

Rational decision making in primates: the bounded and the ecological

Jeffrey R. Stevens

Center for Adaptive Behavior and Cognition
Max Planck Institute for Human Development

A young female rhesus macaque steals furtive glances at the male off to her right. He just arrived to the territory and therefore immediately piques her interest. The alpha male, however, sits a few meters off, basking in the sun. Being in estrus, the young female faces a choice: solicit a mating from the alpha male or follow the unfamiliar male into the brush to sneak a mating with him. Mating with the alpha male almost guarantees 'good genes' for her offspring. But something pushes her towards the unfamiliar male. Mating with him reduces the probability of inbreeding and adds a bit of genetic diversity to her offspring. Additionally, mating with the new male could act as an investment in the future: the current alpha male is getting old, and befriending a prospective alpha male could yield future benefits. In addition, spreading the possibility of paternity may secure protection for the offspring. Yet, this mating also involves risks. Males often vocalize while mating, which attracts the attention of other males. If the unfamiliar male vocalizes, the alpha male may attack the female. Should she take the safe option or risk punishment for possible future payoffs?

Primates constantly face decisions which influence their survival and reproduction. Continue foraging in this tree or move on to another? Expose oneself to a hidden predator by straying from the group or enjoy the safety of having other potential victims nearby? Defend one's territory from invaders or abandon it and seek a new home? In all of these cases, primates must tradeoff the costs and benefits associated with uncertain and delayed decision outcomes. The outcomes of these choices influence survival and reproduction, and natural selection should favor those individuals whose choices lead to the propagation of their genes.

The vast majority of economic analyses of decision making define good or 'rational' decisions as those consistent with a set of mathematical principles. Yet, this ignores the evolutionary pressures on decision making for the sake of mathematical elegance (Kacelnik, 2006; Stevens, 2008). Meanwhile, the standard psychological view of decision making seeks to empirically undermine the economic theory but cannot offer an alternative explanatory theory. Here, I emphasize an evolutionarily informed framework for studying decision making: the bounded and ecological rationality approaches (Gigerenzer et al., 1999). Though these approaches have traditionally focused on human decision making, they are just as relevant for other species, including other primates. To illustrate the relevance of bounded and ecological rationality to the study of primate decision making, I begin by introducing various visions of rationality found in the economic and psychological study of decision making. I then explore how primate studies inform three aspects of decision making: utility, uncertainty, and time. Together, these aspects will guide our understanding of the evolutionary origins of primate and human rationality.

Visions of rationality

Rationality means different things to different people. Kacelnik (2006) proposed that rationality refers to decisions which are either 1) *consistent with expected utility maximization* for economists and psychologists, 2) *consistent within the self* for philosophers, and 3) *consistent with fitness enhancement* for biologists. Of particular relevance here are the economic, psychological, and biological views. A review of these visions of rationality will frame the question, "Are primates rational decision makers?"

Rational choice

Which would you prefer: receiving two bananas with certainty or receiving either one or three bananas with equal probability? Depending on whether you like bananas and your level of hunger, this may be tricky to answer. These questions of decisions under risk mirror fundamental choices that we and other animals frequently face. Very little in life is certain, so all organisms choose between options without knowing the exact consequences.

Economists have approached the question of uncertainty by developing *expected utility theory* (reviewed in Wu et al., 2004). In expected utility theory, three features characterize all options: magnitude, utility, and

probability. Magnitude (x) refers to the amount of the benefit (or cost) associated with the option. Utility (u) is the mapping of magnitude onto some measure of satisfaction or 'goodness'¹. Finally, probability (p) is the chance that the outcome occurs. Thus, if one faces risky options, the expected utility is $EU = p \cdot u(x)$, where utility is some function of magnitude. Von Neumann and Morgenstern (1947) formalized expected utility theory to show that following specific mathematical principles maximizes expected utility. Therefore, expected utility acts as a normative standard for what decision makers should maximize when making risky choices.

Expected utility maximization assumes consistent choice, which requires a number of principles to hold (reviewed in Luce and Raiffa, 1957; Rieskamp et al., 2006b). First, choices must be transitive, meaning that a fixed order of preference exists. If an individual prefers bananas over apples and apples over oranges, then she must prefer bananas over oranges to maintain transitivity. Second, when transitive, choices are also independent from irrelevant alternatives, meaning the relative preference between options should not be affected by the presence or absence of other options. If bananas are preferred to apples, the addition of watermelons to the choice set should not affect the banana/apple preference. Finally, choices must be invariant, meaning that option A is preferred to B regardless of presentation format. If endowed with a banana and an apple and asked to give up one, preferences should be the same as if asked to choose freely between a banana and apple. Given these and a few other principles, one can show that preferences follow expected utility calculations.

Irrational choice?

Not long after von Neumann and Morgenstern (1947) published the principles of expected utility theory, cracks began appearing in the mathematically elegant framework when data showed violations of expected utility predictions (Allais, 1953; Ellsberg, 1961). Soon afterwards, Kahneman and Tversky started a cottage industry of demonstrating violations of the theoretical predictions (Tversky and Kahneman, 1981; Kahneman et al., 1982; Kahneman and Tversky, 2000). In both experimental settings and in real-world decisions, laypeople and experts made 'irrational' choices². Subjects showed intransitivity, irrelevant alternatives changed preferences, and the framing of decision questions greatly influenced preferences. Economists minimized the problem by calling these findings 'anomalies', whereas psychologists emphasized their robustness and labeled them 'biases'. To psychologists, expected utility theory was deeply flawed because it rested purely on mathematically derived principles and not what we know about human behavior.

Kahneman and Tversky (1979) injected a bit of psychological realism into decision theory when they proposed *prospect theory*. Instead of using a utility function characterized over total wealth, prospect theory uses a value function $v(x)$ relative to a reference point. So with 100 bananas in the bank, a decision maker would view a choice between gaining one and losing three bananas as a choice between 101 and 97 bananas from the expected utility perspective but a choice between gaining one and losing three from the prospect theory account. In addition to altering the reference point of the value function, prospect theory integrated data about probability perception into the equation by adding a decision weight function to the probability. Thus, instead of having a linear relationship between the objective probability and the perceived probability, prospect theory assumes that people tend to overestimate low probabilities and underestimate high probabilities. The probability, then, is weighted by the function $\pi(p)$. Thus, prospect theory predicts that preferences depend on which option has the highest value $V = \pi(p) \cdot v(x)$. As a descriptive theory, it nicely fits people's preferences (Kahneman and Tversky, 2000; Wu et al., 2004). However, it does not explain why we have reference points or nonlinear value functions—it takes these as givens and describes how they influence decisions.

Bounded and ecological rationality

Expected utility theory is mathematically elegant but fails to account for many of the data. Prospect theory fits much better descriptively but lacks explanatory power. Parameters in the models are fit to the data, with no a priori predictions about parameter values. Prospect theory therefore offers a slight modification of expected utility theory by patching a few of the holes that data have poked into the theory. But both theories face a more fundamental problem. Namely, neither of these theories adequately addresses two crucial components of decision making: the structure of the mind and the structure of the decision-making environment.

Early in the study of decision making, Simon (1955; 1956) highlighted not only the study of the mind but also the fit between the mind and the environment. He criticized the unrealistic assumption that decision makers have infinite time to decide, full knowledge of the problem, and unlimited computational resources to find an optimal solution to a decision problem. This vision of unbounded rationality contrasts sharply with what we know about human cognition and decision making, so Simon proposed the study of *bounded rationality*—the exploration of decision making given realistic assumptions about cognitive abilities. Real-world decision makers lack knowledge and cannot use optimization processes to make decisions. Thus, much previous research has ignored cognitive processes at work in decision making (but see Payne et al., 1993). The bounded rationality approach calls for realistic models of the decision process based on what we know about cognition rather than on a set of mathematical principles. Knowing the underlying process can help us better understand the decisions. Yet, Simon emphasized that studying the mind exclusively gives you only half of the picture.

