

THE NEUROBIOLOGY OF VISUAL-SACCADIC DECISION MAKING

Paul W. Glimcher

*Center for Neural Science, New York University, New York, New York 10003;
email: glimcher@cns.nyu.edu*

Key Words oculomotor, choice, monkey, human, neuroeconomics

■ **Abstract** Over the past two decades significant progress has been made toward understanding the neural basis of primate decision making, the biological process that combines sensory data with stored information to select and execute behavioral responses. The most striking progress in this area has been made in studies of visual-saccadic decision making, a system that is becoming a model for understanding decision making in general. In this system, theoretical models of efficient decision making developed in the social sciences are beginning to be used to describe the computations the brain must perform when it connects sensation and action. Guided in part by these economic models, neurophysiologists have been able to describe neuronal activity recorded from the brains of awake-behaving primates during actual decision making. These recent studies have examined the neural basis of decisions, ranging from those made in predictable sensorimotor tasks to those unpredictable decisions made when animals are engaged in strategic conflict. All of these experiments seem to describe a surprisingly well-integrated set of physiological mechanisms that can account for a broad range of behavioral phenomena. This review presents many of these recent studies within the emerging neuroeconomic framework for understanding primate decision making.

INTRODUCTION

How do animals select movements from their behavioral repertoires for execution? How are decisions, the computational events that connect sensory data and a stored representation of the structure of the world with behavior, accomplished at a mechanistic level? Deriving an answer to that question, particularly with regard to humans, has been a scholarly goal at least since the time of ancient Greece. Aristotle (Aristotle 1986) argued in *De Anima* that it was the nonmaterial soul that served as the mechanism responsible for purposive human behavior, and it was this view that dominated western and arab thought until the Enlightenment. During the Enlightenment, however, a growing conviction arose that material explanations could be derived for most, if not all, observable phenomena. This shift challenged philosophers to reexamine the scholastic view that all connections between sensation and action were made outside the material universe. Descartes

(1649, 1664), in particular, challenged this view when he suggested a dualist approach to understanding how sensation and action might be connected. He proposed that all of human behavior could be divided into two principle classes that could be viewed as the product of two mechanistically distinct processes. The first class included only those behaviors that are fully deterministic with regard to events in the immediate sensory environment, behaviors that predictably link stereotyped sensory stimuli with simple motor responses. Their simple deterministic nature suggested to Descartes that, for these behaviors, the sensory to motor connection lay within the material body, which made those connections amenable to physiological study. For the second class, behaviors in which no deterministic connection was obvious between sensation and action, he followed Aristotle's lead, identifying the source of these actions as the nonmaterial soul.

Descartes' dualist proposal was a tremendous advance for physiologists. It suggested that at least one of the processes that generated behavior could be studied with purely physiological methodologies. During the nineteenth and early twentieth centuries physiologists made good use of this hypothesis, achieving significant headway toward identifying, at least within the spinal cord, neural mechanisms that could deterministically link sensation and action.

This progress in understanding spinal reflexes represented a critical first step in the physiological study of decision making. Reflex studies provided a model for understanding the simplest possible generative mechanism for behavior, a mechanism with which sensory stimuli could be used to trigger motor responses. The mechanisms that underlie other classes of behavior, however, remained largely unexplored by physiologists until quite recently. Over the course of the past 20 years, this has begun to change, and now three classes of behavioral decision making have come under scrutiny, principally in the visual-saccadic system of awake-behaving primates. The first of these classes, deterministic sensorimotor behaviors of the kind Sherrington examined in the spinal cord a century ago, is an area where particular progress has been made. Perceptual signals that trigger fixed motor responses, for example, have recently been examined at a neuron-by-neuron level. These studies have begun to reveal the outlines of the cerebral architecture that underlies decision making of this type. A second class of deterministic (or predictable) behaviors have also begun to be studied, behaviors that are controlled by a mixture of sensory and nonsensory signals. These are behaviors in which decision making is influenced by factors like the estimated value of a motor response or by probabilistic estimates of the current state of the external world. Finally, stochastic (or unpredictable) decision making has also begun to be studied physiologically. These studies of stochastic behaviors are beginning to suggest that even this class of decision making may soon be understandable, at the mechanistic level, in terms of identifiable cellular-level computations performed within a defined neural architecture.

The past 20 years has seen a tremendous growth in our understanding of each of these processes, and a basic theoretical outline is emerging that may be able to explain how many classes of visual-saccadic decision making are accomplished

by primate brains. This article reviews both the findings and the theories that constitute the neurobiological study of decision making today.

THE HISTORICAL ROOTS OF MODERN STUDIES OF DECISION MAKING

For scholars in the ancient world, understanding how the body of a man reacted to external events was a fairly unitary problem. Aristotle had suggested that the connection between sensation and action was accomplished within the physical confines of the human heart but through a nonphysical process mediated by the soul. Sensory data were, he argued, gathered by the nerves and vessels of the body and passed to the heart, where the human soul used that data to select a course of action to be executed by the neural and vascular systems of the body. For Aristotle the physiological systems of the body could, in this regard, be divided into sensory and motor divisions, which were causally linked by the nonphysical processes that took place in the human heart.

In the second century A.D. all of the ancient debate on this subject was codified, analyzed, and passed on to medieval scholars by the greek physiologist Claudius Galen. In dozens of books, many of which survive today, Galen (e.g., 1968, 1916) analyzed texts and performed critical physiological experiments of his own to better understand the nature of the human body. Galen's writings suggest that during his lifetime two principle theories dominated the debate about how decisions are made. Both theories argued that the nervous system could be fundamentally divided into two principle components, a sensory limb and a motor limb. They argued that these two components were intimately connected with an intermediate nonphysical decision-making process called the soul. They differed in where they placed the point of contact between these physical and nonphysical processes: one placed the soul within the human heart and the other within the ventricles of the human brain. Galen's approach to resolving this uncertainty was largely experimental. He noted that animals rapidly deprived of their hearts during sacrifices briefly retained the ability to move, whereas animals which had their spinal cords severed immediately lost all power of response. From this observation, Galen concluded that it must be the brain to which the sensory half of the nervous system carried data and from which the motor half of the nervous system received data.

Over the course of the ensuing scholastic period, Galen's conclusions came to dominate both western and arab thought on the connection between sensation and action. These conclusions suggested that the sensory and motor nervous systems could, at least in principle, be studied fruitfully by physiologists. But these conclusions also suggested that the actual process by which sensation and action were connected, the process of decision making, lay beyond the purview of anatomical or physiological study because it resided within a fundamentally nonphysical, nonmaterial realm. That changed during the early Enlightenment when the rise of a more materialistic world view led a number of thinkers, most prominently the seventeenth-century French philosopher René Descartes, to challenge this classical view.

To Descartes it seemed clear that the clockwork models with which Enlightenment scientists described the physical world could not account for the entire range of behavior that humans could produce (e.g., Gallistel 1980, Glimcher 2003) (see Figure 1). He argued that behavior, and the decisional processes that produce it, should be subdivided into at least two categories. The first of these categories, reflexes, involved simple deterministic linkages of the type that seventeenth-century mechanical models could be used to explain. Only the second, which accounted for more complicated classes of decision making, required the intervention of the nonphysical soul. Descartes even imported Aristotle's and Galen's division of the nervous system into sensory and motor components, arguing that the simple deterministic responses we call reflexes could be explained by assuming that sensory nerves carried highly specific information to a mechanical interface where sensory



Figure 1 Descartes' model of the reflex. Particles of fire make contact with the skin of the foot, displacing a thread that runs to the cerebral ventricles. The thread, in turn, opens a valve, which allows the ventricular pneuma to pass through a nerve-tube into a particular muscle. This pneumatic force leads to the contraction of the agonist muscle and a yoked relaxation of the antagonist muscle (Descartes 1664).

energy was passed directly to the motor system. In the first half of the seventeenth century Descartes made this proposal very explicitly, suggesting that the sensory and motor nervous systems, when they were engaged in producing a simple deterministic behavior, operated through a largely pneumatic system based in the ventricles of the brain and organized around the centrally located pineal gland.

The historical and practical importance of this proposal cannot be overstated. Descartes argued that human decision making could be viewed as the product of separable mechanisms, the understanding of which required fundamentally different approaches. This critically important insight still forms the core around which our approaches to understanding decision making are built. Nowhere can this be more clear than in the work of Charles Sherrington, whose investigations defined the modern physiological approach to the study of decision making.

Sherrington and Modern Studies of Decision Making and Behavior

Sherrington's work was based upon his acceptance of Descartes' proposal that a relatively simple class of behavior existed, which could be described as a fully determinate material process (Sherrington 1947). His avowed goal was to develop a complete mechanistic model for this class of behavior. He hoped this would reduce a large segment of human and animal behavior to conceptual and physiological tractability. In Sherrington's model, the relatively simple behaviors were those in which an event in the external world triggered a fixed behavioral response. These formed, for both Sherrington and Descartes, the simplest possible decision, a definable and determinate connection between the sensory and motor systems.

Sherrington (1906) built on the work of Aristotle, Galen, and Descartes when he argued that the nervous system responsible for this simplest class of determinate decision making could be viewed as being composed of three critical elements: a selective sensory element, which served as a detector for a restricted set of events in the outside world (a circuit element that he called the afferent limb); a selective motor element (which he called the final common path) that led to the activation of muscles; and a point of contact between these two systems (which he referred to as the integrative element). Sherrington, arguing explicitly from Descartes, suggested that this was the simplest nervous system that could accomplish determinate sensory motor linkages. Although he did not believe that this class of mechanism could underlie more complicated types of decision making, he argued convincingly that understanding reflexes would have to be the starting point for engaging more complicated forms of behavioral decision making.

During the first half of the twentieth century these insights dominated studies of decision making. Neurobiologists focused on understanding how the spinal cord, and to a lesser extent the brain, produced deterministic sensory motor linkages of the kind nearly all scientists and philosophers agreed were tractable to physiological investigation. Though there remained uncertainty in physiological circles about whether indeterminate behaviors actually existed, physiologists generally avoided the study of more complicated decision making.

MODERN STUDIES OF SIMPLE DECISION MAKING: THE SENSORY-MOTOR PROCESS

In the second half of the twentieth century huge advances were made in understanding the structure and function of both the sensory and motor systems of the mammalian brain. In landmark experiments, Mountcastle (1957), Kuffler (1953), Hubel & Wiesel (1959, 1962), and others (see also Lettvin et al. 1959) were able to elucidate the basic physiological architecture of mammalian sensory processing. They found that individual cortical neurons in some brain areas could be viewed as highly specific sensory receptors, each maximally activated by a specific pattern of events in the external world. These individual receptors were organized, in turn, into groups of neurons with closely related sensory properties. These groups of neurons formed topographic maps of the sensory epithelium or of the external world itself. The sensory systems, it seemed clear, were composed of hierarchically organized sets of these areas, which were responsive to more and more complicated and specific patterns of sensory stimulation.

Sherrington had argued that for the purposes of understanding simple determinate behaviors the nervous system could be viewed as interconnected sensory and motor subsystems. His model of the sensory system, the afferent limb, was a highly specific receptor that could act as a behavioral trigger. The suggestion that cortical sensory systems were organized to yield receptive fields that were much more sophisticated versions of the receptive fields Sherrington had encountered in the spinal cord fit well with the reflex model of how more complicated sensory motor behaviors might be implemented in the brain.

In parallel with this work on sensory systems, huge strides were also made in understanding those portions of the mammalian brain involved in movement control. In particular, advances were made in understanding how rapid shifts in the line of sight, eye movements called saccades, were produced. The work of people like Wurtz & Goldberg (1972), Bruce & Goldberg (1985), and others (e.g., Sparks & Mays 1990) suggested that sheets of neurons in areas like the cortical frontal eye fields and the mesencephalic superior colliculus formed topographically mapped representations of all possible saccadic eye movements. Activation of one of these maps at a particular point might be correlated with a 10° rightward movement, whereas activation of an adjacent point in the structure might be correlated with an 11° rightward movement. The complete surfaces of these cortical and subcortical maps, at least in primates, appeared to encode all possible saccadic eye movements. It was even shown that electrical activation of these maps at a given point triggered a saccade that was virtually indistinguishable from the saccade that followed normal activation at that site (cf. Robinson & Fuchs 1969, Robinson 1972). As a result, the saccadic eye movement system became the primary model for understanding movement control and included neurons that constituted what appeared to be a final common path for eye movement control. The frontal eye fields and superior colliculus appeared to be analogues, within the brain, of the final common path Sherrington had proposed in the spinal cord.

In the early 1900s Sherrington had argued that, at least within spinal circuits, simple predictable decision processes could be understood as the product of direct connections between specific sensory receptors and specific motor neurons. The growing knowledge of primary sensory and motor systems developed in the 1950s–1980s raised the possibility that a similar model might be able to describe much more complicated predictable sensory to motor behaviors that involved cortical processing. If the cortical sensory systems included specific receptors for complex sensory events, and the motor systems included final common paths for behaviors as complete as orienting eye movements, then it should be possible to study eye movements, which were the fully predictable product of sensory-perceptual experiences, using the Sherringtonian approach. Accordingly, a number of research groups began to examine decision-making processes of this kind—processes in which a perceptual judgment or stimulus triggered one of two possible saccades (cf. Newsome et al. 1989, Glimcher & Sparks 1992, Hanes et al. 1995). Two groups in particular began that examination at the cortical level: Newsome's group at Stanford University and Schall's group at Vanderbilt University.

Saccadic Decisions About Perceptual Motion

Newsome's group began their study of this class of sensorimotor processes by examining one of the most complicated sensory receptive field structures that had been well defined: the visual-motion sensitive neurons of the primate middle temporal area (area MT). Their hypothesis was that the perceptual experience of determining the direction of visual motion reflected the activity of discrete neuronal groups in area MT. To test this hypothesis, monkeys were trained to perform a simple sensory to motor decision-making task believed to involve both the visual cortices and the saccadic motor control maps of the frontal eye fields and the superior colliculus.

