategies, and the ability of the actor to assess payoffs and choose or learn the best alternative under any given set of circumstances.

Further, since the fitness consequences of different strategies often are

difficult to assess, evolutionary ecologists frequently employ more tractable measures such as energy capture per unit time, or fertility rates (see 2.4.3). The key assumption here is that these "proximate currencies" are highly correlated with fitness—that they are good indices or proxy variables for fitness, to use the language of the social sciences.

In summary, there are two ways in which evolutionary ecology diverges from evolutionary genetics: (1) the phenotype, not the genotype, is taken as the unit of study; (2) some correlated but more readily measurable index of evolutionary success is substituted for fitness. Lacking the mechanistic dynamic inherent in population genetic models of evolution, evolutionary ecology has substituted alternatives suited to its focus on phenotypic strategies and proximate currencies. The most popular approaches have been evolutionary game theory (2.2.3) and simple optimality models (2.4).

2.2.3. Evolutionary Game Theory

Natural selection results from the interaction between a population of organisms and the environment. As long as "environment" is external to the evolving population, selection generally favors those alternatives that confer the highest average fitness on their bearers. But when the relevant aspect of the environment consists of conspecifics (as in the case of social interactions), selection operates in a more complex, reflexive manner: The fittest strategy must do well in competition with (copies of) itself, not just in competition with other strategies. The first situation allows for simple optimization analysis (2.4); the latter requires the strategic analysis of evolutionary game theory.

An analytical framework for strategic contexts was first formalized by mathematical economists under the rubric of "the theory of games" (von Neumann and Morgenstern 1944; Luce and Raiffa 1957) and subsequently given a neo-Darwinian application in the theory of "evolutionarily stable strategies" or ESS (e.g., Maynard Smith 1974a, 1982a; Parker 1984; Parker and Hammerstein 1985). The key insight in ESS theory is this: When the relative payoff of alternative strategies or phenotypic traits depends on what other individuals in the population are doing, the outcome favored by natural selection depends on which alternatives are *unbeatable* rather than on which has the highest average payoff. A strategy that has high payoff when it is rare (and rarely encountered) in the population may have low payoff when common, or vice versa. In strategic contexts the payoff from each strategy must be calculated in light of all the possible strategies that can be played against it, including itself. A standard example of an ESS is the Hawk–Dove game (Box 2.1).

ESS Models and Behavioral Ecology. ESS theory involves complexities beyond the scope of this chapter, but it has two implications that need to be

Box 2.1. *The Hawk–Dove Game.* Imagine a situation in which there are frequent pairwise contests over possession of some resource. Suppose there are just two tactics (possible ways of behaving) in such a contest, Hawk and Dove. The Hawk tactic is to fight aggressively over the contested resource; the Dove tactic involves peaceful bluffing, and yielding as soon as real aggression seems likely. Hence, Hawk always beats Dove. The problem is to determine the ESS, i.e., the ratio of Hawk:Dove that will be unbeatable in terms of average fitness payoffs per player.

Suppose the resource itself is worth 100 fitness points, and that fighting and losing exacts a larger cost, say -300 points, while simply giving up without a fight (the tactic Doves follow when encountering Hawks) yields 0 points (no change in fitness). Suppose further (for simplicity) that when Hawk plays against Hawk, or Dove against Dove, each player has an equal probability (.5) of winning or losing.

Given these assumptions, the average payoffs to each tactic for the three possible types of contests are as follows:

1. Hawk vs. Hawk = $(100/2) + (-300/2) = -100$
2. Hawk vs. Dove: Hawk = 100, Dove = 0
3. Dove vs. Dove = $(100/2) = 50$

It is conventional in both classical and evolutionary game theory to present such a payoff structure in matrix form, with the tactics of one player listed at the top of columns, and the tactics of the other listed by row on the left. For simplicity, payoffs in the cells are those gained by the "row" player only.

For the Hawk–Dove game as described, the matrix is:

	Hawk	Dove
Hawk	-100	100
Dove	0	50

The evolutionary equilibrium is reached when the average payoff to Hawk equals the average to Dove, given the frequency of each strategy in the population. In the present hypothetical case, this equilibrium ratio works out to 1:2. That is, the ESS is to play Hawk one third of the time, and to play Dove in the remainder of encounters; alternatively, one could be Hawk or Dove for life, and the 1:2 ratio would express the ratio of each type of individual in the population. In ESS terminology, the first case is termed a "mixed strategy" (e.g., play Hawk with probability .33, play Dove with probability .67); the second is a balanced polymorphism (analogous to that maintained by heterozygote superiority with sickle-cell vs. normal hemoglobin in a malarial environment).

To see why this ratio is the ESS, consider a population consisting initially of all Dove. A single Hawk mutant will win every contest, because Hawk always beats Dove. Since the prize for winning is defined in fitness gains, Hawks will increase in the population (assuming simple inheritance). But as the Hawk tactic proliferates, it begins playing against itself with appreciable frequency; every time it does, the average

payoff to Hawk is -100 , much worse than Dove gets playing against itself (50) or against Hawk (0). At a certain point (which defines the ESS), the frequency of the Hawk tactic will be such that its average payoff will be no higher than the average payoff to Dove. The same kind of argument applies to a population of all Hawk that is invaded by a mutant playing Dove.

To check that the 1:2 ratio is the ESS, we can compute the average payoffs to each tactic given the expected encounter frequencies. At 1 Hawk:2 Doves, each player plays against a Hawk in $\frac{1}{3}$ contests, and against Dove in the remaining $\frac{2}{3}$. The computations are:

Tactic	Average payoff from		Overall average
	Plays against Hawks	Plays against Doves	
Hawk	$.33(-100) = -33$	$.67(100) = 67$	33
Dove	$.33(0) = 0$	$.67(50) = 33$	33

Using this method, the reader can easily check that any deviation from the 1:2 ratio will result in a player obtaining a lower average payoff. In fact, the general solution to the Hawk–Dove ESS is that the equilibrium Hawk frequency equals V/C , where V is the value (fitness gain) of winning the contest and C the cost of losing a fight (in the present example, $V/C = 100/300 = 0.33$). (If $V \geq C$, the equilibrium frequency of Hawk reaches 1.0, and pure Hawk is the only ESS.)

Sources: The Hawk–Dove game was first described in the seminal paper on ESS theory by Maynard Smith and Price (1973). A detailed review of this and many other evolutionary games is found in Maynard Smith (1982a). Dawkins (1976:Chapter 5) provides a clear and non-technical discussion, from which we have borrowed heavily for this exposition. In classical game theory, a formally identical payoff structure is found in the game of Chicken (Rapoport et al. 1976). Further discussion of Hawk–Dove, and applications to human social behavior, can be found in sections 9.5.3 and 10.3.2.

emphasized. First, to apply ESS theory one asks what strategy or set of strategies will be unbeatable over evolutionary time. A strategy is an ESS if when common in the population it cannot be replaced (via natural selection) by specified alternative strategies. Second, ESS theory shows that conflicts of interest between actors result in evolutionarily stable outcomes that often have lower payoffs than could be achieved if cooperation could be somehow ensured. An ESS is the optimal strategy to follow, but only in the special sense that any actor who deviates from it will be worse off (have lower fitness), even though if *all* deviated in concert they might all be better off. Thus, in the Hawk–Dove game, a “conspiracy of doves” would yield a higher per capita fitness, and indeed higher fitness for every member of the population (Dawkins 1976:77). For the hypothetical payoffs listed in Box 2.1, all-Dove would yield 50 fitness points per individual per interaction,

versus the 33 obtained with the ESS; but an all-Dove population can be invaded by Hawks, who will gain higher fitness than Doves when rare, and thus be favored by selection. As this example illustrates, there is no guarantee at all that the ESS will correspond to the strategy that maximizes average fitness in the population.

