

Decision Making

What Can Evolution Do for Us?

Edward H. Hagen, Nick Chater, Charles R. Gallistel,
Alasdair Houston, Alex Kacelnik, Tobias Kalenscher,
Daniel Nettle, Danny Oppenheimer, and David W. Stephens

Do I contradict myself?
Very well then I contradict myself,
(I am large—I contain multitudes.)
—Walt Whitman, *Song of Myself* (1855)

Abstract

This chapter examines the contributions that evolutionary theory can make to an integrated science of decision making. It begins with a discussion of classical decision theory and analyzes the conceptual and empirical failures of this approach. Mechanistic explanations, which do not explicitly invoke evolutionary arguments, are presented to account for these failures. Thereafter, evolutionary approaches to decision making are examined and the failures revisited in light of evolutionary theory. It is shown that in some cases “irrational” behavior might be adaptive. The chapter concludes by exploring the open questions, levels of analysis, and policy implications that an evolutionary approach can bring to decision making.

The Bicycle

Imagine a group of aliens were to come to Earth, commissioned by their superiors with producing an intelligible account (or explanation) of a class of objects that Earthlings classify as “bicycles.” The aliens could pick a set of such items and measure their thermodynamic properties, their conductivity, and other physical features, but unless they consider the purpose for which the bicycle was designed, the analysis will almost certainly fail to capture

the essential elements needed to produce a coherent explanation of why such things as foldable metal bicycles, rigid plastic bicycles, racing bicycles, and mountain bicycles are all called by the same name.

As this example illustrates, considering the process by which the object of study came about and the criteria for its design (the “purpose” of the object) greatly aids the discovery of answers to certain questions. In any science that uses biological organisms as objects of study, evolution by natural selection is that process. Decision making is a functional property of organisms. Comprehensive accounts of decision processes in organisms as varied as amoebas and humans are thus aided by consideration of evolution by natural selection.

There is an important caveat to this argument for an integration of evolutionary theorizing in the study of decision making and behavior. Returning to our alien narrative, imagine that the aliens’ mission was not to understand “bicycles” but “conductivity.” They may well choose to study bicycles because bicycles tend to contain metal bits, rubber bits, and plastic bits, which might make it a great system for exploring how conductivity operates across multiple materials. They might even succeed in producing a coherent theory of conductivity across materials, and for them the fact that bicycles have a purpose is now irrelevant to the success of their research program. They are not interested in the *bicycleness* of their objects of study.

Within the study of decision making there are topics that are analogous to each of these different kinds of alien missions. A researcher who designs financial products may notice that private investors systematically prefer products that have lesser variance in outcome, and this suggests to her that she can make money by selling products that take advantage of this tendency. For her, the goal is to understand behavior with respect to money, and there are many ways to go about this, not necessarily involving asking why, in an evolutionary sense, investors behave as they do. Forays into *prospect theory* (Kahneman and Tversky 1979) have made advances in identifying regularities by which humans make decisions (e.g., reference points and diminishing returns for both gains and losses) without appealing to evolutionary principles.

Progress in decision science is possible without reference to evolution, but given that many of the questions that decision scientists study involve biological agents (including, of course, human beings), it is greatly aided by it. There are two major advantages of taking evolution into account. First, evolved functionality is a suitable source of candidate hypotheses for decision mechanisms. Evolved functions, which biologists term *adaptations*, are a tiny and special subset of all conceivable functions. Adaptations evolved by natural selection, which means that they must have increased the reproduction of the organism (i.e., increased biological *fitness*). Bicycles, although a highly functional form of human-powered transportation, are *not* adaptations.

For a more pertinent example, it is clear that evolution will not produce decision mechanisms that lead to maximization of lifetime accumulation of

resources per se. Natural selection favors the accumulation of resources only insofar as resources contribute to fitness. If a given resource (say money) has a nonlinear relation to fitness, then there is no reason to expect that decision processes will lead to maximization of lifetime wealth accumulation (McNamara and Houston 1992). This may contradict intuitive expectations of people raised in capitalist societies, but is immediately apparent when we take evolution seriously.

A second important advantage is that since evolution acts on the norm of reaction of organisms as a whole and within specific social and environmental scenarios, the integration of decision studies across sciences such as anthropology, economics, psychology, ecology, and neurophysiology is a natural contribution of evolutionary theorizing.

In this chapter, we examine the contributions that evolutionary theory can make to an integrated science of decision making. We begin with a brief introduction to classical decision theory, followed by a discussion of its conceptual and empirical failures. We then examine some mechanistic explanations for these failures that do not explicitly invoke evolutionary arguments. Thereafter, evolutionary approaches to decision making are introduced. We revisit the failures of axiomatic decision theory in light of evolutionary theory, revealing that in some cases “irrational” behavior might be adaptive. In conclusion, we explore open questions, levels of analysis, and policy implications of an evolutionary approach to decision making.

Axiomatic Decision Theory

Decision theory aims to understand how agents—usually humans and non-human animals, but also microorganisms, plants, and artificial life—pursue goals in the face of options. Examples of goals include maximizing happiness, wealth, or calorie intake, and examples of corresponding options include choosing among different careers, investment opportunities, or berry bushes.

The foundations of decision theory were laid in the 17th century in a series of letters between Blaise Pascal and Pierre Fermat, who discussed the problem of dividing stakes between two gamblers whose game is interrupted before its close. To illustrate the problem and its solution, imagine a game with two players, Peter and Paul, who have staked equal money on being the first to win 3 points by tossing a fair coin. Peter wins a point if the coin lands heads, and Paul a point if it lands tails. How should the stakes be divided if the game is stopped when Peter has two points and Paul one?

Because the coin is fair, the players have an equal chance to win the next point. If Peter won the next point, he would win the entire stakes, so he is entitled to at least half the stakes. If Paul won, the players are tied, so each would have an equal chance to win the entire stakes on the next toss; in that case, each is entitled to half the stakes, meaning Peter is entitled to 3/4 of the stakes

overall. In other words, the value of the unfinished game to each player is the sum over the values of each possible outcome (i.e., the stakes), each multiplied by the chances that it will occur. A key generalization that emerged from this discussion is that agents should maximize *expected value* (EV):

$$EV = \sum p_i x_i, \quad (7.1)$$

where p_i and x_i are the probability and amount of money, respectively, associated with each possible outcome ($i = 1, \dots, n$) of that option.

A Concave Utility Function

However elegant, the prescription to maximize expected value raises the so-called St. Petersburg paradox, famously posed by the mathematician Nicolas Bernoulli. Consider the following game of chance: to play, the player pays a fixed fee up front, x , and then repeatedly tosses a fair coin until the first head appears, receiving $2^n x$ if a head comes up on the n^{th} toss. The expected value of the game is therefore:

$$\begin{aligned} EV &= \frac{1}{2} \cdot 2x + \frac{1}{4} \cdot 4x + \frac{1}{8} \cdot 8x + \frac{1}{16} \cdot 16x + \dots \\ &= x + x + x + x + \dots \\ &= \infty. \end{aligned} \quad (7.2)$$

Hence, players should be willing to pay any amount to play, yet common sense suggests that most people would only pay a few dollars to play. The most influential solution to this paradox was offered by Nicolas' cousin, Daniel Bernoulli, who argued that the value of an outcome should not be judged on monetary value, but instead on a concave function of this value, termed *utility*, which reflects the diminishing marginal utility of money (see Figure 7.1). In Bernoulli's (1738/1954) words:

[T]he determination of the value of an item must not be based on its price, but rather on the utility it yields. The price of the item is dependent only on the thing itself and is equal for everyone; the utility, however, is dependent on the particular circumstances of the person making the estimate. Thus there is no doubt that a gain of one thousand ducats is more significant to a pauper than to a rich man though both gain the same amount.

The diminishing marginal utility of money reduces the values of the later terms of Equation 7.2 so that expected utility converges to a finite number, resolving the St. Petersburg paradox. It is also important to point out that individuals with concave utility functions are risk averse. (There have been many attempts to resolve the St. Petersburg paradox; for a review, see Martin 2008.)




