

2011

Mechanisms for Decisions about the Future

Jeffrey R. Stevens

University of Nebraska-Lincoln, jstevens5@unl.edu

Follow this and additional works at: <http://digitalcommons.unl.edu/psychstevens>

 Part of the [Animal Sciences Commons](#), [Animal Studies Commons](#), [Applied Behavior Analysis Commons](#), [Behavior and Ethology Commons](#), [Biological Psychology Commons](#), [Cognition and Perception Commons](#), [Cognitive Psychology Commons](#), [Communication Commons](#), [Evolution Commons](#), [Other Ecology and Evolutionary Biology Commons](#), and the [Other Psychology Commons](#)

Stevens, Jeffrey R., "Mechanisms for Decisions about the Future" (2011). *Jeffrey Stevens Papers & Publications*. 2.
<http://digitalcommons.unl.edu/psychstevens/2>

This Article is brought to you for free and open access by the Psychology, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in Jeffrey Stevens Papers & Publications by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

7

Mechanisms for Decisions about the Future

Jeffrey R. Stevens

Abstract

Evolutionary and psychological perspectives on decision making remain largely separate endeavors. The bounded rationality approach integrates these two perspectives by focusing on simple, plausible mechanisms of decision making and the cognitive capacities needed to implement these mechanisms. Decisions about the future provide a class of decisions that lend themselves to a bounded rationality approach. Though many different mechanisms may exist for making decisions about the future, only a subset of these mechanisms actually require a representation of the future. The bounded rationality approach helps focus on the cognitive capacities and decision mechanisms that are necessary for a full understanding of decision making about the future.

Introduction

A hungry female chimpanzee spies a termite mound and quickly fashions a branch into a long, thin twig. She then digs to uncover 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 decision mechanism does she use to make this choice? What cognitive capacities does she need to implement this mechanism? This foraging situation raises numerous other questions to 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 largely focus 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, in contrast, explore the mechanistic level,

typically concentrating 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 (e.g., Barkow et al. 1992; Dukas 1998; Hammerstein and Hagen 2005; Kacelnik 2006; McNamara and Houston 2009). This integration should be taken seriously when constructing models of cognitive mechanisms and evolutionary outcomes. Here, I highlight how an integration of evolutionary and psychological approaches is integral to an understanding of decision making. First, I discuss the importance of decision mechanisms and review two general approaches to studying decision making. In particular, the bounded rationality approach proposes simple mechanisms by which decisions are made. This emphasis on the decision mechanisms and the cognitive capacities required for the mechanisms may yield a more realistic understanding of how humans and other animals make decisions. Thereafter, I explore a particular class of decisions that address the future. Specifically, I focus on how individuals make decisions that yield benefits in the future and what kinds of cognitive capacities and representations of the future are needed for these decisions.

Mechanisms of Decision Making

Broadly defined, a decision is the result of an evaluation of possible options. This definition does not commit to a particular process (conscious or otherwise) and can be applied across a wide range of taxa (potentially to plants as well; Kacelnik 2003). Decisions 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? Though I distinguish between inferences and preferences as separate entities, they can interact such that inferences can feed into preference decisions and vice versa.

Approaches to Rationality

The nature of rational decision making has been debated for centuries. Over this time, two perspectives on decision making have emerged: unbounded rationality and bounded rationality.

Unbounded Rationality

Historically, many models of decision making have been based on the *Homo economicus* or “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 typically skirt claims about the process of decision making by stating that agents behave “as if” they are rational (Berg and Gigerenzer 2010). 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” behavior.

The unbounded rationality models 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 2008). Yet, typically these models are feasible only in specific, “small-world” circumstances (Savage 1954), and the generality of their application remains unclear. Moreover, even if organisms possess the ability to use higher-order cognitive skills such as optimal decision making, they do not necessarily do so when simpler solutions will suffice. For instance, cotton-top tamarins (*Saguinus oedipus*) use simpler, more approximate amount-based mechanisms when discriminating different quantities of food, even though they can use more sophisticated and precise number-based mechanisms in other situations (Stevens et al. 2007). Thus, though unboundedly rational models are mathematically tractable and elegant, they do not offer realistic accounts of decision-making mechanisms in complex environments.