The goal of Newsome and his colleagues (1989) was to identify a visual stimulus that produced a perceptual experience of visual motion in human observers. They then used that stimulus to test the hypothesis that a perceptual evaluation, when made by a monkey subject, could be fully accounted for by an analysis of the neuronal activity in area MT. They accomplished this goal by presenting, to a monkey fixating a central target, a circular aperture within which a cloud of white dots moved chaotically against a black background for a period of two seconds (Figure 2). On each frame of this video presentation, a small and systematically variable fraction of the dots were moved in one of two possible directions. As a result the video presented a randomized motion signal, which included a variable amount of visual motion in one of two possible directions.

After 2 s, the dot display was terminated and two secondary targets were illuminated on opposite sides of the aperture, sides corresponding to each of the two possible directions of coherent dot motion. The monkeys were trained to evaluate these displays, reporting the direction of the coherent motion they had observed by making a saccade that shifted their point-of-gaze to one of the two secondary

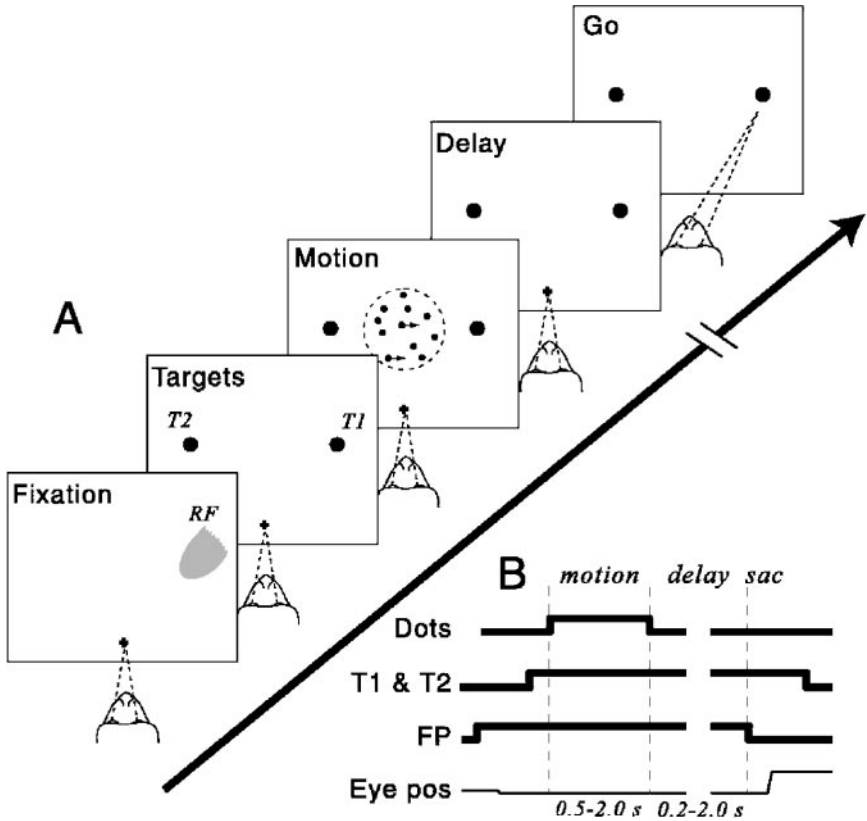


Figure 2 The moving dot task. Monkeys fixate a central point while a display of chaotically moving spots of light are presented within a circular aperture. On any given trial, a small fraction of the dots move in a coherent manner in one of two possible directions. Across trials, the fraction of dots moving in this coherent fashion can be varied systematically, to increase or decrease the strength of the perceived motion signal in either of the two possible directions. After viewing the display for 2 s, monkeys indicate the direction of perceived motion with a saccadic eye movement. Correct responses are reinforced with water or fruit juice. (From Shadlen & Newsome 2001. Reproduced with permission from the *Journal of Neuroscience*.)

targets. The monkeys were then reinforced, with water or fruit juice, if they had correctly identified the direction of dot motion. In essence, the monkeys had been trained so that at the end of each trial they produced a leftward or rightward eye movement, the direction of the eye movement triggered by the motion the animal had perceived.

When these experiments began, it was already known (Zeki 1974, Maunsell & VanEssen 1983a, Albright et al. 1984) that neurons in area MT become active

whenever a stimulus moves across their receptive fields in an idiosyncratic preferred direction. If a moving visual stimulus travels in this direction across the visual receptive field of an MT neuron, the neuron becomes highly active. If, however, the moving stimulus travels in the opposite, or null, direction then the activity of the cell is inhibited. Based on these data and others, it had been argued that MT neurons encoded, in their firing rates, the instantaneous strength of motion in their preferred directions.

Newsome and his colleagues hoped to show that the eye movements made by the animals could be predicted by the activity MT neurons produced while the motion stimulus was being presented. Accordingly, the monkeys were shown moving dot displays in which the fraction of dots moving in a correlated fashion was varied randomly. For any given direction and strength of movement, the probability that the monkey would select right, across many repetitions of each stimulus, was computed. This probability was then compared with the neuronal firing rate. Newsome and his colleagues (1989) found that for an average MT direction-selective neuron, there was an almost perfect correlation between neuronal rate and the probability that the monkey would make the saccade associated with that direction of movement (Figure 3). The likelihood that the monkey would make a rightward saccade was a lawful function of the firing rate of the rightward motion-preferring neurons in area MT. In an extension of this original finding, Salzman and colleagues (1990) were even able to show that electrical activation of neurons in area MT could alter the probability that the animals would produce a particular saccade. Stimulating neurons that preferred rightward movements systematically increased the likelihood that the monkeys would make rightward saccades.

Taken together, these results suggested that the movements of the animals could be predicted from activity in area MT when the monkeys were engaged in a simple sensorimotor behavior, a behavior much more complicated than the kinds of sensorimotor behaviors Sherrington had studied, but one that might well be of the same conceptual class. These results suggested that the sensory neurons in area MT could be viewed as a set of receptors, each capable of identifying a particular stimulus. The outputs of these neurons could then, in principle, be passed to a simple circuit that ultimately led to the activation of either leftward or rightward saccade-encoding.

To further examine the possibility that this simple reflex-like circuit might exist, Shadlen, Newsome, and their colleagues (Shadlen et al. 1996) developed a quantitative model that could account for the behavior of the animal and which outlined the neurobiological processes that would be required to transform these sensory signals originating in area MT into appropriate inputs for the final common paths of the superior colliculus and the frontal eye fields. What Shadlen & Newsome attempted was a formal description of the integrative elements that must intervene between the sensory and motor elements of the circuit they were studying.

Their model (Figure 4) proposed that a group of rightward motion-sensitive neurons in area MT pooled data according to a fully defined algorithm to yield an instantaneous estimate of the current strength of rightward motion in the moving




Figure 3 Relationship between neuronal activity and motion perception. The filled circles and accompanying solid line plot the average response of a single MT neuron when presented with different levels of correlated dot motion. The neuron responds strongly to high levels of correlated dot motion in its preferred direction. (Vertical axis for this plot is not shown.) As the fraction of dots moving in that direction is decreased, firing rate is reduced. As correlated motion increases in the nonpreferred direction, the neuronal firing rate is reduced even further. The open circles and accompanying dashed line plot the probability that the animal will make the eye movement associated with the preferred direction of the neuron, also as a function of dot correlation. Note the precise correspondence between the shapes of the neuronal and behavioral data. (From Newsome et al. 1989. Reproduced with permission from *Nature*.)

dot display. In a similar way, a group of leftward motion-sensitive neurons were hypothesized to extract an estimate of the instantaneous strength of leftward motion. Because the monkeys had viewed the motion stimulus for 2 s, this allowed them to sum 2 s of instantaneous motion information before making a decision. Accordingly, Shadlen and his colleagues proposed that the output of each of the neuronal pools of direction-sensitive neurons was summed, or integrated, over the 2 s display period, to produce an estimate of the average direction of motion presented during the stimulus interval. They next proposed that the neurons that integrated rightward activity should be able to directly activate areas in the colliculus and frontal eye fields that produced rightward movements, and that leftward integrating neurons should be able to activate circuits producing leftward saccades.

In essence, the Shadlen et al. (1996) model instantiated two reflex-like circuits, which connected highly selective sensory receptor systems, in this case the pools

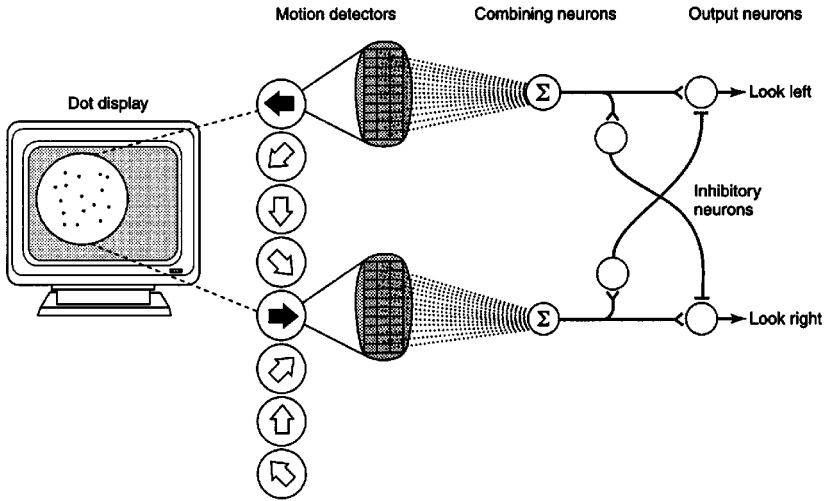


Figure 4 Shadlen and colleagues' (1996) model of a perceptual decision circuit. Pools of neurons in area MT extract the instantaneous strength of visual motion occurring in the display for motion in all possible directions. The instantaneous pooled estimates of motion strength in each of the two possible directions are passed to elements that compute the time integral of that signal to derive an estimate of the average motion signal over a 2-s display interval. These integrative elements project, in turn, to saccade-producing neurons. The integrative elements are postulated to be mutually inhibitory, assuring that only one eye movement is triggered at a time.

of leftward and rightward visual motion sensors in area MT, to a final common path for the generation of eye movements via a novel intermediate stage, which was presumed to perform a mathematical integration of the MT signal.

To make the system capable of decision making, in the sense of making choices, Shadlen and his colleagues recognized that the two possible movements would have to be mutually exclusive. It was essential that the model system not consist of two fully independent sensory motor systems both because animals cannot make rightward and leftward movements simultaneously and because the task the animals were performing required a single motor response. The monkeys had to decide whether to select a rightward or leftward saccade, and this binary nature of the response had to be accounted for by the model. To accomplish the decision, the model employed two inhibitory linkages, which allowed the output of each integrator to inhibit the other integrator's access to its own final common path.¹

¹In fairness to Sherrington (1906), it should be noted that he had proposed a similar strategy for the regulation of mutually incompatible spinal reflexes during walking, but he had probably given little thought to whether neural decision making in the brain might employ a similar strategy

This decisional phase thus consisted of a system of mutual inhibition, which permitted only a single movement to be expressed—the movement associated with the hypothesized integrator that showed the greatest activity.

The Shadlen model was a landmark in studies of how the brain might make decisions. It built upon the work of Descartes and Sherrington to propose a complete cellular and computational model of how a simple sensorimotor decision might be produced, and it made an interesting prediction. It proposed that somewhere between the motion-sensitive neurons of area MT and the hypothesized final common pathways of the frontal eye fields and the superior colliculus lay an element which could integrate the signal produced by MT neurons during the 2-s delay period. This integrative element, the formal model proposed, would need to have three critical properties: *i*) Rightward integrators would need to show a gradual increase in activity during the motion stimulus if the monkey would decide, at the end of the trial, to make a rightward movement; *ii*) the rate at which this activity increased should be a function of the fraction of dots moving into the right in the display—the higher the fraction of dots moving rightward, the faster the activity should grow within the rightward integrative element; and *iii*) in situations in which there was no net motion signal in the dot display, and the monkey was forced to guess, activity in the integrator should still predict the upcoming decision, which it was required to produce, even though in this case the decision was little more than a guess.

In order to test these predictions, Shadlen & Newsome (1996, 2001) recorded the activity of single neurons in the lateral intraparietal area (area LIP) of the posterior parietal cortex, while animals performed the dot motion task. Area LIP was selected as a likely site for the integrative element because it was known that area LIP both projects to the frontal eye fields (Barbus & Meshulam 1981, Asanuma et al. 1985, Lynch et al. 1985) and receives projections from extrastriate visual areas like MT (Maunsell & VanEssen 1983b, Seltzer & Pandya 1984). Further, like neurons in the frontal eye fields and colliculus, individual neurons in area LIP were shown to be most active before movements having particular amplitudes and directions (Gnadt & Andersen 1988). Some LIP neurons were thus associated with, for example, leftward 10° saccades and others with rightward 10° saccades. In short, area LIP already appeared, on anatomical and physiological grounds, to be an interface between the sensory and motor systems of the monkey brain (cf. Colby & Goldberg 1999, Andersen et al. 1997).

What Shadlen & Newsome found (Figure 5) was that on average the firing rates of LIP neurons associated with, for example, rightward saccades grew gradually during the 2-s delay interval if rightward motion was presented and a rightward saccade was later produced. In a similar way the activity of these neurons fell slowly on trials in which leftward motion was presented and a leftward motion was produced. This activity pattern thus validated the first, and most critical, of the model's predictions. Shadlen & Newsome also found that the rate at which the firing rate grew increased as the strength of the motion stimulus grew, validating the second of the model's predictions. Finally, on rare test trials in

which the motion stimulus was completely random (0% correlation), and neither rightward nor leftward motion was presented, the guesses of the monkeys were not only correlated with the strength of LIP activity, but the monkey's decision was also correlated with the strength of LIP activity before the motion stimulus was presented. In other words, it looked as if the LIP neurons were expressing a prestimulus bias toward one direction or another that, under these conditions, was correlated with the behavior of the animals. Again, this is an observation entirely compatible with Shadlen's sensorimotor decision model.