The implications of this disjunction between the evolutionary equilibrium and the possible individual fitness maximum are quite significant. ESS theory shows that selection may move individual behaviors in directions that produce collective consequences that are suboptimal for everyone. It suggests that the evolution of social phenomena is riddled with pervasive or even insoluble contradictions between the interests of each individual and their collective interests as realized through social interaction. The paradigm illustration of this point is the Prisoner’s Dilemma game (Box 2.2).

Although the prisoner’s dilemma structure is one that seems to characterize many collective-action problems, it is certainly not the only possible payoff structure (see 9.5.3 and 10.2.2). Some social interactions are characterized by a more mutualistic structure. But the general argument applies: Social interaction always holds at least some likelihood of strategic conflicts of interest. These require that we analyze the consequences of natural selection and rational choice according to the competitive optimum predicted from game theory.

In summary, ESS theory combines the methods of economic game theory with the explanatory logic of natural selection theory. In the place of the economist’s assumptions of rationality and self-interest, the evolutionary ecologist substitutes evolutionary stability and fitness (Maynard Smith 1982a :2). While not as rigorous as explicit genetic or cultural-transmission models, it provides a convenient and robust approximation when the details of inheritance are unknown and the “phenotypic gambit” (2.2.2) is adopted. ESS models (or their equivalent) are warranted whenever conflicts of interest or frequency dependence effects apply to the characteristics being investigated. In such cases, the optimum is a competitive (evolutionarily stable) rather than simple (average fitness-maximizing) one.

Box 2.2. *The Prisoner’s Dilemma*. In the two-person version of this game, the payoffs have the following pattern:

1. If both cooperate, each player gets c .
2. If both defect (fail to cooperate), each gets b .
3. If one defects and the other cooperates, the Defector gets d and the Cooperator (in this case, altruist) gets a .
4. The payoff values are ordered $a < b < c < d$.

Using the payoff matrix conventions introduced in Box 2.1, we can list Ego’s actions by row (on the left), Alter’s by column (on top), with the payoffs to Ego given in the cells of the matrix:

		Alter's actions:	
		Cooperate	Defect
Ego's actions:	Cooperate	c	a
	Defect	d	b

Since the payoffs will be the same regardless of which player is designated Ego, this is termed a symmetrical game.

Clearly, as long as $2c > (a + d)$, the strategy with the highest average payoff (for Ego and Alter combined) would be to cooperate. But the best strategy in terms of self-interest is to defect, regardless of whether the other actor cooperates. In game-theoretical terminology, we say that in the prisoner's dilemma (PD) game, defection is the "dominant strategy."

As with the Hawk-Dove game (Box 2.1), the ESS can be grasped more easily if we consider whether one strategy (cooperate or defect) can invade a population consisting of the other. The defect strategy can invade a population of Cooperators, because the payoff to defect is greater than that to cooperate both for Defector-Cooperator interactions ($d > a$) and for interactions between Cooperators ($d > c$). But the cooperate strategy cannot invade an all-defect population, because its payoff is lower both against defect ($a < d$) and as compared to the common defect-defect interaction ($a < b$). This is true even though a population of pure Cooperators is characterized by a higher average payoff than a population of all Defectors ($c > b$), and indeed mean fitness declines with each increase in the number of Defectors.

In sum, if there is no way to enforce cooperation, selection (or rational self-interest) will drive the noncooperative strategy to fixation, hardly the best result from anyone's (or everyone's) standpoint. There are circumstances in which a less dismal result could prevail, both in the context of rational-choice models (Taylor 1987) and via natural selection (Axelrod and Hamilton 1981); and of course cooperation may have a non-PD structure. Both points are discussed elsewhere in this volume, particularly in Chapters 9 and 10.

Sources: The PD has a long history in classical game theory (Rapoport 1974; Sugden 1986; Taylor 1987). The earliest application of PD logic to an evolutionary context is probably that of Trivers (1971), although he did not use explicit game theory techniques. Axelrod (1984; Axelrod and Hamilton 1981; Axelrod and Dion 1989) and Hirschleifer (1982; Hirschleifer and Martinez Coll 1988) make extensive use of PD structures for analyzing the evolution of cooperation. Maynard Smith (1982a:Chapter 13) reviews the theory, and Boyd (1988, n.d.) discusses the way in which the PD can be generalized to cover a variety of reciprocal interactions.

2.3. THEORY AND METHOD IN THE SOCIAL SCIENCES

Since theory and method in the social sciences are highly diverse, we must be selective in discussing their relationship to evolutionary ecology.

Decision theory (particularly rational-choice models) has the clearest relevance to and compatibility with evolutionary ecology; hence it is what we focus on here. We concern ourselves with two main issues: the significance of methodological individualism (2.3.1), and the relationship of intentional explanation (of which rational choice accounts are a subset) to explanation in general and selectionist explanation in particular (2.3.2). Richerson and Boyd (3.5.2) discuss some additional linkages between evolutionary ecology and the social sciences, revolving around the effects of cultural transmission.

2.3.1. Methodological Individualism and Social Analysis

Definition. Methodological individualism (MI) holds that the properties of groups (social institutions, populations, societies, economies, etc.) are a result of the actions of its individual members. By "properties of groups" we mean rules, practices, and the like. By "actions" we mean both intentional and unintentional behavior (see 2.3.2). MI stands in opposition to various forms of methodological collectivism, which hold that group properties cannot be reduced to those of its members and their interactions.

The main corollary of MI (and what most people have in mind when they invoke it) is that explanation of the properties of groups should in principle be derived "from the bottom up," in terms of the actions and intentions of individuals. Stated so simply, MI may strike many readers as unproblematic, or even uninteresting. Certainly most of us agree that societies or cultures or bureaucracies do not literally have minds that choose or hands that move. Nevertheless, conventional analyses often proceed as if they did. Hence, if followed consistently MI has consequences that challenge many explanatory practices in the social sciences—including ecological anthropology (Vayda 1986)—in a way that is similar to the subversive effect that individual-level selectionism and ESS theory have had on functionalist explanation in evolutionary biology (see 2.2.2).

The Role of MI in Social Analysis. The primary goal of MI is to provide "microfoundations" or "actor-based accounts" for social phenomena by analyzing the extent to which they are the aggregate outcomes of individual beliefs, preferences, and actions (Homans 1967; Elster 1982; Roemer 1982a). In this sense, MI takes the "black box" of social institutions and processes apart, in order to discover the individual-level mechanisms that provide its workings.

In practice, MI is usually linked to a number of corollary assumptions. One of these is that individual actions result primarily from rational-choice processes. Another is that individuals are basically selfish—guided in their choices by self-interest. Although we will have more to say about these

matters below (see 2.3.2), it is worth noting here that these assumptions are not actually entailed by MI: It is logically defensible to adhere to individualism in a methodological sense without adopting the rational-choice and self-interest corollaries, or any kind of substantive (ethical, political, etc.) individualism (Elster 1982:453). It does not violate MI as such to show that individuals make nonrational decisions, take the welfare of others into consideration, adopt beliefs and preferences from their cultural milieu, or let beliefs about supra-individual entities shape their actions. MI simply holds that in each case these processes originate in and are maintained by the actions of individuals.

In the social sciences MI is most firmly established in neoclassical economics, which postulates individual decision-makers as the fundamental units of action and causality, and views social phenomena like markets and prices as the aggregate outcome of individual choice. Even here, however, MI has proven impractical or inappropriate for some questions, hence the development of macroeconomics and the use of supra-individual actors such as firms, households, and interest groups as analytical shortcuts in economic analysis. Economics, of course, also adheres rather closely to the rational-choice and self-interest assumptions. But many schools of psychological analysis are equally based in MI, yet abjure rational choice as a corollary. Nor does MI necessarily entail orthodox political or economic views: There is a very active branch of Marxist political and economic analysis that is explicitly based in MI and rational-choice/game theory (Elster 1982, 1985; Roemer 1982a, 1982b, 1986).