To fully understand decision making, we must embed the mind in the environment. Gigerenzer and colleagues (1999) have termed this *ecological rationality*—the match between a decision mechanism and the environment. The unbounded rationality approach assumes that expected utility works in all decision making situations—it applies universally. Ecological rationality, however, appeals to the evolutionary idea that adaptations match the environment in which they evolved. Therefore, decision mechanisms should not be universal and domain general but specifically tailored to the environment in which they operate (Barkow et al., 1992). In fact, ecological rationality suggests that we do not possess a single, complex decision-making mechanism used in all contexts. Rather, we have an ‘adaptive toolbox’ of mechanisms (including simple heuristics or rules of thumb) that, when used in the appropriate environment, perform quite well (Gigerenzer et al., 1999). This perspective, then, offers an explanation for the anomalies and biases seen by the experimental economists and psychologists. Rather than being evidence for flawed thinking and irrationality, we are simply putting these decision mechanisms in an unfamiliar and artificial environment—the experimental laboratory. In general, however, our decision mechanisms serve us quite well by exploiting critical aspects of the environment.

Both bounded and ecological rationality offer appealing alternatives to the standard unbounded approach because they rest on realistic evolutionary principles instead of mathematical formalizations. Comparative analyses provide a unique method for testing questions of bounded and ecological rationality because we have great variability across species in their ecological environments. Primates offer an ideal group of species for these investigations because their phylogenetic proximity allows us to test interesting hypotheses about the evolution of human decision making. With this framework in hand, we can now review the bounded and ecological rationality of primate decisions. Note that here I focus on what has traditionally been termed ‘individual decision making’. Though it is likely impossible to completely extract an individual from the social environment, for the purposes of this review I will put aside the exciting and complicated world of social decision making in primates and refer the reader to Maestriperi or Hare and Santos, this volume. Instead, I will focus on three important components of individual decision making³: utility, uncertainty, and time.

Utility and preference

Utility refers to the mapping of the magnitude of a benefit or cost onto some measure of ‘goodness’ or goal achievement (Baron, 2000), and a utility function describes this mapping. Utility itself is difficult to assess, so typically we measure choices to infer preferences. For expected utility theory to work, preferences must follow the principles mentioned previously: transitivity, independence of irrelevant alternatives, and invariance. It is well established that humans violate these principles, deviating from the normative standard (Kahneman and Tversky, 2000). The question remains, “Is this the appropriate standard?” Examining choices in other species can address this question, because if other species also show the biases, then we likely are using the wrong standard. Do other primates violate these standards, suggesting deep evolutionary roots, or do only humans show these violations?

*Transitive preferences*⁴ are not well studied in primates. Though other species seem to show intransitivity (Shafir, 1994; Waite, 2001b), few data exist for primates, so we will not consider transitivity here⁵.

The principle of *independence of irrelevant alternatives* implies that previously available options should not influence the current preferences (Simonson and Tversky, 1992). Waite (2001a), however, found that the

background context does matter for gray jays (*Perisoreus canadensis*). When required to pay a lower relative cost for food in previous choices, the birds reduced their preference for that option later when required to pay a higher relative cost. Tinklepaugh (1928) found a similar result in rhesus macaques (*Macaca mulatta*) and long-tailed macaques (*Macaca fascicularis*). These monkeys observed an experimenter place a banana under one of two cups. The monkeys then immediately chose the correct cup and received the banana. However, in some trials, the experimenter substituted a piece of lettuce under the cup, unbeknownst to the monkeys. When the monkeys lifted the cups and found the lettuce, they rejected the less-preferred food item. Though outside of this context the monkeys readily consumed lettuce, they refused to eat it when expecting a banana. Thus, preferences are not fixed but depend on previous options.

Tinklepaugh's data suggest that the monkey do not have absolute preferences. If they did, they would always consume the lettuce because it is always better than nothing. However, the preference for consuming lettuce is *relative* to the previous availability of the highly preferred banana. The preference is relative to an expectation of other options. There are good, adaptive reasons for avoiding absolute preferences. For instance, it is well known in foraging theory that optimal choices depend on the background environment. If the environment is rich, animals should 'skim the cream' and choose to invest little time in extracting food from patches; if, however, the environment is poor, they should more thoroughly exploit the patches (Stephens and Krebs, 1986). Houston (1997) argued that since the current choice is incorporated into the estimate of the background environment for the next choice point, preferences can change for the same set of options depending on the previous background environment. Thus, relative preferences can result from an ecologically rational mechanism of adaptive decision making.

The classical economic approach to rationality also assumes that decision makers have strictly ordered preferences that are *invariant* to extraneous characteristics of the choice situation. If an agent prefers A to B, it should always prefer A, regardless of whether one is buying or selling A, or whether other choices are made before the A/B choice, etc. A common violation of invariance found in behavioral economics is the endowment effect. This phenomenon occurs when decision makers have a higher preference for an object when they own it (Kahneman et al., 1990). In humans, this is typically demonstrated by showing that subjects require a higher price to sell an object they possess than to buy the same object. Brosnan and colleagues (2007) tested similar effects in chimpanzees (*Pan troglodytes*) by offering them a choice between two different food items and recording their preferences. The experimenters then endowed the chimpanzees with one of the items and measured their willingness to trade for the other item. The choice preferences and trading preferences did not match, suggesting an endowment effect. Though anomalous to economists, an evolutionary perspective provides an explanation for the endowment effect. The question of why owners fight harder than intruders to maintain a resource has a long history in behavioral ecology (Maynard Smith and Parker, 1976; Hammerstein, 1981; Krebs and Davies, 1993). Ownership has its privileges, including additional knowledge about a resource. Even without direct benefits, as an arbitrary rule, respecting ownership can avoid costly conflicts. Thus, owning a resource can increase its value.

Utility building blocks

Simon's (1955; 1956) vision of bounded rationality emphasized realistic assumptions about the cognitive abilities of decision makers. Applying this perspective suggests that we need to consider both the cognitive building blocks or *evolved capacities* needed for exhibiting preferences and the limits placed on these capacities.

Magnitude perception

A critical component of exhibiting preferences is the ability to perceive differences in magnitude between options. Preferring two bananas to one requires discriminating between the amounts two and one. There are many mechanisms used to quantify objects in the world, and many studies have explored these mechanisms in primates (reviewed in Hauser and Spelke, 2004; Brannon, 2005a,b). Brannon (this volume) provides a useful overview of primate quantification, but it is worthwhile to briefly describe two of these mechanisms. The first mechanism discriminates between quantities only approximately via the analog magnitude system. Importantly, the discriminations follow Weber's law: variance in the representation increases with magnitude (Gallistel, 1990). This results in the ratio between quantities rather than the absolute magnitude driving the discrimination. The approximate number system yields a limit to the precision with which individuals can discriminate magnitudes, with larger magnitudes being more difficult than smaller ones (Figure 1). In addition to the approximate system, both humans and other primates seem to have a precise system that tracks individual objects. With this system, individuals can enumerate discrete

quantities but only up to a maximum of three or four objects (Hauser et al., 2000). Therefore, the precision that primates can exhibit in their preferences depends on the magnitude of the amount: they can precisely choose between small amounts and approximately choose between larger amounts.




Figure 1: Weber's law states that the variance around estimates of quantity increases with magnitude. Smaller quantities therefore are easier to discriminate than larger quantities. For instance, there may be little variance for estimations of three and five objects, so discriminating between them is easy. However, the variance dramatically increases for 11 and 13 objects, and estimates greatly overlap for these durations, making them difficult to distinguish.

The standard rationality approach would assume that if an organism has the more precise system, it should use it when making quantity judgments. However, Stevens and colleagues (2007) showed that cotton-top tamarins (*Saguinus oedipus*) used the approximate system in a foraging task, even though this species can use the precise system (Uller et al., 2001; Hauser et al., 2003). The tamarins used the simpler, approximate system as a default mechanism unless the task demanded the more precise system. Thus, different aspects of the decision environment trigger different mechanisms of discrimination.

Valuation

Options differ not only in quantity but also in quality. In many cases, decision makers face choices between qualitatively different reward types, both within and between reward domains. To choose between different types of food or even between different types of reward (food, water, sex, social contact), animals must have a mechanism to evaluate the utility of these reward types; that is, they must have a valuation mechanism that converts different reward types into a common currency. Padoa-Schioppa and colleagues (2006) explored how capuchin monkeys traded off different amounts of various food types to generate a valuation function (for example, one piece of apple may be worth three pieces of carrot). Deaner and colleagues (2005) pitted juice rewards against social information in rhesus macaques. Male monkeys chose between receiving juice and viewing images of either higher-ranking males' faces, lower-ranking males' faces, or female perinea (sexual areas). Interestingly, the valuation functions showed that the monkeys would forego juice to view high-ranking males and female perinea but had to be 'paid' in juice to view low-ranking males. Comparing these kinds of qualitatively different rewards is a critical capacity for decision making, although we do not have good cognitive models for how these tradeoffs occur.