In summary, the Shadlen model was a concrete and quantitative description of how a complete neural circuit within the primate brain might produce a simple decision that was driven by sensory data. It accomplished this by employing a class of neurobiological mechanisms that Sherrington had developed in his studies of the spinal cord. It employed a set of identified neural elements to account for the actual decision-making behavior of conscious primates.

Deciding Which Target Is Different

At the same time that Shadlen, Newsome, and their colleagues were attempting to describe the processes that underlay their sensorimotor decision, Schall and his coworkers were working to understand the next stage in that same process, the mechanism by which sensory, or integrative, signals actually trigger a saccadic eye movement by activating neurons of the frontal eye fields. Schall and his colleagues (Hanes et al. 1995; Hanes & Schall 1996; Thompson et al. 1996, 1997) trained their animals to perform a very simple visual search task shown in Figure 6*a*. In their task, monkeys were presented with an array of eight targets arranged radially around a central stimulus that the animal was fixating. Seven of these targets were presented in a common color but the eighth, an oddball, was presented in a different color. The task of the animal was simply to make a saccade that shifted his point-of-gaze to the oddball as quickly as possible. In many ways this was a task very similar to the one Newsome, Shadlen, and their colleagues were investigating. It differed in that the decision it required could be made much more quickly, allowing Schall and his colleagues to focus on the speed, or reaction times, with which the animals were able to make these decisions.

The logic behind these experiments presumed that the visual signals produced by the eight disks would lead to the activation of specific sets of neurons in the visual cortices, just as the moving dot stimuli led to activation of neurons in area MT. The visual areas activated by the colored disks would, in turn, send projections to both the frontal eye fields and to parietal areas like LIP, which would in turn also send projections to the frontal eye fields. By studying a decision that was being made as quickly as possible, Schall and his colleagues hoped to use the time at which the movement occurred to identify the point in time at which the neural circuitry in the frontal eye fields was irrevocably committed to making a particular movement. They hoped to determine when the decision process Shadlen had identified in area LIP was complete within the final common path of the frontal eye fields.

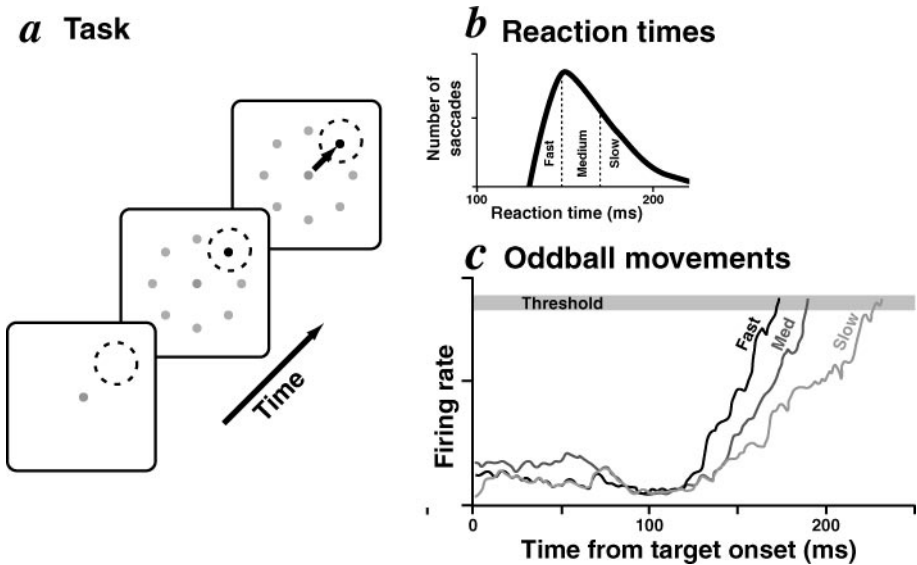


Figure 6 The oddball task. (a) Oddball trials began with the illumination of a central fixation stimulus. After a delay, eight eccentric targets are illuminated. One of those targets is presented, at an unpredictable location, in a unique color. If the monkey looks at the uniquely colored target, he receives a reward. (b) A variability in saccadic reaction times was observed in this task. (c) When the average firing rate was calculated for fast, medium, and slow reaction time trials, all trials reached a threshold level of activity at a fixed interval before saccade onset.

When a monkey shifts gaze toward a visual target that appears suddenly in the response field of a frontal eye field visuomovement neuron, the neuron responds with a fairly stereotyped firing pattern (Bruce & Goldberg 1985). It begins by firing an initial burst of action potentials shortly after target onset; it then fires at a lower rate for a brief period before it produces a gradually building level of activity that culminates in a second burst of action potentials shortly before saccade onset. The critical question Schall hoped to resolve was whether the gradual buildup of activity before the saccade-related burst reflected the completion of a decision-making process like the one Shadlen had studied in area LIP.

Schall and his colleagues answered this question initially by taking advantage of the natural variability in saccadic reaction times that occurs whenever animals perform the oddball task (Figure 6b). On some trials, animals make saccades to the oddball very quickly (~130 ms) and on other trials they respond more slowly (~200 ms). When hundreds of trials are examined, the reaction times for this task are typically distributed in a roughly log-Gaussian fashion around a modal value of about 150 ms.

To use this variability in reaction times to determine whether frontal eye field neurons participate in decision making, Schall and his colleagues had monkeys

perform hundreds of oddball detection trials while the activity of single frontal eye field neurons was studied. For each neuron studied, the trials gathered under these conditions were broken into three groups: those with reaction times shorter than, slightly greater than, or much greater than the average reaction time. Within each of these three groups of trials, Schall and his colleagues plotted the average neuronal firing rate as a function of time under two conditions: under the condition in which the oddball was inside the response field of the neuron they were studying and when one of the seven distractors was inside the response field. They found that on oddball and distractor trials, the neuronal response rose quickly to an early peak at a fixed interval after stimulus onset, but after about 80 ms or more the average firing rate on the oddball and distractor trials diverged. On distractor trials the firing rate fell at this point, whereas on oddball trials the firing rate was maintained or grew. The time at which these two firing rates diverged, they hypothesized, should reflect the time at which the neural circuitry was first beginning to make a decision about whether or not an oddball lay within the neuronal response field.

Many sophisticated models have been proposed in psychological circles to describe the decision-making process under reaction-time conditions like these (e.g., Sternberg 1969a,b; Ratcliff 1978; Luce 1986). Many of them suggest that the internal representations of the oddball and the distractors should continue to diverge until a threshold is crossed, at which point the movement itself should be irrevocably elicited. At a neurophysiological level this suggests that activity in the frontal eye field, perhaps shaped by activity in areas like LIP, might be driven toward a discrete biophysical threshold while the decision process is underway. Crossing that biophysical threshold, which might shift the frontal eye field neurons themselves (or the closely related neurons of the superior colliculus) into a higher firing rate (or burst) mode, might then serve as the physical trigger that executes the selected saccade.

To examine this hypothesis, Schall and his colleagues replotted the data described above, this time comparing the neural activity on fast, medium, and slow trials in which the target lay in the neuronal response field. They found that firing rates increased at different speeds on fast, medium, and slow trials but always seemed to reach a common firing rate, presumably a threshold of some kind, just before saccade onset (Figure 6c). Schall and his colleagues (Hanes & Schall 1995, Hanes et al. 1998) even went a step further, testing whether the threshold, once crossed, irrevocably triggers a saccade. They accomplished this by training animals to abort the saccade to the oddball whenever a visual cue called the stop signal was delivered. By examining the pattern of frontal eye field activity that was present when a stop signal led to a successful abort and comparing this to the activity pattern when a movement was produced anyway, they were able to develop further evidence for the existence of this neurophysiological threshold. Based on all of these data, Schall and his colleagues were able to suggest that the visuomovement neurons of the frontal eye fields reflected the gradual process of neural decision making by which sensory signals give rise to an eye movement.

Visual Saccadic Decision Making in a Sensorimotor Circuit

When these results gathered in area LIP and the frontal eye fields are considered together, a picture of the decision process during predictable sensory-driven decision making begins to emerge. During visual-saccadic decision making of the type required by the moving dot and oddball tasks, neurons of the visual cortices appear to act as receptor systems that perform the function of the afferent arc in the original Sherringtonian model. This sensory signal is, in turn, passed to parietal cortex, which combines information across time. This signal, perhaps via a mutually inhibitory network of some kind, seems to be passed to the frontal eye fields. A movement occurs when this network produces activity in frontal eye field neurons that crosses a biophysical threshold (either in cortex or in the colliculus). Of course, much of this picture is speculative, and it is far from clear that this is all that underlies these decisions. Many other neurons must participate in this process, and for many tasks area LIP may not even participate. But as we understand this class of behavior today, this seems to be the general outline of the decision-making circuit for simple sensorimotor processing in the visual saccadic system.

Variability in Reaction Times

These observations, and the model they suggest, leave many important questions unanswered. Principle among these is how can we account for the differential reaction times that Schall, and others doing similar experiments, routinely observed in these very simple tasks. How is it that a sensorimotor circuit of the type Shadlen & Newsome proposed gives rise to variable reaction times when presented with the same stimulus? Carpenter and his colleagues at Oxford (Carpenter & Williams 1995, Reddi & Carpenter 2000) have addressed this question in studies of humans. Their experiments, in shedding light on that question, may have provided another insight into the general shape of the computations performed by the visual-saccadic decision-making system.

In those experiments, human observers, while fixating, were suddenly presented with one of two possible saccadic targets. Their job was simply to make a saccade that shifted gaze as quickly as possible into alignment with whichever target was presented. Like the oddball task, this is a reaction time sensorimotor task requiring a simple decision. And like the oddball task, when the saccadic target is equally likely to appear at either location, the subject's reaction times show a fixed log-Gaussian distribution. Carpenter and his colleagues examined this reaction-time variability more closely, however, by systematically manipulating the probability that the target would appear at each of the two locations. They found that if a human observer was presented with a block of trials in which the target was more likely to appear on the left than on the right, reaction times to the leftward target were reduced, while reaction times to the rightward target were increased. More precisely, they found that reaction time was a lawful function of the logarithm of the likelihood that the target would appear at a particular location. This was a relationship that suggested that the rate at which activity in the visuomovement

neurons of the frontal eye fields rose to threshold in these human observers might reflect an integrative process that encoded something quite broadly useful. If neurons in areas like LIP or the frontal eye fields were computing the likelihood that the movement they encode will be required, then Shadlen & Newsome's architecture might be capable of even more complicated decision making than had been previously suspected.

The first scientists to examine that possibility explicitly were Gold & Shadlen, and they examined it by revisiting the moving dot task. Gold & Shadlen (2000) wanted to examine a reaction-time version of the task, while studying neurons in the frontal eye fields. They wanted to test the hypothesis that the buildup to the threshold Schall had observed, and which was presumably the product of LIP activity, could be accurately modeled as a calculation of the logarithm of the likelihood that a particular movement would be reinforced. The moving dot task seemed an ideal way to test this hypothesis because the monkey's estimate of the likelihood that a leftward eye movement would be reinforced could be expected to grow whenever the monkey was staring at a leftward moving dot display. And further, the rate at which that likelihood grew, or fell, should be determined by the fraction of dots that the experimenters moved to the left.

To test the hypothesis that the rate at which neurons in the frontal eye fields increase their firing rate toward a threshold is correlated with the rate at which estimates of log-likelihood should be growing, Gold & Shadlen trained monkeys to perform a reaction-time version of the moving dot task, which would allow them to measure both behavioral and neuronal thresholds while varying the fraction of dots moving in a particular direction. They then placed stimulating electrodes in the frontal eye fields of these monkeys. In 1983, Schiller & Sandell showed that when a monkey is passively fixating and a brief electrical stimulation is delivered to a fixed location in the frontal eye fields, a saccade is evoked at short latency. When this electrical stimulation is delivered just before the monkey makes his own saccade, the stimulation-elicited movement is found to be an average of the stimulation-associated movement and the animal's own movement. As the frequency of stimulation is increased or decreased, the endpoint of the saccade can even be shown to shift toward the stimulation-elicited movement or toward the monkey's own movement, respectively. This led Gold & Shadlen to conclude that electrical stimulation might be expected to elicit a movement that would be the average of the stimulation-elicited movement and the movement being planned by the monkey, weighted by the instantaneous likelihood (encoded in the frontal eye field firing rate at the time of stimulation) that the planned movement would be reinforced. What they proposed, in effect, was that the direction of the stimulation-induced movement would shift gradually, during the reaction-time interval, toward the movement that would eventually be made, and that the rate of this shift would reflect the rate at which the neurally encoded ratio of the log likelihoods was shifting.

Remarkably, this is exactly what was found, as shown in Figure 7. The schematic shown in Figure 7A plots, with a grey arrow, the movement that might be produced by stimulation during passive fixation. The grey arrow plots the

movement that might be produced if the monkey was viewing a moving dot pattern in which a large fraction of the dots moved upward. The purple arrow, which reflects the difference between these two movements, would thus presumably reflect the strength of the monkey's plan to make an upward movement. Gold & Shadlen found that the longer the same moving dot pattern was viewed by the monkey before the stimulation was delivered, the more the stimulation-induced movement was found to be shifted. This indicated that the stimulation was combining with a growing signal like the one Schall had observed previously in the frontal eye fields during the oddball task. Perhaps just as impressive was the observation that the rate at which the movement shifted and the extent to which it shifted were found to be a lawful function of the pattern of dots presented in the motion stimulus (Figure 7B). This was exactly what one would expect, at a qualitative level, if the internal signal being combined with the electrical stimulation encoded something like a log-likelihood.

In a final step, Gold & Shadlen (2000, 2001) examined whether the data they had gathered could be formally modeled as the signature of a log-likelihood calculation like the one Carpenter had originally identified behaviorally in his human reaction-time experiments. To test that possibility, Gold & Shadlen developed a model that could estimate likelihood functions for the visual stimuli that the monkeys had been shown and, using an approach similar to Carpenter's, could use these likelihood estimates to predict the amplitudes and directions of the stimulation-elicited movements they had observed. What they found was that early in the delay interval the likelihood computation, more formally the log of the likelihood ratios of the two possible movements, exactly predicted the data that they had obtained, which strongly supported their hypothesis.