Some schools of social science research implicitly or explicitly reject MI in favor of methodological collectivism. They postulate collective entities or goals—such as classes, development of the productive forces, cultural systems of meaning, social or ecological equilibrium, or population pressure—as the prime movers of history and social structure. In fact, underlying much of the theoretical discord in the social sciences is a contest among differing interpretations of the relative power and appropriate realm of application of individualist and collectivist approaches. Whatever their virtues may be, collectivist approaches are much less compatible with evolutionary ecology than those that rely on MI, for reasons we discuss next.

MI and Evolutionary Ecology. The levels-of-selection controversy (2.2.2) resulted in a general consensus among evolutionary biologists that adaptations will rarely be found that increase group persistence or well-being at the expense of its individual members. For instance, the notion that populations evolve so as to regulate their size or maximize ecosystem stability or efficiency seems implausible once one attempts to specify the individual-level mechanisms that would produce such outcomes and to explain why natural selection would favor them. Setting aside technical details, it is fair

to say that something rather close to MI is the working assumption of evolutionary ecology analyses of behavior. Indeed, it has been said that the ultimate goal of evolutionary ecology is to explain the structure and functioning of populations and ecosystems in terms of the properties and evolutionary history of their individual members (e.g., Orians 1973; Pulliam 1976).

The close parallels that often exist between evolutionary ecology and methodological-individualist social science are nicely illustrated by the issue of population regulation. The notion that “population pressure” (imbalance between resources and population size) is something human groups seek to avoid and that many social practices are designed to maintain population–resource balance by either restricting population growth or increasing resource productivity, has been very popular in the history of ecological anthropology and related fields (review in Bates and Lees 1979). Most of these theories assign collectivist notions like “societal homeostasis” the role of agency in sociocultural evolution, and pay little attention to the beliefs and actions of individuals. The MI critique of such views is quite similar to the selectionist one we have sketched above (2.2.2). Both focus on the “free-rider” problem of reproductive restraint, and both caution that group benefits may be only incidental outcomes of individual self-interest. They differ, of course, in the causal mechanisms invoked: Selectionist explanations emphasize population structure and group versus individual mortality and fertility rates, while MI social science points to rational self-interest (i.e., individual incentives to have offspring and disincentives to sacrifice for the collective good).

This example illustrates how, even though starting from rather different premises, both MI and evolutionary ecology make one skeptical of arguments purporting to explain individual characteristics in terms of their group-level benefits, unless it can be shown that special circumstances (centralized coercion, powerful group selection, or highly biased cultural transmission [see 3.4.4 and 3.5.2]) prevail. Similar arguments can be made with respect to conservation of game (Hames 1987b; see 7.5.1), ritual regulation of warfare (Peoples 1982; Vayda 1986; see 10.4.1), or many other group benefit arguments common in the social sciences (a number of which are discussed in other chapters of this volume). The general point, for either MI social scientist or evolutionary ecologist, is that explanation of social phenomena, including group-level benefits, should pay attention to individual-level mechanisms.

2.3.2 *Intentional Behavior and Rational Choice*

Despite the convergence of evolutionary and rational-choice accounts of collective action, there remain some salient differences. One of the most


obvious is that the primary causal force in evolutionary theory is natural selection, while in rational-choice theory it is individual decision-making. One possible way to bridge this gap is to nest one approach within the other. If the mechanisms of individual decision-making are themselves products of evolution, they should be explainable in selectionist terms: Rationality is an evolved adaptation. But the tidiness of this resolution is somewhat deceptive, for it glosses over a number of complex issues. To understand these, we need to place both evolutionary and rational-choice accounts within a more general philosophical framework.

Modes of Explanation. Following Elster (1983), we distinguish three modes of scientific explanation: causal, functional, and intentional. Causal explanation is the scientific "ideal," and the only accepted mode in the modern physical sciences. While the subtleties of this mode are too involved to discuss here, the basic logic is simple enough: Causal explanations specify the physical mechanisms involved when one event or process determinatively leads to another. Causal explanation is typically reductionist, in the sense that it proceeds by explaining the properties of objects or processes in terms of the interaction of component objects or processes—the properties of a gas by its pressure, temperature, and chemical composition, and the properties of an organism by those of the DNA molecules that it inherited and the environment within which it develops.

Functional explanation is the most controversial of the three modes under consideration. It can be outlined succinctly as follows (simplified from Elster 1983:57):

1. There exists one or more actors (A) who possess some characteristic (C).
2. C produces some beneficial consequence (B)—beneficial for A.
3. B is unintended, and perhaps unrecognized, by A.
4. B maintains C by a causal feedback loop passing through A.

To summarize, one might say that functionalism explains characteristics by virtue of their *unintended beneficial consequences*. Following Froemming (1986), we can represent the argument graphically thus:



Functional explanation is not employed in the modern physical sciences, where "beneficial consequence" has no acceptable meaning, but it plays distinctive roles in (evolutionary) biology and in the social sciences. In biology, B properly refers to fitness, with other benefits being relevant only

to the extent that they are good correlates of fitness. The causal feedback loop is simply natural selection: C increases in frequency if it yields a higher B than some alternative C possessed by other A in the same population. In contrast to causal explanation then, biological functionalism accounts for the properties (C) of organisms (A) by their fitness benefits (B). It explains the characteristics of organisms in terms of selective advantage: Phenotypes and their underlying DNA sequences are the way they are because they have evolutionarily outcompeted the alternatives that have been tried on the field of organic evolution. (The use of function here is distinct from its meaning in proximate analysis [1.2.1]. In the one case we are explaining the ultimate evolutionary origins of a phenotypic design, in the other its proximate operation.)

While functional explanation is distinct from causal explanation, it is not independent from it. Specifically, if a functional explanation in biology states that B *maintains* C, this makes sense only if C *causes* B (i.e., the characteristic yields a fitness benefit), and if in turn the causal mechanisms of heredity ensure that C recurs in at least some descendants of A. More generally, we can say that any functional explanation depends for its validation on an underlying causal explanation, and thus is less fundamental than the latter:

We can use functional explanation in biology because we have a causal theory—the theory of evolution by natural selection—that licenses us to explain organic phenomena through consequences that are beneficial in terms of reproductive capacity. Even if in a given case we are unable to tell the full causal story, we may be able to advance an explanation in terms of these beneficial consequences. (Elster 1983:21)

This is precisely the tactic of evolutionary ecology.

In the social sciences, functionalism takes various forms, but the structure sketched above is common to them all. In comparison to biological functionalism, however, in social science there is much greater ambiguity concerning the kinds of beneficial consequences involved, the identification of the beneficiaries, and most importantly the causal feedback loop of proposition (4). Elster (1983:56ff.) suggests that sociofunctionalism occurs in at least two varieties: Strong forms (a la Malinowski) propose that *all* social phenomena have beneficial consequences that explain them; weak forms (a la Merton) propose that *whenever* social phenomena produce unintended benefits, they can be explained by these benefits. Topically, sociofunctionalism covers a wide range. In addition to Malinowski and Merton, prominent examples include British structural functionalism (e.g., Radcliffe-Brown 1952), cybernetic human ecology (e.g., Rappaport 1968), cultural materialism (e.g., Harris 1979b), and historical materialism (e.g., Cohen 1978; cf. Elster 1982).