Inhibitory control

Organisms must not only discriminate the magnitudes of benefits to establish a preference, but they also must favor the larger (positive) outcome. When motivated, this is not a problem for primates (except for very large rewards, see Silberberg et al., 1998). In fact, primates have a very difficult time going against this preference. In a task in which chimpanzees had to point to the smaller of two rewards to receive the larger reward, they failed miserably (Boysen and Berntson, 1995). In addition to the chimpanzees, bonobos (*Pan paniscus*), gorillas (*Gorilla gorilla*), orangutans (*Pongo pygmaeus*), rhesus and Japanese macaques (*Macaca fuscata*), squirrel monkeys (*Saimiri sciureus*), tamarins, and lemurs (*Eulemur fulvus* and *E. macaco*) all fail, at least initially, on this task (Silberberg and Fujita, 1996; Anderson et al., 2000; Kralik et al., 2002; Genty et al., 2004; Murray et al., 2005; Vlamings et al., 2006). Clearly, an unboundedly rational agent would adapt quickly to the contingencies of this task, but the preference for a large reward is so powerful that primates cannot inhibit their propensity to choose this. Of course, evolutionarily, it makes sense to employ the simple

heuristic 'choose the larger'. When in an organism's ecology would they opt for a small reward when a larger is present? This certainly must occur only rarely.

Uncertainty and risk

In a letter to Jean-Baptiste Leroy, Benjamin Franklin stated that "in this world nothing can be said to be certain, except death and taxes" (13 November 1789). That leaves a lot of uncertainty in the world. As agents navigating in this world, we must deal with this uncertainty in an adaptive manner. Knight (1921) posited a useful distinction between uncertainty (not knowing the distribution of possible payoffs) and risk (knowing the distribution of payoffs but not knowing which payoff will be realized). Though we and other animals frequently face uncertainty, this is difficult to study in the laboratory, so we will focus on risk.

Risky gambles

Much of the work undermining classical expected utility theory involved asking human subjects about their preferences in risky gambles. Rather than the banana examples suggested previously, subjects chose between risky reward amounts. Would you prefer a 50% chance of receiving \$100 (and a 50% chance of receiving nothing) or a 100% chance of receiving \$50? Though these two options have equal expected values, most people have a strong preference for the sure thing—they avoid risk. Risk-averse preferences arise with nonlinear utility functions, specifically when utility increases at a slower rate than the magnitude of the benefit (Figure 2). In general, this seems to be the case for human risk preferences for intermediate to large gains (Tversky and Kahneman, 1992).

Interestingly, nonhuman animals seem to show a similar pattern. Animals typically are tested by repeatedly experiencing repeated choices similar to the banana examples mentioned previously. In a review of risk sensitivity in over 25 species, Kacelnik and Bateson (1996) found that most species were either risk averse or risk neutral. Work on primates, however, has provided a more mixed result. Early tests of risk sensitivity in rhesus monkeys showed risk aversion (Behar, 1961). Yet, more recent studies have shown a preference for risky rewards in these macaques (McCoy and Platt, 2005; Hayden and Platt, 2007). Meanwhile, cotton-top tamarins and bonobos (*Pan paniscus*) seem to avoid risk, while common marmosets (*Callithrix jacchus*) ignore risk and chimpanzees prefer risk (Heilbronner et al., 2008, Heilbronner, Stevens, and Hauser, unpublished data). Why does such variation exist?




Figure 2: If utility is a linear function of magnitude (red line), then decision makers should be neutral to risk; they should be indifferent between a guaranteed banana and a 50/50 chance of no bananas or two bananas. If utility increases at a decelerating rate with magnitude (blue line), increments of utility are less valuable at larger magnitudes than they are at smaller magnitudes. Satiation offers an important biological example of diminishing utility because limited gut capacity constrains the utility of excess amounts of food. The additional utility of receiving three over two bananas is high, but the additional utility of receiving 103 over 102 is not as great, primarily because no one can consume 102 bananas. Diminishing utility implies risk aversion because the additional gain of the risky option is valued less than the loss.

Variation in risk preferences

One of the first hypotheses proposed to account for differences in risk sensitivity was the 'energy-budget

rule' (Caraco et al., 1980; Stephens, 1981). This rule suggests that hungry individuals should prefer risks because the safe option will not allow them to survive. Though this seems to work in some situations, there is no evidence for this rule in primates. Hayden and Platt (2007) tested an alternative idea proposed by Rachlin and colleagues (1986). Instead of preferring risky options because of hunger, animals may prefer risky options when they require low costs—specifically, when another choice will arise soon. Repeatedly choosing the risky option guarantees receiving the large payoff at some point. With short time delays between choices, waiting a few more seconds for this jackpot is not that costly. Rhesus macaques seem to follow this rule. The macaques preferred the risky option more when facing shorter delays between choices than with larger delays (Hayden and Platt, 2007).

The ecological rationality approach may also account for some of the patterns of risk preferences seen in primates. In particular, when species experience risk in their natural ecology, they may have decision rules that bias them towards risky options. With this hypothesis in mind, Heilbrunner and colleagues (2008) predicted a species difference in risk preferences between chimpanzees and bonobos. Although their diets overlap quite a bit in their natural habitat, bonobos feed primarily on terrestrial herbaceous vegetation, an abundant and reliable food source, and chimpanzees rely more on fruit, a more temporally and spatially variable food source (Wrangham and Peterson, 1996). Moreover, chimpanzees face risks when they hunt monkeys and other small mammals; bonobos rarely hunt. Interestingly, wild chimpanzees engage in this risky activity more often when fruit is abundant rather than scarce, a direct contrast to the energy budget hypothesis for risk-seeking behavior (Gilby and Wrangham, 2007). Given the generally higher level of risky choice in chimpanzees, Heilbrunner and colleagues predicted that this would select for risk-taking decision mechanisms. As predicted, chimpanzees preferred the risky choice in a laboratory experiment, whereas bonobos preferred the safe option. Therefore, to exploit risky options in their natural environment, natural selection has likely endowed chimpanzees with ecologically rational decision mechanisms, yielding preferences for risky outcomes even in captive laboratory situations.

The framing of risk

Though risk aversion and risk-seeking preferences do not pose a great challenge to expected utility theory, framing effects do challenge the theory. Prospect theory highlights two types of framing effects: reference dependence and loss aversion (Tversky and Kahneman, 1991). Reference dependence refers to viewing choices as gains or losses relative to a reference point rather than as absolute increases or decreases in utility. Thus, a set of outcomes could result in the exact same levels of wealth but be framed as a gain or loss. In the classic Asian disease problem, a medical treatment has a particular effectiveness in combating a disease (Tversky and Kahneman, 1981). However, when the outcome of a treatment is framed as number of people saved (a gain), subjects prefer the risky option more than when framed as number of people that die (a loss). Thus, framing the exact same outcome as either saving or losing lives greatly influences risk preferences. Yet, this reference dependence is not symmetric. People will try to avoid losses more than they will try to obtain gains—a phenomenon termed loss aversion (Tversky and Kahneman, 1981). We have already shown that, in risky gambles, people typically avoid risk over gains. When experiencing a loss, however, (a sure loss of \$50 or a 50% chance of losing \$100), people prefer risks to avoid the guaranteed loss.

Hundreds of studies have documented the effects of reference dependence and loss aversion in humans (Kahneman and Tversky, 2000). If this is truly a bias, then we might expect to find it only in humans. But if framing effects offer an adaptive, ecologically rational advantage, other animal species may exhibit them. Though first demonstrated in European starlings (Marsh and Kacelnik, 2002), Chen and colleagues (2006) explored reference dependence and loss aversion in brown capuchin monkeys (*Cebus apella*). To test reference dependence, the monkeys chose between two risky options. In one option, subjects saw one food reward and either received one or two rewards with equal probability. In the other option, they saw two rewards and received either one or two rewards with equal probability. Though identical in outcome, the reference point (number of initial rewards) varied, resulting in a perceived gain or loss. In this condition, subjects strongly preferred the gain option, showing clear reference dependence. Another experiment tested loss aversion. Here, one option consisted of seeing and then receiving a single reward, and the other option consisted of seeing two rewards but always receiving one. Again, the monkeys faced identical outcomes—a guaranteed one reward—but receiving that one reward could have been neutral or perceived as a loss. Again, the monkeys avoided the loss option, revealing the precursors to loss aversion in nonhuman primates. Thus, we share framing effects with other primates, suggesting deep evolutionary roots for this phenomenon.