Figure 7.1 A typical utility function.

The von Neumann and Morgenstern Axioms

The ideas of Pascal, Fermat, and Bernoulli were axiomatized by von Neumann and Morgenstern (1947) as follows. Agents face mutually exclusive outcomes A_i . A lottery L is:

$$L = \sum p_i A_i, \quad (7.3)$$

where p_i represents probabilities that sum to 1. Agents choose lotteries (not outcomes!). A preference for lottery L over lottery M is denoted $L \succ M$. There are four axioms:

1. **Completeness:** for any two lotteries (L, M) exactly one of the following holds: $L \succ M$, $L \prec M$, or $L = M$. In other words, when comparing two lotteries, agents are never undecided about their preference (although they might be indifferent).
2. **Transitivity:** If $L \preceq M$ and $M \preceq N$ then $L \preceq N$.
3. **Continuity:** If $L \preceq M \preceq N$ then there exists a probability $p \in [0,1]$ such that $pL + (1-p)N = M$.
4. **Independence:** If $L \preceq M$, then for any N and $p \in (0,1]$, $pL + (1-p)N \preceq pM + (1-p)N$. In other words, the preference for M over L is unaffected by the inclusion of N .

In the classical view, widely adopted in economics, agents whose decisions conform to these axioms are, by definition, *rational*.

States and Strategies

Agents' decisions will usually depend on the state of the agent (Mangel and Clark 1986; McNamara and Houston 1986). State, in this context, might include psychological variables representing information that the agent has about the

world (including information about past outcomes); physiological variables such as energy reserves; and morphological variables such as body size.

A strategy is a rule that specifies the action that should be performed in each state. It is also possible that, for a particular state, one of several behavioral options is chosen randomly. The sequence of behavior generated by a strategy might also be stochastic because the organism's state might not be exactly determined by behavior. For example, the strategy in a particular state of energy might be to forage, but the consequences of foraging, and hence the next energetic state, are stochastic.

This definition of strategy does not specify the process that generates strategies, nor does it even contain the idea of maximizing utility or fitness, or otherwise "doing well." What a strategy is good for depends on the process that generated it. In axiomatic decision theory, agent strategies are generated by "rational thought" with the aim of maximizing subjective utility. The mathematical analysis of such strategies is termed *game theory* (von Neumann and Morgenstern 1947).

Mapping Needs to Utility

The Need for a Common Currency

Agents typically have numerous, constantly recurring needs, such as food, water, and safety. Under the axioms of decision theory, there exists a utility function mapping preferences for lotteries to the real numbers. The need for such a common currency arises because decisions must be made whose consequences cannot compensate for the lack of other consequences. For example, no amount of nutrient ingestion will remedy a water shortage nor pass on genes to the next generation; no amount of water ingestion will compensate for a nutrient shortage (starvation) nor pass on genes to the next generation; and no amount of copulation will restore either a water deficit or an energy deficit.

Most models of decision making assume that the decision is based on whether a decision variable of some kind exceeds one threshold or another. The decision variable represents the strength of the evidence or the strength of the need. When evidence is commensurable (i.e., when evidence of one kind can be weighed against evidence of another kind), then combining the evidence into a single decision variable is not problematic. However, evidence of a need for nutrients cannot be weighed against the evidence of a need for water. In animals, including humans, there are many neurons and neural systems involved in decision making, and these neural systems engage in different tasks and may come to different conclusions (Kalenscher et al. 2010). Animals nonetheless must, and do, decide between going to the river and going to the orchard.

One way in which such agents could do so is by mapping all needs into a common currency, usually referred to as *subjective value* or *utility*. Transforming

the different options into utility allows them to be represented on a common scale and rank-ordered according to value/utility. The greater the difference between the values of the available options, the higher the propensity of an agent to choose the option with the highest value. If the analysis of the options each reproducibly (i.e., without noise) generated an output independent of the other options in the common currency (\pm some amount of subjective utility), then simply summing the outputs will yield a real-valued decision variable that is guaranteed to satisfy basic considerations like transitivity because the ordering of the real numbers is transitive.

Despite the theoretical advantages of a common currency, it is far from clear whether this is physically implemented in most animals (cf. Shizgal 1997) or other decision-making organisms.¹

Old and New Approaches to Mapping Needs to Utility in Economics

In economics, there have been several approaches to mapping agents' needs to utility, which we broadly characterize as "old" and "new." "Old" involved the Jeremy Bentham-ite idea that each possible outcome has a hedonic (pain/pleasure) value which has a magnitude and can be weighted by its probability, and with the decision rule being to choose whichever is most positive (pleasure) or least negative (pain) (Bentham 1789). The Benthamite individual thus has some kind of "hedonometer" to convert outcomes into this common currency for comparison.

Because of tremendous difficulties in fitting value functions to the great variety of choices expressed by different actors, Paul Samuelson (1938) proposed to refrain from making any assumptions about the content of utility functions (i.e., what people actually value and to what degree) but to infer individual preferences, whatever they are, from observed choices. Within this "revealed-preferences" approach (which we term "new"), people behave as if they maximized a utility function, whatever this is, as long as they meet certain consistency requirements, including transitivity, independence, and completeness. In short, the idea was to forget about how the individual did the processing and weighing up of options but focus instead on how a rational individual's actual behavior could be represented by substituting "utility" indices for outcomes and weighting these by (subjective) probabilities such that if *A* is observed to be chosen over *B* the indices and weights would give a higher expected utility for *A* than for *B*.

¹ It is even conceivable that "decision making" by plants employs a common currency. There has been a recent explosion of research on "computation" and "intelligence" in plants. Like animals, plants have sophisticated hormonal signaling networks, such as the ethylene, jasmonic acid, and salicylic acid pathways, which regulate growth and development and mediate responses to environmental stressors. These and many other signals must be integrated by the plant to reach "decisions" (Gilroy and Trewavas 2001; Trewavas 2005). It is not out of the question that, e.g., some signal molecule acts as a common currency in plants.

It is impossible, at least within the “new” framework of revealed preferences, to make any statements about utility in someone who does not adhere to these consistency requirements. Hence, in economics, consistency of choice, as indicated by transitivity, completeness, etc., is the gold standard against which the quality of a decision can be evaluated; it is the hallmark of economic rationality.

Irrationality and Exploitation

The worry about decision-making mechanisms that violate basic principles of rationality, such as transitivity, is that they expose the decision maker to potentially catastrophic exploitation. To illustrate, suppose agent X had intransitive preferences $A > B > C > A$, and agent Y possessed A . Agent Y could then sell A to agent X for $B + \varepsilon$, then sell B to agent X for $C + \varepsilon$ and then sell C to agent X for $A + \varepsilon$ (where ε is some small amount). Agent X has now given 3ε to agent Y and received nothing in return, a form of exploitation known as a “Dutch book.”

Failures of the Axiomatic Approach

Despite the mathematical elegance of axiomatic decision theory, and despite the vulnerability to exploitation of agents that violate these axioms, the axiomatic approach fails on both conceptual and empirical grounds. These failures, described next, threaten those behavioral sciences that have grounded their disciplines in axiomatic decision theory, with economics being the prime example. Our principal goal in this chapter is to assess the extent to which evolutionary theory can provide an accurate, formal theory of human (and nonhuman) behavior to replace axiomatic decision theory, a task taken up below (see section on The Evolutionary Theory of Decision Making).

Conceptual Failures

Axiomatic decision theory makes several problematic assumptions about decision-making agents (at least if those agents are taken to be humans or nonhuman organisms) and has important gaps. First, although agents are assumed to have complete, transitive preferences, the theory does not explain why agents have the preferences they do, or even why they have any preferences at all. Second, agents are assumed to be able to maximize their utility functions rapidly under all conditions. The theory thus ignores practical limits on time, information, and computational power.

Axiomatic decision theory also ignores the fact that decisions are typically made in a highly structured environment, and this structure can be exploited by agents to simplify the problems they face, thus allowing them to make “good”

decisions with minimal computation. Numerous experiments reveal decision-making biases that are often interpreted as errors. People are more likely to judge a statement as true if they have heard the statement before, for instance, regardless of the actual truth of the statement, a phenomenon referred to as the *truth effect* or *reiteration effect* (for a review, see Dechêne et al. 2010). Yet it is possible that such biases instead improve decisions by exploiting environmental structure.