The attraction of functional explanation in the social sciences is that it relieves one of the admittedly implausible assertion that all social institutions or practices were created or persist because of consciously perceived beneficial consequences—that (for example) Yanomamo warfare exists because its practitioners recognized its efficacy in regulating population. But because functional explanation relies on unrecognized and unintended relationships, to be convincing it must provide a specified causal feedback mechanism:

A functional explanation can succeed only if there are reasons for believing in a feedback loop from the consequence to the phenomenon to be explained. In the case of functional explanation in biology . . . we have general knowledge—the theory of evolution by natural selection—that ensures the existence of some feedback mechanism, even though in a given case we may be unable to exhibit it. But there is no social-science analogue to the theory of evolution, and therefore social scientists are constrained to show, in each particular case, how the feedback operates. (Elster 1983:61)

Unfortunately, that is rarely done. Existing functional explanations in social science almost invariably fail to postulate any explicit feedback mechanisms, let alone to demonstrate their existence or causal efficacy. The absence of a causal foundation for sociofunctionalism has led to a proliferation of types of beneficial effects, units proposed as beneficiaries (ranging from individuals to ecosystems), and time scales over which the benefits are assessed. This ambiguity gives sociofunctionalism a tremendous degree of license in constructing explanations that are broadly plausible but lack rigorous theoretical foundation or empirical evaluation.

The third mode of explanation is *intentional*. It is virtually limited to the social sciences—we no longer argue that water seeks its own level or light finds the shortest path in any but a metaphorical manner. Some analyses of animal behavior make reference to intentionality (e.g., Griffin 1981, 1982), but ever since Darwin we do not explain most biological characteristics of living things via intentionality: Natural selection, unlike a deity (or an animal breeder), has no intentions.

Intentional explanation accounts for actions in terms of the beliefs and preferences of the actor; both are necessary constituents of intentional explanation (Elster 1983:Chapter 3). We might know that someone prefers intellectual stimulation to high salary, but in order for that to explain why she or he chose a career in anthropology over one in law we must invoke her or his beliefs about a number of things: the intellectual content of each profession, the likelihood of succeeding in either, and so on.

But if beliefs and preferences provide the foundations of intentional explanation, they also define its limits. In order to explain *them*, we must eventually invoke some other explanatory mode—causal (as in the neuro-

physiology or ontogeny of learning) or functional (as in the biological evolution of our cognitive machinery or the cultural evolution of our values). While it is possible to derive some beliefs and preferences from “higher-order” beliefs and preferences, this procedure (like the joke about what the turtle holding up the world is standing on, “It’s turtles all the way down!”) cannot go on indefinitely. As with functional explanation, then, intentional explanation is inherently incomplete.

Rational Choice. The paradigm form of intentional explanation is rational-choice theory. As used here, the term *rational* refers not to ends, but only to the relation between ends and means. Rational actors are those who pursue their ends (whatever they might be) as effectively as possible, and who do not commit logical errors in ordering their preferences. To explain the second point briefly, a rational actor has what is called “transitive preference rankings”: if she or he prefers A to B, and B to C, then she or he will prefer A to C.

This restricted meaning of *rational* is sometimes referred to as “thin rationality,” to emphasize that it does not make any substantive predictions about what it is actors actually value or prefer. In this view, it is not the purview of economics (for example) to explain why some people prefer to maximize pecuniary benefits and others prestige or patriotism or piousness. In principle, then, rational-choice theorists are no more committed to “materialistic” measures of value such as wealth or hedonistic satisfaction than they are to preference rankings that put spiritual enlightenment or service to humanity above all else. But in developing any particular model or hypothesis, one must move beyond thin rationality and posit something substantive about preferences. Preference rankings can be derived inductively (by observing what actors choose, or asking them what they would choose, under specified conditions) or deductively (from some additional assumptions about the cultural or biological determinants of preferences).

Rational choice obviously implies that actors do indeed have a choice. More specifically, rational-choice theory divides the factors determining outcomes into two categories: choices and constraints. Choices refer to the elements over which the actor can exercise intentional control; constraints are everything else (see 2.4.4). Where preferences are known or presumed, rational-choice theory is tested by observing situations where constraints change in a certain manner, then comparing changes in actual choices to those predicted from the theory (given the specified shift in constraints and the original preferences).

In order for choice—rational or otherwise—to operate, an actor must have beliefs as well as preferences; this raises what economists call the problem of information. Rational-choice theory assumes that actors possess (or can gain) “sufficient” information to make an informed choice. The

simplest models assume that actors possess "complete information" (in the sense of knowing at least the probability distribution of outcomes—"the odds") if not "perfect information" (knowing the outcomes with certainty). More complicated models relax this assumption, and allow actors to gain information (at some cost) as they make choices or otherwise sample reality. This allows predictions to be made about the optimal degree of ignorance, based on the marginal costs and benefits of information. While the postulate of complete information may often be unrealistic, it is a useful simplification, in that it allows one to build relatively simple, general, and testable models. In the case of creatures that have elaborate means of obtaining, storing, and transferring information (i.e., humans) it is not such an unreasonable assumption if the relevant sector of the environment does not change too rapidly (see 2.4.4 for further discussion).

In addition to assuming that actors know what they want and how to get it, most rational-choice theorists assume that actors are utility maximizers. In essence, utility is economists' way of talking about the *relative* amount of *satisfaction* derived from *consumption* of any *good*. The four italicized terms in this definition need to be explained.

Like preferences, satisfaction is individually defined, and rational actors are assumed to be able to weigh the satisfaction derived from alternative uses of their time and other resources. Although the classical economists (Adam Smith through Ricardo) thought that utility was something that could be measured on an absolute scale ("cardinal utilities"), neoclassical economics abandoned this assumption in favor of relative ("ordinal") utility. The most behavioristic of contemporary economists simply speak of "revealed preference," that is, the relative preferences for alternative goods implied by the actor's actual choices. "Consumption" and "good" are used in a very broad sense: One "consumes" food, but also sleep, theories, and mating opportunities. All of these (and anything that yields satisfaction, including pursuit of ideals) are "goods."

Rational-choice theory has been repeatedly criticized for taking as given that which needs most to be explained: the preferences and beliefs of individuals, and the social milieu, which determines to a considerable extent the rewards to different choices. In defense of the rational-choice framework, two things bear mentioning. First, it *can* explain in a limited but important sense, by holding preferences and beliefs constant and then allowing one or two other factors to vary. For example, if we know the relationship between wealth and consumption, we can use microeconomic models to predict the effect of changes in prices or income on consumption levels. This can even be done on a historical time scale, as in the attempts of economic historians to explain long-term socioeconomic changes on the basis of changes in technology, population, and market opportunities (e.g., McCloskey 1975b; North 1981). Following the principle of methodological

individualism, social institutions and milieus are analyzed as products of past (and continuing) individual actions, not of some independent force.

Second, some preferences and beliefs can be derived from other, more fundamental ones. For example, preferences for achieving or maintaining social status might lead us to alter our preferences for particular goods or activities as these become more or less fashionable. And of course it may be in the interests of those with exceptional influence or power to manipulate the beliefs and preferences adopted by the less influential or powerful—both in the domestic (familial) and extradomestic (political) domains.

But even accepting these two points, it remains true that rational-choice theory fails to explain much about variation in preferences and beliefs. It is possible that evolutionary ecology can fill at least part of this gap. Preliminary evaluation of this possibility is one important goal of this volume.

Evolutionary Ecology and Intentionality. If evolutionary ecology is to contribute to social science, it will have to come to terms with the role of intentionality. Although a measured discussion of this issue is complex, and made more difficult by the long and acrimonious history of debate on Darwinian analyses of human behavior (1.3.1), it is useful to consider four distinct viewpoints:

1. Intentionality is autonomous and supersedes organic evolution.
2. Intentionality has no explanatory role to play in science.
3. Intentionality is a set of genetically evolved proximate mechanisms.
4. Intentionality is shaped by both cultural and genetic evolution.

