

The Evolutionary Biology of Decision Making

Jeffrey R. Stevens

Center for Adaptive Behavior and Cognition, Max Planck
Institute for Human Development, 14195 Berlin, Germany

Abstract

Evolutionary and psychological approaches to decision making remain largely separate endeavors. Each offers necessary techniques and perspectives which, when integrated, will aid the study of decision making in both humans and nonhuman animals. The evolutionary focus on selection pressures highlights the goals of decisions and the conditions under which different selection processes likely influence decision making. An evolutionary view also suggests that fully rational decision processes do not likely exist in nature. The psychological view proposes that cognition is hierarchically built on lower-level processes. Evolutionary approaches to decision making have not considered the cognitive building blocks necessary to implement decision strategies, thereby making most evolutionary models of behavior psychologically implausible. The synthesis of evolutionary and psychological constraints will generate more plausible models of decision making.

Introduction

A hungry female chimpanzee spies a termite mound and quickly fashions a branch into a long twig. She then uncovers a tunnel in the mound and inserts her twig. Soon, she extracts the twig, revealing a dozen wriggling termites clinging on tightly. The expert angler carefully plucks off and consumes each insect. As she repeats the process, she depletes the soldier termites arriving to defend their nest. When should she leave this hole to either excavate another tunnel or seek a new mound altogether? What cognitive process does she use to make this decision? What cognitive building blocks does she need to implement this process? Of the various processes possible, why does she use this one? This foraging situation raises these and numerous other questions for

biologists and psychologists interested in decision making in both humans and nonhuman animals.

Tinbergen (1963) posited four levels of analysis for why a behavior exists: the phylogenetic, functional, developmental, and mechanistic levels. Evolutionary biologists focus primarily on why behavioral decisions exist from a functional perspective. For example, what benefit exists for leaving the termite hole now versus in ten minutes? Psychologists, by contrast, explore the mechanistic level, concentrating typically on cognitive mechanisms involved in decision making. For instance, what information does the chimpanzee use to decide when to leave, and how does she acquire this information? Regrettably, the functional and mechanistic studies of decision making have remained largely separate endeavors, with many behavioral biologists and psychologists reluctant to cross disciplinary boundaries. Yet, the emergence of cognitive ecology and evolutionary psychology as fields demonstrates a recent push to integrate behavioral function and mechanism across species (see, e.g., Barkow et al. 1992; Dukas 1998; Hammerstein and Hagen 2005; Kacelnik 2006). This integration should be taken seriously when models of cognitive mechanisms and evolutionary outcomes are constructed. Here, I propose ways in which an evolutionary analysis can aid psychologists and ways in which a psychological analysis can aid evolutionary biologists in studying human and nonhuman animal decision making.

Evolutionary theory offers well-developed models of how a process of selection influences a characteristic or trait over time. I begin by briefly reviewing an evolutionary approach of decision making. Then I explore different selection processes and how they can influence a decision mechanism. Considering the kinds of pressures acting on decisions can help evaluate the feasibility of their mechanisms. Psychologists often characterize the cognitive building blocks or psychological capacities (such as memory capacities, learning structures, and attentional abilities) required for higher-level behavior. Cognitive building blocks are often absent from evolutionary models of behavior, leaving the models unconstrained and incomplete. Combining selective pressures and cognitive building blocks will lead to more plausible evolutionary and psychological models of decision making.

Natural Selection and Decision Making

Natural selection is a process by which biological evolution occurs; it results when heritable variation in a trait has differential influences on fitness. We can define fitness as the expected number of descendants produced by an individual, usually measured as survival and reproduction. Individuals that produce more descendants in future generations have higher fitness. A heritable behavior that gives even a slight advantage in survival to one individual relative to others will propagate over evolutionary time, resulting in more genes for that

advantageous behavior in the population. Thus, natural selection favors traits resulting in higher survival and reproduction. With a constant environment and enough genetic variation and time, this type of process results in organisms reaching fitness optima. Many biologists tend to focus on this end state and thereby consider natural selection an optimizing process. Inevitably, a process that favors an increase in fitness results in reaching an optimum, but natural selection as a process does not optimize globally to adapt organisms perfectly to their environments. Instead, natural selection “optimizes under constraints,” and a number of factors can constrain the evolution of behavior. First, the fitness payoffs of a behavior must trade off both the benefits and the costs. For instance, a foraging rule that extends the time an animal stays in a patch of food might increase an individual’s overall food intake, resulting in a fitness advantage over an individual with a shorter patch exploitation time. This rule, however, has consequences outside of the foraging domain that may balance (or exceed) the benefits of longer foraging times. For instance, foraging in a patch for longer time intervals may increase exposure to predators, potentially offsetting the benefits of higher food intake. Natural selection optimizes total net fitness across all domains of an organism’s survival and reproduction, thereby constraining optimization in any single domain. Second, natural selection does not act as a designer, creating traits *de novo*. Rather, it acts as a tinkerer, building on previous traits and accumulating change. The evolutionary history of an organism sets the starting point and thereby constrains potential evolutionary trajectories (see McElreath et al., this volume). Thus, natural selection optimizes behavior given trade-offs with other behaviors and existing evolutionary building blocks.

Many behavioral biologists accept the assumption of natural selection as an optimizing (under constraints) process and use mathematical optimization techniques as a tool to predict end-state behavior given a set of environmental parameters. Numerous behavioral domains use optimization techniques, especially foraging theory (Stephens and Krebs 1986). The termite fishing example highlights one of the classic models of optimal foraging: the patch model. When in a patch of food, foragers receive a particular rate of gain; that is, a particular amount of food per unit time spent foraging. As the forager continues to consume food, the patch may deplete, thereby diminishing the returns to staying in the patch. Thus, foragers must decide when to leave and travel to a new patch versus stay and continue to deplete the current patch. Optimization theory predicts that in some circumstances foragers should leave a patch when the intake rate at that patch drops below the average intake rate at the remaining patches (Charnov 1976). This optimal policy suggests that longer travel times between patches should result in longer patch residence times. Experimental and field data qualitatively support predictions derived from the optimal policy, but the data do not necessarily provide a good quantitative fit. Including more realistic assumptions about the role of physiological state, predation risk, and alternative activities provides a better account of the data (Nonacs 2001).