Schooler and Hertwig (2005) and Reber and Unkelbach (2010) argue that the truth effect is mediated by cognitive fluency, such as retrieval fluency (the ease with which an object is remembered): repeated exposure to a statement increases the ease with which that statement is processed and this, in turn, increases the perception that the statement is true. The latter increase occurs, according to these authors, because statements which are true are more likely to be encountered than those that are false. Hence, cognitive fluency as a cue of truth is epistemically justified (for a review of the effects of fluency on decision making, see Oppenheimer 2008; for an argument that, for gossip statements, reiteration reduces error, thus increasing believability, see Hess and Hagen 2006).

Simon (1956, 1990) dubbed this alternative view of decision making *bounded rationality* and made an analogy with a pair of scissors: one blade represents the cognitive limitations of the decision maker and the other the structure of the environment. Understanding decision making requires understanding both blades (for further examples and a discussion of the relationship between decision making and environmental structure, see Gigerenzer et al. 1999).

Empirical Failures

Perhaps not surprisingly, there is abundant evidence that humans (Tversky 1969; Loomes et al. 1991; Grace 1993; Tversky and Simonson 1993; Kalenscher and Pennartz 2010; Kalenscher et al. 2010) and nonhuman animals (Navarick and Fantino 1972, 1974, 1975; Shafir 1994; Waite 2001; Shafir et al. 2002; Bateson et al. 2003) systematically and predictably violate transitivity and other assumptions of axiomatic decision theory. (Rieskamp et al. 2006 reviews much of this evidence and the extent to which it might be consistent with relaxed consistency assumptions.)

Moreover, the same neural systems that have been implicated in representing economic utility, identified by assuming consistency of choice (a manifestation of the utility function: “as-if becomes as-is”), are also implicated in representing local intransitive value in people making intransitive choices (Kalenscher et al. 2010). This suggests that the neural systems containing value signals do not necessitate transitivity to represent the attractiveness of one commodity over another; they thus do not work according to the requirements of a utility function.

To provide a detailed example of the empirical failure of axiomatic decision theory, we discuss intertemporal choice. We then describe mechanistic and state-dependent explanations of apparent rationality violations, followed by an introduction to evolutionary decision making theory and a reanalysis of these violations in light of it.

Intertemporal Choice

Starting with Paul Samuelson, economists in the 20th century came up with a prescriptive theory on how choices should be made between future outcomes (Samuelson 1937; Koopmans 1960; Lancaster 1963; Fishburn and Rubinstein 1982). This framework, *discounted utility theory*, posited that a decision maker behaves as if she maximized discounted utility (DU), with DU being the sum of the discount-factor-weighted utilities of all possible final states. DU assumed that the discount function, by which the (undiscounted) utilities of the outcomes are multiplied, decreases exponentially with time (Samuelson 1937):

$$f(t) = e^{-rt}, \quad (7.4)$$

where r is the discount rate. One of the important implications of exponential discounting is that the rate by which future rewards are devalued will be constant over time:

$$\frac{f(t + \Delta t)}{f(t)} = \frac{e^{-r(t+\Delta t)}}{e^{-rt}} = e^{-r\Delta t}. \quad (7.5)$$

In other words, if a glass of beer is valued half as much tomorrow as it is today, then, under exponential discounting, it is valued half as much a year and a day from now as it is a year from now.

Constant discounting has important implications for economic rationality and time-consistency of preference. According to DU, it is not irrational or nonoptimal per se to prefer small, short-term over large, long-term rewards, even if the preference for immediacy results in an overall reduced net gain over time. However, DU requires consistency over time. That is, if an individual prefers a small, short-term reward over a large, long-term reward and both rewards are shifted in time by an identical time interval, then the preference for the small, short-term reward should be preserved because both rewards should be discounted by the same rate. For example, it could be perfectly rational for a rock star to live fast and die young if he really accepts this consequence of the early deathbed. In contrast, DU would label behavior as time-inconsistent if a decision maker fails to act in accordance with his long-term interest. For instance, the failure to make appropriate retirement provisions would be irrational if the agent actually wishes to have a good and healthy lifestyle at old age.

Time-constant discounting, or more precisely, the stationarity axiom in DU (Koopmans 1960), predicts that the ranking of preferences between several

future outcomes should be preserved when the choice outcomes are deferred into the future by a fixed interval because the two outcomes should be discounted by the same fraction. A wealth of empirical studies suggests that all species tested, including humans, monkeys, pigeons, rats, mice, leeches and dragonflies, violate the principle of constant discounting and other implications of DU.

In a paradigmatic test of this prediction, human or nonhuman subjects first choose between a smaller, sooner reward (SS) or larger, later reward (LL), such as receiving \$1 today or \$2 tomorrow. A common delay is then added to both options, such as receiving \$1 in seven days or \$2 in eight days. An exceptionally reliable finding across humans and nonhuman animals (Chung and Herrnstein 1967; Rachlin and Green 1972; Ainslie 1974, 1975; Green et al. 1981; Thaler and Shefrin 1981; Logue 1988; Ben Zion et al. 1989; Loewenstein 1992; Green et al. 1994; Kirby and Herrnstein 1995; Green et al. 1997; Bennett 2002; Frederick et al. 2002; McClure et al. 2004; Rohde 2005; McClure et al. 2007) is that the tendency to choose SS declines dramatically with the increasing length of the added delay: although many people choose \$1 today over \$2 tomorrow, these same individuals pick \$2 in eight days over \$1 in seven days. This suggests that the prolongation of the delays resulted in a preference reversal even though the difference in delays remained identical (immediacy effect; Thaler and Shefrin 1981; Ben Zion et al. 1989).

If these experimental results are taken at face value, human and nonhuman agents appear to add extra value to immediate outcomes, a behavior best approximated by nonconstant discount functions, such as hyperbolic (Mazur 1984, 1988) or quasi-hyperbolic (Laibson 1997) functions. A hyperbolic discount function is of the form:

$$f(t) = \frac{1}{1+rt}, \quad (7.6)$$

where r is the discount rate. In contrast to exponential discounting, under hyperbolic discounting the rate by which future rewards are devalued is *not* constant over time. To illustrate this, consider that when $t = 0$ (i.e., the present):

$$\frac{f(t + \Delta t)}{f(t)} = \frac{1/(1+r(t + \Delta t))}{1/(1+rt)} = \frac{1}{1+r\Delta t}. \quad (7.7)$$

On the other hand, when $t \gg \Delta t$ (e.g., the distant future):

$$\frac{f(t + \Delta t)}{f(t)} \sim 1. \quad (7.8)$$

Under hyperbolic discounting, despite the fact that a glass of beer is valued half as much tomorrow as it is today, its value a year and day from now is (almost) equal to its value a year from now.

As noted earlier, because agents who use nonconstant discount functions and exhibit other violations of rationality assumptions are vulnerable to exploitation, and because such violations challenge axiomatic decision theory, considerable effort has been made to explain these violations.

One possibility, of course, is that some such violations are artifacts of the experimental procedures. For instance, experimental procedures might inadvertently induce changes in state, at least in nonhuman animals, producing apparent violations of transitivity. This could happen if training procedures meant to teach the animal about choices also altered feeding rates, and hence the animal's energetic state. Choices, which seem to reveal intransitive preferences, result instead from choices made in different states, which do not violate transitivity, and has been demonstrated in a study of starlings (Schuck-Paim et al. 2004).

Despite the possibility that some violations are actually experimental artifacts, the empirical evidence against axiomatic decision theory is strong enough to compel us to seek an alternative. We turn now to explanations of irrational decision making that do not explicitly invoke evolutionary arguments, such as those involving the mechanisms of decision making, followed by an explicitly evolutionary approach to decision making which, in some instances, identifies an adaptive logic underlying seemingly irrational decisions.

Mechanisms of Decision Making

Some of the observed failures of axiomatic decision theory could be the result of the cognitive architecture of decision making.

