

Decisions, Decisions, Decisions: Choosing a Biological Science of Choice

Review

Paul W. Glimcher¹
Center for Neural Science
New York University
New York, New York 10003

Behavioral ecologists argue that evolution drives animal behavior to efficiently solve the problems animals face in their environmental niches. The ultimate evolutionary causes of decision making, they contend, can be found in economic analyses of organisms and their environments. Neurobiologists interested in how animals make decisions have, in contrast, focused their efforts on understanding the neurobiological hardware that serves as a more proximal cause of that same behavior. Describing the flow of information within the nervous system without regard to these larger goals has been their focus. Recent work in a number of laboratories has begun to suggest that these two approaches are beginning to fuse. It may soon be possible to view the nervous system as a representational process that solves the mathematically defined economic problems animals face by making efficient decisions. These developments in the neurobiological theory of choice, and the new schema they imply, form the subject of this article.

Introduction

In the 1650s, Blaise Pascal puzzled over one of the most complicated decisions faced by enlightenment philosophers: “Either God is or he is not,” he wrote. “But to which view shall we be inclined?” Pascal took what was then a completely novel approach to this problem, arguing that one should “weigh up the gain and loss involved in calling heads that God exists” or tails that he does not (Pascal, 1670).

Over the ensuing three centuries, Pascal’s notion that *gain* and *uncertainty* must interact whenever rational decision making occurs has become the foundation for all rigorous analysis of human choice behavior. Modern economic theory rests almost entirely upon this idea. Some thinkers, however, have criticized Pascal’s approach, and in particular its use to characterize financial decision making, as unsystematic. These scientists have noted that while uncertainty can be rigorously quantified with tools like *Bayesian Estimation*, understanding the subjective value of a financial gain or a loss in mathematical terms may be impossible (for a review of these criticisms, see Glimcher, 2003). Recently, however, a group of biologists has argued that economic approaches applied to biological questions might overcome this limitation by relying on an alternative system of valuation rooted in evolutionary theory. “Paradoxically,” John Maynard Smith wrote in 1982, “it has turned out that game theory [a branch of modern economic theory] is more readily applied to biology than to the field of economic behavior for which it was origi-

nally designed...the theory requires that the values of different outcomes (for example, financial rewards, the risks of death and the pleasures of a clear conscience) might be measured on a single scale. In human applications this measure is provided by ‘utility’—a somewhat artificial and uncomfortable concept: In biology, Darwinian fitness provides a natural and genuinely one-dimensional scale.”

By anchoring economic theory to Darwinian notions of fitness, these behavioral ecologists undertook to explain the *ultimate* causes of human and animal behavior. Whether a moose decides to eat algae or grass on a particular day (Belovsky, 1984), with whom a monkey decides to mate (Dunbar, 1984), or whether a young salmon decides to forage or hide (Bull et al., 1996) can all, these scientists contend, be viewed as decisions which are the products of evolution acting to maximize fitness in the face of uncertainty. Behavioral ecologists explain behavior in evolutionary terms by quantitatively analyzing the decisions animals make and correlating the outcomes of those decisions with theoretical or empirical measurements of fitness. Their goal is to relate a theoretical assessment of what an animal should do with behavioral data about what animals actually do. They try to understand what causes a behavior by understanding the evolutionary constraints under which that behavior evolved.

In contrast, neurobiologists devote their resources to understanding what causes behavior to occur in a more *proximal* sense. Neurobiologists seek to understand the mechanistic cellular processes by which the brain produces specific behavioral acts. They seek to understand how a cue light causes a rat to press a lever for food or how a visual target presented to a monkey elicits an eye movement; neurophysiologists have sought to understand what causes a behavior by studying the sensory-motor architecture of the brain.

The existence of these parallel approaches to understanding behavior raises an interesting and often overlooked question. What is the relationship between ultimate and proximal causation in behavior? While evolution may shape behaviors toward efficient forms, the specific environmental variables which guide evolution may not be represented explicitly by the neural architecture evolution produces. There may be little reason to believe that the computations which influence fitness at an evolutionary level are echoed by the computations made within the brains of individual animals. In the last decade, however, evidence has been accumulating that the brains of complex animals like mammals perform operations which closely correspond to the optimization problems behavioral ecologists describe as the ultimate causes of behavior. These data suggest that the environmental problems animals face may shape not only behavior but also the neural hardware that generates that behavior. If this is widely true, then it may be of tremendous importance to neurophysiologists. The studies of behavioral ecologists may define the computations performed by the primate neurophysiological architecture for decision making just as mathematical studies of im-

¹Correspondence: glimcher@cns.nyu.edu

age encoding shaped our understanding of sensory encoding in the last century. This article reviews some of the growing evidence that economic analyses of behavior may provide neurobiologists with critical information, not just about what problems animals face, but about how their brains solve those problems.

Ultimate and Proximal Causes of Decision: Static Conditions

In 1966, *The American Naturalist* published two papers that, although there were important precedents (cf. Holling, 1959), largely invented quantitative behavioral ecology. J. Merritt Emlen writing alone (Emlen, 1966) and Robert MacArthur writing with Eric Pianka suggested that whenever an animal must make a decision about what food to eat it faces an optimization problem that can be quantified and solved in economic terms. Evolution, they suggested, could be viewed as a force that might well drive organisms not just to solve these problems but to solve them efficiently. "There is a close parallel between the development of theories in economics and population biology. In biology, however, the geometry of the organisms and their environment plays a greater role. Different phenotypes have different abilities at harvesting resources, and the resources are distributed in a patchwork in three dimensions in the environment. In this paper we undertake to determine in which patches a species would feed and which items would form its diet if the species acted in the most economical fashion. Hopefully, natural selection will often have achieved such optimal allocation of time and energy expenditures, but such 'optimum theories' are hypotheses for testing rather than anything certain" (MacArthur and Pianka, 1966).