Assuming that organisms make optimal decisions, what cognitive process do they use? Optimality theorists readily admit the infeasibility of animals using optimal decision processes to arrive at optimal outcomes: animals do not calculate a range of expected fitness consequences and apply calculus to find the optimum. Instead, optimality theorists suggest that natural selection acts as the optimizing selection process, generating decision processes that result in approximately optimal outcomes (Stephens and Krebs 1986; Smith and Winterhalder 1992; Houston and McNamara 1999). Therefore, animals that use decision processes or strategies approaching the optimum outcome transfer more genes to future generations. This focus on strategies is important because natural selection selects for decision processes not the behavior per se, since only the decision process is heritable.

An evolutionary perspective emphasizes three important components of a decision: the decision processes, outcomes, and selection processes. To understand how natural selection can act on decisions, we must first map out the components of the decision mechanism.

The Anatomy of a Decision

Decisions—broadly defined here as the results of an evaluation of possible options—can take a variety of forms, including both inferences and preferences. *Inferences* go beyond the information given to make predictions about the state of the world; for instance, knowing the color of a fruit, can a decision maker infer its ripeness and sugar content? In contrast, *preferences* rank the desirability of options; for instance, would a decision maker prefer to receive a small food item now or a large food item tomorrow? This distinction has not been widely acknowledged (cf. Gigerenzer and Todd 1999), yet it could have important implications for the evolution of decision making.

Figure 13.1 illustrates the primary components of a decision mechanism. This description of a decision is silent on whether each component acts at a conscious or subconscious level, and most of the components can act at either level. Many decisions begin with a goal or task, such as foraging for food. The decision mechanism then gathers and processes information to reach a decision, which may result in an action and a payoff outcome.

Information

A decision maker must first search for information about possible decision options, although some information may be readily available for processing. The search for information can occur at an internal or external level (Fiedler and Juslin 2006). Internal search often refers to searching through memory for relevant information about options, whereas external search refers to perceiving information in the physical and social environment. In our foraging example, the chimpanzee may retrieve internally from memory information on

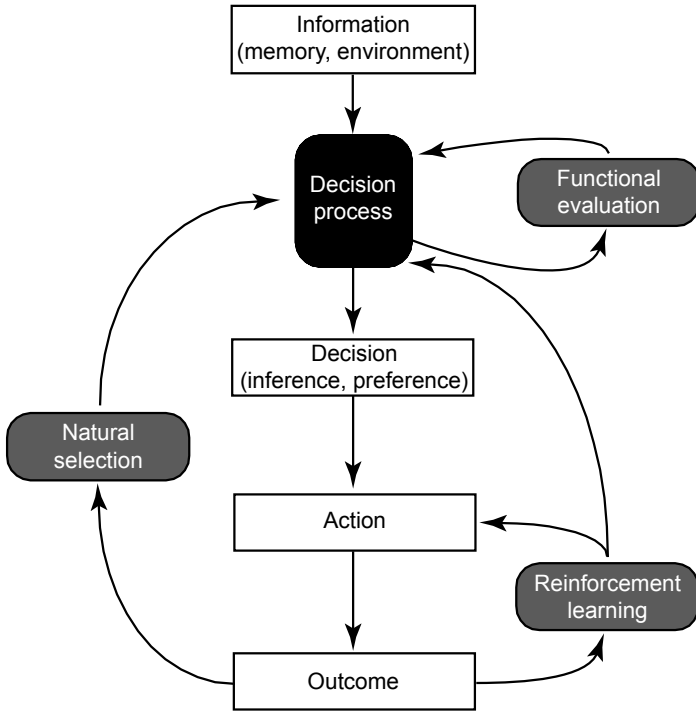


Figure 13.1 The decision mechanism and selection processes. The decision mechanism begins when information from either the internal or external environment feeds into the decision process. The decision process then generates a decision, which possibly leads to an action and outcome. Selection pressures such as natural selection, reinforcement learning, and functional evaluation can alter the decision mechanism by providing feedback about the realized or possible outcomes. Note that reinforcement learning can either influence action directly (learning a behavior) or can influence the decision process (learning a strategy).

the intake rates at other termite mounds and may track externally the gain and foraging time at the current mound when deciding when to leave.

Decision Process

In the decision process, a mechanism processes and integrates information to make a decision. Although other perspectives exist, I focus here on two general views of the decision-making process: unbounded rationality and bounded rationality.

Historically, many models of decision making have been based on the “economic man” perspective, in which decision makers can access all information relevant to a decision and arrive at optimal inferences via rules of logic and

statistics (e.g., Bayes's rule, linear regression) or exhibit optimal preferences via rules of probability (e.g., expected utilities). An unboundedly rational decision maker uses all information available to arrive at the decision producing an optimal outcome. Proponents of unbounded rationality focus on the optimal outcomes and skirt claims about optimal processes by stating that agents behave "as if" they are rational. Nevertheless, any claims of unbounded rationality require that agents possess sophisticated mental inference or preference functions that, when supplied with all relevant information, output the optimal decision. Deviations from the norms of linear regression, Bayes's rule, or expected utility are considered normatively "irrational" (but not necessarily unsophisticated) behavior.

An alternative to the omniscience and unlimited computational power required of economic man is a perspective emphasizing a more realistic view of tools available to decision makers. The bounded rationality approach advocates a plausible notion of the capacities of and constraints on the mind, as well as the interaction of the mind and the decision-making environment (Simon 1956; Gigerenzer and Selten 2001). This bounded rationality approach implies a set of computationally simple heuristics that use only partial information to make good, robust decisions that apply to specific decision-making environments (Payne et al. 1993; Gigerenzer et al. 1999). Rather than having general-purpose statistical devices requiring extensive information and complicated computations, decision makers often succeed by using simple heuristics adapted specifically to their environment. The simple heuristics approach makes explicit predictions about the decision process, the outcomes, and the conditions under which heuristics will work.