Summary

The picture that emerges from these studies of sensory-driven decisions is one that probably would have been acceptable to both Descartes and to Sherrington, although it has turned out to be a good deal richer and more complicated than the reflexes they studied. In visual-saccadic decision making of this type, it appears that sensory elements, acting as detectors, gather data and pass it to what both Shadlen and Sherrington would call integrative elements. Our current evidence leads us to suspect that these integrative elements in areas like LIP fire at a rate correlated with something like the likelihood that the movement they encode will be reinforced. The output of this likelihood computation, in turn, appears to be passed to the frontal eye fields where the firing rates of neurons rise toward a threshold, perhaps instantiated as a biophysical state change in collicular or frontal eye field neurons which, when crossed, initiates a high frequency burst of action potentials that triggers a saccade.

Of course, many critical questions remain to be answered. For example, we do not know what other areas participate in this process and how the many areas that might be expected to collaborate in decision making interact or are synchronized.

Also, it must be acknowledged that much of this story remains speculation; but while other interpretations are certainly possible, this does seem to be a plausible account of how a visual stimulus might be used to trigger an eye movement. It is a simple model of a simple decision, and it reflects the best estimate to date of how a brain circuit might make a choice, albeit one based entirely on accumulating sensory data. If this is the case, then how might this circuit, or others like it, manage more complicated decisions? How, for example, might neuronal decision-making circuits incorporate nonsensory data into the decision-making process? Would this require entirely different neuronal circuits or might the neuronal process studied by Newsome, Schall, Shadlen, and their colleagues be generalized to account for decision making of that type as well?

MORE COMPLICATED DECISIONS: USING STORED INFORMATION

The simplest possible class of decision making, the one Descartes and Sherrington had both described at a mechanistic level, occurs when a stereotyped sensory event triggers a fixed motor response. Nearly all of the behaviors that Sherrington studied were of this type. These were behaviors that could be modeled as the product of direct connections between sensory and motor elements of the nervous system—direct connections that existed within the nervous system before the triggering stimulus was delivered. Both the moving dot task and the oddball task are far more complicated than any behavior Sherrington examined, and they undoubtedly engage much more complicated neural circuits than those Sherrington studied, but they are also behaviors that can be modeled as the product of preexisting direct connections between sensory and motor systems.

Consider a variant of the moving dot task in which rightward motion is more likely to occur than leftward motion, much like the behavioral experiments of Carpenter and his colleagues (Carpenter & Williams 1995, Reddi & Carpenter 2000). This would be a situation in which the pretrial likelihood, or more formally the prior probability, that a rightward saccade will be rewarded is greater than the prior probability that a leftward saccade will be rewarded. We know from the Carpenter experiments that if one response is more likely to be correct than another, subjects produce that response more quickly. If an animal performing the moving dot task knew that a rightward movement was twice as likely to be reinforced as a leftward movement, we might expect these prior probabilities to systematically bias the reaction times, and perhaps even the decisions, of the animals. In a similar way, if the animal subject knew that a rightward movement, if correct, would yield twice as much reward as a leftward movement, then we might also expect this to influence the decision-making process. How would such a bias, which reflects nonsensory influences on decision making, be instantiated in neural circuitry, and how might this circuitry interact with the hardware employed for more direct sensory to motor decision making? Would decision making of this type involve a completely distinct process, as Descartes might have supposed, or would it involve

the same systems Newsome, Schall, Shadlen, and others have examined? In the mid-1990s three groups began to examine these questions by studying the activity of neurons in the superior colliculus and the posterior parietal cortex while animals engaged in decision-making tasks that involved nonsensory data.

The Influence of Changing Prior Probabilities on Collicular Processing

Basso & Wurtz (1997) were the first to explicitly examine, while recording the activity of collicular neurons, the neural representation of changing prior probabilities in the visual-saccadic decision-making system. Prior to their work, it had been shown that in the superior colliculus, as in the frontal eye fields, neurons produce a premovement prelude of activity, a gradual increase in firing rate that occurs before a saccade and that rises toward what appears to be a biophysical threshold mediated by local excitatory circuitry (Wurtz & Goldberg 1972, Sparks et al. 1976, Pettit et al. 1999). Once the firing rate crosses this threshold, the neurons begin to produce a very high frequency burst of activity that is followed by a saccade. Like neurons of the frontal eye fields, individual neurons of the superior colliculus encoded particular movements (Wurtz & Goldberg 1972), and Glimcher & Sparks (1992) had used this property to show that the collicular prelude of activity could be used to predict, under many conditions, the saccade that the animal would later produce. They had suggested that the prelude reflected an impending plan to produce a particular movement.

Basso & Wurtz hypothesized that the collicular prelude, or buildup, of activity might encode not simply the movement that would be made, but rather the likelihood that the movement encoded by that neuron would be made. To test this hypothesis they presented animals with one, two, four, or eight possible targets, while recording from single neurons in the superior colliculus. After a delay, all but one of the possible targets, which was selected at random, were extinguished, and the animal was rewarded for shifting gaze to the remaining target. Critically, in their experiment, one of the targets being presented was always placed within the response field of the neuron they were studying. This meant that the prior probability that the stimulus associated with their neuron would be the reinforced target varied from 100% down to a probability of only 12.5%. Basso & Wurtz found that during the delay interval the magnitude of the collicular prelude was a function of the number of potential targets that were visible. The prior probability that the stimulus within the neuronal response field would serve as a saccadic goal seemed to be controlling the magnitude of the premovement prelude.

In a similar way, Dorris & Munoz (1998) explored the activity of collicular neurons while prior probability was manipulated, under conditions similar to those Carpenter and his colleagues had examined behaviorally. In Dorris & Munoz's experiment, animals were presented with blocks of trials in which a single visual target appeared at one of two possible locations, one of which lay within the response field of a collicular neuron under study. They then manipulated the

probability, within a block of trials, that the stimulus would appear at the location within the response field. They found that the rate at which collicular activity grew during the prelude was a function of the likelihood that the movement would be reinforced. Furthermore, they found that the reaction times of the animals were tightly correlated with this rate of growth.

Both of these experiments showed that in the superior colliculus, a principle target of the frontal eye fields, activity levels can reflect the likelihood of a saccade being required. The activity of the collicular neurons was shown, in these experiments, to reflect something more than activity along a sensory to motor pathway. This was a critical observation because it suggested that the sensory to motor circuit that had been studied with the moving dot task and with the oddball task might not be a specialized circuit for purely sensory to motor decision making. Instead, their results suggested that this circuit might also be able to participate in at least two kinds of decision making: sensorimotor decision making and non-sensory decision making. To begin to examine that possibility in more detail, Platt & Glimcher (1999) undertook a series of experiments in parietal area LIP aimed at determining how changes in prior probabilities and in the value of particular saccades influenced the activity of saccade-related neurons in area LIP.

Prior Probabilities and Reward Values in Parietal Cortex

Platt & Glimcher (1999) trained monkeys to perform the cued saccade task (Figure 8*a*). While the animals fixated straight ahead, a pair of target stimuli were illuminated, one red and one green. After an unpredictable delay, the fixation target changed color to either red or green to indicate which of the two stimuli served as the future saccadic goal. After a second unpredictable delay, the fixation stimulus was extinguished, and the animal was reinforced if it shifted gaze to the saccadic goal within 500 ms.

The goal of the experiment was to determine whether neurons in area LIP encoded the prior probability that a given stimulus would serve as a saccadic goal or the amount of reward that looking at that stimulus would produce, independent of any of the simple sensory and motor variables that had been previously examined in these neurons. To achieve this goal Platt & Glimcher developed an experiment in which the animals would be presented with exactly the same visual stimulus and would have to make exactly the same motor response, while either the prior probability or the reward magnitude associated with the movement encoded by the neuron under study was manipulated. One feature of the cued saccade task that they exploited to accomplish this is that on all trials in which the red target is presented at the same location and in which the fixation stimulus turns red, a fixed sensory motor correspondence is produced. The same stimulus is presented, and the same movement is made. If under these conditions the relative likelihoods of red and green trials are systematically manipulated across blocks of trials, then the prior probability that a particular movement will be reinforced can be varied, while the sensory and motor correspondences within each kind of trial are held constant.




Figure 8 The cued saccade task. (a) While fixating a central stimulus, monkeys are presented with two eccentric targets, one inside the response field of the LIP neuron under study. After a delay the central stimulus changes color, identifying one of the eccentric targets as a movement goal and the other as a distractor. In sequential blocks of trials, the likelihood that the central stimulus will turn red or green is manipulated. (b) When trials with identical sensory and motor properties, but from different blocks, were compared, LIP firing rates were found to be strongly influenced by the likelihood that the movement they encode will be required. Solid lines plot average firing rates. Rows of tickmarks show spike times during individual trials. (c) When blocks having many different likelihoods were examined at six sequential epochs during the trials, firing rate was found to be a lawful function of the instantaneous probability that a movement would be required.

The first cued saccade experiment was aimed in this way at determining whether neurons in area LIP encoded prior probabilities independent of either the sensory or motor aspects of the decision the animal was making. In that experiment, animals were presented with, for example, a block of 100 cued saccade trials within which the fixation stimulus was 80% likely to turn red (identifying the rightward movement as rewarded) and 20% likely to turn green (identifying the leftward movement as rewarded). That might be followed by a block in which the fixation stimulus was 20% likely to turn red and 80% likely to turn green. Single neurons were studied this way with five to seven blocks of trials in which the prior probability that the fixation stimulus turned red was systematically varied.

To meet the goal of holding all of the sensory and motor properties of the decision-making system constant, only those trials in which the final color of the

fixation stimulus instructed the animal to make a gaze shift into the response field of the neuron under study were analyzed. Critically, on all of these preselected trials the monkey would have been presented with exactly the same stimulus, and the monkey would have made exactly the same response. The only thing that would have varied across blocks would be the probability that, at the end of the trial, the movement into the response field would be required.

Figure 8*b* shows how a typical neuron in area LIP responded under these conditions. Both the thick black line and the thick grey line plot the average firing rate of the neuron on trials that are identical. The trials plotted in black were, however, drawn from a block in which the central stimulus was 80% likely to turn red, while the grey line plots data from a block of trials in which the central stimulus was only 20% likely to turn red. Note that even though the sensory and motor properties of the trials were identical, the neuron responded quite differently when an underlying nonsensory decision variable was changed.

In fact, there appears to be significant structure in the way that the response evolves over the timeline of these trials. Early in the trials, when the eccentric stimuli are first illuminated, there is a very large difference between the neuronal firing rates observed under the two conditions. The firing rates then converge after the fixation stimulus turns red, an event identifying the gaze shift encoded by this neuron as the reinforced response on all of these trials. In a sense then, this LIP neuron appears to encode not only the prior probability that the movement will be reinforced across blocks, but also seems to carry information related to the instantaneous likelihood that the movement will be reinforced. Early in the trial there is either a 20% or 80% probability of reinforcement associated with the movement encoded by this neuron. Later in the trial there is (for both blocks) a 100% probability of reinforcement, and at this time the firing rates in both blocks converge. Figure 8*c* extends this analysis, plotting the average firing rate as a function of prior probability during six sequential epochs of time. The first graph shows this correlation shortly after the two eccentric stimuli were presented but before the fixation light had changed color, and the last graph shows this correlation just after the movement was produced. Note the strong correlation for this neuron between the prior probability and the firing rate early in the trial. Platt & Glimcher found that this was fairly typical of neurons in area LIP; 75% of the neurons they studied showed a statistically significant correlation between prior probability and firing rate early in the trial, before the fixation stimulus indicated to the animal that on all of these trials the same movement would be required.

To examine the closely related possibility that neurons in area LIP might encode information about the value of a movement to the animal, Platt & Glimcher performed a second experiment. Animals were again presented with blocks of 100 cued saccade trials, but for this experiment the likelihood that the fixation stimulus would turn red or green was always fixed at 50%. Across each block it was the amount of reward that the animal received for correctly performing trials on which the fixation stimulus turned red or green that was varied. In the first block the

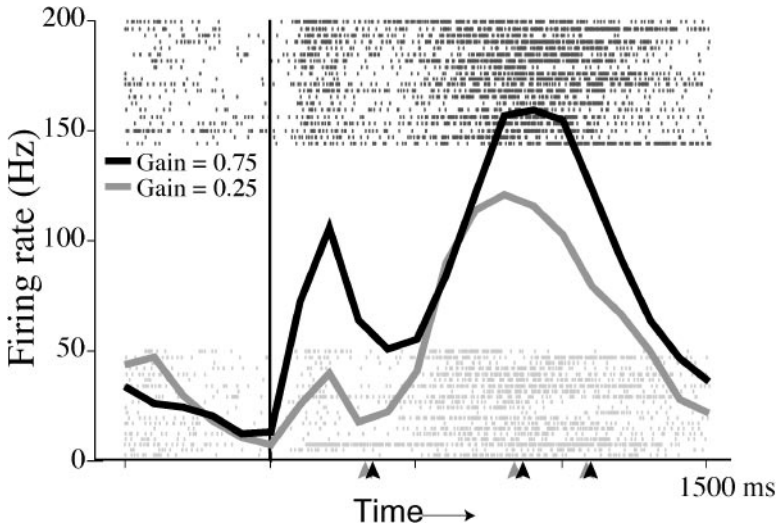


Figure 9 When the magnitude of the reward associated with a movement was varied in the cued saccade task, LIP firing rate was shown to be a function of this variable. Solid lines plot average firing rates. Rows of tickmarks show spike times during individual trials.

animal might receive 0.2 ml of fruit juice as a reward on red trials and 0.1 ml of juice on green trials. In a second block he might receive 0.1 ml on red trials and 0.2 ml on green trials.