The area of behavioral ecology which developed directly from these two papers, *foraging theory*, has made significant advances since 1966, and today, at least six independent optimization problems have been identified, modeled, and tested. One class of foraging problem that has received significant attention, and one of particular relevance to neurobiologists, is the study of *Prey Selection*. As a monkey roams the savannah, it encounters patches of food of different types and qualities, each with a different frequency. Each food type, a plant with nutritious roots, a colony of insects, or a piece of fruit, occurs with a certain probability and a certain caloric density. Each takes a different amount of time and energy to obtain and consume, and so each has a different value, or *profitability*, to the monkey. What happens when a monkey like this encounters eight food patches, each with a different set of sensory properties? Those sensory properties allow the animal to estimate the value of each patch, the cost of harvesting each patch, and the frequency with which a patch of that type is encountered. Prey selection models use these variables to predict which patch, if any, the monkey will attempt to acquire and consume. In complete mathematical form, the prey model was therefore developed (Charnov, 1973; Stephens and Krebs, 1986) to define a method for computing the most efficient predatory strategy for any animal, in any environment. It assumes that optimal predation is achieved whenever a forager achieves a maximum rate of energy intake with a minimal expenditure of effort in a random and unpredictable world.

The Prey Model

The first step in looking for food is to begin *searching*; defined as any activity that takes time and during which the forager may encounter prey. Searching is presumed to cost s units of energy per minute, and animals engage in search for a total time labeled T_s . After a period of searching, it is assumed that the animal encounters a prey item. The forager then has to make the decision around which the prey model is structured, whether to use both time and energy to attempt to capture and eat the prey item or whether to pass it up and continue searching. The process of predating is thus a cycle: search, encounter, decide, search, encounter, decide...

The goal of the model is to characterize the decision-making phase, for which the animal must know (1) the energy gained from prey of each type, (2) the average *handling time* required to catch and consume the prey, (3) the cost, in energy spent, of the handling process, and (4) the rate, in encounters per unit time, at which a prey of each type is detected.

We can characterize the rate of net energy intake in any environment and for any possible prey attack strategy in the following way. First, we determine the profitability of each prey type by multiplying the probability that the forager will attack that prey type, P (the variable controlled by the forager), by the frequency with which that prey is encountered, λ , to determine how often an attack occurs. Then multiply that frequency by the net energy gained from the prey. (The value of the prey minus the energy lost during handling.) This calculation tells us how much energy the forager can expect to gain (per unit time) for adopting this probability of attack with this particular type of prey.

$$\begin{aligned} & \text{average gain per prey type per unit time} \\ & = P \times \lambda \times (\text{energy gained} - \text{energy lost}) \quad (1) \end{aligned}$$

Next, one needs to know the cost of attacking each type of prey in terms of time diverted away from searching for other, potentially better, prey items. Multiply the probability of an attack by the frequency of an encounter and by the total handling time for that prey type.

$$\begin{aligned} & \text{average handling time per prey type} \\ & = P \times \lambda \times (\text{handling time}) \quad (2) \end{aligned}$$

Finally, one sums the first calculation across every possible prey type and multiplies it by the total time spent searching, then subtracts from that the total cost of searching, s , and divides the sum by the time spent searching plus handling, yielding a measure of how much energy is gained, for a given set of attack strategies, per unit time.

$$R = \frac{T_s [\sum \text{average gain per prey type} - s]}{T_s + \sum \text{average handling time}} \quad (3)$$

To figure out directly what specific attack strategy maximizes the rate of energy intake, one differentiates Equation 3 with regard to P , creating a new equation that allows us to compute the set of attack strategies that maximizes R (for details, see Stephens and Krebs, 1986). The prey model thus defines an optimal strategy for

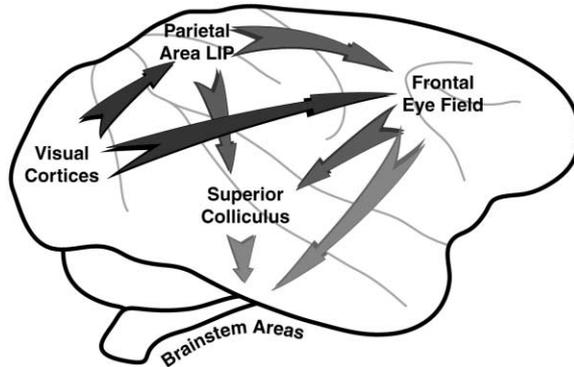


Figure 1. Sensory-to-Motor Connectivity in Visual-Saccadic Decision Making

Visual signals originating in the retina pass via the lateral geniculate nucleus of the thalamus to the visual cortices, the primary visual cortex, and then throughout the extrastriate visual areas. These signals influence activity in saccadic control areas via the lateral intraparietal area, the frontal eye field, and even the superior colliculus. The two principal saccadic control areas, the superior colliculus and the frontal eye fields, project in turn to areas in the brainstem reticular formation that control the speed and position of the rotating eyeball. (For a review of saccadic anatomy, see Glimcher, 1999.)

predation, a behavioral pattern toward which evolution is presumed to drive real species.

Neurophysiological Studies of Decision Making

While neurophysiologists have started to examine the proximal causes of behaviors very much like those produced by a foraging monkey, they have tended to focus almost exclusively on the neural pathways that link sensory stimuli, like the color of a food patch, with movement control pathways, like those that activate the muscles of the arm. Very little attention has been paid to the variables behavioral ecologists study. In part, this reflects the current state of our neurobiological sophistication; both sensory systems and movement control systems are far better understood than the systems which assess the values, likelihoods, and profitabilities of outcomes. But as a result, many classical physiological studies have almost entirely ignored the variables behavioral ecologists identify as critical at the level of ultimate causation. If the vertebrate brain does represent the environmental variables that define strategies for maximizing evolutionary fitness, then these more classical studies of sensory-to-motor connections may only reveal the most superficial properties of the neural architecture for decision making (Figure 1).

In order to better understand how physiological studies of decision making differ from ecological studies of decision making, consider the following set of landmark experiments by Jeffrey Schall and his colleagues (cf. Hanes and Schall, 1996; Schall and Thompson, 1999). For these experiments, thirsty rhesus monkeys were trained to stare straight ahead at a centrally located fixation stimulus. Shortly after the monkey began staring straight ahead, eight secondary targets appeared, arranged radially around the central fixation stimulus. Seven of those targets appeared in a common color, and one appeared in a different color, an *oddball*. If the animal looked at any of the seven common color targets,

the trial ended. If he looked at the oddball, he received a drop of fruit juice as a reward.