Decisions, Actions, and Outcomes

The decision process, of course, results in a decision: either an inference or a preference. Though I distinguish between inferences and preferences as separate entities, they can interact such that inferences can feed into preference decisions and vice versa. However, as internal constructs, inference and preference decisions are invisible to selection pressures such as natural selection and reinforcement learning. A decision maker must translate a decision into action (even if the action is not to act) to experience external consequences (Rachlin 1989). A decision process that generates a decision not connected to an action cannot be selected for via natural selection or learning.

In summary, the decision mechanism inputs information from the internal and external environment into a decision process, which generates a decision. The decision maker chooses an action based on this decision and receives an outcome. Stopping here, however, raises the question: How do we know whether a decision is good or bad? An evolutionary analysis provides an answer.

Selection Criteria and Processes

To determine whether a decision is good or bad, a selection process must evaluate the outcome relative to some criterion. Hammond (2000, 2007) described two types of selection criteria: correspondence and coherence. *Correspondence* refers to the degree to which decisions achieve empirical accuracy; that is, whether they reflect the true state of the world. For instance, we can evaluate an inference about how fruit color relates to sugar content based on how well this inference corresponds with the true relationship between color and sugar content. Alternatively, *coherence* refers to following norms, usually rational norms such as Bayesian reasoning and expected utility theory. An inference about fruit color and sugar content cannot only correspond to the state of the world; it can also cohere to a Bayesian analysis of an individual's prior experience with color and sugar content.

Inferences and preferences differ in which types of criteria apply to them. Both criteria can apply to inferences: how accurately they correspond to the state of the world and how consistently they cohere to a norm (as demonstrated by the fruit color/sugar example). Preferences, however, have no correspondence criteria, since they simply reflect an internal ranking and therefore rely only on coherence criteria. When deciding between a smaller, sooner versus larger, later reward, for example, there is no "correct" choice based on the state of the world. Rather, the choice depends on the relevant norm: Are you maximizing expected utility or minimizing the time to your next reward? The lack of a universal criterion makes matching a decision to a coherence criterion quite difficult. Most studies of inference assess whether performance coheres to predictions of logical analysis. For example, in the Wason selection task (Figure 13.2; see also Schooler, this volume), most subjects violate logical analysis of the situation. However, the pattern of data fits the assumption that decision makers are maximizing information gain rather than following logic (Oaksford and Chater 1994). Additionally, when the situation is framed as a social contract rather than an abstract logic problem, subjects provide the logically correct solution much more often, possibly triggering an evolved cheater detection mechanism (Cosmides and Tooby 1992; Gigerenzer and Hug 1992). Given that multiple coherence criteria exist and different criteria seem to govern behavior in different contexts, it may prove difficult to predict a priori which coherence criterion is relevant to a particular decision.

Selection processes evaluate how well outcomes match the selective criteria, and different selection processes require different criteria. As one class of processes, Skinner (1981) proposed selection by consequences—processes which select or reject behaviors based on their realized outcomes. The selection-by-consequences view emphasizes that the process of selection works on previous consequences. In decision making, two selection processes clearly use selection by consequences (natural selection and reinforcement learning), whereas a third (functional evaluation) does not. This is not an exhaustive list




Figure 13.2 Wason selection task. Subjects must decide which of four cards to turn over to test the rule: if there is a vowel on one side, then there is an even number on the other. This is logically equivalent to the rule *if p, then q*. Most subjects choose A or A and 2. The logically correct solution is A and 7; that is, following the rules of logic, subjects should investigate *p* and *not-q*.

of selection processes for human or other animals, but it offers a starting point for exploring feedback mechanisms that shape decisions.

Natural Selection

Natural selection provides the ultimate example of selection by consequences. When an individual with specific inferences and preferences produces more offspring than other individuals, the genes influencing those inferences and preferences will proliferate in the population. Specifically, natural selection favors genes for decision processes that result in good decisions, actions, and outcomes from a fitness perspective. Note that with natural selection no direct pressure exists for accuracy or correspondence per se, just pressure to produce descendents. Accuracy is only valuable when tied to fitness. For instance, it makes no sense evolutionarily for a strict carnivore to make inferences about fruit color and sugar content, because it faces no fitness consequences for making the inference. Similarly, selection cannot act at the decision stage alone because without behavior, no differential fitness exists. Inferences and preferences without action have no consequences upon which selection can act. Even if an action is taken, selection does not act on the action itself because multiple decision processes can yield the same action. Natural selection can only act on heritable decision processes that generate actions and produce outcomes.

The slow pace at which natural selection can track a changing environment means that genetically coding all decision mechanisms would leave organisms poorly adapted to their environment. Dawkins (1976) pointed out that our brains are not slaves to our genes; they allow us to process information and execute actions flexibly. Dawkins compared our genes to policy makers and our brains to executives. Stanovich (2004) described decisions under direct influence of genes as “short-leashed” and more flexible behavior controlled by brains rather than genes as “long-leashed.” Natural selection has equipped decision makers with longer-leashed decision mechanisms by creating other selection processes that respond more flexibly to the environment: reinforcement learning and functional evaluation.

Reinforcement Learning

Learning, of course, does not require evolutionary time or the differential survival of individuals. Instead, over the lifetime of an individual (or potentially much shorter time periods), actions yielding good outcomes occur more often than actions yielding poor outcomes (Skinner 1938). Therefore, learning selects based on the coherence criteria of reinforcer maximization. Why then is one outcome more reinforcing than another? Because natural selection built learning mechanisms to allow flexible decision making, the value of reinforcement is tied directly to evolutionarily relevant commodities. More food, water, and sex, along with less pain, typically result in higher fitness, though not always. A mechanism exploiting this correlation will tend to serve the organism well; substituting reinforcer maximization as a proxy for fitness maximization allows a more flexible proximate mechanism to achieve a good evolutionary result.