(In all cases, by *intentionality* we mean conscious elements of decision-making—beliefs and preferences—as these are used in intentional explanations.) Let us briefly consider each of these positions.

Position (1) is widespread in the social science literature, but takes various forms. Some see the autonomous status of human intentionality in a historical dimension (e.g., Bock 1980; Slobodkin 1978), an intellectual tradition that can be traced back to Marx. Others emphasize the social matrix in which intentions are formed and enacted, and see this social process as contradicting or eluding Darwinian reductionism (e.g., Sahlins 1976; Ingold 1986). Still others do not appear to deny Darwinism a limited explanatory role, but see the causal distance between genes and human action as a major impediment to evolutionary analysis (e.g., Reynolds 1976; Kitcher 1985).

At the other extreme is position (2), which denies explanatory validity to concepts such as intentionality, decision-making, choice, and rationality (e.g., Dunnell 1980; Rindos 1985). In this view, evolutionary change takes place independently of the vagaries of human intentions; to grant intentions a causal role is to succumb to teleological thinking. Thus, whatever force

intentionality might have in a psychological sense, to proponents of (2) it cannot explain the trajectory of social evolution any more than mutation can explain the outcomes of organic evolution. Indeed, Rindos (1989) explicitly compares intentional phenomena to the random-variation-generating mechanisms organic evolution (mutation, recombination, etc.). This view has historical precedent in the superorganicism of Kroeber (1917) and the cultural-evolutionary philosophy of Leslie White (1949), both of whom ardently denied that the direction of cultural evolution could be explained by human choices.

The simplest way in which evolutionary ecology could give explanatory depth to rational choice theory would be to derive fundamental beliefs and preferences from neo-Darwinian principles. This is position (3), the program of human sociobiology (e.g., Alexander 1979; Chagnon and Irons 1979; Harpending et al. 1987). A common way of expressing this position is to refer to intentionality—or more precisely, the cognitive mechanisms underlying intentional phenomena—as proximate mechanisms (1.2.1). Like other proximate factors, these mechanisms are evolved characters; they can be analyzed in and of themselves (like a physiologist studies the digestive system, for example) or they can be analyzed as adaptations that have been shaped by natural selection to serve certain (fitness-enhancing) purposes.

This sociobiological view of human action is particularly instructive with regard to explaining preferences. While economics or psychology can derive second-order preferences from more basic ones (as noted above), there comes a point at which fundamental preferences or goals (conscious or unconscious) must be taken as given in order to anchor the entire analysis. Intentionality, like learning, is a *derived* force (Boyd and Richerson 1985; Campbell 1965; Rosenberg 1980). In principle, evolutionary biology offers a way to identify fundamental preferences that are likely to be favored by natural selection. A completely convincing account would require demonstration of the proximate cognitive mechanisms and their specific genetic bases (Kitcher 1985; Tooby and Cosmides 1989; Symons 1989). But many feel that a plausible initial case can be based on the fit between adaptive predictions and the empirical facts of human action (Alexander 1990; Irons 1990; Smith 1987c; Turke 1990).

The sociobiological view of human decision-making offers a special refinement of the concept of self-interest. As we noted above, the “thin” version of rational-choice theory does not specify the content of self-interest. It proposes simply that actors are goal-seeking, whatever their individual goals may be. In practice, social scientists invoke various kinds of self-interest (well-being, wealth, power, status, etc.). But they rarely agree on the exact list involved or the rank order of its elements, perhaps because the goals are not derived from any underlying theoretical principles. In contrast, sociobiology specifies a single underlying maximand: inclusive fitness. As

described earlier (2.2.2), inclusive-fitness interests sometimes are “selfish” in the colloquial sense, but can also be self-sacrificing (to offspring and other close kin, and in the short run even to unrelated reciprocators) in specified evolutionary contexts (9.4, 10.4). Evolutionary theory can even be used to predict when it would be fitness enhancing to violate the canons of rationality and be swayed by benevolent or malevolent emotions (Hirshleifer 1987; Frank 1988). It offers explicit predictions about the exact form and content of “self-interest,” and of evolved preferences and inclinations in general, in different environmental circumstances. In this it goes beyond the thin rationality of neoclassical economics or decision theory.

We have already argued that this viewpoint does not necessarily entail the narrow sort of genetic determinism that most social scientists rightly reject (1.3). It does, however, remain problematic, since neither humans nor other animals actually seem to hold fitness maximization as a goal. Rather, they possess various psychological mechanisms that lead them to learn some things easily, other things with difficulty or not at all, and to invest outcomes with specific positive or negative valences. In other words, actors may be utility maximizers, but they are fitness maximizers only to the extent that the utility functions defined by their evolved cognitive and emotional machinery are correlated with fitness. Given the logic of Darwinian evolution, high correlation of fitness and utility is the most reasonable expectation, but even ardent sociobiologists recognize that rapid environmental change and perhaps other factors can reduce or even destroy this correlation in specific cases (e.g., Symons 1990; Tooby and Cosmides 1990).

The final viewpoint on intentionality, listed above as position (4), holds that human preferences are best explained as the *joint* product of genetic and cultural evolution. In contrast to (3), this position holds sociobiological theory to be incomplete, and cultural inheritance to be an independent determinant of human action. In contrast to (1), and to the received view in contemporary social science, it adopts an explicitly Darwinian approach to understanding cultural variation, including the preferences and beliefs that inform intentional action (for further discussion, see Chapter 3).

In summary, both functional and intentional explanation of human behavior either require evolutionary analysis or at minimum are complementary with it. In the case of functional explanation, natural selection is currently the only process capable of providing the causal feedback element necessary for a convincing analysis. In the case of intentional explanation, natural selection is needed to underwrite a theory of preference formation. In neither case is there presently a well-established or compelling alternative to selection theory. However the problems of reconciling evolutionary biology with intentionality and culture are eventually resolved, the evidence summarized in this volume suggests strongly that some significant portion of the preferences and beliefs exhibited by human beings in diverse times and

places have been shaped directly or indirectly by natural selection. While it is certainly too early to draw firm conclusions about the explanatory limits or potential of evolutionary ecology in the human sphere, we are confident that it has an important role to play in explaining human social and behavioral patterns.

2.4. OPTIMIZATION ANALYSIS

Optimization theory is one of the most frequently used analytical tools in evolutionary ecology. It is also one of the more controversial. In this section, we offer a brief summary of the explanatory logic behind the optimization approach, its relation to natural selection and to human decision-making, and the general elements common to any optimization model. Although optimization underlies ESS analyses of selection in a strategic context (2.2.3), the term itself is more closely associated with analyses in a parametric context (as defined in 1.2.1).

2.4.1. *The Logic of Optimization Analysis*

Elements. While optimization analysis is used in many diverse settings, from decision theory and economics to engineering, all optimality models share certain basic elements: (1) an actor that chooses or exhibits alternative strategies or states; (2) a strategy set defining the range of options available to the actor; (3) a currency in which the costs and benefits of alternatives are measured; and (4) a set of constraints that determine the feasible strategies and the payoffs associated with each.

In keeping with methodological individualism (2.3.1) and the presumed prevalence of individual-level selection (2.2.2), in evolutionary ecology actors are usually defined as individual organisms. Each of the other three elements is discussed in some detail in sections 2.4.2–2.4.4, following consideration of some general issues.

Epistemology. What is the relationship between optimization analysis and the logic of explanation? Optimization is not some general principle of nature, nor are optimality models intended as realistic descriptions of the behavior of individual actors or the process of adaptation by natural selection. Instead, research using optimality models tests particular hypotheses, each of which shares the elements and structure common to optimization analysis.