Under conditions like these, we know quite a lot about both the sensory and motor processes that must become active in the monkey's brain. When the targets illuminate, we know that eight locations in the visual cortices become active. These signals propagate through the visual system toward saccadic eye movement control centers like the frontal eye fields (FEF) and the superior colliculus (SC). Only one of the eight locations, however, represents the oddball and ultimately leads to activation of the eye movement control circuitry in those areas. So how is the translation from eight visual signals to one motor command actually accomplished? To answer that question, Schall and his colleagues recorded the activity of neurons in the saccadic movement maps of the FEF while monkeys performed this oddball detection task.

Whenever monkeys executed hundreds of these oddball detection trials, Schall noticed that there was a natural variability in the speed with which the monkeys produced their saccades. Based on this behavioral variability, trials could be divided into *fast*, *medium*, and *slow* groups. Examining the activity of FEF neurons on these trials, they found that each neuron was most active before a movement in a particular direction, that neuron's *best movement*. Taking into account both the behavioral and neurophysiological data, trials could be subdivided into those in which the movement acquiring the oddball was the *best movement* for the neuron being studied and those trials on which the oddball elicited a different movement.

Schall found that FEF neurons (Figure 2) rose to an early peak firing rate shortly after stimulus onset, but only after about 80 ms was there evidence, in these neurons, of an underlying decision process. At that point, neuronal firing rates continued to grow if the *best movement* for the studied neuron was required to look at the oddball; otherwise, firing rates dropped back toward baseline levels. Importantly, regardless of the rate of increase in neuronal activity, the movement occurred at a roughly fixed interval after the firing rate reached a specific level. This led Schall to suggest the existence of a decisional threshold which the rising activity had to cross in order for a movement to be produced.

Ecological Studies of Decision Making

While these experiments do tell us something important about the relationship between sensation and action in the primate nervous system, what is most striking to a behavioral ecologist is that almost none of the variables which guide decision making were manipulated in those studies. The values, costs, profitability, and likelihoods of reward associated with stimulus and movement were not varied; instead, all of those variables were held constant under all conditions. As a result, to a behavioral ecologist these experiments seem more about movement production than about decision making per se.

To more completely understand this point of view, consider a classic ecological study of decision making by John Krebs and his colleagues (Krebs et al., 1977). In Krebs' experiment (Figure 3), at the beginning of an experimental session, hungry titmice, a small European bird in the chickadee family, were placed in a 1 m³ cage. The floor of the cage was opaque except for a 2.5 inch

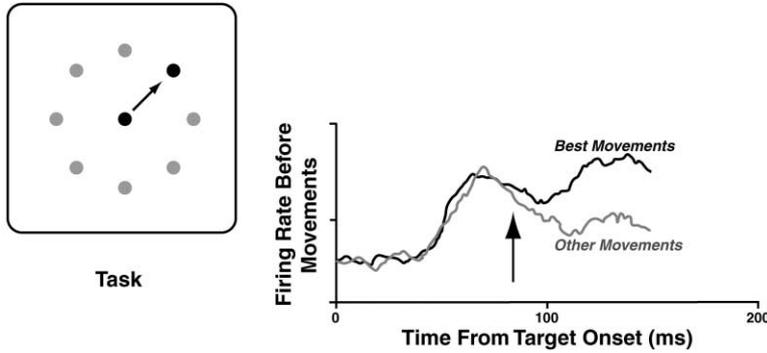


Figure 2. Oddball Task with Schall Activity Plot
In Schall's oddball task, eight radially arranged targets appear simultaneously, one in an *oddball* color. The animal receives a reward if he looks at the oddball immediately after it appears. When the oddball elicits a movement encoded by the frontal eye field neuron under study (the best movement), firing rates are different than when any other movement is elicited. The difference between these two firing patterns is apparent about 80 ms after target onset. Current evidence suggests that this difference reflects a neuronal decision about what movement to produce.

hole, placed beneath a perch, which overlooked a black rubber conveyor belt sliding past the opening at a speed of 5 inches per second. Out of sight of the bird, the experimenters could place mealworms (a favorite food of titmice) on the belt, which would present the worms to the birds through the 2.5 inch hole for about 0.5 s. The birds had to decide whether to grab a worm as it went by or whether to wait for the next worm to appear.

In order to characterize the serial decision problem that the birds faced in terms of the classical prey selection model (Charnov, 1973; see Stephens and Krebs, 1986, for a full treatment of the model), three classes of variables were systematically manipulated in this study. First, the caloric value of the mealworms was controlled. Worms were presented in two sizes which differed in value by a factor of two. Second, the handling time (and thus the cost) associated with each worm type was measured for each individual bird and, in the case of the small worms, was manipulated. This was accomplished by attaching pieces of paper tape to the small worms, which the birds had to remove. By measuring the time it took each individual to handle and consume both large and small worms, it was possible to determine the relative *profitability* of the two worm types to each

individual. Because some birds were able to strip the paper tape off more quickly than others, the profitability of the small worms varied significantly from individual to individual. Third, the experimenters systematically manipulated the frequency with which, in a given experimental session, worms of each type were encountered. In each session, large and small worms were placed on the conveyor belt according to a predetermined pseudo-random sequence yielding a specific mean encounter frequency for that prey type. Varying the encounter frequency was selected because under these circumstances the prey selection model makes a critical prediction. It predicts that for a given bird the average rate at which large worms are encountered should set a minimum threshold for profitability. Any prey that falls beneath this threshold should never be selected, regardless of the frequency with which it is encountered (cf. Stephens and Krebs, 1986).

When the experiment was run, the actual choices made by each bird were well predicted by the prey selection model. For birds who could strip the tape away quickly, the model predicted that the small worms should always be above the threshold of profitability. For the birds that handled the small worms slowly, however, the prey model made a different prediction. At one or more of the rates at which large worms were presented, each of these birds should have decided to select only the large worms and to ignore the small worms. What Krebs and his colleagues found was that the choice behavior of the birds did reveal the existence of this profitability threshold and at a level that was almost identical to the level predicted by the model. They also, however, observed one significant deviation of behavior from the model. Although the birds showed strong preferences for the large worms at the right times, they did not show an absolute preference. When they should have been ignoring the small worms completely, they were still sampling those worms about 10% of the time, a result which has been widely repeated but never entirely explained.