Some types of learning differ from natural selection in the decision components used. For instance, in operant conditioning, a response is associated with a reinforcer; this maps onto the action and outcome components of Figure 13.1. Therefore, operant conditioning does not influence the decision process itself but simply the frequency of an action given its outcome. In contrast, Rieskamp and Otto (2006) suggested that reinforcement learning can also work at the level of the decision process or strategy. Decision makers can learn to implement different decision processes based on feedback from the outcomes received. Reinforcement learning can then adapt either actions or the decision process itself to the reinforcement contingencies in the environment.

Functional Evaluation

Natural selection favors appropriate decision processes over evolutionary time, and reinforcement learning selects behavior over the lifetime. Both processes select actions based on previously experienced outcomes. Perhaps natural selection has allowed even more flexibility in decision making by further lengthening the leash. Novel decision-making environments may arise, making hard-wired rules obsolete and trial-and-error learning too costly or slow. Perhaps natural selection has generated a selection process that selects not on previous outcomes but on behavior to achieve a desired end state or goal. Using this selection process—here termed functional evaluation—a decision maker mentally evaluates the decision options available and chooses the one with the potential to maximize the relevant selection criterion. Instead of maximizing fitness or reinforcer value, functional evaluation assesses potential actions and outcomes relative to a decision goal.

Because functional evaluation considers the fit between outcomes and the selective criteria directly, selection is embedded completely within the decision process phase of the decision mechanism. Unlike natural selection and

learning, functional evaluation selection occurs before making an action. Actions do not die out and are not reinforced, but can be entertained mentally and compared to the decision goal. Functional evaluation requires an internal representation of the coherence criteria, such as a mental utility function. Also, unlike natural selection and learning, there exists no single selective criterion for functional evaluation. Individuals can maximize utility, predictive accuracy, sampling opportunities, or many other criteria. There is no single criterion because decision makers face multiple decision goals simultaneously.

Functional evaluation should not be confused with a decision process. It does not refer to the conscious processing of information. In fact, functional evaluation can occur at either the conscious or subconscious level. It simply refers to the selection process in which outcomes are evaluated before they are experienced.

Optimal Processes and Outcomes

The “anatomical” classification of decision making emphasizes two processes at work: the selection process and the decision process. Recall that the selection process evaluates the outcomes with the correspondence or coherence criteria, whereas the decision process gathers and processes information to arrive at a decision. In natural selection, the selection process works directly on the decision process, which then generates the decision outcome. This hierarchy helps assess which (if any) aspects of decision making are optimal. I contend that an optimal selection process is required to create an optimal decision process, which is required to produce universally optimal outcomes. I define optimal outcomes as the best possible payoffs, optimal decision processes as processes that *always* produce optimal outcomes, and optimal selection processes as processes that *always* produce optimal decision processes.

As previously stated, natural selection is not an unboundedly optimal selection process, rather it optimizes under constraints. Because natural selection optimizes under constraints, it cannot produce universally optimal decision processes and therefore cannot yield optimal outcomes across all decision-making domains. It can nevertheless result in optimal outcomes for specific circumstances. To take an example from foraging, natural selection will not generate a mechanism that obtains all information about all possible foraging opportunities, trades the lost opportunity for other fitness-enhancing endeavors, and then calculates the optimal foraging choice. This would represent a universally optimal decision process that always yields the best fitness outcomes. Yet, a mechanism may evolve that can sample from two foraging patches and choose the one that yields the highest intake rate, resulting in an optimal outcome. The difference lies in the scale: optimality can evolve at small scales for particular circumstances but not at large scales across all domains. As Houston et al. (2007, p. 1532) explain: “We cannot expect natural selection, having no

foresight, to shape organisms to act rationally in all circumstances, but only in those circumstances which it encounters in its natural setting.”

The limitations of natural selection’s optimization under constraints not only influence the decision process and outcomes, but also trickle down into the other selection processes. Because both reinforcement learning and functional evaluation selection mechanisms result from evolution by natural selection, they cannot be unboundedly optimal selection processes; no evolved selection or decision process can be unboundedly optimal. A selection process would have to emerge outside of natural selection to be truly optimal.

In summary, the selection process view makes two contributions to the study of decision making. First, it requires an analysis of the pressures that shape decision making. Natural selection, reinforcement learning, and functional evaluation influence the decision mechanism. Natural selection and reinforcement learning select decision processes or actions based on previous realized consequences, whereas functional evaluation selects decisions in the absence of direct feedback from an outcome.

Second, the selection perspective suggests that the constrained optimization of natural selection will limit the degree of optimality for evolved selection processes and decision processes. It seems implausible for a constrained optimization selection process to generate a universally optimal decision process following a rational norm. Instead, natural selection likely creates decision processes that increase fitness rather than cohere to the rules of logic and probability theory (Hammond 2007). Therefore, the frequent labeling of behaviors as “irrational,” “anomalies,” or “biases” assumes a particular perspective of rational norms. A broader, evolutionary perspective, however, cautions against using these labels, emphasizing instead an understanding of the decision goals, selection pressures, and decision-making environment.

The Evolution of the Decision Process

Given that natural selection cannot generate universally optimal decisions, what kinds of decision processes are feasible? As mentioned previously, many studies of decision making assume, at least implicitly, unbounded rationality and thereby imply an implausibly omniscient, temporally unconstrained, and computationally unlimited decision maker. There are, however, examples in which agents seem to make unboundedly rational decisions (Glimcher 2003; Glöckner, this volume). These decisions, however, occur in specific circumstances, and the generality of their application remains unexplored. Moreover, even if organisms possess the ability to use higher-order cognitive skills like optimal decision making, they do not necessarily do so when simpler abilities will suffice. For instance, cotton-top tamarins (*Saguinus oedipus*) use simpler, amount-based mechanisms when discriminating different quantities of food, even though they can use more sophisticated, number-based mechanisms in other situations (Stevens et al. 2007).

A more computationally and evolutionarily plausible perspective is the bounded rationality approach, which assumes that decision makers use simple heuristics and satisficing rules to make decisions with minimal information over short time periods (reviewed in Hutchinson and Gigerenzer 2005). Although many behavioral biologists ignore the process of decision making, they work under the assumption that animals do not use optimal decision processes. Instead, biologists often assume that animals use rules of thumb (heuristics) that approach optimal outcomes.