Uncertainty building blocks

Which evolved capacities does an organism need to cope with uncertainty and risk? When given a choice between risky gambles, a decision maker must compare the probabilities of each outcome. In the human risk literature, subjects typically choose based on written descriptions of probabilities (e.g., a 50% chance) and/or on visual displays (e.g., a pie chart with half of the circle colored in). These techniques allow fairly accurate discriminations between probabilities. In the animal risk literature (and in some human experiments, see Hertwig et al., 2004), the subjects repeatedly experience the outcomes to gauge the level of risk. Therefore, to choose between gambles, animals must discriminate probabilities based on experienced outcomes. Though few studies have explicitly tested this, Weber's law may describe probability discrimination. Both Herrnstein and Loveland (1975) and Bailey and Mazur (1990) showed that pigeons' choices for the less risky option increased as the ratio between the small to large probability decreased (probabilities became less similar). In addition, Bailey and Mazur and Krebs et al. (1978) showed that pigeons and starlings (respectively) took more time to stabilize their preferences when the ratio between probabilities increased, further suggesting difficulty in discriminating similar probabilities. Thus, like numerical magnitude, probability discrimination likely follows Weber's law: individuals can discriminate a 10% from a 20% chance better than an 80% from a 90% chance. This has important implications for how animals deal with risk. When facing unlikely events, animals may discriminate probabilities well and therefore respond appropriately to risk. For more likely events, however, animals may ignore the probabilities and simply focus on the payoffs to determine choice.

Time

All decisions have a temporal component, from choosing to search for predators instead of searching for prey to delaying reproduction until the next breeding season. Delayed payoffs often have both benefits and costs (Stevens and Stephens, in press). They may be beneficial when investing time in obtaining resources allows for the extraction of more resources. For instance, the more time chimpanzees spend fishing for termites, the more termites they will extract. Yet, delayed rewards often come with a cost.

Temporal discounting

When facing options with smaller, sooner payoffs and larger, later payoffs, animals must make an *intertemporal choice*; that is, they must tradeoff the magnitude of rewards with the delay to receiving them (reviewed in Read, 2004). In some cases, ignoring the temporal component and choosing based on magnitude is best, but in other cases, a long delay may prove too costly. How should animals deal with this tradeoff? They may discount or devalue delayed rewards because the future is uncertain. The risk of not collecting a reward grows with delay because some event may interrupt its collection. For instance, a predator may interrupt an extended courtship or a bank may collapse before an investment matures. Economists have modified the expected utility models to create a *discounted utility model* of delayed benefits (Samuelson, 1937). This model replaces p from the expected utility model with a discounting function that includes a constant rate of interruption λ per unit time. Thus, for a reward amount A delayed for t time units, $DU = e^{-\lambda t} \cdot u(A)$. Again, utility is difficult to assess, so most versions of this model drop utility and just discount the absolute reward amount: $V = A \cdot e^{-\lambda t}$. Because the value of a reward decays exponentially with time, this is called the *exponential model* of discounting (Figure 3).

Though intuitively appealing, the data do not support the exponential model. Humans, pigeons, and rats violate predictions of this model in self-control experiments (Figure 4) by choosing between a smaller, immediate reward and a larger, delayed reward (Ainslie and Herrnstein, 1981; Mazur, 1987; Richards et al., 1997; Frederick et al., 2002). In fact, when choosing between immediately receiving two pieces of food and waiting for six, rats and pigeons only wait a few seconds for three times as much food. Animals would have to face extraordinarily (and unrealistically) high interruption rates for discounting by interruptions to account for this level of impulsivity.

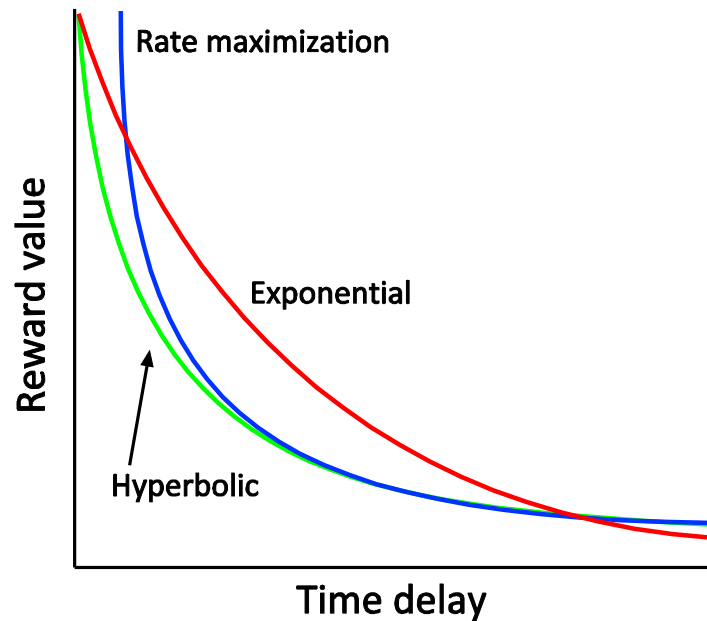


Figure 3: Models of intertemporal choice differ in their predictions about how the value of a reward decreases with the time delay to receiving the reward. Exponential discounting (red line) predicts a constant rate of decrease over time. Hyperbolic discounting (green line) predicts a decreasing rate of decrease over time, such that decision makers exhibit high discount rates at short-term delays but lower discount rates at longer delays. Rate maximization (blue line) predicts similar patterns as hyperbolic, albeit with strange behavior at very small time delays. The advantages of rate maximization models are that they include repeated choices and have biologically relevant parameters.

Psychologists proposed an alternative model that captures the data much better: the *hyperbolic discounting model* (Ainslie, 1975; Mazur, 1987). In the hyperbolic model⁶, $V = A/(1+kt)$, where k represents a fitted parameter that describes the steepness of discounting. Rather than predict a constant rate of discounting over time, this model predicts a decelerating rate over time—decision makers steeply discount at short time delays, and the rate of discounting declines at longer delays (Figure 3). This model nicely matches the data in humans, rats, and pigeons (Bickel and Johnson, 2003; Green et al., 2004). A number of studies have tested intertemporal choice in primate species, including cotton-top tamarins, common marmosets, brown capuchins, long-tailed and rhesus macaques, chimpanzees, and bonobos (Tobin et al., 1996; Szalda-Petree et al., 2004; Stevens et al., 2005; Ramseyer et al., 2006; Dufour et al., 2007; Rosati et al., 2007). Unfortunately, we have too few quantitative data on primates to distinguish between these two models. Nevertheless, in cotton-top tamarins and common marmosets, the rate of discounting slows with time, contradicting predictions of the exponential model (Hallinan et al., unpublished manuscript).




Figure 4: In self-control experiments, animals face a choice between a small reward (A_s) after a short delay (t_s) and a large reward (A_L) after a long delay (t_L). The animals wait for an intertrial interval, choose one option, wait the specified delay, consume their food, then begin another intertrial interval.

Though the hyperbolic model has achieved great empirical success, it suffers from a critical disadvantage: like prospect theory, it is a purely descriptive model. It cannot make a priori predictions about intertemporal choices and thus lacks explanatory power (Stevens and Stephens, in press). In addition, both exponential and hyperbolic models of discounting rest on the assumption of one-shot choice. In self-control

experiments, however, animals face repeated choices between the same options, and the repeated nature of these experiments has important implications for models of intertemporal choice (Kacelnik, 2003; Stevens and Stephens, in press).

Rate

The evolutionary approach to optimal foraging theory takes as its major assumption the fact that animals maximize their intake rate; that is, the amount of food gained per unit time (Stephens and Krebs, 1986). In the classic patch-choice scenario, an animal travels among the many patches of food in her environment and extracts resources from each patch. The question is: at what point should she stop extracting resources from the current patch and move on to search for a new patch? This scenario sounds quite similar to the self-control situation tested in the laboratory. Animals must choose between staying in a patch for a short time to extract a small gain or staying for longer time to extract a larger gain. Optimal foraging theory makes predictions about how long to stay in a patch: stay until the intake rate drops below the background intake rate in the environment (Stephens and Krebs, 1986).