Bounded Rationality

An alternative to the omniscience and unlimited computational power required of *Homo economicus* is a perspective that emphasizes 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 (Gigerenzer and Selten 2001; Simon 1956). 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 and Gaissmaier 2010). That is, rather than having general-purpose statistical devices that require extensive information and complicated computations, decision makers often succeed by

using less information and simple heuristics specifically adapted to their environment. The simple-heuristics approach makes explicit predictions about the decision process, the outcomes, and the conditions under which heuristics will work.

An evolutionary perspective on decision making highlights the gap between unbounded and bounded rationality. Many models of decision making in animals use optimization to find the best solution to a decision problem. However, despite using unboundedly rational models, behavioral biologists do not suggest that animals use optimal decision mechanisms. Optimization models are used only because natural selection approximates an optimizing process under constraints. Instead, biologists often assume that animals use rules of thumb (heuristics) that approach optimal outcomes. Consequently, the evolutionary perspective on decision making distinguishes between optimal outcomes and feasible mechanisms that can approach those outcomes.

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; Stevens and King 2011). As an example, biologists have investigated the use of simple rules in the “patch-choice” model of foraging (Stephens and Krebs 1986; Wilke et al. 2009) described in the chimpanzee termite-fishing example. Recall that in this scenario, foragers 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 may be computationally difficult in complex environments. A number of researchers have proposed simple patch-leaving rules that avoid some of the complicated computations (Figure 7.1). 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 even simpler rules dispense with the requirement of directly monitoring the current intake rate and instead indirectly estimate this rate. Animals using these rules may leave after consuming a certain number of prey items (fixed number rule), after a certain time period after arriving to a patch (fixed time rule), or after a certain time period of unsuccessful foraging (giving-up time rule). Empirical evidence suggests that different 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 animals is not 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 argue that in certain environments heuristics can achieve good outcomes. Given that the human brain has been built by evolution through natural selection, we might suppose that the costs of decision computations weigh heavily in the evolution of decision mechanisms, and mechanisms with simple