Evolving Simple Heuristics

The simple heuristics approach emphasizes the presence of rules for ordering information search, stopping search, and making a decision (Gigerenzer et al. 1999). These heuristics require little computation and use minimal information to make decisions. The rules of thumb approach in biology has coincidentally focused on the fast and frugal nature of rules of thumb without explicitly considering the search, stop, and decision rules of the heuristics.

Animals use rules of thumb in a number of important decision-making contexts, ranging from navigation to nest construction (Marsh 2002; Hutchinson and Gigerenzer 2005). Rules of thumb are particularly well studied in optimal foraging (Stephens and Krebs 1986), which may seem like a rather unlikely area for their application given the intense focus on optimization models. Biologists, however, distinguish between optimal outcomes and feasible mechanisms that can approach those outcomes.

Biologists have investigated the use of simple rules in the patch choice model of foraging described in the termite fishing example. Recall that in this scenario, foraging chimpanzees must decide when to leave a patch and move on to another. The optimal policy recommends leaving when the intake rate at the current patch equals the average intake rate for the remaining patches under this policy. Calculating or estimating this average intake rate in the environment (including travel times) is not computationally trivial. A number of researchers have proposed simple patch-leaving rules that avoid some of the complicated computations (reviewed in Stephens and Krebs 1986). For instance, rather than comparing the current intake rate to the average rate, animals may just leave a patch when the current intake rate drops below a critical threshold. Other simpler rules even dispense with the requirement of directly monitoring the current intake rate and instead estimate this rate indirectly. Animals using these rules may leave after consuming a certain number of prey items (fixed number rule), after a certain time period after arriving at a patch (fixed time rule), or after a certain time period of unsuccessful foraging (giving-up time rule). Empirical evidence suggests that different animal species use these various rules in different foraging situations (Stephens and Krebs 1986; van Alphen et al. 2003; Wajnberg et al. 2003).

The use of simple rules and heuristics by nonhuman animals is not terribly surprising. This perspective, however, has stimulated more controversy when applied to human decision making (see Todd and Gigerenzer 2000 and subsequent commentaries). Do humans use simple heuristics for important decisions? Gigerenzer and colleagues have argued that in certain environments heuristics can achieve good outcomes. Given that the human brain has been built by evolution through natural selection, we can infer that the costs of decision computations weigh heavily in the evolution of decision processes and that processes with simple decision rules tend to prevail over complex computations when yielding similar outcomes.

Evolving Satisficing Rules

When using another mechanism of bounded rationality—satisficing—decision makers search through options and select the first one that surpasses some threshold or aspiration level (Simon 1956). Economists and psychologists frequently investigate satisficing in decision making from the individual to the institutional level. Biologists, however, have not warmed to the concept of satisficing. Skepticism in biology derives from at least two sources (Stephens and Krebs 1986; Smith and Winterhalder 1992). First, satisficing theory provides no a priori justifications for how thresholds are set: What is “good enough”? Second, if natural selection is an optimizing process, then satisficing is not evolutionarily stable. An individual who optimizes and consistently exceeds another individual’s decision threshold (e.g., makes better than “good enough” decisions) will have higher fitness, and this optimizing trait will spread relative to the satisficing trait. Stephens and Krebs (1986) suggested that if the effect of performance on fitness is a step function, a satisficing rule could evolve, but they argue that this rigid function is uncommon in nature. Instead, it is likely that fitness increases continuously with performance.

Despite skepticism towards satisficing, some biologists champion it as a reasonable decision process (Janetos and Cole 1981; Ward 1992), and recent theoretical investigations have modeled satisficing strategies in foraging and mate choice contexts (Todd and Miller 1999; Carmel and Ben-Haim 2005). In these cases, satisficing can produce nearly optimal outcomes with appropriate thresholds, depending on the costs of information acquisition and the level of environmental variation. Additionally, these satisficing mechanisms yield decisions that are more robust to uncertainty (Todd and Gigerenzer 2000; Carmel and Ben-Haim 2005), yielding better than “good enough” results.

Using limited information and simplified decision rules can produce optimal or near optimal outcomes, and thus heuristics and satisficing can offer plausible processes for decision making. These decision processes act independently of the selection process. Heuristics and satisficing can be genetically hardwired, subject directly to natural selection; they can be learned based on reinforcer value; or they can be functionally evaluated relative to a decision goal.

Evolving the Mechanisms for Decision Making

An evolutionary perspective can clarify the role of selection pressures and limit the set of plausible decision processes. Thus far, however, the evolutionary perspective has largely ignored the cognitive building blocks needed to implement decisions. For instance, the rational-actor approach in game theory assumes that all individuals have perfect knowledge of all other individuals' knowledge and beliefs, meaning that I know that you know that I know that you know and on and on. This implies that game theoretic decision making requires the cognitive building block of theory of mind (McCabe and Smith 2000; Hedden and Zhang 2002; McCabe and Singer, this volume), so decision makers can represent the knowledge and beliefs of other agents. The building block approach allows us to predict which types of decision makers utilize which types of decision processes: organisms lacking specific cognitive abilities cannot implement certain decision processes. My colleagues and I have explored this idea in the realm of animal cooperative interactions (Stevens and Hauser 2004; Stevens, Cushman et al. 2005).