It is certainly possible to analyze biological diversity without employing optimization analysis. Consider the example of hunter–gatherer prey

choice. We could describe each instance of prey choice in terms particular to the specific group being studied. We might note that the prey chosen varied systematically, and perhaps isolate a statistical relationship between prey abundances and the size or caloric value of chosen prey. But these findings would be specific to that case and lack any broader relevance. We would have an observation or correlation in search of an explanation.

In contrast, by using optimization analysis we can apply a general methodological framework to any particular case, and thereby test and refine models that have widespread applicability (cf. 1.2.1). For example, the prey choice model (6.2.3) can be applied whenever some basic assumptions about prey distributions and currency are met. It yields a set of simple yet powerful predictions about prey choice that might apply equally to seed-gathering Paiutes and seal-hunting Inuit. Although the assumptions and predictions of this optimization model might not precisely fit any particular case, it appears to capture the basic elements of diverse cases well enough to have received substantial empirical support in studies of a considerable number of species and situations (see section 6.2.5 for data on humans; Stephens and Krebs 1986 for nonhumans). Furthermore, it can be modified by altering the currency or constraint assumptions to make it more realistic for any particular case under investigation (as discussed in 6.4 and 6.5).

The value of the fine-grained prey choice model is not that it provides realistic descriptions of the cognitive processes or behavioral tactics underlying prey choice in bees, birds, and humans; these doubtless are quite diverse. Rather, it is due to the simplicity and generality of its formulation, its resulting testability, and its potential for analytical manipulation. If empirically successful, this implies that the model has correctly identified the adaptive goals involved in foraging behavior in diverse species.

In sum, optimization is not a theory (in the usual sense of an explanatory framework consisting of substantive propositions about the real world), but a method. Put succinctly, the epistemological role of optimization analysis is to provide a systematic means of generating hypotheses about the structure and function of living things. In other words, “the role of optimization theories in biology is not to demonstrate that organisms optimize [but] . . . to understand the diversity of life” (Maynard Smith 1978:52). Thus, for evolutionary ecology the optimization approach serves as a bridge between the abstract principles of natural selection theory and the diverse empirical facts of any real-life case. But the legitimacy of recasting evolutionary processes and outcomes in terms of optimization analysis is itself subject to debate, an issue we take up next.

Optimization and Evolution. Simplifying somewhat, three views on the link between optimization analysis and evolutionary theory can be found in the literature: (1) evolution via natural selection is an optimizing process; (2) optimization analysis has little or no valid relation to evolutionary theory; (3)

optimization analysis is a convenient heuristic tool or simplification for analyzing evolutionary outcomes.

The first view holds that optimality models merely formalize the process underlying Darwinian evolution (e.g., Cody 1974). This view was propounded during the early period of ecological optimization research, but has since been qualified in significant ways. First, natural selection is not the only evolutionary force, so even if selection is an optimizing process, other forces (drift, pleiotropy, etc.) and constraints (developmental, genetic, etc.) may produce nonoptimal outcomes. Second, the identification of selection with optimization is inexact. Selection favors existing variants with higher fitness, not necessarily the best possible or imaginable variant; it trades in relative rather than maximum advantage.

The second view holds that the criticisms just outlined identify crippling or even fatal weaknesses in the analytical partnership of evolution and optimization (e.g., Sahlins 1976; Gould and Lewontin 1979; Lewontin 1987). One version of this critique argues that evolutionary optimization analysis proceeds from a "Panglossian" assumption that all is for the best in this best of all possible worlds. This assumption is then easily demolished by showing that selection may instead lead to inefficiency, waste, conflict, and even extinction. Other critics have pointed out that individual organisms face a variety of adaptive problems, which involve trade-offs or compromise. For example, the most efficient foraging strategy might constrain the organism to suboptimal predator avoidance.

The solution offered by natural-selection theory to this latter criticism is in principle quite straightforward. Since other traits possessed by an organism are part of the environment of any one trait, the optimal strategy for any given trait is defined as that member of the feasible set that contributes the greatest fitness to the organism, *given the other traits that the organism possesses and given that each of these is optimal in the same context*. But in practice, this principle brings an intractable complexity to the analysis of any particular adaptive problem, for to understand one problem/trait/strategy properly, one would have to understand all of them.

Although important work is being done in analyzing optimal trade-offs between two major strategy sets (e.g., predation and foraging), evolutionary ecologists have not shown much interest in tackling whole-organism adaptation with formal optimality models. Indeed, it is not clear that this is even possible within existing theoretical and practical constraints. Instead, analyses have typically focused on a single optimization problem at a time, as if it were isolated from other (potentially interacting) problems. Some critics charge that this amounts to an illegitimate "atomizing" of organisms into component traits, each removed from the integrated matrix of the organism and its historical changes (e.g., Lewontin 1979; Gould and Lewontin 1979).

Whenever these single-trait analyses fail and the analyst cites competing adaptive goals as a possible explanation, the critics see this as a cover-up aimed at obscuring failure of the enterprise by post hoc rationalization.

This is a serious criticism, which has engendered an extended discussion (e.g., Maynard Smith 1978; Dawkins 1982; Oster and Wilson 1984; Krebs and McCleery 1984; Kitcher 1985:Chapter 7; Williams 1985; Dupré 1987; and articles by Gould and Lewontin, cited above). Perhaps the best that we can add here is that it remains a rather hypothetical problem. We rarely know in fact what degree of contradiction and compromise exists among the various traits of an organism; the most direct way to find out is to take the atomistic or "piecemeal" approach and test models of trait optimization. Given the numerous reasons why particular optimality hypotheses might fail, simple models have been empirically successful to a surprising degree.

We believe the third view listed above underlies the great majority of optimization analyses in evolutionary ecology. Optimality models are useful tools for ecological analysis even though they are not "true" in any simple sense (1.2.2). It is especially important not to confuse the process of selection with its results. In the narrowest sense, the process is one of relative advantage only. But selection is persistent and cumulative. Given sufficient genetic variation and consistency of selection pressures, it is plausible that one of its cumulative results will be a trajectory of improvement in designs. The result may be a design that can fairly be characterized as optimal with respect to the fitness currency, the design problem, and the relevant constraints.

A skeptic might still object that environmental change is ubiquitous, hence consistency of selective pressures unlikely. But this is a matter for empirical resolution (Jochim 1983). A more interesting reply is that selection can favor, and demonstrably has favored, the evolution of capacities for phenotypic adjustment to rapidly shifting environmental conditions. Such abilities are central to behavioral aspects of phenotype, including human learning abilities, rationality, and cultural transmission—all arguably products of selection for optimal (fitness-maximizing) phenotypic design (2.3.2, 3.3).

When evolutionary ecologists invoke optimization arguments then, they are not granting omniscient directionality to natural selection. Rather, they are postulating the evolution of phenotypes whose attributes are "optimal" in the specific sense expressed in the hypothesis under test (e.g., a foraging strategy that maximizes energy efficiency, presuming that the latter is a robust correlate of fitness). Viewed this way, optimization models are simply shortcuts to understanding the outcome of evolutionary history, "a tactical tool for making educated guesses about evolutionary trends" (Oster and Wilson 1984:284).

2.4.2. The Strategy Set

A strategy set (sometimes termed "phenotype set" or "feasible set") consists of the options (choices, decisions, alternative states) available to the actor. The content of the strategy set may be limited to a very few alternatives (e.g., Hawk vs. Dove, see Box 2.1), or it may be more complex (e.g., the various combinations of prey types a forager might harvest, which for just 5 possible prey types amounts to 31 alternative combinations). Specification of the strategy set is obviously a critical step in optimization analysis.