Figure 9 plots the behavior of an LIP neuron under these conditions. Again, only trials that were identical in their sensory and motor properties were selected for use in this figure. The trials differ only in the magnitude of reward that the monkey could expect to receive for making this same movement into the response field of the neuron we study. Note that the neuron fires more strongly during trials in which the animal could expect to receive a large reward (average firing rate shown by the black line) and more weakly during trials in which the animal could expect to receive a small reward (grey line). As with prior probability, the data indicate that neurons in area LIP encode the values of movements even when the sensory and motor properties of the task are held constant.

A second interesting finding that Platt & Glimcher made (M.L. Platt and P.W. Glimcher, unpublished observations) was that little or no correlation was observed between reward magnitude and firing rate when only a single target was presented. The importance of this observation became particularly clear when they noted that it was the relative magnitudes of the leftward and rightward rewards (Platt & Glimcher 1999) that were linearly correlated with LIP firing rates. This suggests that what neurons in area LIP were encoding in this experiment was something

more like the value of each movement, relative to the available alternatives, than the absolute value of each movement.²

Taken together, all of these data suggest that when animals must take into account either the relative value of each movement or the prior probability that each movement will be reinforced, neurons in area LIP and in the superior colliculus appear to reflect this fact. Thus the data presented so far suggest, as a preliminary hypothesis, that sensory signals are analyzed by areas like the extrastriate visual cortices and are then combined with nonsensory signals that encode the prior probabilities and values of outcomes to an animal. These combined variables appear to be represented in posterior parietal cortex and in its targets, where those signals can produce saccades by driving some classes of neurons above a threshold.

This outline, however, raises two critical sets of questions that have to be answered if an analysis of neural decision making is to proceed. First, how, at a computational level, might we expect all of these sensory and nonsensory variables to be combined by a biological decision making system? Second, where are the nonsensory variables, encoding the prior probability and value of alternative outcomes, stored and updated?

Formal Models of Sensory and Nonsensory Decision Making

Unlike physiologists, social scientists have a long history of thinking about both sensory and nonsensory decision making in computational terms, and almost all of the models within that tradition can be traced to the work of the Enlightenment mathematician Blaise Pascal (Arnauld & Nicole 1662, Pascal 1670; for a review of this subject see also Kreps 1990, Glimcher 2003). To understand Pascal's approach to the problem of how nonsensory factors should influence decision making, consider deciding how to play the game of roulette. Imagine being offered the opportunity to bet \$50 either on the color red or on the number 17. How, Pascal wondered, should one decide efficiently between these two alternatives? To make that decision Pascal argued that one had to know both the likelihood that you would win each of the two wagers and the amount that each wager was worth if you did win. Given these two pieces of information, Pascal argued that one could determine the value of each wager on average by multiplying these two quantities—a computation which yields what modern economists call the expected value of the wager. A rational decision maker, Pascal argued, then simply chooses the wager with the higher expected value. In an ideal casino the wager resolves this way:

$$\begin{aligned} \text{Bet red; probability of winning } 18/36 = 0.5; \text{ gain} = 100; \\ \text{expected value} = 0.5 \times 100 = 50; \end{aligned} \quad 1.$$

²It should be noted that psychologists have often argued that it is the relative value of the options available to a decision maker, rather than the absolute value, that should influence decision-making behaviors.

Bet 17; probability of winning $1/36 = .0277$; gain = 1800;
expected value = $0.0277 \times 1800 = 50$. 2.

In this case, the two wagers have identical expected values. This means if you play this game of roulette ten thousand times, regardless of whether you wager red or 17, you stand to win an essentially identical amount of money. Therefore, Pascal would have argued that an efficient decision maker is one who is indifferent between betting on red or on 17. Pascal's insight is critical because it suggests that the process of efficient decision making, or what economists call rational decision making, can be modeled at a computational level as a mechanism that combines the likelihood and magnitude of a gain to determine the value of a course of action.

In the middle of the eighteenth century two other critical advances occurred that revised Pascal's original approach and laid the foundations for modern economic theories of choice. Whether the economic analysis is used in behavioral ecology to compute expected evolutionary fitness or in market analysis to determine the worth of a corporation, all economic analyses at least have their roots in these advances. These two advances were the development of expected utility theory and the development of the Bayesian theorem.

Expected utility derives from a simple observation that was made by a number of mathematicians over the course of the eighteenth century. In the roulette example described above, Pascal would have argued that humans, if rational, should be indifferent between the wager on red and the wager on 17. Empirically, however, eighteenth century scientists observed that under conditions like these humans were not indifferent. When asked to choose between two bets of equal expected value but with different probabilities of winning, humans routinely select the lower risk bet.³ This aversion to risk was formally characterized within decision theory by the Swiss mathematician Daniel Bernoulli (1954 [1738]) using a concept he called utility. Rational decision makers are, Bernoulli argued, naturally prudent. Whenever they have to choose between two options of equal expected value that present different levels of risk (and hence must offer different levels of rewards if they are to have the same expected value), decision makers always prefer the lower risk option. It was as if, when computing expected value, the higher rewards (necessarily associated with higher risks) had less influence on decision making than expected, as if the value of a gain to the decision maker grew more slowly than the numerical value of the gain. Placing this within the framework of Pascal's expected value theory, Bernoulli chose to include risk sensitivity in models of decision making by suggesting that humans do not multiply likelihood with gain directly but instead multiply likelihood with a concave function of gain called utility. Thus utility, the subjective value, which Bernoulli argued guides choice behavior, is a function of value as shown in Figure 10. It is the product of this quantity and likelihood, a product known as expected utility, that Bernoulli argued

³As do animals.

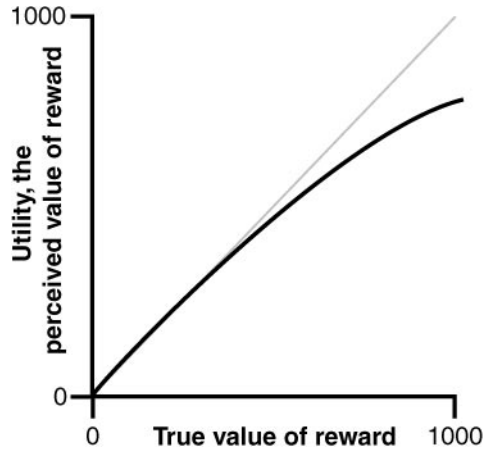


Figure 10 A utility curve. The black line plots the relationship between the objective values of rewards and the subjective utilities of those rewards. Utility grows more slowly than value, an idea consonant with the assumption that rational decision makers are averse to risk.

represents the decision variable employed when efficient choices are made. (See Kreps 1990 for an overview of expected utility theory with regard to economics or Glimcher 2003 for an overview of the theory with regard to neurobiology.)

The second critical advance in the development of formal models of decision making was Bayes' theorem for estimating likelihoods. The Bayesian theorem (Bayes 1958 [1763]; see also Duda et al. 2000 for a review of the Bayesian theorem) is a mathematical tool for accurately estimating the likelihood of any event. It does this by allowing us to ask the following question formally: Given my prior knowledge of how often I have observed that the world appeared to be in state x , and given my knowledge of how well correlated my current sensory data is with the actual world state x , then precisely how likely is it that the world is in state x ?

Imagine a monkey trained to fixate a spot of light while two eccentric spots of light are also illuminated, as in the cued saccade experiment. In this experiment the central fixation light changes color to indicate which of the two eccentric target lights, the left one or the right one, will serve as the goal on this trial. However, the color of the central fixation light (or more precisely the wavelength of the light emitted by the central stimulus) can be any one of a hundred different hues (or wavelengths). A Bayesian description of this task begins by observing that there are two possible world states from the point of view of the decision maker.

In mathematical notation we designate these two world states as w_1 and w_2 . State w_1 exists when a leftward saccade will be rewarded, and state w_2 exists when a rightward saccade will be rewarded. After observing 100 trials, the subject notes that on 25% of trials a leftward movement was rewarded, irrespective of the color of the

fixation light, and on 75% of trials the rightward movement was rewarded. Based upon this observation it can be said that the prior probability that world state w_1 will occur [known as $P(w_1)$] is 0.25, and the prior probability of world state w_2 is 0.75.

To make these prior probabilities more accurate estimates of the state of the world, one next has to take into account the color of the central fixation stimulus and the correlation of that stimulus color with each of the world states. To do that one needs a graph that plots the probability of encountering a particular stimulus wavelength (which we will call λ) when the world is in state w_1 . Figure 11A plots an example of such a probability density function showing the likelihood of each value of λ when the world is in state w_1 and when in state w_2 . This is referred to as the conditional probability density function for λ in world state w_1 , or $P(\lambda | w_1)$.

Next, in order to get the two graphs in Figure 11A to indicate how likely it is that a given λ was observed and that the world is in a given state, these graphs have to be corrected for the overall likelihood that the world is in either state w_1 or state w_2 . To do that, one multiplies each point on the graphs by the prior probability of that world state. The graph on the left thus becomes: $P(\lambda | w_1)P(w_1)$, where $P(w_1)$ is the prior probability for world state w_1 as described above. Note that in Figure 11B this has the effect of rescaling the graphs that appeared in Figure 11A. Finally, one has to determine how likely it is that any given value of λ will occur regardless of world state. To do this, one needs to count up all the times that λ has been observed at a specific value and then plot the probability density function for all values of λ (irrespective of which movement was rewarded) as shown in Figure 11C.

With these data, it becomes possible to ask, when a given wavelength of light is observed, what the likelihood is that on this trial a rightward movement will be rewarded (that the world is in state w_1) and what the likelihood is that a leftward movement will be rewarded (world state w_2). To compute these likelihoods one divides the curves shown in Figure 11B by the curve shown in Figure 11C. This essentially corrects the likelihood that one would see a particular λ in a particular world state for the overall likelihood that one would ever have seen wavelength λ , a likelihood shown in Figure 11D. This is the essence of the Bayesian theorem given by the equation:

$$\text{Probability of } w_1 \text{ given the current value of } \lambda = P(\lambda | w_1)P(w_1) / P(\lambda). \quad 3.$$

To restate this one could say: The best possible estimate of the probability that a rightward movement will be rewarded is equal to the probability that the central stimulus would be this color on a rightward trial multiplied by the overall probability of a rightward trial divided by the probability that this particular color would ever be observed. The result is usually referred to as a posterior probability, and it reports, in principle, the best estimate that you can derive for this likelihood. Therein lies the absolute beauty of Bayes' theorem. Bayes' theorem provides a mechanical tool that can report the best possible estimate of the likelihood of an event, and it computes this likelihood by combining a nonsensory estimate of prior probability with current sensory data. The Bayesian theorem is a critical advance because no decision process that must estimate the likelihood of an uncertain

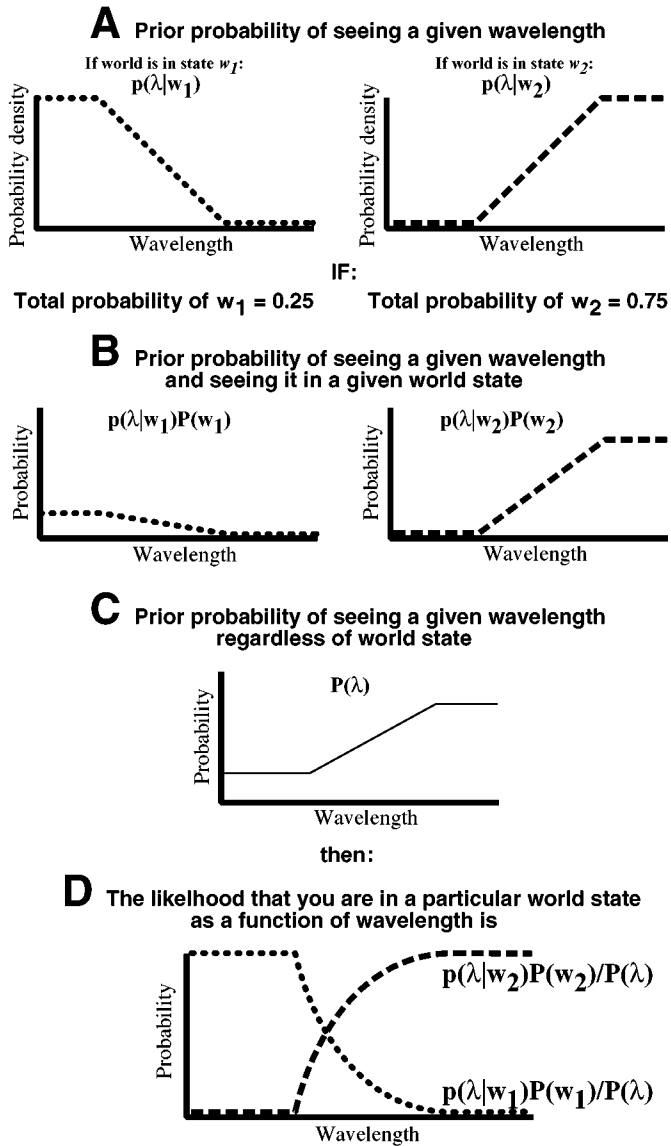


Figure 11 The Bayesian theorem. Graphs present an example of using the Bayesian theorem to compute likelihoods. See text for details.

outcome can ever do better than a Bayesian estimate of that probability based on both prior probabilities and current sensory evidence.

At a theoretical-computational level these two insights, a description of how likelihoods can be estimated and a description of how choices can be made efficiently, suggest a model for how we might think about decision making within the nervous system. It suggests that sensory data and prior probabilities must be combined, at a site not yet identified, to yield an estimate of posterior probabilities. Next, that posterior probability has to be combined with an estimate of the value of each outcome to derive an expected utility for each behavioral response, which can then guide choice.