A Case Study in Cooperation

One of the most-studied games in biology is the Prisoner's Dilemma (Figure 13.3), a game in which individuals can either cooperate or defect. Cooperation maximizes the total payoff to everyone involved in the interaction (mutual cooperation provides higher benefits than mutual defection); however, any individual receives higher personal payoffs by defecting, resulting in a sizable temptation to cheat. Unilateral cooperation is not evolutionarily stable in this game; therefore, cooperation is altruistic because cooperators pay a fitness cost by foregoing the temptation to cheat. The possibility of altruistic cooperation in the Prisoner's Dilemma intrigues biologists, because the presence of altruistic behavior violates the standard principles of natural selection. A defector will have a higher payoff than an altruist and therefore will contribute more genes to the next generation. Biologists have proposed a number of solutions to the Prisoner's Dilemma that may allow cooperation to exist, including kin selection, reciprocal altruism, punishment, and reputation formation (Dugatkin 1997). Reciprocal altruism, in which individuals pay a short-term cost of cooperation for the future benefit of a social partner's reciprocated cooperation (Trivers 1971), has probably gained the most attention. Axelrod and Hamilton (1981) investigated an evolutionary strategy for reciprocal altruism called Tit-for-Tat (TFT), in which a player begins by cooperating and copies the opponent's previous behavior. TFT can outperform pure defection when individuals repeatedly engage in a Prisoner's Dilemma with the same partners (although TFT is not evolutionarily stable; Selten and Hammerstein 1984; Boyd and Lorberbaum 1987). Therefore, a rule such as TFT can, in theory, maintain cooperation in the face of defection in the Prisoner's Dilemma. Yet, despite

	Cooperate	Defect
Cooperate	R	S
Defect	T	P

Figure 13.3 Prisoner's Dilemma payoff matrix. In the Prisoner's Dilemma, players must decide to cooperate or defect, and the payoffs are ranked such that $T > R > P > S$. Mutual cooperation pays more than mutual defection ($R > P$), but defection always pays more than cooperation for an individual ($T > R$ and $P > S$).

theoretical evidence supporting the viability of reciprocal altruism, it is not well supported by empirical evidence in animals (Clements and Stephens 1995; Noë and Hammerstein 1995; Pusey and Packer 1997; Hammerstein 2003).

Why is reciprocal altruism rare in animals when TFT is such a simple heuristic? Perhaps, despite numerous theoretical investigations of TFT, no models have included the cognitive building blocks required to implement reciprocal strategies (Hammerstein 2003; Stevens and Hauser 2004; Stevens, Cushman et al. 2005). When Trivers (1971) first introduced the concept of reciprocal altruism, he outlined necessary prerequisites, including one cognitive building block: the ability to detect cheaters (individual recognition). This single requirement does not likely capture the cognitive sophistication required for utilizing reciprocal strategies. In particular, the delay between the cost of a cooperative act and the benefit of reciprocated cooperation introduces a number of cognitive challenges. Minimally, this delay interacts with memory, inhibitory control, and impulsivity processes.

Memory

Limitations in memory decay, interference, and capacity can constrain the frequency of reciprocal altruism. Models of forgetting predict exponential or power functions (Wixted 2004) that decay rapidly over time. Therefore, longer time intervals between cooperative acts may make reciprocal altruism more difficult. Even with short time delays between cooperative interactions, previous memories can interfere with recall, and every potential new partner increases the computational load of tracking debts owed and favors given. Tracking reciprocal obligations with multiple individuals may place a computationally intensive burden on memory systems.

Inhibitory Control

Reciprocal interactions begin with an inhibitory problem: Can an individual inhibit the choice of the large benefit (the temptation to defect)? Animals often fail to inhibit their preferences for larger rewards. In the reversed-contingency task, subjects must reach toward the smaller of two visible rewards to receive the larger reward (Boysen and Berntson 1995). Numerous primate species

tested in this paradigm failed repeatedly to learn the contingency, because they could not inhibit their response to choose the larger of two options: the inhibitory system appears to be too weak. Subjects can learn to solve this task, but the solution requires hundreds to thousands of trials, correction procedures, or symbolic representation of the quantities (Boysen et al. 1996; Silberberg and Fujita 1996; Murray et al. 2005). This strong preference for larger rewards suggests that choosing altruistic actions (those providing smaller benefits when a larger benefit is available) may also prove difficult. Therefore, strong inhibitory control capacities are required to implement reciprocal strategies.

Self-control

Reciprocal altruism not only requires choosing a smaller benefit but also waiting for the reciprocated benefit, thus, requiring self-control. In reality, however, animals behave quite impulsively. When given a choice between a smaller, immediate reward and a larger, delayed reward, animals show a strong preference for immediacy (Mazur 1987). For instance, when offered a choice between an immediate option and a delayed option, marmoset monkeys (*Callithrix jacchus*) refuse to wait more than 20 seconds for three times as much food (Stevens, Hallinan et al. 2005). Most other species tested so far show a similar lack of self-control, with the exception of the apes, which can wait for minutes instead of seconds (Rosati et al. 2007). This intense preference for immediacy suggests that animals cannot forego the instant benefits of defection in lieu of the delayed benefits from reciprocal cooperation. An operant experiment with blue jays (*Cyanocitta cristata*) demonstrated that some subjects cooperated in a Prisoner's Dilemma more often when they played an opponent using TFT and their impulsivity was experimentally reduced than when playing a defecting opponent or when impulsivity was not reduced (Stephens et al. 2002). Reciprocal altruism requires patience.

A simple heuristic like TFT may be easy for humans to implement, but it requires cognitive building blocks—such as high memory capacity, inhibitory control, and patience—beyond the abilities of many animal species. This perspective raises the question: Are cognitive building blocks constraints on the evolution of behavior or are they evolutionary adaptations tuned for specific decision contexts? For instance, researchers often refer to memory as a constraint, emphasizing memory capacities and loads. Psychologists have begun exploring the adaptive nature of memory, showing how memory processes may track the usefulness of information in the environment (Anderson and Schooler 1991 and Schooler, this volume). Yet, building blocks such as memory may be adaptive only in a specific context. It remains unclear under what conditions they constrain or facilitate higher-level decision making.