The rate-based approach to intertemporal choice has two key advantages over the standard hyperbolic approach. First, it is by definition a model of repeated choice. Each decision faced by animals is embedded in a series of decisions. Thus, the rate-based approach models choices in situations very similar to the self-control experiments. Second, rate models provide an explanation of the pattern of temporal preferences seen in animals rather than just a description. Animals that maximize their intake rates should survive and reproduce better than their counterparts that do not – rate models are built on a foundation of evolutionary fitness.

We have sound theoretical reasons to favor rate models, but how do they fare empirically? Actually, quite well. The short-term rate model makes similar predictions as the hyperbolic model (Figure 3). Bateson and Kacelnik (1996) first demonstrated the effectiveness of the short-term rate model in describing choices by European starlings in a self-control experiment⁷. Stevens and colleagues (2005) then tested a variant of this model: $A/(t+h)$, where h represents the time required to handle the food, an important component of the foraging timeline (Rosati et al., 2006). Cotton-top tamarins chose the option that maximized this intake rate. Like rats and pigeons, the tamarins waited only 10-15 seconds for three times as much food. However, the results do not appear so 'impulsive' because we have an explanation for waiting such short delays that does not require unrealistic rates of interruption (Stevens and Stephens, in press).

Data on bonobos also agree with rate-maximization predictions, but this time long-term rather than short-term rate predictions (Rosati et al., 2007). The long-term rate $A/(\tau+t+h)$ includes τ , the time between trials (or between patches in the natural scenario). The long-term rate includes all of the relevant time intervals, and therefore this is the standard model used in optimal foraging theory. Why would some species ignore the intertrial interval while other species include it? An ecological rationality approach may help answer this question.

Ecological rationality of intertemporal choice

Tamarins and bonobos match predictions of rate-based models of intertemporal choice. Yet, bonobos wait over a minute for three times as much food, whereas tamarins wait less than 15 seconds. Why might the tamarins use short-term rates when making these types of choices? An answer lies at the heart of ecological rationality. Namely, the natural decision environment strongly shapes the decision mechanisms. Though the self-control experiments appear similar to natural patch-foraging scenarios, a key difference arises: animals rarely face simultaneous choices in nature but often face sequential choices in patch situations. Rather than facing two binary options, animals regularly choose to stay or leave a patch. Stephens and Anderson (2001) argued that this represents the natural foraging decision that animals typically confront, and therefore this should be the situation for which decision mechanisms evolved. Interestingly, a decision rule that assesses short-term intake rates in a patch situation will automatically result in long-term rate maximization because they are logically equivalent (Stephens and Anderson, 2001).

If short- and long-term rules are logically equivalent, this then may explain why bonobos differ from the tamarins and marmosets. Both represent 'adaptive peaks', so it does not matter on which peak a particular species rests. Each rule, however, has advantages and disadvantages. The short-term rule has the advantage of simplicity and possibly higher accuracy. It is simpler because it ignores information, namely

the intertrial interval. This may also increase accuracy because, by ignoring this time interval, animals can estimate much smaller intervals, which results in more accurate estimates (Stephens et al., 2004). Despite the benefits of the short-term rule, using the long-term rule provides advantages as well, primarily because it applies broadly and should result in the 'correct' outcome in most cases. Unfortunately, this is not true for the short-term rule. Though the short-term rule works well in patch situations, this does not directly carry over to the self-control situation (Stephens and Anderson, 2001). In fact, as shown by the tamarins, the short-term rule results in 'impulsive' choice in self-control situations, so the animals do not achieve the optimal long-term rate of gain. Thus, tamarins use the short-term rule because it works well in a natural foraging task—patch exploitation. When placed in an artificial situation, the rule 'misfires'. Bonobos, in contrast, may possess more acute estimation abilities and therefore use the more powerful long-term rule.

Specialized diets also raise interesting questions about ecologically rational intertemporal choice. Common marmosets offer an intriguing case because they specialize on a rather unique food source: tree exudates. Marmosets have specialized teeth, as well as digestive morphology and physiology, adapted to gouging holes in tree bark and eating the sap and gum that exude from the holes (Stevenson and Rylands, 1988; Rylands and de Faria, 1993). Because this foraging strategy requires waiting for the sap to exude, Stevens and colleagues (2005) proposed an ecologically rational response to this type of foraging strategy: a decision mechanism that is biased towards waiting for longer delays. Compared to the more insectivorous tamarins, marmosets should show stronger preferences for delayed rewards. As predicted, the marmosets waited longer than tamarins in a self-control situation, suggesting that the natural foraging ecology shapes the decision mechanisms (Stevens et al., 2005).

Intertemporal building blocks

Time poses unique cognitive challenges for organisms. Intertemporal choice combines establishing preferences over rewards with tracking these payoffs over time. Now, we explore what kind of cognitive building blocks might be required to make intertemporal choices.

Time perception

To cope with delayed rewards, an organism must perceive the delay. Unfortunately, we know little about primate time perception. We do, however, know a lot about time perception in rats, which likely applies to primates. Like magnitude judgments, time interval judgments seem to follow Weber's law. In fact, Gibbon (1977) showed that variance scaled with magnitude in time perception before applying it to quantification. Given that both primate numerical judgments (see Brannon, this volume) and human time perception (Allan, 1998) follow Weber's law, likely nonhuman primate time perception does as well. This has important implications for the study of intertemporal choice. First, logarithmic time perception (resulting from Weber's law) may result in the hyperbolic pattern of intertemporal choice typically seen in both humans and other animals (Takahashi, 2005; Takahashi et al., 2008). This finding could rescue the exponential model of discounting by overlaying a subjective time perception function. Second, as we have already seen, long time delays make discriminating the delays difficult. If time delays are viewed as equivalent, decision makers should opt for the larger (and therefore later) reward, resulting in more patient choice.

Delayed gratification

In addition to perceiving time, delayed rewards require a motivational ability to wait—decision makers must delay gratification. Mischel and colleagues pioneered the study of delayed gratification by measuring how long children would wait for delayed rewards (Mischel and Ebbsen, 1970; Mischel et al., 1989). They recorded the ability to wait at different ages and showed that delayed gratification at a young age strongly predicted intelligence, academic success, standardized test scores, and drug use much later in life.

Beran and colleagues have tested primates in a similar paradigm by offering chimpanzees and rhesus macaques a stream of food rewards (Beran et al., 1999; Evans and Beran, 2007b). The experimenter placed rewards in front of the subjects one by one at a particular rate, say one every 10 seconds. If the subject began consuming the food, however, the experimenter stopped the flow of food. Chimpanzees performed quite well on this task and some waited for over 10 minutes for the stream of food to be completed (Beran and Evans, 2006). This corroborates findings in the self-control task in which chimpanzees wait longer than any other species tested so far, including bonobos (Dufour et al., 2007; Rosati et al., 2007). Interestingly, this ability can be influenced by attentional factors. Evans and Beran (2007a) found that offering chimpanzees a distraction during the waiting period could significantly increase their abilities to delay gratification. Thus,

delayed gratification could provide a key building block required for patient choice, but it too is mediated by other factors such as attention.

Conclusion

Bounded rationality and ecological rationality are both firmly grounded in an evolutionary perspective on decision making. Natural selection places limitations on cognitive capacities and tailors cognitive mechanisms to the environment in which they are used (Barkow et al., 1992; Stevens, 2008). Darwin (1871; 1872) emphasized a continuity between the 'mental powers' of humans and other animals, and the analysis provided here confirms a similar continuity for decision making. Nonhuman primates show many of the so-called biases or anomalies demonstrated by human decision makers. The human decision-making literature has highlighted these results as failures to achieve a normative outcome. The evolutionary approach suggests that the bias lies not in the behavior but in the normative criteria used. Organisms did not evolve to follow a mathematically tractable set of principles—rather, natural selection favored decision strategies that resulted in greater survival and reproduction. In some cases, the evolutionary and normative perspectives may overlap. However, the normative perspective often fails to properly account for the role of the environment in decision making, whereas this is a critical part of the evolutionary view. Natural selection shapes decision strategies to match the environment.