In 1966, MacArthur and Pianka wrote that "optimum theories" [which describe behavior as the product of evolution and the environment] are hypotheses for testing rather than anything certain." Recent tests of these hypotheses, like the one conducted by Krebs and his colleagues, suggest that economic approaches may often describe both behavior and the environmental variables that shape it quite accurately (Krebs and Davies, 1997). Neurophysiological studies of decision making,

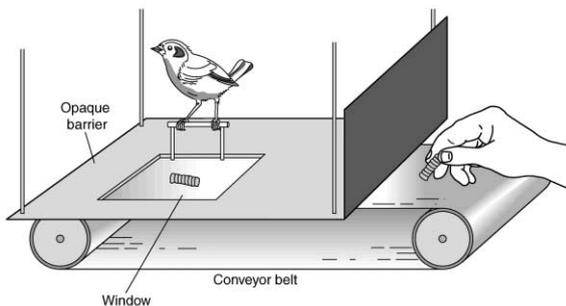


Figure 3. Krebs Experiment: Bird Foraging over Conveyor Belt

In Krebs' experiment, hungry birds of the species *parus major* stand over a conveyor belt. An experimenter places mealworm segments of two sizes on the belt in a pseudorandom sequence. The bird faces a serial decision problem; it must decide which segments to eat and which to ignore. The decisions the bird makes are influenced by the mean rates at which both prey types are encountered, the difficulty of capturing and eating the segments, and the relative values of the two different size pieces. Charnov's prey model predicts the quantitative pattern of decisions that birds make with significant precision. Illustration from Glimcher, 2003, courtesy of MIT Press.

however, have until recently almost completely ignored these variables as tools for understanding the nervous system. Of course, the notion that optimum theories of the type MacArthur and Pianka pioneered can be used to study the neural architecture for decision making is more a hypothesis for testing rather than anything certain. But a number of laboratories have begun to perform those tests, and the results are encouraging.

Decision Variables and Neural Architectures

C. Randy Gallistel and his colleagues have spent over 20 years studying the decisions rats make when they work for brain stimulation reward, the direct electrical activation of neural structures like the medial forebrain bundle (MFB) with surgically implanted electrodes (cf. Gallistel et al., 1981; Gallistel, 1994). Unlike ecological studies in which rewards are experienced through intact perceptual systems, experiments with brain stimulation reward bypass at least some of the sensory component of standard sensory-motor decision making. These studies therefore permit one to test the hypothesis that variables which serve as the ultimate causes of behavior can govern the decisions made by a reduced segment of the neural architecture, shifting the focus of the economic approach toward a study of the proximal causes of behavior.

Gallistel's most recent work along these lines (cf. Leon and Gallistel, 1992; Mark and Gallistel, 1993, 1994) has focused on the study of rats who must choose between two different patterns of stimulation. In those experiments, a rat is placed in a cage that contains two levers. If the rat depresses the right lever, the MFB is activated by a brief train of electrical pulses delivered at a fixed frequency and current. If the rat depresses the left lever, the MFB is also activated but at a different frequency and current. The rat must decide between these two options, allowing the experimenter to assess the relationship between the subjective value of the stimulation and the frequency and current of MFB activation. The only drawback to this approach toward understanding the value of MFB stimulation, however, is that one can only rank order the values of the left and right levers.

In order to overcome that limitation, Gallistel and his colleagues adjusted the dynamic structure of the rewards in the standard two-lever choice task in order to extract, from the animal's behavior, a precise estimate of the exact value of any pattern of stimulation. To do this, they adopted a strategy first pioneered for the study of choice behavior in pigeons by Richard Herrnstein (1961), in which each of the rat's two levers was configured so that the likelihood it would release a stimulation train increased in proportion to the time since the lever was last pressed (a *variable interval schedule*). Under this regime, the left lever might, for example, become enabled with an additive probability of 0.5 per second, while the right lever might become enabled on any given second with a lower additive probability of 0.3. Importantly, once enabled, a lever remained enabled until the reward was harvested by the rat. In a configuration like this, during the first second of an experiment in which the two levers yield rewards of equal value but with these different probabilities, the left lever is clearly more profitable, and the rats should respond on it exclusively. But after 3 s spent pressing the left lever, the right lever reaches a point at which it is now more likely to present

a reward than the left lever. Under these conditions, the animal should now switch to the right lever.

In the more general case, where the values of the reward produced by the two levers varies, foraging theories predict that response rates on each lever should be proportional to their relative profitabilities. If one of the two levers presents a standard reward and the other presents a reward that is varied across blocks of trials, then it should be possible to determine the precise value of any physiological stimulation to the rat with regard to a fixed standard by observing the fraction of time she budgets to each lever.

When Gallistel and his colleagues examined this two-lever variable interval schedule with self-stimulation as the reward, they were able to apply an economic-style analysis to the choices rats made in order to describe the precise subjective profitability of any pattern of MFB activation. This allowed them to derive an equation defining the value of stimulation as a function of the current, frequency, and duration of MFB stimulation. Gallistel's results are critical because they were among the first to suggest that economic-style approaches could be used to study neurobiological phenomena. They show that economic approaches can be used to characterize decision making not just at the level of the whole organism, but even when the sensory-perceptual systems by which the animal assesses natural rewards are replaced by the direct electrical activation of an internal neural structure. If the approaches behavioral ecologists employ for studying the ultimate causes of behavior can be used to study electrical stimulation of the MFB, can these approaches also be used in neurophysiological studies of decision making?