Conclusion

A truly integrative study of decision making must synthesize evolutionary and psychological approaches. Though the emerging fields of cognitive ecology and evolutionary psychology have begun this integration, much work remains. Considering the selective pressures on decisions refines which kinds of correspondence and coherence criteria are feasible for decisions. Natural selection does not favor coherence to rational norms, but increases fitness relative to others in the population. Considering selection also emphasizes that natural selection is a process of optimization under constraints. Because a constrained optimization process cannot generate a universally optimal process, unboundedly optimal decision mechanisms cannot exist in nature. Therefore, studying decision making with an eye on evolution can aid in understanding the goals of decision and thereby explain (or dispel) notions of irrational choice. Despite the advantages of accounting for natural selection in decision making, an entirely evolutionary, outcome-based approach overlooks the limitations that cognitive abilities impose on decision processes; certain cognitive building blocks must exist to implement decision processes. Many decisions can be made with a set of simple building blocks, whereas some require more sophisticated cognitive abilities. Thus, a complete understanding of decision making rests on the appropriate integration of ultimate goals and evolutionary pressures with the psychological mechanisms of choice.

Acknowledgments

I am especially grateful to Peter Carruthers, Gerd Gigerenzer, Peter Hammerstein, John Hutchinson, and David Stephens for comments on an early version of the chapter and for thought-provoking discussions about evolution and cognition.

References

- Anderson, J. R., and L. J. Schooler. 1991. Reflections of the environment in memory. *Psychol. Sci.* **2**:396–408.
- Axelrod, R., and W. D. Hamilton. 1981. The evolution of cooperation. *Science* **211**:1390–1396.
- Barkow, J. H., L. Cosmides, and J. Tooby. 1992. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford Univ. Press.
- Boyd, R., and J. P. Lorberbaum. 1987. No pure strategy is evolutionarily stable in the repeated Prisoner's Dilemma game. *Nature* **327**:58–59.
- Boysen, S. T., and G. G. Berntson. 1995. Responses to quantity: Perceptual versus cognitive mechanisms in chimpanzees (*Pan troglodytes*). *J. Exp. Psych.: Anim. Behav.* **21**:82–86.
- Boysen, S. T., G. G. Berntson, M. B. Hannan, and J. T. Cacioppo. 1996. Quantity-based interference and symbolic representations in chimpanzees (*Pan troglodytes*). *J. Exp. Psych.: Anim. Behav.* **22**:76–86.

- Carmel, Y., and Y. Ben-Haim. 2005. Info-gap robust-satisficing model of foraging behavior: Do foragers optimize or satisfice? *Am. Nat.* **166**:633–641.
- Charnov, E. L. 1976. Optimal foraging: The marginal value theorem. *Theor. Pop. Biol.* **9**:129–136.
- Clements, K. C., and D. W. Stephens. 1995. Testing models of non-kin cooperation: Mutualism and the Prisoner's Dilemma. *Anim. Behav.* **50**:527–535.
- Cosmides, L., and J. Tooby. 1992. Cognitive adaptations for social exchange. In: *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, ed. J. H. Barkow, L. Cosmides, and J. Tooby, pp. 163–228. Oxford: Oxford Univ. Press.
- Dawkins, R. 1976. *The Selfish Gene*. Oxford: Oxford Univ. Press.
- Dugatkin, L. A. 1997. *Cooperation among Animals: An Evolutionary Perspective*. New York: Oxford Univ. Press.
- Dukas, R. 1998. *Cognitive Ecology: The Evolutionary Ecology of Information Processing and Decision Making*. Chicago: Univ. Chicago Press.
- Fiedler, K., and P. Juslin. 2006. Taking the interface between mind and environment seriously. In: *Information Sampling and Adaptive Cognition*, ed. K. Fiedler and P. Juslin, pp. 3–29. Cambridge: Cambridge Univ. Press.
- Gigerenzer, G., and K. Hug. 1992. Domain-specific reasoning: Social contracts, cheating, and perspective change. *Cognition* **43**:127–171.
- Gigerenzer, G., and R. Selten. 2001. *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.
- Gigerenzer, G., and P. M. Todd. 1999. Fast and frugal heuristics: The adaptive toolbox. In: *Simple Heuristics That Make Us Smart*, ed. G. Gigerenzer, P. M. Todd, and the ABC Research Group, pp. 3–34. Oxford: Oxford Univ. Press.
- Gigerenzer, G., P. M. Todd, and the ABC Research Group. 1999. *Simple Heuristics That Make Us Smart*. Oxford: Oxford Univ. Press.
- Glimcher, P. W. 2003. *Decisions, Uncertainty, and the Brain*. Cambridge, MA: MIT Press.
- Hammerstein, P. 2003. Why is reciprocity so rare in social animals? A protestant appeal. In: *Genetic and Cultural Evolution of Cooperation*, ed. P. Hammerstein, pp. 83–94. Cambridge, MA: MIT Press.
- Hammerstein, P., and E. H. Hagen. 2005. The second wave of evolutionary economics in biology. *Trends Ecol. Evol.* **20**:604–609.
- Hammond, K. R. 2000. Coherence and correspondence theories in judgment and decision making. In: *Judgment and Decision Making: An Interdisciplinary Reader*, ed. T. Connolly, H. R. Arkes, and K. R. Hammond, pp. 53–65. Cambridge: Cambridge Univ. Press.
- Hammond, K. R. 2007. *Beyond Rationality: The Search for Wisdom in a Troubled Time*. Oxford: Oxford Univ. Press.
- Hedden, T., and J. Zhang. 2002. What do you think I think you think? Strategic reasoning in matrix games. *Cognition* **85**:1–36.
- Houston, A. I., and J. M. McNamara. 1999. *Models of Adaptive Behaviour: An Approach Based on State*. Cambridge: Cambridge Univ. Press.
- Houston, A. I., J. M. McNamara, and M. D. Steer. 2007. Do we expect natural selection to produce rational behaviour? *Phil. Trans. Roy. Soc. Lond. B* **362**:1531–1543.
- Hutchinson, J. M. C., and G. Gigerenzer. 2005. Simple heuristics and rules of thumb: Where psychologists and behavioural biologists might meet. *Behav. Proc.* **69**:97–124.
- Janetos, A. C., and B. J. Cole. 1981. Imperfectly optimal animals. *Behav. Ecol. Sociobiol.* **9**:203–209.