The available data does seem to suggest that the primate nervous system may actually be performing a calculation of this general type; all of these values have already been demonstrated in the activity of neurons in, or near, posterior parietal area cortex. Prior probabilities have been shown to influence activity in area LIP and in the superior colliculus (Basso & Wurtz 1997, Dorris & Munoz 1998, Platt & Glimcher 1999). Likelihoods based on sensory data, a quantity which is closely related to posterior probability, have also been shown to be encoded in LIP activity and to most likely reflect computations performed in extrastriate visual neurons (Shadlen & Newsome 1996, 2001). Finally, the relative value of rewards, perhaps their relative expected utility to the animal, has been shown to be encoded in area LIP (Platt & Glimcher 1999).

Of course the hypothesis that these neurons participate in an expected utility computation raises other critical questions. Where are prior probabilities (or less formally, memories that guide decision making) encoded and updated and how are the values of outcomes encoded or updated? Human neurological studies, human functional imaging studies, and physiological studies of monkeys evaluating rewards are beginning to provide answers to these questions.

Brain Imaging and Studies of Decision Making

The first broadly influential studies to evaluate the neurological basis of human decision making in economic terms were those by Bechara and colleagues (1994, 1997). Those studies focused on describing the capabilities of humans with damage to ventromedial frontal cortex in a decision making task that relied on stored, nonsensory information. Human subjects, some with injuries to the ventromedial prefrontal cortex, were presented with four decks of playing cards and given a loan of play money. They were instructed to turn over cards, one at a time, from any of the four decks. The backs of the cards were printed with a positive or negative dollar value and subjects won or lost the value, in play money, that was printed on each card they turned over. The task of the subjects was simply to identify the most profitable deck of cards and to continue selecting cards from that deck until the game was terminated, at an unpredictable time, by the experimenter.

In each game, two of the four decks presented to each subject contained a majority of cards that presented a \$100 win and occasional large losses that produced a net loss of \$250 per 10 cards turned over. The other two decks mostly

included cards that yielded a gain of \$50 and an occasional card that imposed a small loss, a net gain of \$250 per 10 cards. The expected value of two of the decks was thus $-\$250$, and the expected value of the other two decks was $+\$250$. Bechara and his colleagues found that normal human subjects approach this task by sampling all four decks (usually switching after each loss) for a brief period. After an initial sampling of all four decks, normal subjects transiently preferred the risky decks with large gains and a negative expected value before settling appropriately on the decks with a positive expected value. Patients with bilateral lesions of the ventromedial frontal cortex, however, seemed unable to adequately assess the actual expected values of the two decks, continuing to prefer the high-risk negative expected value deck and never identifying the positive expected value deck as preferable. Although the exact role of the ventromedial frontal cortex is not specified entirely by these experiments, they do suggest this area plays some role in nonsensory decision making. It may even be that this area is important for the storage or updating of nonsensory data used to estimate expected utility.

In a closely related brain imaging study, Elliott and colleagues (2000) examined a slightly different card-based decision task and found that in addition to the role of ventromedial frontal cortex demonstrated by the Bechara studies, midbrain and ventral striatal regions are active when rewards are obtained, whereas regions of the hippocampus become active in response to losses. High reward levels were also found to activate the globus pallidus, thalamus, and subgenual cingulate cortex. (For closely related results see also London et al. 2000, O'Doherty et al. 2000, Rolls 2000).

These studies, among others, have led to the proposal that the frontal cortex and a cluster of nuclei in the basal ganglia play a critical role in nonsensory human decision making. Breiter and his colleagues (2001) tested this hypothesis explicitly by measuring brain activity in human subjects playing a series of lotteries for actual money. In those experiments, activity in the orbital gyrus of frontal cortex tracked the expected utility of each lottery as expressed behaviorally by the subjects, and neural responses in the nucleus accumbens of the basal ganglia were monotonically related to the values of outcomes to the subjects, a result supported by the findings of Knutson et al. 2001.

In summary, these studies suggest that areas of the frontal cortex and of the basal ganglia seem to be involved in nonsensory aspects of decision making. Damage to some of these areas has been shown to reduce the efficiency of human decision makers in nonsensory decision-making tasks. Activity in some of these areas has been correlated with fluctuations in expected utility. Taken together, these data raise the very preliminary possibility that the frontal cortices and the basal ganglia may play a role in nonsensory decision making that is analogous to the role played by the visual cortices in sensory decision making. These areas may provide nonsensory signals, as decision making inputs, to regions like area LIP. This is a conclusion compatible with recent studies of neuronal activity in monkey prefrontal cortex made while animals performed the moving dot task (Kim & Shadlen 1999). These data suggested that signals in cortical areas 8 and 46, measured during simple sensory to motor decision-making tasks, cannot be construed as uniquely sensory

or motor but rather as some compound of the two. However, the strongest support for the hypothesis that the frontal cortices and basal ganglia store or update data about the prior probability or magnitude of reward associated with a particular movement comes from new work with monkeys on the role of the basal ganglia in motivated behavior.

Studies of the Primate Basal Ganglia

These recent physiological studies in monkeys have begun to suggest that estimating the value of a reward, and determining how the value of an obtained reward differs from the reward that might have been expected, may be a critical function of the basal ganglia. In one of these experiments Schultz and his colleagues (1997) randomly presented thirsty monkeys with one of three different trial types, none of which required any behavioral response by the subject. In the first trial type, an auditory tone was presented and about a second later a fluid reinforcement was delivered to the thirsty monkey. In the second trial type, which occurred rarely, the fluid was delivered without a preceding tone. In the third trial type, also presented rarely, the tone was presented but no fluid was delivered. Schultz and his colleagues found that when the fluid was delivered to experienced monkeys without the preceding tone, dopaminergic neurons of the substantia nigra pars compacta and the ventral tegmental area, both nuclei of the basal ganglia, increased their discharge rate 100–200 ms after the unsignaled reinforcement was delivered. In contrast, if the reinforcement was delivered after the predictive tone, no neuronal response was observed after fluid delivery. Instead, the tone itself elicited an increase in neuronal firing rate with a delay of about 100–200 ms. Finally, on trials in which the tone was presented without reinforcement, the neurons still produced an increase in firing rate 100–200 ms after tone onset but then produced a decreased firing rate 100–200 ms after the reinforcement would have been delivered, had it been a reinforced trial. Based on these data, and other studies of substantia nigra pars compacta activity (Schultz & Romo 1990a,b; Ljungberg et al. 1992; Schultz et al. 1993), Schultz (Schultz et al. 1997, Schultz 1998) proposed that neurons of the pars compacta carry a signal encoding the difference in value between the reinforcement expected by a subject and the reinforcement actually received.

In a similar way, accumulating data from studies of the caudate nucleus and the substantia nigra pars reticulata, which lie within the basal ganglia and are reciprocally connected with the substantia nigra pars compacta, also suggest that basal ganglia may play a role in estimating and updating knowledge of both reward likelihoods and values. Studies of saccade-related neurons in the caudate nucleus, for example, now clearly suggest that the activity of these neurons is strongly influenced by the likelihood that a reward will be obtained (Hikosaka & Sakamoto 1986; Hikosaka et al. 1989a,b,c; Kawagoe et al. 1998; Takikawa et al. 2002; Lauwereyns et al. 2002). Likewise, neurons in the substantia nigra pars reticulata, a nucleus that directly modulates the activity of collicular saccade-related neurons (cf. Rinvik et al. 1976; Hopkins & Niessan 1976; Hikosaka & Wurtz 1983a,b,c,d, 1985a,b), are also preferentially modulated before movements that are likely to

yield a reward (Handel & Glimcher 2000). Although we still have very little data about how these patterns of activity may be related to decision making, it does seem clear that these areas carry information of a type critical to efficient decision making—information about the value and likelihood of behavioral responses.

Summary

All of these studies hint at an extended model of the primate decision-making system in which frontal and basal ganglia circuits store and update estimates of prior probabilities and reward value, which are used as nonsensory inputs for decision making by the parietal, frontal, and collicular circuits. Current experiments in a number of laboratories seek to better understand the role of each of these elements of the decision-making process. Perhaps the most important observation that arises from these studies is that the same neural circuits that appear to generate simple sensorimotor behaviors, of the kind Descartes and Sherrington studied, also seem to be able to account for more complicated nonsensory decision making. The parietal decision-related circuits, for example, seem to combine sensory and nonsensory signals as necessary to achieve efficient decision making under a range of circumstances.

What this cannot tell us, however, is how the richest kinds of behavior, the stochastic and unpredictable decision making we often refer to as volitional, are produced. To begin to ask how behaviors of that kind are generated, and to understand how the neural circuits that produce it may be related to the circuits we have already examined, neurobiologists are turning to an additional theoretical framework from economics, the theory of games.

STOCHASTIC DECISION MAKING AND THE THEORY OF GAMES

Descartes argued that all of human behavior could be sorted into two separable categories. His first category was made up of simple motor responses that were reliably triggered by specific stimuli. The second category was made up of behaviors that were fundamentally unpredictable, behaviors he considered voluntary. It was his proposal that the first category of behaviors could be mechanistically explained by simple sensory to motor linkages. Recent studies of this class of behavior in awake-behaving primates now seem to support a modified form of this proposal. Actions that depend almost exclusively on currently available sensory data do seem to be mediated by an integrative element that links sensory and motor components of the nervous system.

As the preceding section of this review indicates, over the past 20 years neurobiologists have begun to examine a third class of behaviors that lies, in complexity, somewhere between the two classes Descartes had identified. This third class involves nonsensory data; it is a class of decision influenced by the values or likelihoods of available rewards. Studies of these nonsensory decision-making

TABLE 1 Von Neumann & Morgenstern's H Matrix (from Von Neumann & Morgenstern 1944)

	1: Moriarity goes to Dover to intercept Holmes at ferry	2: Moriarity goes to Canterbury to intercept Holmes on train
1: Holmes gets off train in Dover to escape to the continent	Holmes: -100 Moriarity: 100	Holmes: 50 Moriarity: -50
2: Holmes gets off train in Canterbury to escape Moriarity	Holmes: 0 Moriarity: 0	Holmes: -100 Moriarity: 100

processes are beginning to suggest that the same neural circuits that mediate direct sensory to motor linkages in the visual-saccadic system may also mediate these less direct and more complex decisions.

If we hope to extend the neurophysiological study of decision making to include the study of stochastic behaviors like the ones Descartes argued were the product of the human soul, then an even richer conceptual framework is required; this is because expected utility theory cannot account for stochastic, or unpredictable, patterns of behavior. To make it clear why expected utility theory fails to adequately describe the stochastic patterns of decision making that we can observe under many conditions, consider the following example developed by the mathematician John Von Neumann and the economist Oskar Morgenstern in the 1930s and 1940s (Von Neumann & Morgenstern 1944) (Table 1).

Sherlock Holmes desires to proceed from London to Canterbury and thence from Canterbury to Dover and thence to the Continent in order to escape from Professor Moriarity who pursues him. Having boarded the train he observes, as the train pulls out, the appearance of Professor Moriarity on the platform. Sherlock Holmes takes it for granted—and in this he is assumed to be fully justified—that his adversary, who has seen him, might secure a special train and overtake him. Sherlock Holmes is faced with the alternative of going to Dover or of leaving the train at Canterbury, the only intermediate station. His adversary—whose intelligence is assumed to be fully adequate to visualize these possibilities—has the same choice. Both opponents must choose the place of their detrainment in ignorance of the other's corresponding decision. If, as a result of these measures, they should find themselves, *in fine*, on the same platform, Sherlock Holmes may with certainty expect to be killed by Moriarity. If Holmes reaches Dover unharmed he can make good his escape.

...

Let us now consider the H matrix of [the table above]. The fields (1,1) and (2,2) correspond to Professor Moriarity catching Sherlock Holmes, which it is reasonable to describe by a very high value of the corresponding matrix element, say -100. The field (2,1) signifies that Sherlock Holmes successfully

escaped to Dover, while Moriarity stopped at Canterbury. This is Moriarity's defeat as far as the present action is concerned, and should be described by a big negative value of the matrix element [for Moriarity]—in the order of magnitude but smaller than the positive value mentioned above—say, -50 . The field (1,2) signifies that Sherlock Holmes escapes Moriarity at the intermediate station, but fails to reach the Continent. This is best viewed as a tie, and assigned the matrix element 0.

...

[Von Neumann & Morgenstern then present a mathematical analysis of this matrix and conclude that only one efficient strategy exists.] Moriarity should go to Dover with a probability of 60% while Sherlock Holmes should stop at the intermediate station with a probability of 60%—the remaining 40% being left in each case for the other alternative. (From Von Neumann & Morgenstern 1944)

The critical point that Von Neumann & Morgenstern meant for their mathematical analysis to convey is that neither Holmes nor Moriarity can compute, in isolation, a single determined optimal strategy of the kind that classical expected utility theory was designed to model. If Holmes simply concludes, based upon an analysis of his situation, that he should detrain at Canterbury, then his determinate action, because it is determinate, makes him easy prey for Moriarity. Holmes can only minimize his loss against Moriarity by behaving stochastically. Both Holmes and Moriarity know this. To maximize their mutual expected utilities, they must each produce behaviors that cannot be reduced to certainty by their opponent. They must each effectively roll a set of dice as their trains pull into Canterbury station and let their actions be guided by this randomizing event. This is, at an absolute level, the best strategy that they can adopt. No other approach is better. Mixed strategy games like these are those in which, by definition, behaving probabilistically is an optimal solution.

A number of researchers have recently begun to examine how the theory of games might be used to analyze the neural architecture active when competitive stochastic behaviors are produced, a class of unpredictable behaviors that appears in every respect to be voluntary. 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. In their experiment, subjects played a two-person game called trust and reciprocity, which 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 two, then player two must decide between taking all of a larger 405-cent gain for herself or keeping only 225 cents and returning 180 cents 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 cents for herself.

Cooperating with player one offers 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 the more classical economic examples Pascal and others investigated, 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 on assumptions both about the likelihood of encountering the same player and about how one's own behavior will influence the behavior of the opponent. This is a property that it shares with the Holmes-Moriarity game and which defies explanation with nongame-theoretic tools.

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.