To begin to answer that question, Platt and Glimcher (1999) developed a neurophysiological experiment based loosely on foraging approaches like the ones Krebs and Gallistel employed but in this case designed to examine the neural substrate for visual-saccadic decision making in rhesus monkeys (Figure 4). In that experiment, thirsty monkeys were trained to stare straight ahead at a central visual stimulus. After a delay, two eccentric targets were illuminated, and the monkeys had to choose whether to look at the left target, the right target, or to abort the trial. The critical manipulation was that on sequential blocks of 100 trials the amount of juice that the monkeys would earn for each of the leftward and rightward movements was systematically manipulated, and the animals' decisions about which movement to make were recorded. Finally, while the monkeys made decisions under these varying conditions, the activity of single neurons in parietal cortex was studied to test the hypothesis that the relative profitabilities of the leftward and rightward responses were being represented within the neural architecture itself. The posterior parietal cortex was selected for examination, at least in part, because it was a major source of input to the frontal eye fields (FEF).

At a theoretical level, the problem that the monkeys faced at the beginning of each block of 100 trials was first to determine the relative profitability of both responses. Once that had been accomplished, animals might be expected to adopt an efficient foraging strategy. They should, therefore, have produced each response several times at the beginning of a new block

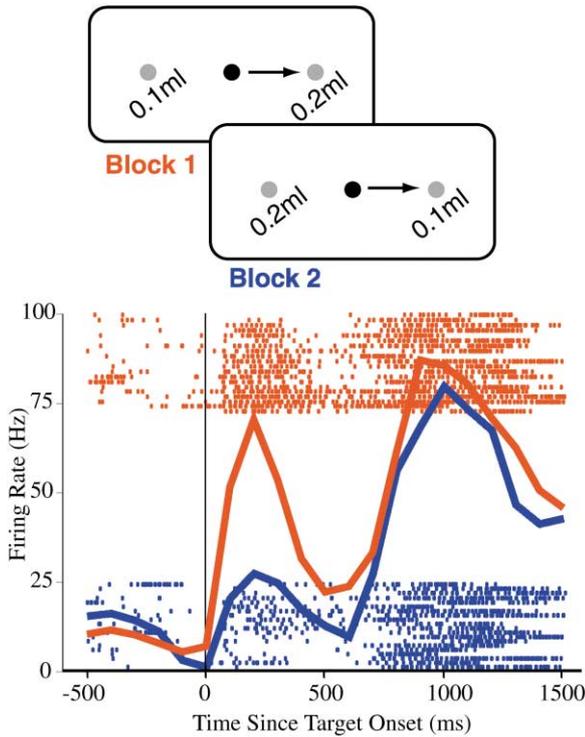


Figure 4. Platt and Glimcher Task and Unit Data

In Platt's experiment, monkeys run blocks of trials in which they must decide whether to look left or right in order to obtain a fruit juice reward. In sequential blocks, the relative values of the two movements are manipulated. Neuronal firing rates in the lateral intraparietal area are strongly influenced by the relative value, to the animal, of the movement they encode.

and then settled on a strategy of producing only the more profitable response. Instead, Platt and Glimcher found that the monkeys typically matched their rates of responding during each block to the relative profitability of the two responses, as the rats had in Gallistel's experiments. In 1972, the behavioral ecologist Stephen Fretwell developed a model for foraging under these circumstances that might provide some insight into the ultimate cause of this apparently suboptimal behavior. Fretwell noted that whenever animals forage in groups and have to compete with each other for access to rewards of different values, probability matching of this type is always an optimal strategy. This raised the possibility that monkeys may probability match under many circumstances because the neural computations that they perform reflect an evolutionary assumption that they are competing with other monkeys.

In any case, while the animals were engaged in this matching behavior, the activity of neurons in the eye movement control area of posterior parietal cortex was assessed. Neurons associated with rightward movements were, under these circumstances, found to carry a signal which was highly correlated with the relative profitability of the rightward movement. Similarly, neurons associated with leftward movements seemed to encode the relative profitability of the leftward movement. Essentially, the ratio of these two activities was correlated not only with the relative profitabilities that

controlled behavior in Fretwell's model of multi-animal foraging, but also with the actual moment-by-moment probability matching behavior of the animals in which these neurons resided. To a first approximation then, the neuronal data seemed to suggest that computations performed by the neural architecture for decision making were at least related to computations that should be the ultimate causes of decision-making behavior.

More recently, Gold and Shadlen (2000) have come to a similar conclusion in a study of activity in the FEF of monkeys during another kind of visual-saccadic decision-making task. In that experiment, monkeys were trained to stare at a field of chaotically moving spots of light. A small fraction of those spots, however, moved coherently in a single direction, either to the right or to the left. If the monkey reacted to this display by looking in the direction that the small fraction of coherently drifting spots were moving, he earned a juice reward. Gold and Shadlen (2001) applied a formal decision theoretic analysis to the problem the monkeys faced. They reasoned that when a large fraction of the spots were moving coherently the monkeys ought to be able to assess the likelihood that a rightward movement would be rewarded quite quickly. When only a tiny fraction of the dots were moving coherently, the monkeys ought to maximize the likelihood that they had correctly identified the more profitable movement. If the decision-making architecture reflects this calculation explicitly, then it should produce neural activity with a temporal profile which reflects a moment-by-moment estimate of the mathematical likelihood that a movement would be rewarded.

To test this hypothesis, Gold and Shadlen (2000) used electrical stimulation in the FEF to trigger a saccade at a variable interval after the moving spot display began. Under normal conditions, stimulation of the FEF elicits a saccade having a fixed amplitude and direction, and it was hoped that the temporal profile of the saccadic decision process would be revealed as a systematic deviation in the endpoint of the stimulation-induced movements. What they found was that the stimulation-induced movements were indeed systematically biased and in a way that was correlated, at each point in time during the decision-making interval, with the theoretically derived estimate of the likelihood that a given movement would be rewarded. Once again, the neurons seemed to be encoding a signal closely related to the variables that an economically based model indicated should be the ultimate cause of the behavior.

Summary

Over the course of the last several decades, behavioral ecologists have repeatedly demonstrated that animals often select between uncertain options of variable value in ways that are highly efficient; evolution appears to push animals toward efficient decision making within their evolved niches. Behavioral ecologists have been able to show that models of optimal decision making rooted in economic theory do a surprisingly good job of describing the computations that animals perform. More recently, neurobiologists have begun to appropriate this approach, using economic tools developed for studying the ultimate evolutionary causes of behavior for the examination of the neural architecture which

serves as the proximal cause of that behavior. The studies presented here and literally dozens of other closely related studies (see the other articles in this issue of *Neuron*) have begun to suggest that the explicit computations modeled by behavioral ecologists can be dissected at a neuro-computational level.