This lesson about the limitations of the normative approach is an important one because many fields, especially neuroscience, use tools borrowed from economics to greatly expand the study of decision making. Though this could prove a fruitful enterprise, the data on decision making in other animals cautions against relying exclusively on the normative models of decision making. As a biological science, neuroscience should heed the warning of Dobzhansky (1964): "nothing makes sense in biology except in light of evolution" (p. 449). Of course, the emphasis on evolutionarily plausible models is nothing new to the field of neuroethology, and this tradition should continue in the study of decision making. Rather than looking for expected or discounted utility in the brain, perhaps we should test hypotheses about the bounded and ecological rationality of decision making. The bounded rationality approach makes clear predictions about what cognitive building blocks might be recruited and neuroscience can help test this. In addition, neuroscience can allow us to explore what aspects of the environment are relevant for triggering specific decision rules. The neuroethological approach offers an ideal interface between the evolutionary and mechanistic approaches to decision making.

The study of primate decision making has much to offer the larger field of decision making, and prospects are bright for future contributions. To advance the field, we must to begin exploring primate decisions in the wild or at least in more naturalistic situations. Currently, most studies of primate decisions (and animal decisions more broadly) occur in the laboratory with rather artificial scenarios. Though these studies provide valuable information on primate decision making, they may also lead us astray. If natural selection tailors decision mechanisms to the environment in which they are adapted, then the artificial nature of the laboratory might not trigger the appropriate mechanism. In this case, behavior seen in both humans and other animals may simply be spurious results elicited by an unnatural environment (Houston et al., 2007b; Stevens and Stephens, in press). Thus, exploring natural behaviors in natural environments underlies our understanding of primate decisions.

Another important advance that primate researchers may offer is the development and testing of process models of decision making. David Marr (1982) introduced the idea of three levels of information processing. The computational level emphasizes the goal of the system, the algorithmic level emphasizes the processes used, and the implementational level emphasizes the neural circuitry required to process information. Most models of decision making in animals act at the computational level—that is, level of evolutionary function—in the field of behavioral ecology or the implementational level in the field of neurobiology. The algorithmic level has been greatly neglected as a relevant level of analysis. Yet, adding an analysis of the relevant cognitive processes and algorithms could constrain and improve evolutionary models of decision making. Only an integrative approach across these levels of analysis—from the evolutionary to the psychological to the neurobiological—will yield satisfying answers to questions of the nature of decision making.

Notes

- Utility does not necessarily increase linearly with magnitude. The difference between consuming one and three bananas is not the same as the difference between consuming eleven and thirteen bananas. A difference of two bananas is much more relevant when there are fewer to begin with.
- ² Kahneman and Tversky do not use the term 'irrational'. They describe deviations from the normative theory as 'biases', 'fallacies', or 'cognitive illusions'. Nevertheless, the emphasis on putative errors highlights the irrational nature of human decision making (Rieskamp et al., 2006a).
- ³ Most work on decision making in animals, including primates, involves food as the reward, though there are important exceptions (Deaner et al., 2005). Therefore, in this review, I will also focus on food as the reward domain. Nevertheless, many of the principles mentioned here apply to other reward domains, and the study of animal decision making needs more in depth exploration of other reward types.
- ⁴ Though we have few data on transitive *preference*, transitive *inference* is well studied in primates. Transitive preference refers to an ordered preference over choices. Transitive inference refers to the ability to infer a transitive relation between objects (Vasconcelos, 2008). For instance, if individual A is dominant to B and B is dominant to C, can an individual infer that A is dominant to C without actually seeing A and C interact? Primates are quite good at these tasks, both with arbitrary objects in artificial laboratory tasks (reviewed in Tomasello and Call, 1997) and with social agents in the more naturalistic social inferences (Cheney and Seyfarth, 1990).
- ⁵ Despite the sparse data on animal intransitivity, Houston and colleagues (Houston, 1997; Houston et al., 2007a) have developed a series of models demonstrating that intransitivity might be a perfectly adaptive strategy when a decision maker is either updating her assessment of the environment or betting on a changing environment. Schuck-Paim and Kacelnik (2004) also assert that intransitive choices (particularly when food is the reward) can be perfectly biologically rational (Kacelnik, 2006) when the internal state changes (i.e., decision makers become satiated).
- ⁶ Though $V=A/(1+kt)$ is called the hyperbolic model, it represents only a single instance from a class of hyperbolic models. In fact, most nonexponential models (including the rate models) are hyperbolic. Nevertheless, the term hyperbolic discounting typically refers to Mazur's (1987) version.
- ⁷ The key difference between the short-term and long-term rate models is that the short-term model ignores the travel time or intertrial interval. Thus, animals using the short-term rate only focus on the times between choice presentation and reward consumption.

References

- Ainslie G (1975) Specious reward: a behavioral theory of impulsiveness and impulse control. *Psychol Bull* 82:463-496.
- Ainslie G, Herrnstein RJ (1981) Preference reversal and delayed reinforcement. *Anim Learn Behav* 9:476-482.
- Allais PM (1953) The behavior of rational man in risk situations: a critique of the axioms and postulates of the American School. *Econometrica* 21:503-546.
- Allan LG (1998) The influence of the scalar timing model on human timing research. *Behav Proc* 44:101-117.
- Anderson JR, Awazu S, Kazuo F (2000) Can squirrel monkeys (*Saimiri sciureus*) learn self-control? A study using food array selection tests and reverse-reward contingency. *J Exp Psych Anim B* 26:87-97.
- Bailey JT, Mazur JE (1990) Choice behavior in transition: development of preference for the higher probability of reinforcement. *J Exp Anal Behav* 53:409-422.
- Barkow JH, Cosmides L, Tooby J (1992) *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. New York: Oxford University Press.
- Baron J (2000) *Thinking and Deciding*. Cambridge, UK: Cambridge University Press.
- Bateson M, Kacelnik A (1996) Rate currencies and the foraging starling: the fallacy of the averages revisited. *Behav Ecol* 7:341-352.
- Behar I (1961) Learned avoidance of nonreward. *Psychol Rep* 9:43-52.
- Beran MJ, Evans TA (2006) Maintenance of delay of gratification by four chimpanzees (*Pan troglodytes*): the effects of delayed reward visibility, experimenter presence, and extended delay intervals. *Behav Proc* 73:315-324.
- Beran MJ, Savage-Rumbaugh E, Pate JL, Rumbaugh DM (1999) Delay of gratification in chimpanzees (*Pan troglodytes*). *Dev Psychobiol* 34:119-127.
- Bickel WK, Johnson MW (2003) Delay discounting: a fundamental behavioral process of drug dependence. In: *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice* (Loewenstein G, Read D, Baumeister RF, eds), pp 419-440. New York: Russell Sage Foundation.
- Boysen ST, Berntson GG (1995) Responses to quantity: perceptual versus cognitive mechanisms in chimpanzees (*Pan troglodytes*). *J Exp Psych Anim B* 21:82-86.
- Brannon EM (2005a) Quantitative thinking: from monkey to human and human infant to adult. In: *From Monkey Brain to Human Brain* (Dehaene S, Duhamel J-R, Hauser MD, Rizzolatti G, eds), pp 97-116. Cambridge, MA: MIT Press.
- Brannon EM (2005b) The numerical ability of animals. In: *Handbook of Mathematical Cognition* (Campbell JID, ed), pp