Voting Can Produce Intransitivity

One suggestion is that intransitivities are the consequence of a system where multiple independent decision mechanisms in the brain "vote" for a choice, and the winning choice becomes the agent's decision. Condorcet's voting paradox tells us that when decisions are made by aggregating over votes (rather than, e.g., summing real-valued outputs), and where there are three or more voters (in this case, decision-making mechanisms), the revealed preferences of the system as a whole may very well be intransitive, even though every voter has a transitive ordering of the options. When choosing among two of the three decision options by voting, the preferences (see Table 7.1) will yield $A > B > C > A$.

Kenneth Arrow (1950) generalized Condorcet's paradox in his *impossibility theorem*, which proved that no voting system can be designed that satisfies three fairness criteria:

Table 7.1 Preferences illustrating Condorcet’s paradox, where *A*, *B*, and *C* are decision options; I, II, and III are decision-making agents; and 1st, 2nd, and 3rd represent the ordinal preferences of each agent.

	I	II	III
<i>A</i>	1st	3rd	2nd
<i>B</i>	2nd	1st	3rd
<i>C</i>	3rd	2nd	1st

1. If every voter prefers alternative *X* over alternative *Y*, then the group prefers *X* over *Y*.
2. If every voter’s preference between *X* and *Y* remains unchanged, then the group’s preference between *X* and *Y* will also remain unchanged (even if voters’ preferences between other pairs like *X* and *Z*, *Y* and *Z*, or *Z* and *W* change).
3. No single voter possesses the power to always determine the group’s preference.

Hence, some behavioral tendencies might simply be by-products of the neural architecture (for a discussion, see Schneider et al. 2007).

Representation

Constraints on representation may be another central way in which cognitive architecture shapes decision making, including exhibiting intransitive preferences and other irrational behaviors. The concept of representation encompasses different phenomena according to different authors, so it is worth prefacing any discussion or example of its role in decision making with a minimum attempt at semantic precision.

By representation, we mean a system of symbols that is homomorphic² to another system. A representation has three essential properties (Gallistel 2001):

1. Reference: the symbols in the representing system must be causally connected to that to which they refer.
2. Homomorphism: operations on the symbols in the representing system must be homomorphic to processes and relations that obtain between their referents.
3. Functionality: the operations on the symbols in the representing system must at least occasionally direct behavior in ways consistent with the state of that aspect of the represented system to which the symbols refer.

Neurons involved in yaw correction in the housefly provide an example of all three properties (yaw is rotation around an axis perpendicular to the fly’s wings). First, large-field image motion (which under natural circumstances is a

² A homomorphism is a “structure-preserving” transformation.

measure of yaw) alters spike trains in two axons in the visual system of the fly, establishing causality. Second, a linear operation on the spike train data recovers an accurate representation of the yaw waveform, establishing homomorphism (Rieke et al. 1997). Third, there is good evidence that the spike trains in these neurons drive the yaw correcting responses of the flying fly, establishing behavioral function.

Whether “thoughts” (higher-level cognitive processes), in general, are best understood as computations over representations, or as emergent phenomena of the simultaneous operation of large numbers of simple, interconnected processing units (neurons, presumably, or groups of neurons) is, of course, a classic debate in the fields of cognitive science and the philosophy of mind. If the symbolic view of cognition is correct, then it has important implications for decision making in humans. Insofar as information-processing systems in non-human organisms, including microorganisms and plants, are also symbolic, the implications we describe next apply to them as well; in fact, some of the most compelling examples come from insects.

A decision is made between alternatives: between alternative possible states of the world or between alternative actions. The representation then specifies the set of alternatives. Thus, it constrains the decision. No decision can be taken on an alternative not specified in the representation. Consider, first, decisions as to which, among a range of possible states of the world, is in fact the currently prevailing state, a problem often solved with Bayesian analysis. A Bayesian decision maker has a representation of the possible states of some aspect of the world, which precedes and constrains the outcome of the decision process. Consider, second, decisions as to which action to take. Again, the animal cannot choose an action that it cannot represent. In reinforcement learning, the values that are computed attach to the symbols that represent the possible actions. No decision can be taken on an action that has not been symbolically represented.

This is a very general definition of the concept of representation as it applies to choices, one that encompasses a broad range of views regarding the veridicality of representations. At the “minimalist” extreme, a representation is simply a member of a set over which a nonzero probability distribution has been defined. The set and its members need not have much, if any, structure. The primary requirement is that, in the environments in which an agent typically finds itself, the set of representations helps generate functional behavior.

At the same time, this definition of representation can also accommodate highly structured representations of the world. If an agent solves a problem like finding its way in space, then it must somehow represent space, in the sense that it must compute solutions using “quantities” that represent attributes of the way-finding problem. For an agent that navigates along a chemical gradient, the representation of space could be as simple as “higher concentration” and “lower concentration.” At the other extreme, an agent that must achieve pinpoint accuracy in order to navigate long distances, such as a migrating bird,

might have a representation of space whose structure closely approximates the structure of Euclidean space (i.e., \mathbb{R}^3 with the standard metric). Much ink has been spilled on the question of whether nonhuman animals possess such a “cognitive map,” and although there are many threads to this story—from “short-cut style experiments” to the analysis of so-called place cells in the mammalian hippocampus—there does not seem to have been a satisfying resolution of this question (see, e.g., Collett and Graham 2004; Cruse and Wehner 2011; Dyer 1991; Gallistel and Cramer 1996; Gould 1986; McNaughton et al. 2006; Menzel et al. 2005; Tolman 1948).

State-Dependent Valuation Learning

The foregoing theoretical arguments make it clear that agents’ representations could shape their decision making, but they do not address the issue of which aspects of experience are first encoded and then recalled to drive choices. Given a certain experience (say an action followed some time later by access to a food item), the subject can encode perceptual information about the food item or the time that has lapsed between action and outcome, as well as encoding internally generated information such as changes in energetic state consequent on ingestion. It is possible to explore these issues experimentally.

Consider a subject that experiences two stimuli: red and blue. Responding to red leads to outcome A whereas responding to blue leads to outcome B . In one experiment, these stimuli are randomly assigned to two internal states, so that, for example, red is experienced when the individual is hungry (H) and blue when it is not hungry (N). We start with a case where the sizes of A and B are equal. Once the subject is acquainted with each stimulus and its consequences, choices between them are introduced in both states (H and N) and preference is measured. The questions involved are: Which is preferred (if any), and what kind of representation may explain such preference? The problem is described in Figure 7.2 (Aw et al. 2011).

Focusing on Figure 7.2b, let us consider several putative criteria for choice. The subject may:

1. Use representations of outcome sizes, or equivalently of the change they cause in energetic state, shown as ΔR_A and ΔR_B . Since these are equal, the subject would be indifferent.
2. Form distorted representations so that A is encoded as being bigger than B . This is shown in the cartoon representing the bird’s memory, where ΔR_A appears larger than ΔR_B . This may happen because the significance of the change in state is greater when the animal is hungrier. If this happens, the bird will prefer A (or the red stimulus that causes it in our example).
3. Use representations of the significance, or value, consequent of the changes, shown as ΔH_A and ΔH_B , leading again to preference for the