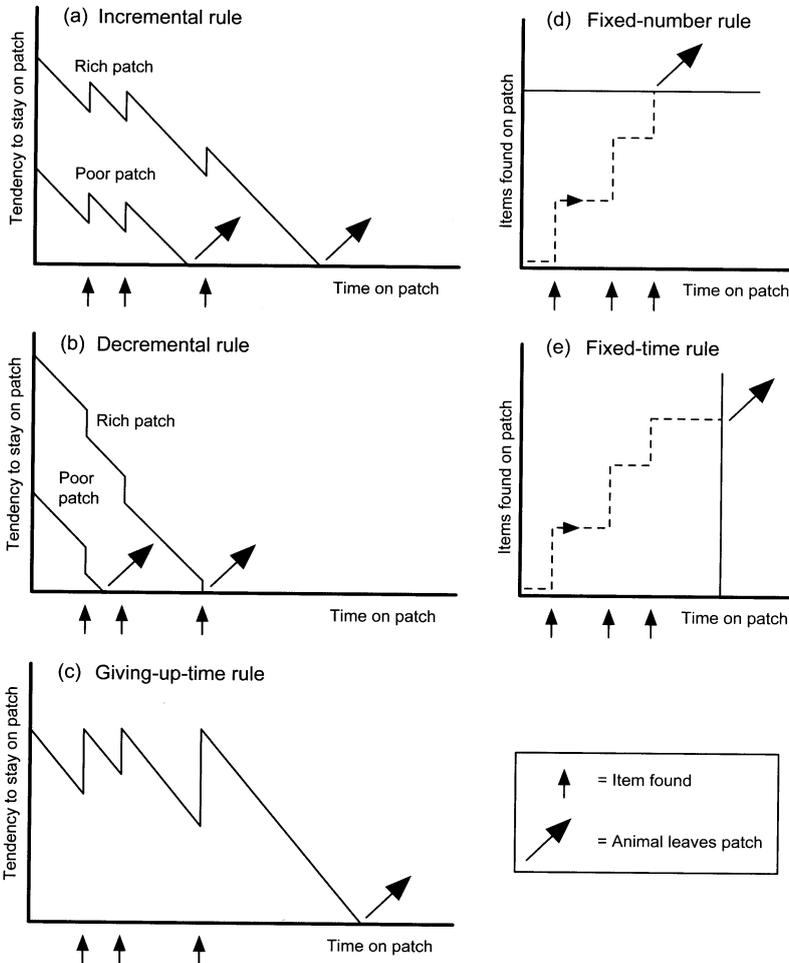


Figure 7.1 Biologists have tested a number of patch-leaving rules. (a) With an incremental rule for deciding when to leave a patch, each resource capture (indicated by small arrows) increases the probability of staying in a patch. (b) With a decremental rule, each resource capture reduces the probability of staying. (c) With a giving-up time rule, the tendency to stay in the patch declines with unsuccessful search and is reset to a maximum with each resource found. (d) With a fixed-number rule, a patch is left after a fixed number of items have been found. (e) With a fixed-time rule, the patch is left independent of the number of food items found. Reprinted with permission of the Cognitive Science Society from Wilke et al. (2009).

decision rules tend to prevail over complex computations when yielding similar outcomes.

The bounded rationality approach involves not only an exploration of heuristics and other decision mechanisms but also an investigation of the cognitive

capacities that underlie these mechanisms. This is particularly useful when studying the evolution of decision making, because species differ in their underlying capacities and therefore may differ in which decision mechanisms they can implement. Consequently, carefully outlining the required capacities is critical for studying the mechanism. As an example, though tit-for-tat and similar strategies have been promoted as simple decision rules that can generate cooperation (Axelrod and Hamilton 1981; Nowak 2006), the underlying capacities needed to implement these strategies have not been well studied (Stevens and Hauser 2004). When measuring memory capacity, for instance, it seems unlikely that even humans have the memory structure and accuracy required to implement tit-for-tat (Stevens et al. 2011). Thus, consideration of the underlying cognition needed for decision rules will help constrain the possible list of feasible mechanisms.

To summarize, the unbounded rationality approach to decision making focuses on optimal outcomes, whereas the bounded rationality approach emphasizes the cognitive mechanisms of decision making. In addition to testing decision heuristics, the bounded rationality approach highlights the importance of investigating the underlying cognitive capacities needed for decision mechanisms. With this general overview of bounded rationality in hand, we can now focus on a more specific class of decisions.

Making Decisions about the Future

Most of the decisions made by animals, humans included, involve some aspects of the future. Individuals must make inferences about the future (e.g., predicting the presence of a predator at a goal destination) as well as preferences about the future (e.g., investing in social partners to achieve future rewards). Here, I focus on a subset of preferences about the future known as *intertemporal choices*; that is, choices between options with future rewards (Read 2004; Stevens 2010b).

The termite-fishing chimpanzee introduced earlier in this chapter faces an intertemporal choice. Should she continue fishing in the current termite mound to extract more food or move on to another mound? From searching for food and mates to investing in territory, offspring, and social partners, intertemporal choices are ubiquitous in animal decision making. Researchers have studied these kinds of questions under a host of different names: delayed gratification, impulsivity, patience, self-control, temporal discounting. The key feature of these decisions is that animals act now to influence their future state, sometimes at a cost to their current state. Can animals forgo immediate benefits for delayed ones?