- Kacelnik, A. 2006. Meanings of rationality. In: *Rational Animals?*, ed. S. Hurley and M. Nudds, pp. 87–106. Oxford: Oxford Univ. Press.
- Marsh, B. 2002. Do animals use heuristics? *J. Bioecon.* **4**:49–56.
- Mazur, J. E. 1987. An adjusting procedure for studying delayed reinforcement. In: *Quantitative Analyses of Behavior: The Effect of Delay and of Intervening Events on Reinforcement Value*, ed. M. L. Commons, J. E. Mazur, J. A. Nevin et al., pp. 55–73. Hillsdale, NJ: Lawrence Erlbaum.
- McCabe, K. A., and V. L. Smith. 2000. A comparison of naive and sophisticated subject behavior with game theoretic predictions. *Proc. Natl. Acad. Sci.* **97**:3777–3781.
- Murray, E. A., J. D. Kralik, and S. P. Wise. 2005. Learning to inhibit prepotent responses: Successful performance by rhesus macaques, *Macaca mulatta*, on the reversed-contingency task. *Anim. Behav.* **69**:991–998.
- Noë, R., and P. Hammerstein. 1995. Biological markets. *Trends Ecol. Evol.* **10**:336–339.
- Nonacs, P. 2001. State dependent behavior and the Marginal Value Theorem. *Behav. Ecol.* **12**:71–83.
- Oaksford, M., and N. Chater. 1994. A rational analysis of the selection task as optimal data selection. *Psychol. Rev.* **101**:608–631.
- Payne, J. W., J. R. Bettman, and E. J. Johnson. 1993. *The Adaptive Decision Maker*. Cambridge: Cambridge Univ. Press.
- Pusey, A. E., and C. Packer. 1997. The ecology of relationships. In: *Behavioural Ecology: An Evolutionary Approach*, ed. J. R. Krebs and N. B. Davies, pp. 254–283. Oxford: Blackwell Science.
- Rachlin, H. 1989. *Judgment, Decision, and Choice: A Cognitive/Behavioral Synthesis*. New York: W. H. Freeman.
- Rieskamp, J., and P. E. Otto. 2006. SSL: A theory of how people learn to select strategies. *J. Exp. Psychol.: Gen.* **135**:207–236.
- Rosati, A. G., J. R. Stevens, B. Hare, and M. D. Hauser. 2007. The evolutionary origins of human patience: Temporal preferences in chimpanzees, bonobos, and adult humans. *Curr. Biol.* **17**:1663–1668.
- Selten, R., and P. Hammerstein. 1984. Gaps in Harley's argument on evolutionarily stable learning rules and in the logic of "tit for tat." *Behav. Brain Sci.* **7**:115–116.
- Silberberg, A., and K. Fujita. 1996. Pointing at smaller food amounts in an analogue of Boysen and Berntson's (1995) procedure. *J. Exp. Anal. Behav.* **66**:143–147.
- Simon, H. A. 1956. Rational choice and the structure of the environment. *Psychol. Rev.* **63**:129–138.
- Skinner, B. F. 1938. *The Behavior of Organisms: An Experimental Analysis*. Englewood Cliffs, NJ: Prentice Hall.
- Skinner, B. F. 1981. Selection by consequences. *Science* **213**:501–504.
- Smith, E. A., and B. Winterhalder. 1992. Natural selection and decision making. In: *Evolutionary Ecology and Human Behavior*, ed. E. A. Smith and B. Winterhalder, pp. 25–60. New Brunswick, NJ: Transaction Press.
- Stanovich, K. E. 2004. *The Robot's Rebellion: Finding Meaning in the Age of Darwin*. Chicago: Univ. Chicago Press.
- Stephens, D. W., and J. R. Krebs. 1986. *Foraging Theory*. Princeton: Princeton Univ. Press.
- Stephens, D. W., C. M. McLinn, and J. R. Stevens. 2002. Discounting and reciprocity in an iterated Prisoner's Dilemma. *Science* **298**:2216–2218.
- Stevens, J. R., F. A. Cushman, and M. D. Hauser. 2005. Evolving the psychological mechanisms for cooperation. *Ann. Rev. Ecol. Evol. Syst.* **36**:499–518.

- Stevens, J. R., E. V. Hallinan, and M. D. Hauser. 2005. The ecology and evolution of patience in two New World monkeys. *Biol. Lett.* **1**:223–226.
- Stevens, J. R., and M. D. Hauser. 2004. Why be nice? Psychological constraints on the evolution of cooperation. *Trends Cog. Sci.* **8**:60–65.
- Stevens, J. R., J. N. Wood, and M. D. Hauser. 2007. When quantity trumps number: Discrimination experiments in cotton-top tamarins (*Saguinus oedipus*) and common marmosets (*Callithrix jacchus*). *Anim. Cog.* **10**:429–437.
- Tinbergen, N. 1963. On aims and methods of ethology. *Z. Tierpsychol.* **20**:410–433.
- Todd, P. M., and G. Gigerenzer. 2000. Précis of Simple Heuristics That Make Us Smart. *Behav. Brain Sci.* **23**:727–741.
- Todd, P. M., and G. F. Miller. 1999. From pride and prejudice to persuasion: Satisficing in mate search. In: Simple Heuristics That Make Us Smart, ed. G. Gigerenzer, P. M. Todd, and the ABC Research Group, pp. 287–308. Oxford: Oxford Univ. Press.
- Trivers, R. L. 1971. The evolution of reciprocal altruism. *Q. Rev. Biol.* **46**:35–57.
- van Alphen, J. J. M., C. Bernstein, and G. Driessen. 2003. Information acquisition and time allocation in insect parasitoids. *Trends Ecol. Evol.* **18**:81–87.
- Wajnberg, E., P. A. Gonsard, E. Tabone, C. Curty, N. Lezcano et al. 2003. A comparative analysis of patch-leaving decision rules in a parasitoid family. *J. Anim. Ecol.* **72**:618–626.
- Ward, D. 1992. The role of satisficing in foraging theory. *Oikos* **63**:312–317.
- Wixted, J. T. 2004. The psychology and neuroscience of forgetting. *Ann. Rev. Psych.* **55**:235–269.