McCabe and his colleagues found when studying the brains of their subjects under these conditions that whenever a subject chose to cooperate, a region in the prefrontal cortex was more active than when they decided to act rationally against the computer. To a very rough first approximation, this suggests that frontal areas do respond differently when subjects make different evaluative choices under game theoretic conditions. However, though this represents an important first step, it tells us little about how the decision process itself occurs at a neurophysiological level. One reason is that it is difficult to unequivocally identify an optimal solution to the decision-making problem faced by choosers under the strategic situation McCabe and colleagues examined. Identifying games for which optimal solutions can be unambiguously defined may therefore be an important part of using game theory to understand the neurophysiology of complex decision making. Luckily, such a definitional system exists for a broad category of strategic interactions that have been well studied by experimental and theoretical economists.

Nash Equilibria

Consider the classic automotive game of chicken. Smith and Jones sit in cars parked on opposite ends of a bridge. At a signal they drive toward each other as fast as they can. At the instant before the two cars collide, Smith and Jones each decide whether or not to swerve to the right. To formalize this game, a payoff matrix is shown below.

If neither Smith nor Jones swerves, both suffer a significant loss. If one player swerves, that represents a loss to the swerver of 10 and a gain to the winner of 50. Finally, we characterize a draw as a negligible gain to both players.

TABLE 2 Chicken payoff matrix

	Smith continues	Smith swerves
Jones continues	Jones: -100 Smith: -100	Jones: 50 Smith: -10
Jones swerves	Jones: -10 Smith: 50	Jones: 1 Smith: 1

In the late 1940s, the mathematician John Nash recognized (1950a,b; 1951) that, given a matrix of this type, if we think of Jones and Smith as playing the game of chicken again and again (neglecting the possibility of their deaths for the purposes of the initial mathematical analysis) the Jones-Smith interaction must come to rest at some sort of equilibrium point—a point at which the losses and gains between the two players are equally balanced and neither player has an incentive to swerve more or less often than this equilibrium level (Table 2). Where would that point be? Is there a way to calculate the location of an equilibrium point in a mixed strategy game of this type?

Nash solved this problem in the following way. First, begin by calling the probability that Smith will swerve: *P*smith, swerves. Because this is a probability, it will be a number ranging from 0 (he never swerves) to 1 (he always swerves). Because Smith either swerves or goes straight, the sum of the probabilities of swerving and going straight must equal 1.

$$P_{\text{smith, straight}} + P_{\text{smith, swerve}} = 1; \quad 4.$$

or put another way,

$$P_{\text{smith, straight}} = 1 - P_{\text{smith, swerve}}. \quad 5.$$

By analogy,

$$P_{\text{jones, swerve}} = 1 - P_{\text{jones, straight}}. \quad 6.$$

We can think of a variable like *P*smith, straight as the likelihood that Smith will decide to go straight. Starting with this variable we can figure out what the expected utility is, for Jones, of the situation in which Jones swerves and Smith goes straight. The probability of Smith going straight is *P*smith, straight and the loss to Jones under these conditions is 10. Therefore, if Smith goes straight,

$$\text{gain to Jones for swerving} = P_{\text{smith, straight}} \times -10. \quad 7.$$

In exactly the same way we can use a probabilistic approach to see the gain to Jones if Smith decides to swerve. Therefore, if Smith swerves

$$\text{gain to Jones for swerving} = P_{\text{smith, swerves}} \times 1. \quad 8.$$

So in total, if Jones makes a decision to swerve,

$$\begin{aligned} \text{gain to Jones for swerving} &= (\text{Psmith, straight} \times -10) \\ &+ (\text{Psmith, swerves} \times 1); \end{aligned} \quad 9.$$

or we can say equivalently

$$\begin{aligned} \text{gain to Jones for swerving} &= [(1 - \text{Psmith, swerves}) \times -10] \\ &+ (\text{Psmith, swerves} \times 1). \end{aligned} \quad 10.$$

Nash's critical insight was that a specific course of action for Jones would be an equilibrium point if and only if no other course of action was better for Jones. It could only be an equilibrium point if Jones was indifferent between swerving and going straight because swerving and going straight had equal expected utility for him. So if the

$$\begin{aligned} \text{gain to Jones for continuing} &= [(1 - \text{Psmith, swerves}) \times -100] \\ &+ (\text{Psmith, swerves} \times 50), \end{aligned} \quad 11.$$

and the

$$\begin{aligned} \text{gain to Jones for swerving} &= [1 - \text{Psmith, swerves}) \times -10] \\ &+ (\text{Psmith, swerves} \times 1) \end{aligned} \quad 12.$$

then these equations are equal when

$$\begin{aligned} &[(1 - \text{Psmith, swerves}) \times -100] + (\text{Psmith, swerves} \times 50) \\ &= [(1 - \text{Psmith, swerves}) \times -10] + (\text{Psmith, swerves} \times 1). \end{aligned} \quad 13.$$

To solve this equation algebraically,

$$90 = 139 * \text{Psmith, swerves}, \quad 14.$$

and

$$.647 = \text{Psmith, swerves}. \quad 15.$$

As long as there is a 64.7% chance that Smith will swerve, swerving and not swerving are of equal expected utility to Jones. As long as Smith will swerve 64.7% of the time, Jones has no incentive to care what he (Jones) does; his two choices are equally good (or bad).

The absolutely critical idea here is that Smith's best possible plan is to swerve 64.7% of the time; this is the point at which Jones will be indifferent to his own actions. As soon as either player behaves in any other manner (as in Jones asserting that he will go straight no matter what and meaning it), then the other player faces a standard deterministic decision-making problem of the kind we have already encountered: There is one rational strategy, and one connection between the sensory and motor systems is appropriate under these conditions. As long as both players are seeking an optimal solution, however, they must reach this stochastic equilibrium point if they behave efficiently. The equilibrium point computed in this way defines the only behavioral pattern for which neither the option of swerving

TABLE 3 Inspection game payoff matrix, employee side only

	Employer inspects	Employer does not inspect
Employee works	Employee: Wage-Cost of Work, W-I	Employee: Wage-Cost of Work, W-I
Employee shirks	Employee: 0	Employee: Wage, W

nor the option of continuing is a demonstrably better plan. It is at this indifference point, where no pure action is any better than any other pure action, that the players come into equilibrium. The Nash equilibrium computation thus identifies a rational strategy for a chooser and also provides us with a critical computational insight. Nash equilibria are, definitionally, situations in which the expected utility of all available options is equivalent, even when the situations require that the response probabilities associated with each choice are not.

The Neurobiology of Voluntary Behaviors at Nash Equilibrium

To examine the activity of parietal neurons during decision making of this type, Dorris & Glimcher (2001) developed an animal model for strategic decision making using a conflict known as the inspection game (Table 3). In that conflict, two players, the employer and the employee, compete to maximize their earnings by each making a single decision on each trial. Employees must decide whether to attend work to be assured that they will earn their wage or whether to shirk. An employee who shirks successfully (without being detected) earns both his wage and a bonus, whereas he earns nothing if detected. In a similar way, the employer must decide whether or not to pay an inspection fee, which allows him to check up on his employee. These two sets of decisions yield a 2×2 payoff matrix shown above.

According to Nash's theory, an equilibrium point will be reached in this game when the expected utility (EU) for each choice is equal for both players. Thus at Nash equilibrium for the employee,

$$EU_{\text{working}} = EU_{\text{shirking}}. \quad 16.$$

This, given the payoff matrix, expands to

$$[P_1 \times 0] + [(1 - P_1) \times W] = [P_1 \times (W - C)] + [(1 - P_1) \times (W - C)], \quad 17.$$

and, solving for P_1 ,

$$P_1 = C/W, \quad 18.$$

where EU_{working} is the expected utility for choosing to work, EU_{shirking} is the expected utility for choosing to shirk, P_1 is the probability of the boss inspecting and $1 - P_1$ is the probability of the boss not inspecting, and C is the cost of work to the employee. In a similar way, Nash equilibrium for the employer is

$$EU_I = EU_N. \quad 19.$$

This, given the payoff matrix, expands to

$$\begin{aligned} & [P_S \times -I] + [(1 - P_S) \times (V - I - W)] \\ & = [P_S \times -W] + [(1 - P_S) \times (V - W)], \end{aligned} \quad 20.$$

and, solving for P_S ,

$$P_S = I/W, \quad 21.$$

where EU_I is the expected utility of choosing to inspect, EU_N is the expected utility for choosing to not inspect, P_S is the probability of the employee shirking and $(1 - P_S)$ is the probability of the employee working, V is the utility of the product to the boss, I is the cost of inspection to the boss, and W is the wage paid by the boss to the employee.

This suggests that the employee is indifferent between his two courses of action when the employer sets the probability of inspection equal to C/W . In a similar way, the employer becomes indifferent when the employee sets the probability of shirking to I/W . The strategy that should thus evolve as the system comes into dynamic equilibrium is for the employee to behave stochastically—what most of us would probably call volitionally—by unpredictably shirking with a probability determined by the ratio of the cost of inspections to the employer and the value of his wage. If either player deviates significantly from this strategy, the behavior of his opponent will drive him back to the dynamic equilibrium point.

To understand how actual primates play games of this kind and to determine how the visual-saccadic decision circuit might participate in decision making of this type, Dorris & Glimcher (2001) had human volunteers play the game of work or shirk. In that experiment, one human was assigned the role of employer and the other the role of worker. The subjects then completed several hundred plays of work or shirk over a one- to two-hour period. At the beginning of each play, they independently decided what to do on that iteration: whether to work or to shirk, whether to inspect or not to inspect. After they had made their independent decisions, a computer informed them of what they had earned for that play (a monetary gain they actually received at the end of the session). In a typical experiment, a pair of subjects played a series of blocks of 150 trials under each of several different payoff matrix conditions. In each block, a new rate of shirking and inspecting was the Nash equilibrium rate.

Figure 12a shows what a typical human employee did over the course of about 300 plays of the game. In black is plotted the percentage at which the employee decided to shirk, represented as a running average over 20 sequential plays. In grey is plotted the Nash equilibrium solutions during the two sequential blocks of trials presented on this day. Note that over the course of the first 50 plays, the employee comes into fairly precise accord with the Nash equilibrium. The employee's behavior then remains at or near this stable equilibrium until the payoff matrix is changed, after which it converges on the new equilibrium rate. A second

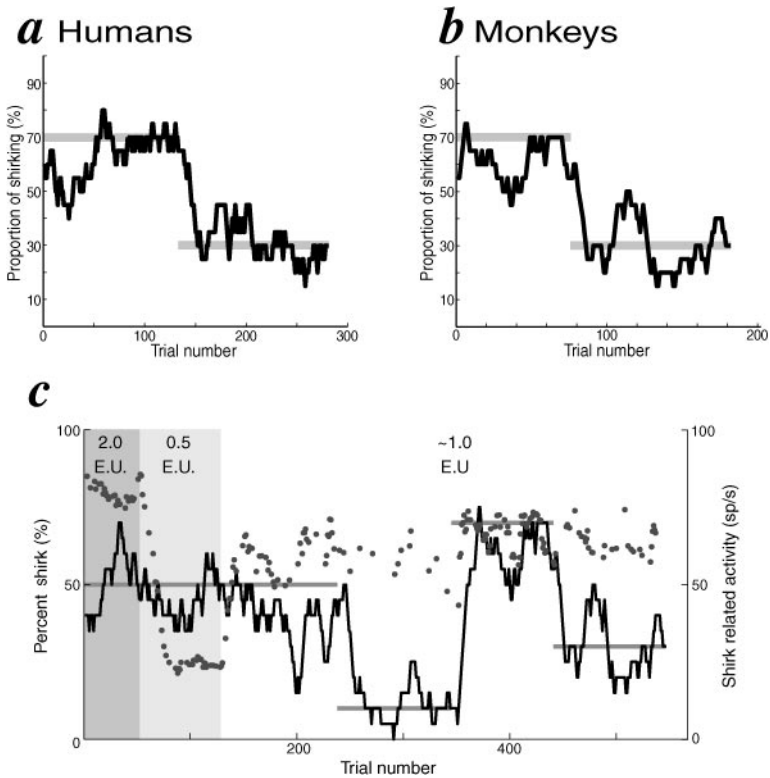


Figure 12 The inspection game. (a) The behavior of a typical human employee during several hundred trials of the inspection game. The jagged black lines plot a 20-trial running average of the rate of shirking produced by this subject. The grey lines plot optimal Nash equilibrium strategies. (b) A monkey playing the inspection game. (c) Behavior of a monkey, plotted with black line, during seven blocks of trials. In the first two blocks the monkey is instructed which movement to make. In the first block, movement into the response field of the neuron under study yields twice as much reward as the other movement (relative expected value = 2). In the second block the relative expected value of this same movement is one half of the value of the other movement. The next five blocks examine different Nash equilibrium conditions in which the animal is free to choose in competition with an opponent, but in which the relative expected utility of the two movements is equal. The dots plot the firing rate of an LIP neuron on each of these trials. The LIP neuron tracks expected utility.

set of experiments extended this result, showing that the behavior of the human employees was unchanged if the human employer was replaced with an intelligent adaptive computer program also seeking to maximize its gain.

Next, monkeys were trained to play this same game against a computer employer but in a visual-saccadic form to earn a fruit juice reward rather than money. Each

trial began when a central fixation stimulus was illuminated, which the monkey had to stare at in order to begin the trial. The computer employer then secretly selected one of two possible trials to present to the monkey, the Inspect trial or the No-Inspect trial. In both trials, two eccentric visual targets were next illuminated, one to the left and one to the right of the central stimulus. After a brief delay the central target was extinguished, and the monkey had to decide whether to look left or to look right. Looking to the left counted as working and looking to the right counted as shirking. Based on what kind of trial the computer opponent had selected, and what movement the monkey had selected, the monkey received an appropriate juice reward and the next round of play began. Dorris & Glimcher found that, like human players, the monkeys seemed to quickly figure out that there is an optimal rate (Figure 12*b*).