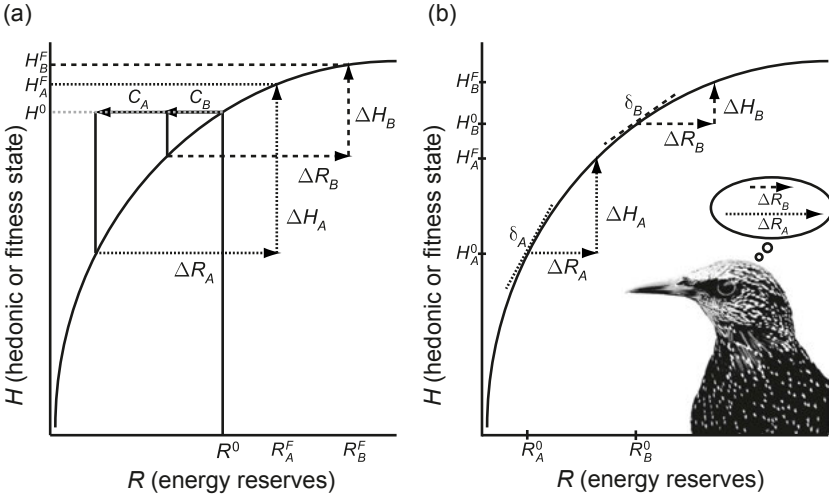


Figure 7.2 A nonlinear relation between state and a measure of subjective value. (a) Depiction of a situation where an animal in a given state, R^0 , pays different costs C_A and C_B for equally sized rewards, A and B , which then cause increases in reserves, ΔR_A and ΔR_B of equal magnitude, but with different consequences, ΔH_A and ΔH_B . (b) Depiction of a situation where the encounters with A and B occur when the subject's state is manipulated so that A is encountered in R_A^0 and B in R_B^0 , where the former refers to stronger hunger. In both cases, because of the concave shape of the relation between state and consequences, A and B differ in value gain despite being equal in size. After Aw et al. (2011).

stimulus causing A . This differs from second hypothesis because percepts are encoded here veridically; their different affective consequences are encoded and then used. Hypotheses 2 and 3 can be differentiated by designing experiments so that the subject is induced to reproduce its memory for each outcome rather than choosing between them.

4. Use representations of the slope of the value versus state function at the time of encounter with each stimulus (“red gives more value per unit of reward”). This appears as δA and δB . In this case, it would also lead to preference for A , but could be differentiated from the other hypotheses by using protocols where the magnitudes of A and B are not equal. Under hypothesis 4 the magnitude of the outcomes exerts no influence.

As we see, one hypothesis leads to indifference and three to preference for red; namely, the stimulus associated to greater hunger.

Such experiments were conducted in starlings, fish, and locusts (Aw et al. 2009, 2011; Marsh et al. 2004; Pompilio et al. 2006; Pompilio and Kacelnik 2005). In all three species, the subjects preferred the stimuli associated with a state of greater need, refuting the possibility that direct representations of the metrics of the outcomes, without any additional encoding of value, may be sufficient to explain choice.

Further experiments permit some level of differentiation between the surviving three hypotheses. If the fourth hypothesis were true, the subjects would be indifferent to the size of the outcome. We can discard this in the starlings because when the outcome metrics are manipulated, so that the stimulus associated with hunger takes values below that of the alternative, preferences for it decline and eventually reverse. This sensitivity to outcome magnitude eliminates (at least for the only species where it was tested) hypothesis 4.

Hypotheses 2 and 3 can also be differentiated for starlings. In experiments where outcomes *A* and *B* differed in the delay in delivery of food, it was shown that the birds' preferences were independent of having accurate representations for the delays in both options (Pompilio and Kacelnik 2005). They sometimes preferred responding to receive outcome *A* in spite of knowing accurately that its metrics were worse (longer delay) than in outcome *B*. For locusts, however, preexistent neurobiological evidence indicates that it is likely that the gain of receptors is adjusted, so that they would perceive the stimulus associated with hunger as being more salient than the alternative. Admittedly this last matter is not as categorically sorted as the former; however the point here is to illustrate how the nature of the representation and its effect on choice can be unfounded with suitable experimentation.

From an evolutionary standpoint, these studies provide an example of convergence to the same choice mechanism (i.e., state-dependent valuation learning) in a set of very distant organisms, raising the question of its adaptive significance (for an evolutionary model, see McNamara et al. 2012).

Multiple Attribute Problem

Choosing between options with many attributes is a fundamental problem of decision making and provides another illustration of the impact of representation. Broadly speaking, there are two approaches. One approach attempts to measure each option directly through a common currency. Suppose that we are attempting to decide which of two apples to eat, and each varies in both size and taste. Very crudely, according to this approach, we attempt to value the apple based on, say, size and taste, as having a certain utility (say 23 vs. 25 "utils"). The apple that is assigned the highest value is preferred. If people could stably associate such complex objects with values using such an internal currency, then choices would be highly stable and transitive. However, the problem with such "object-based" comparisons is that it is extremely difficult to know how to map complex objects onto utilities; indeed, when people are asked to do this explicitly using evaluations, the results are notoriously unstable (e.g., Hsee and Rottenstreich 2004).

Attribute-based comparison is an alternative approach. Here, the agent does not evaluate each apple separately. Instead, the apples are compared according to size (e.g., apple *A* is slightly bigger than apple *B*) as well as taste (apple *A* is a very much less tasty variety than apple *B*). Each attribute thus provides a

potential argument for or against the choice of apple *A* or *B*. The agent then attempts to weigh these “arguments”: the big difference in taste may overwhelm the small difference in size, and result in apple being chosen.

Almost all versions of attribute-based decision making will not be equivalent to any object-based approach. For example, if one object is valued slightly better on one dimension, and yet identical on all others, it will be chosen nearly always. From the attribute-based view, this is easily exploitable, because there is just one “argument” for one object, and none against it. By contrast, the choice between two objects, such that one is better on attribute *f* while the other is better on attribute *g*, will be highly unstable, because the agent has two almost equally good arguments to trade off (Loomes et al. 2012). Moreover, if the object has three attributes, which are appropriately arranged, then Condorcet’s paradox may arise (see earlier discussion).

Internal Conflict

Some decidedly irrational behaviors in humans and other animals appear to reflect “conflict” in the decision-making machinery. Rats offered both food and shock at the end of an alley, for instance, oscillate at a certain distance from them; this behavior could indicate conflict between approach and avoidance mechanisms. Human alcoholics willingly take drugs that will make them sick when they next consume alcohol. Livnat and Pippenger (2006) argue that if behavior is generated by a decision-making system that is subject to computational constraints, if conflict is defined in a particular way in terms of utility functions, and if parts of the decision-making machinery can be assigned utility functions based on information theoretic considerations, then even an optimal system designed for a single purpose can comprise agents that are in conflict with one another and may occasionally produce maladaptive behavior. In the rat case, this would involve approach and avoidance mechanisms that “selfishly” optimize their respective goals, with computational constraints preventing an optimal trade-off between the two goals; this results in maladaptive oscillations under certain parameters of food and shock.

The Evolutionary Theory of Decision Making

The synthesis of evolutionary theory and decision theory is achieved primarily by using *fitness*,³ rather than subjective utility, as the common currency to compare options (e.g., McNamara and Houston 1986), with fitness optimized by natural selection rather than rational thought. What evolves are strategies (not preferences!) that maximize fitness across the statistically variable sets of environments encountered by members of a population over evolutionary time. The

³ The number of descendants left far into the future.

mathematical analysis of such strategies is termed *evolutionary game theory* (Maynard Smith and Price 1973; Maynard Smith 1982; for reviews of the strategy concept in biology, see Hammerstein et al. 2006; Laubichler et al. 2005).

The substitution of fitness and natural selection for subjective utility and rational thought, respectively, overcomes many of the conceptual limitations of axiomatic decision theory. First, fitness is a core concept in the theory of natural selection. Unlike subjective utility, it has an exceptionally strong theoretical justification and applies to all forms of life, including viruses, bacteria, plants, and animals. Second, because fitness is optimized by natural selection in a population over many generations, and not by cognitive processes within the agent, agents can behave optimally even with extremely limited computational abilities. In fact, the strategy concept can be applied not only to behavior but also to ontogeny and morphology (e.g., Hagen and Hammerstein 2005). Third, because natural selection acts to adapt agents in a population to their environment, there is a fundamental relationship between decision making and environmental structure (for a review of the increasingly productive interaction between evolutionary biology, economics, and other social sciences, see Hammerstein and Hagen 2005).

This synthesis has implications that are not as widely recognized as they should be. In axiomatic decision theory, rational agents are characterized by their complete transitive preferences. The strategic behavior of agents in a particular context is the result of optimizing computations in the agents' brains, which takes place more-or-less in real time. In evolutionary game theory, in contrast, the agent is a strategy executor rather than a utility optimizer. In particular, evolved agents do not have a fitness utility function coded in their brains and do not compute how to optimize their fitness (the distinction between agents that optimize over preferences vs. those that simply execute strategies might be critical for interpreting the seemingly anomalous results of some behavioral economics experiments; Hagen and Hammerstein 2006).