In some cases, optimization theory may provide considerable guidance in defining the strategy set. For example, in the prey choice problem, the fine-grained prey choice model (see 6.2.3) tells us to rank the prey types by return rates, reducing the strategy set to just 5 alternatives (take only the highest-ranked type; take the 2 highest ranked; etc.). But there are many cases in which the specification of the strategy set is an empirical or inductive problem, for which theory (including optimization theory) offers little assistance.

Satisficing versus Maximizing. If the analyst faces difficulties in defining and characterizing the strategy set, so might the actor. Does the actor know all of the possible solutions to a given decision problem and their payoffs? Can she or her calculate the optimal one? Mindful of these questions, some propose an alternative to optimization analysis, termed *satisficing*. Rather than expect the actor to be an optimizer (an impossible task, in this view), we might expect him or her to choose any alternative that is "good enough" (Simon 1955; Winter 1964; Elster 1983:74ff.).

Satisficing has two major analytical limitations, however. First, there is no general criterion by which the analyst or the actor can decide what is "good enough." Thus, satisficing remains an ad hoc method; the goal (satisficing) must be determined empirically by observing an actor assumed to have that goal. If satisficing is then used to explain the actor's behavior, the circularity is evident. Second, satisficing has no clear meaning in an evolutionary context (Krebs and McCleery 1984:119ff.; Foley 1985:224). A "satisfaction threshold" is irrelevant to selection, which in parametric contexts favors the best (fitness-maximizing) alternative among those available in a population. And in strategic contexts, satisficing can never be an ESS; because a strategy that is closer to the optimum will invade and outcompete the satisficing strategy. There is no "good enough" in the eyes of selection, only "better than."

Satisficing is sometimes advocated by analysts who wish to acknowledge the actor's incomplete information or cognitive limitations. But in these cases, satisficing is best replaced by an optimizing approach with appropriate information constraints (see below). Proponents of satisficing might reply

that without knowing the costs of acquiring information, the optimal amount of information remains unknowable, causing an "infinite regress" problem in determining how much information to gather to determine the value of information (Winter 1964; Elster 1983:139ff.). This is a serious problem for single actors using rational choice in a novel situation. It may be less of a problem for actors who pool their experience (through observation and cultural transmission). And it is certainly less problematic if the optimizing force is natural selection rather than choice, for then selection "does the work" of determining which of the existing heritable strategies in a population actually comes closest to the optimum.

2.4.3. Currency

The currency in an optimization model calibrates the costs and benefits of alternatives in the strategy set, in order to rank preferences and to determine the optimal solution. Once a currency is chosen, the analyst can manipulate the optimization model to produce a set of predictions. For example, the currency in the fine-grained prey choice model (see 6.2.3) is some measure of food value (typically calories) obtained per unit foraging time. The different costs (in foraging time) and benefits (in food value) obtained from various prey choice combinations define the net benefit.

Ideally a currency would be characterized by maximal generality, precision, and realism. But this combination is rarely possible (see 1.2.2). The most general currencies commonly used in optimization models are fitness and utility. But these present problems, particularly in obtaining precise measurements of choice outcomes.

Fitness offers the strongest deductive basis for ranking different outcomes in terms of selective value. But because it is a lifetime measure summing the effects of many different phenotypic characters, it is generally impractical as an empirical currency. In addition, it can be argued that fitness is an inappropriate empirical measure for behavioral choices. As noted above (2.3.2), organisms do not rank alternatives by fitness consequences per se, but rather by various cognitively defined costs and benefits. Furthermore, fitness is a probabilistic measure (2.2.1), and the issue of evolutionary interest is really *design* for fitness rather than current fitness outcomes themselves (Williams 1966a:102). To study the costs and benefits of short-term behavioral choices in an evolutionary framework then, we need a proxy for fitness, a measure that we have reason to believe is correlated with fitness (or has been so in the evolutionary past) but that is more suitable for empirical research.

Utility has proven to be a useful theoretical construct in the development of economics and decision theory. Actors exhibit preferences that reflect psychological valuations of the consequences of alternative choices. Since

the psychological mechanisms underlying these evaluations are at least in part genetically evolved, it is a reasonable presumption that utility is a fairly robust correlate of fitness. Of course, this is less likely when conditions are evolutionarily novel (2.3.2), or when the relevant preferences are strongly shaped by certain forms of cultural evolution (3.4). But if utility is a good proxy for fitness on theoretical grounds, it is not much help on methodological ones. Utility (like revealed preferences or satisficing) cannot be established independently of the actual choices made by an actor. It is a descriptive term for inferred preferences. To say that rational actors maximize utility is thus true by definition (Krebs and McCleery 1984:94).

In sum, to advance the goal of understanding organisms as products of evolutionary design we need currencies that are less general and more operationally useful than fitness or utility. Depending on the problem under analysis, the currency of choice might be reproductive success, mating frequency, survival frequency, resource harvest rate, and so on (as detailed in subsequent chapters of this volume). No single currency is best for all analyses, but some have proved useful for large domains of evolutionary ecology.

2.4.4. Constraints

The elements of any optimization model can be divided into two categories: variables that are subject to choice on the part of the actor, and those that are not. The latter are termed constraints. Stephens and Krebs (1986:9) suggest the following classification of constraints:

1. Extrinsic
2. Intrinsic
 - a. of abilities
 - b. of requirements

Extrinsic constraints are exogenous to the actor, features of the social or natural environment that are beyond the actor's control—at least under the *ceteris paribus* ("all else being equal") assumptions of a particular optimization model. For example, in the Hawk–Dove model (Box 2.1), the value of the contested resource is an extrinsic constraint.

Intrinsic constraints are those endogenous to the actor's phenotype. They include abilities (behavioral, cognitive, and the like) and requirements (physiological, nutritional, etc.). The simplest and most general models assume few intrinsic constraints. For example, in the Hawk–Dove game, the analysis is considerably simplified if we can assume that all players employing the same tactic have equal ability, so that any Hawk has an equal chance of winning against another Hawk, or any Dove against another Dove. This simplification obviously exacts some price in terms of reduced realism or empirical accuracy.

The specification of constraints involves a balance among the competing goals of realism and generality, a compromise between our desire to be faithful to the facts and the demands of analytical tractability (the ease of mathematically manipulating and comprehending the model). In general, an increase in realism is obtained by adding constraints. For example, in the Hawk–Dove game we might include intrinsic constraints like differences in fighting ability (e.g., Maynard Smith and Parker 1976), or such extrinsic constraints as nondivisibility of the resource so that Doves cannot share (e.g., Maynard Smith 1982a:Chapter 3). But increased realism often comes at the price of diminished generality, testability, and analytical comprehension (see 1.2.2). The most efficacious balance must be judged relative to specific research goals and current theoretical and empirical understanding.

Cognitive Constraints. One class of intrinsic constraints of particular interest to social scientists is that imposed by limitations in the cognitive mechanisms and information that actors possess. There are two distinct issues here: the relationship between cognition and selection, and the way in which limited information affects decision-making.

Many optimization models assume the actors possess complete information in the relevant domain. If one seeks to understand what information an actor actually possesses, how it is acquired and updated, and its role in decision processes, then the complete-information assumption is clearly inappropriate. Instead, one needs to build and test models that incorporate information constraints. Bayesian decision theory, models of cultural transmission (3.3.1), and sampling models are appropriate avenues for this research. Some of these analyses treat the acquisition of information as an optimization problem in which there is a trade-off between the value of increased knowledge and the cost of acquiring it (Stephens and Krebs 1986:Chapter 4).

To return to the prey choice problem, it might be that prey that superficially resemble each other (e.g., belong to a single species) in fact differ significantly in their expected return rates. In some cases, gathering the information needed to discriminate prey types might have low cost (e.g., stranded vs. free-swimming whales). In others, the cost of obtaining the requisite information might outweigh any possible benefits (e.g., if one needs actually to capture the prey item before assigning it to one or another type).