The Psychology of Intertemporal Choice

Some of the most amazing instances of waiting for delayed rewards in animals come from species that cache food for the winter. Nutcrackers (*Nucifraga*

columbiana) can bury 33,000 pine seeds each winter and wait months to recover them (Vander Wall and Balda 1977). Every year, these birds make thousands of intertemporal choices in which they choose between an immediate, smaller (relative) payoff versus a larger (relative) payoff in the future. How do they and the myriad other caching species make these decisions? Do they have a concept of the harsh conditions facing them in the coming months and plan accordingly? Or are there other ways to solve this problem? Animals can make intertemporal choices in a variety of ways. Indeed, many kinds of intertemporal choices do not require a representation of the future or any kind of planning ability. Though there are likely more, below I explore four types of mechanisms that may yield intertemporal choices: simple rules, reinforcement decay, reinforcement rate, and temporal discounting.

Simple Rules

In many cases, animals may use simple rules to make intertemporal choices. These rules do not require a representation of the future, and they may not even need any estimates of time. This is likely the case in many instances of caching. Rather than anticipate the future dearth of food, caching species implement simple strategies that respond to salient environmental variables. For caching, this probably entails a propensity to cache modulated by hormonal variations that result from changes in day length. Indeed, in the laboratory, experimenters elicit caching behavior in seasonal caching species by reducing the daily light:dark ratio (e.g., Pravosudov et al. 2010). This is not to say that caching is not flexible. On the contrary, much of the evidence of caching in corvids suggests remarkable flexibility in their caching decisions, depending on social context and the caching environment (Clayton et al. 2005; Emery and Clayton 2001).

As another example, parasitoid wasps (*Leptopilina heterotoma*) lay more eggs in lower-quality hosts when an impending storm is coming (Roitberg et al. 1992). Rather than assessing the future uncertainty of the storm and dumping eggs in anticipation of possibly losing the opportunity to lay, the wasps respond directly to manipulations of barometric pressure. Thus, various species use rather simple rules to convert environmental input into intertemporal choices. In these situations, the animals do not represent the future or any elements of time or reward magnitude. Instead, they respond rather directly to environmental cues.

Reinforcement Decay

In their natural habitats, animals continually face intertemporal choices. Yet, most of the work on intertemporal choice in animals occurs in the laboratory under operant conditions (e.g., Green and Myerson 2004). These studies often

present two stimuli to signal the options. For instance, choosing a blue circle results in a smaller food amount available sooner, and choosing a yellow square yields a larger food amount available later. Laboratory studies offer the advantages of tight control over the reward magnitudes and delays, allowing precise manipulation of relevant factors for intertemporal choices. The artificial nature of the task, however, allows alternative accounts of the phenomenon of interest. Rather than choosing between future rewards, the animals may simply be choosing between two reinforcement decay rates; that is, the strength of the association between the stimuli and the reward decays with the time since last reinforcement. Therefore, choices may be based on the past reinforcement history (decaying stimulus strength) instead of the future payoffs. Using operant paradigms does not necessitate the decay explanation, but it does offer an alternative that does not consider future states.

One solution is to replace the arbitrary stimuli with the actual rewards, so that subjects choose between the rewards rather than stimuli (Stevens et al. 2005c; Rosati et al. 2007). An additional solution is to switch from a purely temporal task to a spatial task. In these tasks, animals choose between smaller, closer rewards and larger, more distant ones (Stevens et al. 2005c; Mühloff et al. 2011). Therefore, the animals can see both the rewards and the cost required to obtain the rewards (the spatial distance). These kinds of spatial tasks mimic natural foraging problems (Janson 2007; Noser and Byrne 2007) and can mitigate reinforcement decay explanations of intertemporal choice.