Ecological, or Substantive, Rationality

We do not wish to infer that, in evolutionary models, agents never optimize in real or ontogenetic time. When they do, however, it is in a circumscribed socio-ecological domain (e.g., foraging or mate choice) that was important to fitness over many generations. This requires a mechanism that makes some (possibly erroneous) commitments to the nature of the reality about which decisions and choices are being made, an approach to decision making dubbed *ecological rationality* (Simon 1956; Gigerenzer et al. 1999) or *substantive rationality* (Chater, this volume).

The problem of characterizing the environments in which decision mechanisms evolved, and the environments in which they operate (and the former could well differ from the latter), comes up in many different guises. For example, in understanding the performance of a heuristic, one clearly needs to

know something about the range of situations in which the heuristic is applied. Moreover, a common and reasonable explanation of why we fail to observe globally optimal decisions is that the decision mechanism in question evolved to make decisions in a specific environment, and we would therefore hypothesize that this mechanism should perform at a high level in that environment even though it performs poorly in other environments—a problem often referred to as an “evolutionary mismatch.” For instance, a popular hypothesis holds that human obesity occurs because humans evolved in conditions of relative scarcity, which, the argument goes, selected for a tendency to favor foods high in fat and carbohydrate, leading to overeating in the modern environment where these foods are plentiful (e.g., Nesse and Williams 1996). In such cases, the behavioral anomalies studied provide insight into the actual evolutionary pressures to which the decision mechanism was adapted.

Despite its importance, characterizing the environments in which a mechanism evolved versus those in which it operates is often challenging and very seldom even attempted. Indeed, in most decision-making studies, consideration of “the environment” never goes beyond the specific situation presented to subjects.

The research program suggested here has three components: (a) selecting the environment to characterize, (b) choosing the relevant dimensions of the environment to measure, and (c) evaluating and interpreting the performance of a hypothesized mechanism in the now characterized environment. Each presents a different set of challenges.

First, there is the problem of selecting the appropriate environment to characterize. Should we study the current environment in which the animal lives or the environment it experienced in its selective past? From an evolutionary perspective, one is obviously most interested in the environment in which selection has acted, but identifying this could be quite difficult. In most cases, one would have to make some type of stationarity assumption (e.g., that the current environment fairly represents the “adaptation-relevant environment”). Stationarity could, of course, fail in several ways. A given lineage may have experienced a cyclic environment, for example, or its distribution might have covered a large geographic distribution such that the lineage has sampled a wide range of conditions.

Second, it is difficult to choose which attributes of the environment one needs to characterize. To make headway, we probably have to focus on a particular class of decisions. The best example stems from the study of the visual environment. In these studies, investigators have collected thousands of images of the human visual environment and characterized the properties, such as brightness and color distributions (Olshausen and Field 2000), allowing vision scientists to interpret the neural mechanisms of visual process in terms of the statistical regularities of the visual environment. This is the best developed example that we know, and we suggest that it should be a model for future studies of the decision environment. We should, perhaps, point out that this

“best example” was developed to study the mechanisms of visual perception rather than decision making.

A hypothetical example that is directly applicable to decision making is the problem of the foraging environment. Suppose we consider a classic foraging problem like patch exploitation (Charnov 1976; Stephens and Krebs 1986). We would want to know the probability distribution of patch types, the gain functions associated with these types, and distributions of travel times associated with these types. In general, we would not expect these properties to be independent of each other, so we would actually want to know, for example, the joint distribution of gain function properties and travel times. Moreover, it is clear that regularities, in the form of correlations between economic properties, are exactly the type of thing that might influence the evolved form of decision rules.

Assuming that one can solve the problems described above (and these problems should not be minimized), the problem of evaluating a decision mechanism, given a well-characterized selective environment, seems relatively straightforward. Yet, even here one must make decisions about the appropriate level of analysis. Crudely using this information to understand a hypothesized neural mechanism requires a different tool box than evaluating the performance of a simple patch-leaving rule.

Even when there are mechanisms, such as reinforcement learning, that solve many different classes of problems, there is a need for domain-specific adaptations. Reinforcement learning can only be deployed where selection has made relevant end-states rewarding, and prepared constraints on the types of actions or stimuli which can become associated with those rewards.

Clear examples of the role of preparedness come from conditioned taste aversions in omnivores (Garcia and Koelling 1966; Rozin and Kalat 1971; Holder et al. 1988; see also Hammerstein and Stevens, this volume). The relevant learning mechanisms make strong species-specific commitments to the nature of the cues that are likely to prove reliable in predicting the nutritive and pathogenic properties of the different foods that they must decide to ingest or not to ingest. Specifically, animals readily form associations between novel flavors and gastrointestinal malaise but fail to do so when it the color of the food predicts ingestibility.

Interestingly, vampire bats do not readily form conditioned taste aversions (Ratcliffe et al. 2003), although they are presumably able to learn associatively in other domains. Given that they do not encounter variation in the level of putrefaction of possible foods (since they ingest blood directly from healthy animals), this illustrates the role of natural selection in maintaining particular associative learning pathways for specific domains in particular lineages.

In another well-studied example, many animals, including insects, use the Sun as a directional referent. To do this, they must learn the solar ephemeris, the compass direction of the Sun as a function of the time of day. The

mechanism that learns this makes strong a priori commitments to the form of this function (Dyer and Dickinson 1994).

In studies of humans, economists routinely take for granted that people making economic decisions should and do assume that their environment contains other decision makers that have accurate knowledge of the relevant aspects of the world; decision makers who reason and make decisions in accord with principles of rationality identical to the principles that the decision maker herself employs.

Irrationality Revisited

Evolutionary decision theory provides powerful tools to help resolve the many apparent violations of rationality assumptions in humans and nonhuman animals. Perhaps foremost among these tools is the comparative approach: decision making can be studied in numerous, often distant, branches of the tree of life, including, in principle, organisms without nervous systems, such as plants and bacteria. If a violation of rationality assumptions, such as hyperbolic discounting, occurs systematically across taxa as well as across ecological contexts within taxa, this provides a vital clue to the nature of evolved decision mechanisms.

Below we revisit some apparent violations of rationality assumptions discussed earlier in light of the evolutionary theory of decision making. In some cases, we find that behavior which violates assumptions of axiomatic decision theory is, from an evolutionary perspective, adaptive.

Intertemporal Choice

How is it possible that evolution has shaped a behavior—hyperbolic discounting—that is so ubiquitous across domains and species, yet so apparently and dramatically maladaptive? Optimal foraging theory (Stephens and Krebs 1986) might provide some insight.

Optimal foraging theory makes predictions about the behavior of an animal foraging in an environment with patches of food that vary in location, density, quality, and other variables. The animal needs to decide whether to put its background activity—foraging—on hold to enter a food patch when encountered. Once the animal enters a patch, the marginal rate of energy return decreases with time because the patch is progressively depleted. The animal, therefore, needs to decide when to leave to seek the next patch.

When making such a decision, the animal has to take into account the current (decelerating) energy gain rate, and thus the time between energy unit consumptions in the current patch and the time (and effort and risk) needed to reach the next patch. Within the classical optimal foraging framework, animals have to trade off, among other things, delays and reward magnitudes. Hence, they face an intertemporal choice.

Optimal foraging theory proposes that animals maximizing Darwinian fitness should use foraging strategies that maximize the net energy gain per time unit over the long term. In formal terms, it is assumed that organisms maximize, at least in the long run, the ratio of food intake and the time needed to obtain and consume the food (Stephens and Krebs 1986). In a choice between large, delayed rewards and small, short-term rewards, rate maximization predicts that animals prefer large rewards when the ratio of reward amount per time unit is higher for the large than for the small reward.