In all of these cases, however, animals are decision makers who must select and execute a rational course of action in a passive world. The world is conceived of as presenting a fixed problem that the animal must solve. While clearly valuable, studies of this kind may fail to engage the richest and most complicated kind of decision making, the unpredictable or stochastic decisions that humans and animals make when faced with more complicated environmental situations.

Dynamic Conditions: The Theory of Games

In the middle of the twentieth century, the mathematician John Von Neumann and the economist Oskar Morgenstern (Von Neumann and Morgenstern, 1944) became interested in understanding when and how stochastic behavior, behavior in which humans behave unpredictably, might be described as an efficient strategy for maximizing wealth. Von Neumann recognized that most economic problems which had been well described at that time involved straightforward optimization of the kind foraging theorists would later study. In those problems, the likelihoods and values of all possible future outcomes are considered static variables insensitive to the actions of the decision maker. What Von Neumann realized was that problems of this type fail to capture situations in which the profitability of a course of action is influenced both by the actions of the decision maker and by the actions of intelligent opponents who may themselves be influenced by the decision maker.

Consider a famous experiment in behavioral ecology conducted by D.G.C. Harper in 1982 when a flock of ducks wintered on the main pond at the Cambridge University Botanical Garden (Harper, 1982). Each morning, two experimenters would walk to different banks of the pond, each with a sack of 5 g breadballs. At a signal, they would both begin throwing the balls onto the ground at a fixed rate: experimenter one throwing one ball every 10 s and experimenter two throwing one ball every 20 s. To characterize this as a standard foraging problem, imagine that 20 ducks are fixed in position in front of experimenter one and nine ducks are fixed in position in front of experimenter two. A single free duck who must choose between walking toward experimenter one or experimenter two should compute the profitability, in grams of bread per minute, of each action. Experimenter one provides a profitability of $30 \text{ g/min}/20 \text{ ducks} = 1.5 \text{ g/m/duck}$. Experimenter two provides a profitability of $15 \text{ g/min}/9 \text{ ducks} = 1.67 \text{ g/m/duck}$. The rational strategy under these conditions would be for the free duck to walk toward experimenter two.

What Von Neumann recognized when thinking about human decision making under similar conditions was that the situation becomes much more complicated when all the subjects (in our case, ducks) are free to make this same decision at the same time. Imagine a situation like the one above but in which each duck is free to move. Under these conditions, each of the 20

ducks in front of experimenter one realizes that it would be more profitable to switch to experimenter two. Of course, if all of them switch to experimenter two, then any duck remaining in front of experimenter one will profit enormously; as ducks shift toward experimenter two, their actions alter the profitability of standing in front of experimenter one. Under these conditions, the dynamic interactions of the ducks influence profitability as much as does the rate at which breadballs are thrown. What then is the optimal response under these conditions?

The modern solution to this class of problem was developed by John Forbes Nash in the 1950s. Nash recognized that under conditions like these the population as a group could be viewed as a dynamic system which would ultimately reach a stable group solution, or *equilibrium*, when the expected value of each resource patch was equivalent. The ducks would reach a group equilibrium when and only when ten ducks stood before experimenter two and 20 ducks stood before experimenter one, rendering the profitability of both patches 1.5 g/m/duck (Nash, 1950a, 1950b; for a derivation specifically with regard to foraging, see Fretwell, 1972).

To make Nash's insight clear, consider a case in which 21 ducks stand momentarily in front of experimenter one and nine stand in front of experimenter two. The population will reach a stable point most efficiently (and each duck will be guaranteed a maximal return assuming all the other ducks behave rationally) if and only if in the next moment each duck standing before experimenter one shows a 1 in 21 chance of shifting to experimenter two. Each duck should behave stochastically but in a manner constrained by the ecological problem that the ducks face.

Harper's experiment was critical because it tested both the idea that the ducks could reach this kind of stable Nash equilibrium and the notion that the ducks could behave unpredictably. When the experiment was performed, it was found that the flock of 33 ducks assorted themselves within as little as 90 s at a Nash equilibrium solution, precisely matching their behavioral distribution to the relative profitabilities of the two experimenters as predicted by Nash's (and Fretwell's) equations. If the experimenters then changed to new rates of breadball throwing, the ducks would immediately resort themselves, assuming the new equilibrium distribution in as little as 90 s. One thing that was particularly striking about this result was the speed with which the ducks achieved this assortment. After 90 s of breadball throwing, as few as ten breadballs have been dispersed. Long before half the ducks have obtained even a single breadball, they have produced a precise equilibrium solution.

Harper also tested the game theoretic hypothesis derived from other work (cf. Maynard Smith, 1982) that each duck should behave unpredictably on a moment-by-moment basis. Even when the flock was at a stable equilibrium, Harper found that individual ducks were constantly in motion. When the equilibria required that one third of the ducks stand in front of experimenter one, it was observed that each duck spent a random one third of its time standing in front of experimenter one. The behavior of individual ducks as they solved this sensory motor problem was stochastic and unpre-

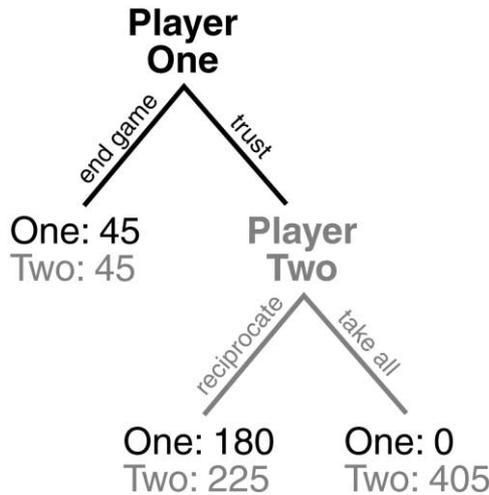


Figure 5. Extensive Form Game Tree for Trust and Reciprocity

In McCabe's Trust and Reciprocity game, a round begins when player one decides whether to end the game immediately or to pass control of the game to player two. If player one ends the game, both she and player two receive 45 cents. If player one elects to trust player two, then player two must decide how to divide a much larger gain. She can elect either to keep 405 cents of it for herself or to reciprocate player one's trust by turning 180 cents back to player one. McCabe and his colleagues found that humans were much more likely to trust or reciprocate with other humans than with computer programs. Activity in an area of prefrontal cortex seemed to be correlated with trusting and reciprocating.

dictable at a local level but maintained the Nash equilibrium globally.