- 85-108. New York: Psychology Press.
- Brosnan SF, Jones OD, Lambeth SP, Mareno MC, Richardson AS, Schapiro SJ (2007) Endowment effects in chimpanzees. *Curr Biol* 17:1704-1707.
- Caraco T, Martindale S, Whittam TS (1980) An empirical demonstration of risk-sensitive foraging preferences. *Anim Behav* 28:820-830.
- Chen MK, Lakshminarayanan V, Santos LR (2006) How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *J Polit Econ* 114:517-537.
- Cheney DL, Seyfarth RM (1990) *How Monkeys See the World: Inside the Mind of Another Species*. Chicago: University of Chicago Press.
- Darwin C (1871) *The Descent of Man and Selection in Relation to Sex*. London: Murray.
- Darwin C (1872) *The Expression of the Emotions in Man and Animals*. London: Murray.
- Deaner RO, Khera AV, Platt ML (2005) Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Curr Biol* 15:543-548.
- Dobzhansky T (1964) Biology, molecular and organismic. *Amer Zool* 4:443-452.
- Dufour V, Pele M, Sterck EHM, Thierry B (2007) Chimpanzee (*Pan troglodytes*) anticipation of food return: coping with waiting time in an exchange task. *J Comp Psychol* 121:145-155.
- Ellsberg D (1961) Risk, ambiguity, and the Savage axioms. *Q J Econ* 75:643-669.
- Evans TA, Beran MJ (2007a) Chimpanzees use self-distraction to cope with impulsivity. *Biol Lett* 3:599-602.
- Evans TA, Beran MJ (2007b) Delay of gratification and delay maintenance by rhesus macaques (*Macaca mulatta*). *J Gen Psychol* 134:199-216.
- Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: a critical review. *J Econ Lit* 40:351-401.
- Gallistel CR (1990) *The Organization of Learning*. Cambridge, MA: MIT Press.
- Genty E, Palmier C, Roeder J-J (2004) Learning to suppress responses to the larger of two rewards in two species of lemurs, *Eulemur fulvus* and *E. macaco*. *Anim Behav* 67:925-932.
- Gibbon J (1977) Scalar expectancy theory and Weber's law in animal timing. *Psychol Rev* 84:279-325.
- Gigerenzer G, Todd PM (1999) Fast and frugal heuristics: the adaptive toolbox. In: *Simple Heuristics that Make Us Smart* (Gigerenzer G, Todd PM, eds), pp 3-34. Oxford: Oxford University Press.
- Gigerenzer G, Todd PM, the ABC Research Group (1999) *Simple Heuristics that Make Us Smart*. Oxford: Oxford University Press.
- Gilby IC, Wrangham RW (2007) Risk-prone hunting by chimpanzees (*Pan troglodytes schweinfurthii*) increases during periods of high diet quality. *Behav Ecol Sociobiol* 61:1771-1779.
- Green L, Myerson J, Holt DD, Slevin JR, Estle SJ (2004) Discounting of delayed food rewards in pigeons and rats: is there a magnitude effect? *J Exp Anal Behav* 81:39-50.
- Hallinan EV, Stevens JR, Hauser MD (unpublished manuscript) Rate maximization in common marmosets (*Callithrix jacchus*) and cotton-top tamarins (*Saguinus oedipus*).
- Hammerstein P (1981) The role of asymmetries in animal contests. *Anim Behav* 29:193-205.
- Hauser MD, Spelke ES (2004) Evolutionary and developmental foundations of human knowledge: a case study of mathematics. In: *The Cognitive Neurosciences III* (Gazzaniga M, ed), pp 853-864. Cambridge, MA: MIT Press.
- Hauser MD, Carey S, Hauser LB (2000) Spontaneous number representation in semi-free-ranging rhesus monkeys. *Proc R Soc Lond B* 267:829-833.
- Hauser MD, Tsao F, Garcia P, Spelke ES (2003) Evolutionary foundations of number: spontaneous representation of numerical magnitudes by cotton-top tamarins. *Proc R Soc Lond B* 270:1441-1446.
- Hayden BY, Platt ML (2007) Temporal discounting predicts risk sensitivity in rhesus macaques. *Curr Biol* 17:49-53.
- Heilbrunner SR, Rosati AG, Stevens JR, Hare B, Hauser MD (2008) A fruit in the hand or two in the bush? Divergent risk preferences in chimpanzees and bonobos. *Biol Lett* 4:246-249.
- Herrnstein RJ, Loveland DH (1975) Maximizing and matching on concurrent ratio schedules. *J Exp Anal Behav* 24:107-116.
- Hertwig R, Barron G, Weber EU, Erev I (2004) Decisions from experience and the effect of rare events in risky choice. *Psychol Sci* 15:534-539.
- Houston AI (1997) Natural selection and context-dependent values. *Proc R Soc Lond B* 264:1539-1541.
- Houston AI, McNamara JM, Steer MD (2007a) Violations of transitivity under fitness maximization. *Biol Lett* 3:365-367.
- Houston AI, McNamara JM, Steer MD (2007b) Do we expect natural selection to produce rational behaviour? *Phil Trans R Soc B* 362:1531-1543.
- Kacelnik A (2003) The evolution of patience. In: *Time and Decision: Economic and Psychological Perspectives on Intertemporal Choice* (Loewenstein G, Read D, Baumeister RF, eds), pp 115-138. New York: Russell Sage Foundation.
- Kacelnik A (2006) Meanings of rationality. In: *Rational Animals?* (Hurley S, Nudds M, eds), pp 87-106. Oxford: Oxford University Press.
- Kacelnik A, Bateson M (1996) Risky theories: the effects of variance on foraging decisions. *Amer Zool* 36:402-434.
- Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47:263-292.
- Kahneman D, Tversky A (2000) *Choices, Values, and Frames*. New York: Russell Sage Foundation; Cambridge University Press.
- Kahneman D, Slovic P, Tversky A (1982) *Judgment under Uncertainty: Heuristics and Biases*. New York: Cambridge University Press.

- Kahneman D, Knetsch JL, Thaler RH (1990) Experimental tests of the endowment effect and the Coase theorem. *J Polit Econ* 98:1325-1348.
- Knight FH (1921) *Risk, Uncertainty, and Profit*. Boston, MA: Houghton Mifflin.
- Kralik JD, Hauser MD, Zimlicki R (2002) The relationship between problem solving and inhibitory control: cotton-top tamarin (*Saguinus oedipus*) performance on a reversed contingency task. *J Comp Psychol* 116:39-50.
- Krebs JR, Davies NB (1993) *An Introduction to Behavioural Ecology*. Oxford: Blackwell Science.
- Krebs JR, Kacelnik A, Taylor P (1978) Test of optimal sampling by foraging great tits. *Nature* 275:27-31.
- Luce RD, Raiffa H (1957) *Games and Decisions: Introduction and Critical Survey*. New York: Wiley.
- Marr D (1982) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: W. H. Freeman and Co.
- Marsh B, Kacelnik A (2002) Framing effects and risky decisions in starlings. *Proc Nat Acad Sci USA* 99:3352-3355.
- Maynard Smith J, Parker GA (1976) The logic of asymmetric contests. *Anim Behav* 24:159-175.
- Mazur JE (1987) An adjusting procedure for studying delayed reinforcement. In: *Quantitative Analyses of Behavior: The Effect of Delay and of Intervening Events on Reinforcement Value* (Commons ML, Mazur JE, Nevin JA, Rachlin H, eds), pp 55-73. Hillsdale, NJ: Lawrence Erlbaum Associates.
- McCoy AN, Platt ML (2005) Risk-sensitive neurons in macaque posterior cingulate cortex. *Nat Neurosci* 8:1220-1227.
- Mischel W, Ebbesen EB (1970) Attention in delay of gratification. *J Pers Soc Psychol* 16:329-337.
- Mischel W, Shoda Y, Rodriguez ML (1989) Delay of gratification in children. *Science* 244:933-938.
- Murray EA, Kralik JD, Wise SP (2005) Learning to inhibit prepotent responses: successful performance by rhesus macaques, *Macaca mulatta*, on the reversed-contingency task. *Anim Behav* 69:991-998.
- Padoa-Schioppa C, Jandolo L, Visalberghi E (2006) Multi-stage mental process for economic choice in capuchins. *Cognition* 99:B1-B13.
- Payne JW, Bettman JR, Johnson EJ (1993) *The Adaptive Decision Maker*. Cambridge: Cambridge University Press.
- Rachlin H, Logue AW, Gibbon J, Frankel M (1986) Cognition and behavior in studies of choice. *Psychol Rev* 93:33-45.
- Ramseyer A, Pele M, Dufour V, Chauvin C, Thierry B (2006) Accepting loss: the temporal limits of reciprocity in brown capuchin monkeys. *Proc R Soc Lond B* 273:179-184.
- Read D (2004) Intertemporal choice. In: *Blackwell Handbook of Judgment and Decision Making* (Koehler D, Harvey N, eds), pp 424-443. Oxford: Blackwell.
- Richards JB, Mitchell SH, de Wit H, Seiden LS (1997) Determination of discount functions in rats with an adjusting-amount procedure. *J Exp Anal Behav* 67:353-366.
- Rieskamp J, Hertwig R, Todd PM (2006a) Bounded rationality: two interpretations from psychology. In: *Handbook of Contemporary Behavioral Economics* (Altman M, ed), pp 218-236. Armonk, NY: M.E. Sharpe.
- Rieskamp J, Busemeyer JR, Mellers BA (2006b) Extending the bounds of rationality: evidence and theories of preferential choice. *J Econ Lit* 44:631-661.
- Rosati AG, Stevens JR, Hauser MD (2006) The effect of handling time on temporal discounting in two New World primates. *Anim Behav* 71:1379-1387.
- Rosati AG, Stevens JR, Hare B, Hauser MD (2007) The evolutionary origins of human patience: temporal preferences in chimpanzees, bonobos, and adult humans. *Curr Biol* 17:1663-1668.
- Rylands AB, de Faria DS (1993) Habitats, feeding ecology and range size in the genus *Callithrix*. In: *Marmosets and Tamarins: Systematics, Behaviour, and Ecology* (Rylands AB, ed), pp 262-272. Oxford: Oxford University Press.
- Samuelson PA (1937) A note on measurement of utility. *Rev Econ Stud* 4:155-161.
- Schuck-Paim C, Pompilio L, Kacelnik A (2004) State-dependent decisions cause apparent violations of rationality in animal choice. *PLoS Biology* 2:e402.
- Shafir S (1994) Intransitivity of preferences in honey-bees: support for comparative-evaluation of foraging options. *Anim Behav* 48:55-67.
- Silberberg A, Fujita K (1996) Pointing at smaller food amounts in an analogue of Boysen and Berntson's (1995) procedure. *J Exp Anal Behav* 66:143-147.
- Silberberg A, Widholm JJ, Bresler D, Fujita K, Anderson JR (1998) Natural choice in nonhuman primates. *J Exp Psych Anim B* 24:215-228.
- Simon HA (1955) A behavioral model of rational choice. *Q J Econ* 69:99-118.
- Simon HA (1956) Rational choice and the structure of the environment. *Psychol Rev* 63:129-138.
- Simonson I, Tversky A (1992) Choice in context: tradeoff contrast and extremeness aversion. *J Marketing Res* 29:281-295.
- Stephens DW (1981) The logic of risk-sensitive foraging preferences. *Anim Behav* 29:628-629.
- Stephens DW, Krebs JR (1986) *Foraging Theory*. Princeton: Princeton University Press.
- Stephens DW, Anderson D (2001) The adaptive value of preference for immediacy: when shortsighted rules have farsighted consequences. *Behav Ecol* 12:330-339.
- Stephens DW, Kerr B, Fernandez-Juricic E (2004) Impulsiveness without discounting: the ecological rationality hypothesis. *Proc R Soc Lond B* 271:2459-2465.
- Stevens JR (2008) The evolutionary biology of decision making. In: *Better than Conscious? Decision Making, the Human Mind, and Implications For Institutions* (Engel C, Singer W, eds), pp 285-304. Cambridge, MA: MIT Press.
- Stevens JR, Stephens DW (in press) The adaptive nature of impulsivity. In: *Impulsivity: Theory, Science, and Neuroscience of Discounting* (Madden GJ, Bickel WK, eds). Washington, DC: American Psychological Association.
- Stevens JR, Hallinan EV, Hauser MD (2005) The ecology and evolution of patience in two New World monkeys. *Biol Lett* 1:223-226.