To illustrate the implications for putative consistency violations, assume an animal chooses between (a) two food items delivered after 2 seconds (rate: one item per second) and (b) four food items delivered after 8 seconds (rate: 1/2 item per second); hence (a) > (b). Let us now play the economic game presented above and add a common time interval to both options: if both rewards were delayed by 10 seconds, the energy rate for option (a) would change to 0.17 items per second (two food items after 12 seconds) and for option (b) to 0.22 items per second (four food items after 18 seconds). Thus, in contrast to DU, optimal foraging theory predicts a preference reversal: the animal should prefer (a) over (b), but (b') over (a'). The same behavior that is labeled irrational in economics may be considered well-adapted in biology (e.g., Fawcett et al. 2011).

Experimental work by Stephens corroborates this notion. He has shown that the same animals that perform poorly in traditional intertemporal choice tasks, similar to the ones described above, perform remarkably well in economically equivalent tasks whose structure, however, resembles more the patch-like environment under which the intertemporal choice policies presumably evolved (Stephens and Anderson 2001; Stephens 2002, 2008; Stephens et al. 2004).

Apparent Violations of Transitivity

Houston et al. (2007b) use the foregoing evolutionary strategic framework to obtain *apparent violations of transitivity*. The term “apparent” is used because the violation appears if an observer takes the organism to be choosing between two options, ignoring the possibility that the organism is able to make *repeated* choices between options. The model involves a single state-dependent rule for choosing between two options, given that these options will always be available to the organism in the future.

Specifically, it is adaptive for a foraging animal to consider the alternative food sources when determining the value of a given source. If the value of one food source depends on the nature of the alternative food sources available, intransitive choices can occur. Assume an animal is foraging in an environment containing food sources that can be characterized, among others things, by reward magnitude, reward probability, and predation risk. For example, a starving animal may prefer a rich, but risky (high predation risk) food source (source *A*) over a poorer, but safer food source (source *B*), and may prefer *B*

over a rich, safer source that has, however, a low reward probability (source C ; thus preferring $A > B$ and $B > C$). Under certain circumstances, this animal may prefer the rich, safer source C with the low reward probability over the rich, but risky source A , thus preferring $C > A$, because the availability of the high-risk source A acts as an insurance against the outcome variance of option C . The animal could first try out the low-risk patch C , and, if no reward is found, it could still opt for source A , which yields a large reward with a high probability, albeit at a high predation risk.

At the level of these strategies there is no intransitivity; for any pair of options, there is a single best state-dependent strategy. Even if experiments do not involve repeated choices (as in this model), animals might use rules that evolved to cope with environments in which the options that are available to a decision maker persist into the future. Migrating birds are seen to use current food availability to predict future food availability (Houston 1997), for example, and humans might also behave as if interactions will be repeated (e.g., Hagen and Hammerstein 2006).

Multiple Attribute Problem

Russo and Doshier (1983) maintain that, in multiattribute choice (e.g., a decision between gambles varying in reward probability and magnitude), an agent comparing the levels of the attributes separately will fail to treat each option as an integrated whole, but will evaluate the available options much faster and with fewer errors than an agent integrating probability- and reward-representations into a single utility representation for each option. Assuming that the world is structured in such a way that attribute-based comparisons generally result in recommendations similar to utility-based comparisons, then attribute-based policies will be favored by evolution because of their supremacy over utility-based policies in terms of processing speed and accuracy, which, however, comes at the (possibly small) cost of occasional intransitivities.

These are not the only optimality accounts of intransitive choice (Houston 1991, 1997). They do, however, illustrate the gist of the idea that evolution may not have favored choice consistency, but the development of mental policies that were adapted to the environment in which they evolved.

Two Open Questions

Specificity versus Generality of Decision Making

The *heuristics-based approach* (known in behavioral ecology as the rules-of-thumb approach) correctly and usefully separates the functional from mechanistic aspects of decision processes, yet it tempts theorists to propose overly specific heuristics. For instance, there have been a number of heuristics proposed

for whether people find a message persuasive. These include the “audience response heuristic” (Axson et al. 1987), the “consensus heuristic” (Giner-Sorella and Chaiken 1997), and the “endorsement heuristic” (Forehand et al. 2004), all of which suggest that the more positively the message is received by others, the more positively it will be received by the target (but which vary based on the specific nature of who the “others” are). There is also the “likability heuristic” (Chaiken 1980)—people are more persuaded by likable speakers—and the “expertise heuristic” (Ratneshwar et al. 1987) which suggests that people are persuaded more by experts.

Instead of having a repertoire of extremely specific heuristics, it might be that there is a more general strategy that the highest validity, positively valenced information available in a given context is what determines persuasion. More generally, the most valid information to a given task is used to solve that task. Indeed, Shah and Oppenheimer (2008) have argued that there may be only five general forms of heuristics, and that other heuristics can be created as specific instantiations of a combination of these elemental heuristics.

Kacelnik (this volume) suggests that one robust, heritable, broad-domain algorithm might be when facing a novel problem, display an inherited broad-domain behavior, and then use outcomes to modify the response in the appropriate direction (basically, a loose version of reinforcement learning). As an example, he considers a puppy that is trying to reach a moving ball; initially it runs directly to the ball, thus having to correct its direction constantly and being highly inefficient. With experience, the puppy starts to anticipate these corrections and may converge to a constant angle of gaze control in its chases. The point here is that there was no need to postulate that a heuristic for constant gaze was selected through evolution. What was selected in this case was the very broad “learning through consequences” mechanism, which the puppy uses for such different situations as identifying which human is more likely to deliver a treat versus a kick, which plants sting, what size of puddle is jumpable, or how to catch a moving object.

By developing more general principles rather than a laundry list of heuristics, we have more broadly predictive theories as well as more plausible targets for natural selection.

Substantive versus Formal Rationality

Two types of mechanisms that might be present in the cognitive systems of humans and nonhuman animals can be distinguished: those which embody *substantive* and *formal* rationality. Mechanisms involving substantive rationality assist the organism in reasoning or decision making, by embodying contingent aspects of the world or fitness-relevant goals. For example, as we mentioned earlier, Olshausen and Field (2000) measured spatial statistics of natural images, and argued that the receptive fields in the early cortical visual processing area V1 may arise because they are optimized to encode the image using the

minimum number of active neurons. (Substantive rationality is similar to ecological rationality; we use the former term because we want to remain agnostic about some of the theoretical commitments of Gigerenzer and colleagues, with whom the latter term is closely associated.)

Substantive information may be more abstract, such as the apparently strong constraints on how bees learn to use the movement of the Sun in navigation (Wehner and Rossel 1985). It also may include procedures, such as nest-building or information concerning fitness-relevant goals of the organism, such as food and sex. This substantive information contributes to the fitness of the organism, in the context of its actual environment and ecological niche.

In contrast, there might be cognitive mechanisms that embody formal rationality; that is, mechanisms that do not themselves make commitments to the nature of reality but instead impose a priori consistency constraints that are critical to the adaptiveness of decisions. For example, virtually all decision making must take account of the uncertain (probabilistic) nature of the input on which the decision is based. Bayesian inference, the normative form of probabilistic inference, provides the formal rationality constraints for this problem and therefore yields optimal results. Insofar as one believes that evolution optimizes the properties of critical mechanisms, a biological decision-making machine should have mechanisms that implement Bayesian inference.

For Bayesian reasoning, the problem specificity, or substantive content, is found in the “support” (in the statistical sense of the word); that is, in the representation of the possibilities. Prior probabilities (or probability densities) and likelihoods are defined over this support. For example, in deciding where one is, the support is the vector space of one’s possible locations. In deciding on what course to follow, the support is the points on the unit circle (the possible compass directions in which one could head). In deciding on whether one is within the radius around one’s goal within which one should stop the straight run toward the goal and initiate the search pattern (Wehner and Srinivasan 1981), the support is the possible distances from the goal. In deciding whether to respond to a cue that may or may not predict the time of occurrence of an event of interest, such as the onset of food availability, the support is the measure of how well the onset of the predictor predicts the onset of the predictee. This support is the interval on the entropy dimension between 0 and the source entropy (because the mutual information between the predictor and the predictee cannot be negative and cannot be greater than the source entropy, the amount of uncertainty about when the next predictee might occur; Balsam and Gallistel 2009).