Harper's experiment suggests that, even when animals are dynamic and unpredictable, their behavior may still be described as the product of an evolutionary process that optimizes decision making. Thus, even stochastic behavior may be ultimately caused by the environmental constraints that drive evolution. Indeed, a growing number of other studies support this conclusion. For example, Craig Packer and his colleagues (cf. Packer and Ruttan, 1988; Packer et al., 1990) have shown that the strategic behavior of lions who compete within a pride can be well described using game theoretic approaches. Of course, the question that this raises is whether game theoretic approaches which define the ultimate causes of competitive and stochastic behavioral decisions can also be used to better understand the neural substrates that serve as the proximal causes of these unpredictable behaviors.

Game Theory and Neural Architectures

A number of researchers in economics and neuroscience have recently begun to examine how the theory of games might be used to analyze the neural architecture active when competitive and stochastic behaviors are produced. Kevin McCabe and colleagues (2001) pioneered this approach when they examined the brains of human subjects engaged in a strategic game using functional magnetic resonance imaging (fMRI). In their experiment, subjects played a game called *trust and reciprocity* (Figure 5), in which two players, only one of which is inside the scanner, work sequentially and interactively to earn money. The trust and reciprocity

game begins with the first player, who must decide whether to terminate the game immediately, in which case both players earn a 45 cent cash payoff, or whether to turn control of the game over to player two. If control passes to player 2, then player 2 must decide how to split a much larger gain, 405 cents. Player two must decide between taking all 405 cents for herself or keeping only 225 and returning 180 to player one.

For a game theorist, this conflict is particularly interesting when subjects face a new opponent on each trial. Under those circumstances, if player two is perfectly rational, given the chance, she will always take all 405 for herself. Cooperating with player one offers her no advantage. Player one knows this and should, therefore, always be compelled to end the game on the first play, which guarantees her a small but at least positive outcome. When players encounter each other repeatedly, however, a different optimal strategy can emerge. The two players can cooperate in fear of future retribution, electing to trust one another in order to reach the 180/225 outcome on each play.

Like more classical foraging examples, the trust and reciprocity game examines a situation in which subjects must decide between one of two possible responses. But unlike more classical examples, the optimal solution depends upon assumptions both about the likelihood of encountering the same player again and assumptions about how one's own behavior will influence the behavior of the opponent. This is a property that defies explanation with non-game theoretic tools and makes it similar in many ways to Harper's duck experiment.

What McCabe and his colleagues found was that a typical subject was very likely to cooperate with a human opponent, even when she was told that she would face a different opponent on subsequent trials. Humans turned out to be more cooperative with other humans than was strictly rational, almost as if their brains were performing a computation that assumed this opponent would, sooner or later, be encountered again. However, when subjects were told that they faced a computer opponent, they often took a different and more purely rational approach. They almost never cooperated. What McCabe and his colleagues found when studying the brains of their subjects under these conditions was that whenever a subject chose to cooperate with a human opponent a specific region in prefrontal cortex was more active than when they decided to act rationally against the computer. While this does not tell us too much about how cooperativity is computed neurally, the McCabe experiment is an important first step because it demonstrates that game theoretic approaches can be used to study the neurobiological basis of stochastic decision making.

Recognizing the significance of this strategy, two other groups have also begun to explore game theoretic techniques for identifying variables that might be encoded in the neural decision-making architecture. Dorris and Glimcher (M.C. Dorris and P.G. Glimcher, 2001, Soc. Neurosci., abstract) have trained monkeys to participate in a classic strategic conflict called the *inspection game*, which is based loosely on Harper's studies of foraging ducks. Ongoing studies of single neurons in the brains of monkeys playing this game may well provide insight into the computational architecture involved in the pro-

duction of stochastic behavior. In a similar vein, Berns and colleagues (G.S. Berns et al., 2001, Soc. Neurosci., abstract) have begun to examine, using fMRI, the brains of pairs of humans engaged in strategic interactions. In those experiments, two humans in two fMRI scanners play a game classically called *matching pennies* while both are simultaneously scanned. These simultaneous studies may soon provide insight into the moment-by-moment interplays of neural activity that characterize stochastic decision making between pairs of subjects.

Summary

In the 1950s and 60s, a number of neurophysiologists became interested in understanding how the sensory systems of the brain encoded information about the outside world. One approach to this problem was to derive an estimate of how an optimally efficient sensory system *would* operate. Behavioral tests then sought to determine the sensory efficiency of human and animal subjects with regard to these theoretically defined estimates. Neurophysiological experiments sought to extend this approach, searching for evidence that the neural architecture actually employed such strategies. Horace Barlow encapsulated the argument for this strategy in 1961 when he wrote “[The tendency of sensory systems to respond only when a stimulus changes but not when a stimulus remains constant] may be regarded as a mechanism for compressing sensory messages into fewer impulses by decreasing the tendency for serial correlations in the values of impulse intervals. There must be many occasions when neighboring sense organs are subjected to the same stimulus, and this will lead to correlations in the values of impulses in neighboring fibers. This offers scope for further economy of impulses, and one might look for a mechanism to perform the appropriate recoding. [A form of recoding that] would...diminish the correlations between impulse intervals spatially separated...[and thus would, according to information theory, achieve a more nearly optimal encoding of the sensory event]” (Barlow, 1961).