Actors are also constrained by their cognitive and cultural capabilities for processing information. Formal optimization models define solutions using geometric tangents, partial derivatives, algebraic inequalities, or the like. These mathematical techniques provide convenient and elegant means of arriving at general and precise solutions, but they certainly do not replicate the everyday decision processes of actors. Instead, real actors (including people) are likely to use rules of thumb, less robust but far simpler ways of

comparing outcomes and finding the (approximately) best solution. For example, instead of ranking prey types by net caloric return per unit handling time (6.2.3), foragers might use cruder but fairly effective rules involving the size and fat content of available prey (Jochim 1983).

While the relationship between cognitive constraints and evolutionary forces is complex, two general points can be made here. First, widespread rules of thumb are probably products of cultural or genetic evolution that have had robust selective advantages. If so, their primary adaptive value may be that they offer low-cost solutions to problems that are difficult or even insoluble when tackled by rational choice or individual trial-and-error learning. Actors may not replicate the procedures of an optimization model, but simple rules of thumb or cognitive algorithms provided by natural or cultural selection may allow them to approach the solution quite closely under conditions approximating the environments in which these "short-cuts" evolved.

Second, human actors in particular may increase their probability of attaining the optimum in complex decision problems by drawing on the accumulated experiences of others, as transmitted through conversation, lore, and culturally acquired beliefs. Indeed, this may underlie the evolution of the elaborate system of cultural transmission in the hominid lineage (3.3).

Individually Variable Constraints. We usually think of optimization models as prescribing a single best solution to a given problem. From this perspective, it is reasonable to average together the actions of a set of individuals, and to view individual variation in choices as random "noise" due to errors in measurement, decision-making, or both. But the process of selection and the logic of methodological individualism both suggest otherwise. Since individuals (and classes of individuals) can be expected to differ in their intrinsic constraints (due to their age, experience, etc.) we should not expect them to share the same optimum for a given decision problem, even in the same environmental situation.

There are at least three distinct explanations for individual variation, besides the observational and decision errors mentioned above. First, individuals exhibiting what appear to be suboptimal responses may be making "the best of a bad job" (Dawkins 1980). That is, such individuals may temporarily or permanently lack the abilities or resources to achieve the outcomes with highest fitness payoffs. That would be the case, for example, with a forager who could not efficiently pursue a normally high-ranked prey type because of some physical handicap, and therefore wisely chose to ignore it. To average his or her response in with those of nonhandicapped foragers, or treat it as nonoptimal, would be to ignore the relevant difference in constraints.

Second, individuals (or classes of individuals) may differ in their constraints, and hence in their optima, in ways that cannot be ranked as "better

versus worse" or "more versus less constraints." A good example of this is constraint difference between males and females, which can affect optima for such decision categories as foraging (Chapter 6), time allocation (Chapter 7), or use of space (Chapter 8), as well as the more obvious domains of mating and parenting strategies (Chapter 11).

Finally, even individuals who do not differ in their constraints may display different but equally optimal phenotypes. This will be the case whenever the evolutionary equilibrium includes two or more tactics, and tactic differences are either a facultative response of individuals (a mixed strategy) or heritable (i.e. polymorphism; see 2.2.3). In such a situation, it makes no sense to say that one phenotype is superior (optimal) compared to another. Because "superiority" is frequency dependent, at equilibrium the payoffs to each will be equal. Prominent examples include "ideal-free" habitat distributions (8.3.2), evolutionarily stable group sizes (10.2.1), and female choice of polygyny versus monogamy as a function of variable male resources (11.3.2). To characterize one strategy (living in the "richest" habitat, foraging alone, being monogamous, etc.) as the optimal one is to overlook the differences in constraints created by the actions and characteristics of other actors.

2.5. CHAPTER SUMMARY

This chapter has surveyed the central ideas of evolutionary biology and social theory relevant to human behavioral ecology. The main points are as follows:

1. Natural selection requires three conditions: individual variation, heritability, and differential fitness.
2. Natural selection favors variants with highest *fitness* or *evolutionary stability*. Fitness refers to the statistical propensity of a variant to leave descendants, rather than to actual numbers of descendants of any given individual.
3. When fitness is frequency dependent, as is true in many social interactions, natural selection favors an *evolutionarily stable strategy* (ESS). An ESS is the variant (or combination of variants) that, when characteristic of most members of the population, cannot be replaced by another variant through natural selection. There is no reason to expect that an ESS will maximize average population fitness.
4. Evolutionary biologists distinguish between individual selection, group selection, and kin selection. Individual selection favors variants with maximal surviving offspring; kin selection favors variants with maximal surviving genetic relatives (including nondescendant kin); group selection favors variants producing maximal survival or colonizing success of groups

(trait groups, demes, etc.). The balance among these forms of selection is currently debated, but most evolutionary ecology models emphasize individual selection.

5. Because the exact mode of inheritance of behavioral phenotypes is usually unknown, evolutionary ecologists often adopt the "phenotypic gambit." That is, they assume that inheritance is simple, phenotypes are highly flexible, and proxy currencies are robust fitness correlates.

6. Methodological individualism analyzes social phenomena in terms of the actions and interests of individuals; it is adopted in several areas of social science, and is generally compatible with the analytical approach of evolutionary ecology.

7. Explanations can be classed as causal, functional, or intentional. Evolutionary explanations are often functional in form, but full explanation requires that they incorporate the causal process of natural selection. Intentional explanation has a complex relation to evolutionary analysis. Some see intentions as overriding evolutionary analysis, others claim they have no explanatory role in science, and still others incorporate them as a genetically or culturally evolved set of mechanisms with fitness consequences.

8. Optimality models are very commonly used in evolutionary ecology. Their elements include a decision-maker (actor), a set of alternatives (strategy set), a currency (specifying the variable maximized), and a set of constraints (variables outside the actor's control). Optimization models can be powerful tools for analyzing adaptations, even though they are not fully realistic descriptions of the evolutionary process.

ACKNOWLEDGMENTS

We thank the following individuals for providing detailed comments on earlier versions of this chapter: Robert Bettinger, Rob Boyd, Monique Borgerhoff Mulder, Raymond Hames, Kristen Hawkes, Sarah Blaffer Hrdy, and Michael Taylor.

Cultural Inheritance and Evolutionary Ecology

Peter J. Richerson and Robert Boyd

3

3.1. INTRODUCTION

Humans are uniquely reliant on culture as a means of adaptation. To be sure, some human variation results from genetic adaptation. Like most other animals, we adapt using individual learning. Unlike other creatures, however, humans acquire a great deal of adaptive information from other conspecifics by imitation, teaching, and other forms of "cultural transmission."

It is important to distinguish between culture and individual learning. Culture is often lumped with ordinary individual learning and other environmental effects under the heading of "nurture," to be contrasted with genes—"nature." This way of thinking is responsible for much confused thinking about the evolution of human behavior. Culture differs from individual learning because variations are acquired from other individuals. For the most part, humans do not learn their language, occupational skills, or forms of social behavior for themselves, they learn them from parents, teachers, peers, and others. Cultural variants are more like genes than are ordinary learned variants. Like genes, they are inherited and transmitted in a potentially endless chain, while variants acquired by individual learning are lost with the death of the learner.

Because culture is transmitted, it can be studied using the same Darwinian *methods* used to study genetic evolution. Human populations transmit a pool of cultural variation that is cumulatively modified to produce evolutionary change, much as they transmit an evolving gene pool. To understand cultural change we must keep track of all the processes in the lives of individuals that increase the frequency of some cultural variants and decrease the frequency of others. Of course, these processes do differ substantially from the processes of genetic evolution. Most important perhaps,