The prior probability distribution specifies the probabilities (or probability densities) of the different possibilities in the light of both analytic considerations and evidence that has already been processed. An appeal of the Bayesian formulation is that it naturally melds information instilled in the genes through evolutionary time (McNamara and Houston 1980) with information acquired by the animal. The likelihood function specifies the likelihood

of the possibilities in the light of the latest input to the processor (latest signal, newest data). The support for the likelihood function is the same as the support for the prior distribution; namely, the representation of the possibilities. The posterior likelihood function is the point-by-point product of the prior distribution and the likelihood function. It specifies the relative likelihood of different possible states of the world “all considered.” When normalized, it is the posterior probability distribution.

Some cognitive scientists and neuroscientists have suggested that computational mechanisms which carry out (some approximation to) Bayesian inference may be implicated across perception, language processing, and inference (e.g., Chater et al. 2006; Knill and Pouget 2004). Here we illustrate how formal Bayesian inference procedures could decide on behavioral outcomes. A given behavioral option will have various outcomes associated with it. Each outcome can be associated with a value which may be subjective or, in the case of evolutionary explanations of behavior, is its reproductive value. Once the prior information is combined with available information to form the posterior distribution, the posterior probability of each outcome can be computed. The expected (mean) value of the behavioral option can then be computed. This is the average value of the outcome, where the average is formed by weighting each value by its corresponding (posterior) probability of occurrence. When there are a range of possible behavioral options, the decision process will select the option with the greatest expected value.

Another foundational mechanism that operates across all domains is a system of the arithmetic processing of quantities; it enables the elaboration of the representations upon which decision making is based (Gallistel 2011). That is, arithmetic manipulation is necessary for constructing the model that mediates the decision, no matter what the substance of the model (Gallistel 2011). Navigational computations use arithmetic to establish vector space representations from which ranges (distances) and bearings (directions) may be computed. Computations of the mutual information between events distributed in time use arithmetic as well, but they do not use it to create vector spaces. Instead, they use it to compute entropies (amounts of information). In both domains, arithmetic is used to carry out the Bayesian inference (i.e., to take the product of the prior distribution and the likelihood function to marginalize or integrate the result) to take the ratio of different likelihoods (i.e., to form Bayes factors) and to multiply Bayes factors by prior odds.

Similarly, consistency conditions for linking beliefs, values, and actions are provided by various types of decision theory. Indeed, statistical decision theory has proved to be a powerful framework for understanding perceptual motor control and simple perceptual choice tasks (e.g., Bogacz et al. 2006).

The distinction between substantive and formal rationality raises a number of questions: Which organisms have mechanisms that capture different aspects of substantive and formal rationality? Which, if any, organisms embody mechanisms for formal rationality at all? How might the “tinkering” of evolution

(Jacob 1977), and perhaps also learning mechanisms, apply substantive or formal mechanisms outside their original domain. Could there be general principles underlying the inferential “machinery,” and their implementation in neural circuitry, which might be carried across many aspects of brain function? When should we attribute an organism with specific substantive rationality (e.g., concerning particular algorithms for navigation), and when is it appropriate to postulate, in addition, principles of formal rationality concerning, for example, the principles of Euclidean geometry, which might underpin such navigational strategies? Is it necessary that we postulate that substantive or formal information is represented? What neural or behavioral evidence might help in answering such questions?

Policy Implications

We have argued that consistency is not the only standard against which to evaluate the quality of a decision since what is irrational, from the classical point of view, could, from a biological perspective, be adaptive. However, we cannot escape the conclusion that time-inconsistent preferences are a problem for the individual and society, even if they make sense from an evolutionary perspective: if everyone were slave to their present bias, and there is evidence that a significant fraction of the population is, we would live in a society facing severe problems such as old-age poverty (as a result of the failure to take appropriate retirement provisions), severe health issues (as a consequence of the failure to pay regular health insurance premiums when not being ill), eating disorders (as a consequence of succumbing to the lure of tempting, yet unhealthy food), and financial illiteracy (as a consequence of the inability to deal with loans and credit cards).

Hence, to generate policy recommendations, society needs a normative standard against which the quality of a decision can be evaluated. The problem is that a theory which performs well in explaining, describing, and predicting behavior cannot tell us what people ought to do; evolutionary theory does not make policy suggestions. Thus, although we propose that all disciplines concerned with the explanation, description, and prediction of decision making (e.g., economics, psychology, behavioral ecology, and ethology) should eventually converge to a unifying framework on how decisions are made, this would not displace the need for a modern normative standard. A bold and certainly controversial proposition would be that the normatively flavored approaches in neoclassical economics, such as revealed preferences, are nothing but normative (i.e., they have no descriptive, explanatory or predictive value), and that other, more evolutionary flavored approaches replace them. If we adhered to this proposition, we would have two sets of theories: one that does well in describing, explaining, and predicting human (and animal) behavior,

and one that prescribes choice by telling us what we ought to do. Any discrepancy could then be used for policy intervention.

Caveats

To truly advance the science of decision making, an evolutionary approach will need to do a number of things. First, it must offer more than the concept of optimization, which is already broadly used in psychology, but without the evolutionary constraints (e.g., Anderson 1990). Second it must offer more than comparison with other species, which again, already influences decision theory. Third, it must go beyond simply documenting between-species differences in decision making by making testable predictions about both cognitive mechanisms and behavior that apply to a particular species, especially humans. Fourth, it needs to provide constraints that are generally instructive. An important theme of this chapter is that decision-making mechanisms were shaped by, and might take advantage of, the structure of the environments in which they evolved. When it comes to humans, some aspects of the ancestral environment are certain. For instance, women got pregnant and men did not—a fact that undoubtedly influenced the evolution of mating and parenting psychology. Other aspects of the ancestral human environment, such as sex differences in social status, are unknown and might never be known. A related challenge is that it is often difficult to distinguish between behaviors or cognitive patterns and strategies that have been learned versus those that evolved by natural selection.

Finally, even if evolutionary perspectives advance psychological theorizing, it is important to note that there are many other approaches that can also be generative. As such, we would not want evolutionary theorizing to replace or subsume other, useful approaches, but rather it should add to the theorist's toolbox.

Concluding Remarks: Levels of Analysis

As in any science, we seek a compact account of decision-making processes in evolved systems that is as simple as possible, but no simpler—one that is parsimonious, yet able to capture everything that is systematic in biological decision making. Such a theory would enable us to understand and anticipate the properties of decision-making mechanisms not yet investigated. It would enable us to understand the many seemingly paradoxical aspects of human decision making in personal, interpersonal, societal, and economic contexts.

The account we seek should extend across the levels of analyses of cognitive processes delineated by Marr (1982). It should describe the representations on which decisions are based, because, by specifying the alternatives or options between which the decision is to be made, they strongly constrain the results

of the decision. It should offer guidance on what constitutes an appropriate and functionally adequate representation of the alternatives. It should describe how the representations are computed and the computations that mediate the decision itself. This is Marr's algorithmic level of analysis. Finally, it should include a specification of the (mostly neural) mechanisms and processes that implement these algorithms.

We advocate an evolutionary account of decision making. Evolutionary thinking emphasizes function: many organismal structures are best explained as effective means to some end. Evolved decision-making mechanisms would therefore be designed to optimize some resource or outcome that was important to fitness (e.g., energy, mates, and offspring number and quality). Functional thinking, in turn, requires a focus on the structure of the organism's environment: the mechanisms which have been optimized by natural selection to function in particular environments. For decision-making mechanisms, this environmental structure would include the risks, time delays, and threats typical of the environment in which the organism evolved. Evolutionary explanation also involves the gradual modification of a biological lineage, and conserved structures from this lineage, from reflexes to computational patterns in neural circuits, might potentially have profound effects for modern human decision making.

Such a theory will provide a conceptual unification, bringing together the currently very disparate bodies of work in the many disciplines that study decision-making processes and mechanisms.

Acknowledgment

This chapter benefited greatly from discussions with Graham Loomes.