It may be that a similar logic can be applied to studies of decision making. Over the past several decades, behavioral ecologists have made significant advances in their theoretical studies of decision making. Economic approaches have allowed them to build sophisticated models of the environmental constraints which define efficient behavioral decision making in evolutionary terms and to thereby define optimal strategies for decision making. Empirical studies have begun to determine the extent to which the decisions of real animals approximate these ideals and the extent to which animal decision making may fail to approximate these ideals. (A behavioral failure human economists refer to as *irrational*.)

The available results from behavioral ecology suggest that, in practice, models of optimal decision making with regard to fitness often do a good job of characterizing the real decision making processes of animals. Indeed, these analyses of animal behavior as a tool for maximizing fitness are beginning to reveal a more global tendency toward rationality than human economists have identified when characterizing human behavior as a tool for maximizing wealth. Using the evolutionary and economic framework of behavioral ecology, it is beginning to appear that the ultimate causes of behavior can be

examined and in much the same way that Horace Barlow used information theory to describe efficient sensory encoding.

As neurobiologists begin to study the proximal causes of decision making, it seems imperative that these economic approaches to behavior be employed as tools to bridge the gap between ultimate and proximal causes of behavior. Just as the sensory physiologists of the last century used models that were specifically designed to describe efficient sensory encoding and discrimination, neurophysiologists interested in decision making must employ economic models specifically designed to describe the decision-making process. If the success of the sensory physiologists during the last 50 years is any indication, models of decision making rooted in economic theory should provide powerful insights into brain function over the next half century.

References

- Barlow, H.B. (1961). The coding of sensory images. In *Current Problems in Animal Behavior*, W.H. Thorpe and O.L. Zangwill, eds. (Cambridge: Cambridge University Press), pp. 331–360.
- Belovsky, G.E. (1984). Herbivore optimal foraging: a comparative test of three models. *Am. Nat.* 124, 97–115.
- Bull, C.D., Metcalfe, N.B., and Mangel, M. (1996). Seasonal matching of foraging to anticipated energy requirements in anorexic juvenile salmon. *Proc. R. Soc. Lond. B. Biol. Sci.* 263, 13–18.
- Charnov, E.L. (1973). *Optimal foraging: some theoretical explorations*. PhD dissertation, University of Washington, Seattle, Washington.
- Dunbar, R.I.M. (1984). *Reproductive Decisions. An Economic Analysis of Gelada Baboon Social Strategies* (Princeton: Princeton University Press).
- Emlen, J.M. (1966). The role of time and energy in food preference. *Am. Nat.* 100, 611–617.
- Fretwell, S.D. (1972). *Populations In a Seasonal Environment* (Princeton: Princeton University Press).
- Gallistel, C.R. (1994). Foraging for brain stimulation: toward a neurobiology of computation. *Cognition* 50, 151–170.
- Gallistel, C.R., Shizgal, P., and Yeomans, J. (1981). A portrait of the substrate for self-stimulation. *Psychological Reviews* 108, 228–273.
- Glimcher, P.W. (1999). Eye Movements. In *Fundamental Neuroscience*, M.J. Zigmond, F.E. Bloom, S.C. Landis, and L.R. Squire, eds. (New York: Academic Press), pp. 993–1010.
- Glimcher, P.W. (2003). *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics* (Cambridge, MA: MIT Press).
- Gold, J.I., and Shadlen, M.N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature* 404, 390–394.
- Gold, J.I., and Shadlen, M.N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16.
- Hanes, D.P., and Schall, J.D. (1996). Neural control of voluntary movement initiation. *Science* 247, 427–430.
- Harper, D.G.C. (1982). Competitive foraging in mallards: “ideal free” ducks. *Anim. Behav.* 30, 575–584.
- Herrnstein, R.J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exp. Anal. Behav.* 4, 267–272.
- Holling, C.S. (1959). Some characteristics of simple types of predation and parasitism. *Canadian Entomology* 91, 385–398.
- Krebs, J.R., and Davies, N.B. (1997). *Behavioural Ecology*, Fourth edition (Oxford: Blackwell Scientific Publications).
- Krebs, J.R., Erichsen, J.T., Webber, M.L., and Charnov, E.L. (1977). Optimal prey selection in the great tit. *Anim. Behav.* 25, 30–38.
- Leon, M., and Gallistel, C.R. (1992). The function relating the subject-

- tive magnitude of brain stimulation reward to stimulation strength varies with site of stimulation. *Behav. Brain Res.* *52*, 183–193.
- MacArthur, R.H., and Pianka, E.R. (1966). On optimal use of a patchy environment. *Am. Nat.* *100*, 603–610.
- Mark, T.A., and Gallistel, C.R. (1993). Subjective reward magnitude of medial forebrain stimulation as a function of train duration and pulse frequency. *Behav. Neurosci.* *107*, 389–401.
- Mark, T.A., and Gallistel, C.R. (1994). Kinetics of matching. *J. Exp. Psychol. Anim. Behav. Process.* *20*, 79–95.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games* (Cambridge: Cambridge University Press).
- McCabe, K., Houser, D., Ryan, L., Smith, V., and Trouard, T. (2001). A functional imaging study of cooperation in two-person reciprocal exchange. *Proc. Natl. Acad. Sci. USA* *98*, 11832–11835.
- Nash, J.F. (1950a). Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* *36*, 48–49.
- Nash, J.F. (1950b). The bargaining problem. *Econometrica* *18*, 155–162.
- Packer, C., and Ruttan, L.M. (1988). The evolution of cooperative hunting. *Am. Nat.* *132*, 159–198.
- Packer, C., Scheel, D., and Pusey, A.E. (1990). Why lions form groups: food is not enough. *Am. Nat.* *136*, 1–19.
- Pascal, B. (1670). *Pensees*, A.J. Krailsheimer, trans. (New York: Penguin Books).
- Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* *400*, 233–238.
- Schall, J.D., and Thompson, K.G. (1999). Neural selection and control of visually guided eye movements. *Annu. Rev. Neurosci.* *22*, 241–259.
- Stephens, D.W., and Krebs, J.R. (1986). *Foraging Theory* (Princeton: Princeton University Press).
- Von Neumann, J.V., and Morgenstern, O. (1944). *Theory of Games and Economic Behavior* (Princeton: Princeton University Press).