- Stevens JR, Wood JN, Hauser MD (2007) When quantity trumps number: discrimination experiments in cotton-top tamarins (*Saguinus oedipus*) and common marmosets (*Callithrix jacchus*). *Anim Cog* 10:429-437.
- Stevenson MF, Rylands AB (1988) The marmosets, genus *Callithrix*. In: *Ecology and Behavior of Neotropical Primates* (Mittermeier RA, Rylands AB, Coimbra-Filho AF, Fonseca GAB, eds), pp 131-222. Washington, DC: World Wildlife Fund.
- Szalda-Petree AD, Craft BB, Martin LM, Deditius-Island HK (2004) Self-control in rhesus macaques (*Macaca mulatta*): controlling for differential stimulus exposure. *Percept Motor Skill* 98:141-146.
- Takahashi T (2005) Loss of self-control in intertemporal choice may be attributable to logarithmic time-perception. *Med Hypotheses* 65:691-693.
- Takahashi T, Oono H, Radford MHB (2008) Psychophysics of time perception and intertemporal choice models. *Physica A: Statistical Mechanics and its Applications* 387:2066-2074.
- Tinklepaugh OL (1928) An experimental study of representative factors in monkeys. *J Comp Psychol* 8:197-236.
- Tobin H, Logue AW, Chelonis JJ, Ackerman KT (1996) Self-control in the monkey *Macaca fascicularis*. *Anim Learn Behav* 24:168-174.
- Tomasello M, Call J (1997) *Primate Cognition*. Oxford: Oxford University Press.
- Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice. *Science* 211:453-458.
- Tversky A, Kahneman D (1991) Loss aversion in riskless choice: a reference-dependent model. *Q J Econ* 106:1039-1061.
- Tversky A, Kahneman D (1992) Advances in prospect theory: cumulative representation of uncertainty. *J Risk Uncertainty* 5:297-323.
- Uller C, Hauser MD, Carey S (2001) Spontaneous representation of number in cotton-top tamarins (*Saguinus oedipus*). *J Comp Psychol* 115:248-257.
- Vasconcelos M (2008) Transitive inference in non-human animals: an empirical and theoretical analysis. *Behav Proc* 78:313-334.
- Vlamings P, Uher J, Call J (2006) How the great apes (*Pan troglodytes*, *Pongo pygmaeus*, *Pan paniscus*, and *Gorilla gorilla*) perform on the reversed contingency task: the effects of food quantity and food visibility. *J Exp Psych Anim B* 32:60-70.
- von Neumann J, Morgenstern O (1947) *Theory of Games and Economic Behavior*, 2nd Edition. Princeton: Princeton University Press.
- Waite TA (2001a) Background context and decision making in hoarding gray jays. *Behav Ecol* 12:318-324.
- Waite TA (2001b) Intransitive preferences in hoarding gray jays (*Perisoreus canadensis*). *Behav Ecol Sociobiol* 50:116-121.
- Wrangham RW, Peterson D (1996) *Demonic Males: Apes and the Origins of Human Violence*. Cambridge, MA: Harvard University Press.
- Wu G, Zhang J, Gonzalez R (2004) Decision under risk. In: *Blackwell Handbook of Judgment and Decision Making* (Koehler D, Harvey N, eds), pp 399-423. Oxford: Blackwell.

¹ Utility does not necessarily increase linearly with magnitude. The difference between consuming one and three bananas is not the same as the difference between consuming eleven and thirteen bananas. A difference of two bananas is much more relevant when there are fewer to begin with.

² Kahneman and Tversky do not use the term 'irrational'. They describe deviations from the normative theory as 'biases', 'fallacies', or 'cognitive illusions'. Nevertheless, the emphasis on putative errors highlights the irrational nature of human decision making (Rieskamp et al., 2006a).

³ Most work on decision making in animals, including primates, involves food as the reward, though there are important exceptions (Deaner et al., 2005). Therefore, in this review, I will also focus on food as the reward domain. Nevertheless, many of the principles mentioned here apply to other reward domains, and the study of animal decision making needs more in depth exploration of other reward types.

⁴ Though we have few data on transitive *preference*, transitive *inference* is well studied in primates. Transitive preference refers to an ordered preference over choices. Transitive inference refers to the ability to infer a transitive relation between objects (Vasconcelos, 2008). For instance, if individual A is dominant to B and B is dominant to C, can an individual infer that A is dominant to C without actually seeing A and C interact? Primates are quite good at these tasks, both with arbitrary objects in artificial laboratory tasks (reviewed in Tomasello and Call, 1997) and with social agents in the more naturalistic social inferences (Cheney and Seyfarth, 1990).

⁵ Despite the sparse data on animal intransitivity, Houston and colleagues (Houston, 1997; Houston et al., 2007a) have developed a series of models demonstrating that intransitivity might be a perfectly adaptive strategy when a decision maker is either updating her assessment of the environment or betting on a changing environment. Schuck-Paim and Kacelnik (2004) also assert that intransitive choices (particularly when food is the reward) can be perfectly biologically rational (Kacelnik, 2006) when the internal state changes (i.e., decision makers become satiated).

⁶ Though $V=A/(1+kt)$ is called the hyperbolic model, it represents only a single instance from a class of hyperbolic models. In fact, most nonexponential models (including the rate models) are hyperbolic. Nevertheless, the term hyperbolic discounting typically refers to Mazur's (1987) version.

⁷ The key difference between the short-term and long-term rate models is that the short-term model ignores the travel time or intertrial interval. Thus, animals using the short-term rate only focus on the times between choice presentation and reward consumption.