Reinforcement Rate

As an alternative to the reinforcement decay explanation, animals may be attending to reinforcement rate or intake rate; that is, the number of rewards per unit time. To use these rates, individuals must have some estimate of the reward magnitudes and time delays. Research in numerical competence and timing suggests that animals can estimate both quantity and time (Brannon 2006; Gibbon 1977; Haun et al. 2010). Moreover, combining quantity and time into a rate is a core principle of behavioral ecological studies of foraging (Stephens and Krebs 1986) and psychological accounts of learning (Skinner 1938; Gallistel 1990).

Animals can use at least two types of reinforcement rates. Short-term rates focus only on the delay from choice to reward acquisition, whereas long-term rates include the time it takes to consume rewards as well as the time between trials. Intertemporal choice studies using blue jays (*Cyanocitta cristata*) and cotton-top tamarins (*S. oedipus*) are consistent with using short-term rates (Stephens and Anderson 2001; Stevens et al. 2005b), whereas choices of bonobos (*Pan paniscus*) may be consistent with long-term rates (Rosati et al. 2007).

After the rates are estimated, two different decision rules can be employed to make a choice: maximizing and matching. Maximizing predicts that individuals will compare the rates of available options and exclusively choose the

option with the higher rate. Matching, in contrast, predicts the distribution of choices in proportion to the relative rates across the options. If, for example, one option offered a reinforcement rate that is twice as high as that from another option, the higher rate option would be chosen in approximately 67% of the choices (Herrnstein 1961). An extended debate on the importance of matching versus maximizing remains unresolved (Commons et al. 1982; Logue et al. 1990) and has not been properly separated from the short- versus long-term rate issue.

Temporal Discounting

Probably the most commonly studied explanation of intertemporal choice is temporal discounting (Frederick et al. 2002). The discounting approach suggests that the present value of a future reward decreases with the delay to receiving that reward. The discounted value function requires estimates of both reward magnitude and future delay, so this mechanism requires an understanding of the future. In fact, some argue that temporal discounting and future planning or prospective memory are intertwined (Critchfield and Madden 2007).

Much of the work on animal intertemporal choice attributes the choice to discounting without testing whether a discounting mechanism is at work (Green and Myerson 2004). Indeed, the work on human intertemporal choice also relies on the discounting explanation. When making a choice between, say, \$100 today and \$105 in 3 weeks, economists suggest that people discount the value of the future reward. The discounting approach assesses how the present value of the delayed reward decreases with the time to receiving the reward. In both humans and other animals, this value appears to decrease with time delay in a hyperbolic way; that is, the rate of discounting decreases as the delay increases (Figure 7.2).

This hyperbolic pattern of choices, though consistent with a discounting explanation, is not unique to discounting. In fact, the reinforcement rate approach also produces behavior consistent with a hyperbolic decrease in value (Figure 7.2). In addition, a simple rule can account for this pattern, a rule that humans could be using. Rubinstein (2003) and Leland (2002) suggest that rather than making intertemporal choices based on discounting, humans may compare the two reward amounts or the two time delays and assess their similarity. If one attribute (amount or delay) is similar but the other not, then the decision maker would ignore the similar attribute and just focus on the other one. In the monetary example, for instance, one might consider \$100 and \$105 similar, thereby using the delays to decide and choosing the sooner option. When this rule can be used, it can outperform the discounting models in some situations (Stevens 2009), suggesting that simple rules may account for some instances of discounting-like behavior.

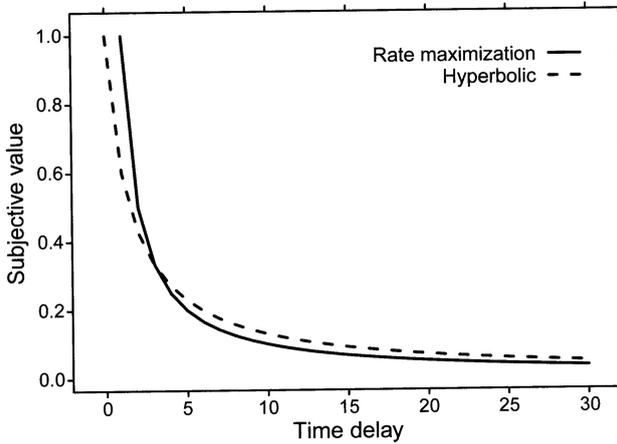


Figure 7.2 Temporal discounting explanations assume that individuals devalue future rewards. Though the hyperbolic discounting model fits animal data quite well, other mechanisms of intertemporal choice, such as rate maximization, also show a hyperbolic pattern.

Intertemporal Choice Mechanisms

Numerous mechanisms exist that allow organisms to make intertemporal choices, and I have provided by no means an exhaustive list. The critical point is that considering the mechanism has important implications for how we study intertemporal choice. First, it influences what kinds of models are relevant. The bulk of work on intertemporal choice uses temporal discounting models to investigate this behavior, both for humans and other animals. In the experimental paradigms of intertemporal choice, however, alternative nondiscounting mechanisms may account for these choices. When given binary forced choices between two options, humans may use similarity rather than discounting. Similarly, animals may use reinforcement rate or decay to make their choices in operant experiments. Again, though most studies of intertemporal choice assume temporal discounting, few spell out and test a clear discounting mechanism.

A mechanistic approach also highlights the underlying cognitive capacities needed to implement the mechanisms. Simple rules used by the parasitoid wasp require only a barometer. The similarity rule proposed for human intertemporal choices needs the ability to categorize amounts and times as similar or different. Temporal discounting involves an understanding of the future, but the type of understanding is not well studied. How do individuals assess the future? How is the future incorporated into the decision-making mechanism? Which types of future cognition are required and which types are optional? For instance, though mental time travel might be a useful capacity to have when implementing temporal discounting decisions, it may not be necessary

for them to be made. However, other theories of intertemporal choice frame the choices as a problem of “multiple selves” in which one must mentally travel in time to determine which future self would be best (Frederick et al. 2002). Clearly, these models require a different set of cognitive abilities.

Though formulating different decision mechanisms is useful, empirically testing between these mechanisms may be difficult. In many cases, they make similar predictions at the behavioral level. Thus, purely behavioral outcome measures may not suffice to discriminate among them. Instead, measures of process are needed to test between mechanisms (Schulte-Mecklenbeck et al. 2010). These measures may include reaction times, eye tracking, active information search, and physiological and neurological measures. Shapiro, Siller, and Kacelnik (2008), for instance, propose a model of intertemporal choice that successfully uses process data to predict choices in starlings (*Sturnus vulgaris*). Their sequential choice model takes reaction times (response latencies) to a single option in the absence of choice to predict choices when multiple options are presented simultaneously. At the moment, few intertemporal choice models are well specified enough at the mechanistic level to provide appropriate process predictions. The sequential choice model provides a nice example of the importance of incorporating decision mechanisms into models of intertemporal choice.

Summary

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. The bounded rationality approach offers a promising perspective that highlights the importance of studying simple mechanisms of decision making. In contrast to unbounded rationality’s assumptions of omniscient agents with limitless time and computational ability, bounded rationality begins with reasonable approximations of how the mind works. Boundedly rational agents often use rather simple heuristics when they perform well. Moreover, this approach emphasizes the cognitive capacities that must be in place to implement various decision mechanisms.

The bounded rationality approach can assist us in understanding decisions about the future. In particular, intertemporal choices involve future benefits, and a number of explanations have been proposed to account for these decisions. Yet, these explanations have not been fully worked out at the mechanistic level, and each mechanism requires different cognitive capacities and representations of the future. Temporal discounting seems to require some kind of understanding of the future, but other explanations can account for many instances of intertemporal choices without the need for a representation of the

future. More specific mechanistic models are needed to better understand how organisms make decisions about the future.

Acknowledgment

Parts of this chapter are based on work from an earlier Ernst Strüngmann Forum (Stevens 2008).