In the cued saccade experiment described in an earlier section of this article (Platt & Glimcher 1999) the probability or value of an outcome was shown to be encoded in the firing rates of LIP neurons. These results suggested that neurons in area LIP might encode something like relative expected utility when animals are engaging in either sensory or nonsensory decision making. To confirm that hypothesis Dorris & Glimcher began each day by examining the activity of each neuron with the cued saccade task. In a first block of trials, shown against a dark grey background in Figure 12*c*, the movement into the response field, when it was required, had a high relative expected utility. They followed this with a block, shown against a light grey background, in which the movement into the response field had a low relative expected utility. As in the original Platt & Glimcher experiment, Dorris & Glimcher found that the neuronal firing rate was strongly modulated by this change in relative expected utilities.

Next, Dorris & Glimcher shifted to the inspection game protocol, in which the animals made their own choices, sampling sequentially a set of five Nash equilibrium conditions. These five conditions required that the animal, if he was to behave efficiently, show five different likelihoods of producing each possible movement; however, because they were all Nash equilibrium conditions, each block of trials reflected a condition in which both movements had equal relative expected utilities. Under these conditions LIP neurons adopted a firing rate intermediate between the rate observed on the preceding two blocks and at a roughly constant level across all five Nash equilibrium conditions. The firing rates of LIP neurons were correlated with relative expected utility and in a way that was independent of the likelihood that the monkey would choose to make any particular movement. This is, of course, exactly what would be expected if neurons in area LIP encoded relative expected utility, a quantity which should guide choice behavior.

An efficient decision maker should produce whichever movement has the higher expected utility. Only when the relative expected utility of the movements is equivalent, when the relative expected utility of each movement is one, should behavior become stochastic. This reasoning, interestingly, raises an additional possibility. One might expect that small fluctuations in LIP rate at Nash equilibrium would be correlated with the small fluctuations in expected utility that serve to keep the

chooser's behavior accurately at the Nash solution. Dorris & Glimcher examined this possibility by deriving a trial-by-trial estimate of expected value (which is correlated with expected utility) and correlating that estimate with neuronal rates in LIP on a trial-by-trial level. They found that these two values were significantly correlated, exactly as predicted by theory.

The Dorris & Glimcher experiment thus leads us toward two conclusions. The first is that neurons in area LIP, and presumably in other closely related areas, do seem to encode relative expected utility, the central element in most economic theories of decision making. The second conclusion these data suggest is that the neural circuit for predictable sensory and nonsensory decision making also seems adequate to account for some stochastic decision making, a class of decision making that, in humans, we call volitional. Even under conditions of strategic conflict that yield stochastic behavior, firing rates in area LIP are correlated with a decision variable appropriate for guiding choice behavior. Indeed, it even seems possible that the random stochasticity observed under these conditions could reflect unpredictable fluctuations in LIP firing rate, fluctuations often referred to as neuronal noise. Though many more questions remain to be answered, these economic approaches to studying the neural basis of decision making do seem to be valuable, and it is clear that they have yielded significant insights into the neuronal basis of decision making. The growing field of neuroeconomics from which these economic approaches to neuroscience stem is clearly beginning to provide powerful conceptual tools for studying the neural basis of choice.

CONCLUSION

Aristotle and Galen argued that the causal force for generating all behavior was the human soul. It served to link events in the outside world with the actions animals take. Descartes challenged this assertion, arguing that behavior could be usefully divided into two categories, each of which could be attributed to a different causal process. Current studies of visual-saccadic decision making in primates suggest an alternative view: that at a conceptual level, behavior can be divided into at least three categories. These studies also suggest, however, that, at a physiological level, decision making in each of these categories may often be the product of a single mechanistic process that combines sensory and nonsensory data to produce efficient behavioral courses of action. Even stochastic decision making seems to follow this model.

Unfortunately, we know very little about how these insights into visual-saccadic decision making will generalize. We know almost nothing about how other classes of motor responses are organized with regard to this architecture or how the many different sensory and motor systems of the brain interact during decision making. It is clear, however, that in the past two decades we have made significant progress toward a neurobiological answer to the question that engaged Aristotle in *De Anima*: How is it that the intellectual faculty organizes behavior?

The *Annual Review of Neuroscience* is online at <http://neuro.annualreviews.org>

LITERATURE CITED

- Albright TD, Desimone R, Gross CG. 1984. Columnar organization of directionally selective cells in visual area MT of the macaque. *J. Neurophysiol.* 51:16–31
- Andersen RA, Snyder LH, Bradley DC, Xing J. 1997. Multimodal representation of space in the posterior parietal cortex and its use in planning movements. *Annu. Rev. Neurosci.* 20:303–30
- Aristotle A. 1986. *De Anima: On the Soul*. New York: Penguin
- Arnauld A, Nicole P. 1662. *Logic or the Art of Thinking*. Transl. JV Buroker, 1996. Cambridge, UK: Cambridge Univ. Press (From French)
- Asanuma C, Andersen RA, Cowan WM. 1985. The thalamic relations of the caudal inferior parietal lobule and the lateral prefrontal cortex in monkeys: divergent cortical projections from cell clusters in the medial pulvinar nucleus. *J. Comp. Neurol.* 241:357–81
- Barbus H, Meshulam M-M. 1981. Organization of afferent input to subdivisions of area 8 in the rhesus monkey. *J. Comp. Neurol.* 200:407–31
- Basso MA, Wurtz RH. 1997. Modulation of neuronal activity by target uncertainty. *Nature* 389:66–69
- Bayes T. 1958 [1763]. An essay toward solving a problem in the doctrine of chances. Originally published in *Phil. Trans. R. Soc.* 53:370–418. Reprinted in *Biometrika* 45:293–315
- Bechara A, Damasio AR, Damasio H, Anderson SW. 1994. Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition* 50:7–15
- Bechara A, Damasio H, Tranel D, Damasio AR. 1997. Deciding advantageously before knowing the advantageous strategy. *Science* 275:1293–95
- Bernoulli D. 1954 [1738]. Exposition of a new theory on the measurement of risk. *Econometrica* 22:23–36
- Breiter HC, Aharon I, Kahneman D, Dale A, Shizgal P. 2001. Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30:619–39
- Bruce CJ, Goldberg ME. 1985. Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* 53:603–35
- Carpenter RH, Williams ML. 1995. Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62
- Colby CL, Goldberg ME. 1999. Space and attention in parietal cortex. *Annu. Rev. Neurosci.* 23:319–49
- Descartes R. 1649. *Passions De L'Ame: the Passions of the Soul*. Transl. S Voss, 1989. Indianapolis, IN: Hackett (From French)
- Descartes R. 1664. *L'Homme: Treatise on Man*. Transl. TS Hall, 1972. Cambridge, MA: Harvard Univ. Press (From French)
- Dorris MC, Munoz DP. 1998. Saccadic probability influences motor preparation signals and time to saccadic initiation. *J. Neurosci.* 18:7015–26
- Dorris MC, Glimcher PG. 2001. A game theoretic approach for studying the neurobiology of decision-making. *Soc. Neurosci. Abstr.* 58.10
- Duda RO, Hart PE, Stork DG. 2000. *Pattern Classification*. New York: Wiley-Interscience
- Elliott R, Friston KJ, Dolan RJ. 2000. Dissociable neural responses in human reward systems. *J. Neurosci.* 20:6159–65
- Galen C. 1916. *On the Natural Faculties*. Transl. AJ Brock. Cambridge, MA: Loeb Class. Libr. (From Greek)
- Galen C. 1968. *On the Usefulness of Parts*. Transl. MT May. Ithaca, NY: Cornell Univ. Press (From Greek)
- Gallistel CR. 1980. *The Organization of Action: a New Synthesis*. Hillsdale, NJ: Erlbaum
- Glimcher PW. 2003. *Decisions, Uncertainty*

- and the Brain: the Science of Neuroeconomics. Cambridge, MA: Mass. Inst. Technol. Press
- Glimcher PW, Sparks DL. 1992. Movement selection in advance of action in the superior colliculus. *Nature* 355:542–45
- Gnadt JW, Andersen RA. 1988. Memory related motor planning activity in posterior parietal cortex of macaque. *Exp. Brain Res.* 70:216–20
- Gold JJ, Shadlen MN. 2000. Representation of a perceptual decision in developing oculomotor commands. *Nature* 404:390–94
- Gold JJ, Shadlen MN. 2001. Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5:10–16
- Handel A, Glimcher PW. 2000. Role of the substantia nigra pars reticulata in saccade generation. *J. Neurophysiol.* 83:3042–48
- Hanes DP, Schall JD. 1995. Countermanding saccades in macaque. *Vis. Neurosci.* 12:929–37
- Hanes DP, Schall JD. 1996. Neural control of voluntary movement initiation. *Science* 247:427–30
- Hanes DP, Patterson WF, Schall JD. 1998. The role of frontal eye field in countermanding saccades: visual, movement and fixation activity. *J. Neurophysiol.* 79:817–34
- Hanes DP, Thompson KG, Schall JD. 1995. Relationship of presaccadic activity in frontal eye field and supplementary eye field to saccade initiation in macaque: poisson spike train analysis. *Exp. Brain Res.* 103:85–96
- Hikosaka O, Sakamoto M. 1986. Cell activity in monkey caudate nucleus preceding saccadic eye movements. *Exp. Brain Res.* 63:659–62
- Hikosaka O, Sakamoto M, Usui S. 1989a. Functional properties of monkey caudate neurons. I. Activities related to saccadic eye movements. *J. Neurophysiol.* 61:780–98
- Hikosaka O, Sakamoto M, Usui S. 1989b. Functional properties of monkey caudate neurons. II. Visual and auditory responses. *J. Neurophysiol.* 61:799–813
- Hikosaka O, Sakamoto M, Usui S. 1989c. Functional properties of monkey caudate neurons. III. Activities related to expectation of target and reward. *J. Neurophysiol.* 61:814–32
- Hikosaka O, Wurtz RH. 1983a. Visual and oculomotor functions of monkey substantia nigra pars reticulata. I. Relation of visual and auditory responses to saccades. *J. Neurophysiol.* 49:1230–53
- Hikosaka O, Wurtz RH. 1983b. Visual and oculomotor functions of monkey substantia nigra pars reticulata. II. Visual responses related to fixation of gaze. *J. Neurophysiol.* 49:1254–67
- Hikosaka O, Wurtz RH. 1983c. Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. *J. Neurophysiol.* 49:1268–84
- Hikosaka O, Wurtz RH. 1983d. Visual and oculomotor functions of monkey substantia nigra pars reticulata. IV. Relation of substantia nigra to superior colliculus. *J. Neurophysiol.* 49:1285–301
- Hikosaka O, Wurtz RH. 1985a. Modification of saccadic eye movements by GABA-related substances. I. Effects of muscimol and bicuculline in the monkey superior colliculus. *J. Neurophysiol.* 53:266–91
- Hikosaka O, Wurtz RH. 1985b. Modification of saccadic eye movements by GABA-related substances. I. Effects of muscimol in the monkey substantia nigra pars reticulata. *J. Neurophysiol.* 53:292–308
- Hopkins DA, Niessen LW. 1976. Substantia nigra projections to the reticular formation, superior colliculus and central gray in the rat, cat, and monkey. *Neurosci. Lett.* 2:253–59
- Hubel DH, Wiesel TN. 1959. Receptive fields of single neurons in the cat's visual cortex. *J. Physiol. (Lond.)* 148:574–91
- Hubel DH, Wiesel TN. 1962. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond.)* 160:106–54
- Kawagoe R, Takikawa Y, Hikosaka O. 1998. Expectation of reward modulates cognitive signals in the basal ganglia. *Nat. Neurosci.* 1:411–16

- Kim J-N, Shadlen MN. 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2(2):176–85
- Knutson B, Adams CM, Fong GW, Hommer D. 2001. Anticipation of increasing monetary reward selectively recruits nucleus accumbens. *J. Neurosci.* 21:RC159
- Kreps DM. 1990. *A Course in Microeconomic Theory*. Princeton, NJ: Princeton Univ. Press
- Kuffler SW. 1953. Discharge patterns and functional organization of the mammalian retina. *J. Neurophysiol.* 16:37–68
- Lauwereyns J, Watanabe K, Coe B, Hikosaka O. 2002. A neural correlate of response bias in monkey caudate nucleus. *Nature* 418:413–17
- Letzvin JY, Maturana HR, McCulloch WS, Pitts WH. 1959. What the frog's eye tells the frog's brain. *Proc. Inst. Radio Eng.* 47:1940–51
- Ljungberg T, Apicella P, Schultz W. 1992. Responses of monkey dopamine neurons during learning of behavioral reactions. *J. Neurophysiol.* 67:145–63
- London ED, Ernst M, Grant S, Bonson K, Weinstein A. 2000. Orbitofrontal cortex and human drug abuse: functional imaging. *Cereb. Cortex* 10:334–42
- Luce RD. 1986. *Response Times*. New York: Oxford
- Lynch JC, Graybiel AM, Lobeck LJ. 1985. The differential projection of two cytoarchitectural subregions of the inferior parietal lobule of macaque upon the deep layers of the superior colliculus. *J. Comp. Neurol.* 235:242–54
- Maunsell JHR, VanEssen DC. 1983a. Functional properties of neurons in the middle temporal visual area (MT) of the macaque monkey. I. Selectivity for stimulus direction, speed and orientation. *J. Neurophysiol.* 49:1127–47
- Maunsell JHR, VanEssen DC. 1983b. The connections of the middle temporal visual area (MT) and their relationship to a cortical hierarchy in the macaque monkey. *J. Neurosci.* 3:2563–86
- McCabe K, Houser D, Ryan L, Smith V, Trouard T. 2001. A functional imaging study of cooperation in two-person reciprocal exchange. *Proc. Natl. Acad. Sci. USA* 98:11,832–35
- Mountcastle VB. 1957. Modality and topographic properties of single neurons of cat's somatosensory cortex. *J. Neurophysiol.* 20:408–34
- Nash JF. 1950a. Equilibrium points in n-person games. *Proc. Natl. Acad. Sci. USA* 36:48–49
- Nash JF. 1950b. The bargaining problem. *Econometrica* 18:155–62
- Nash JF. 1951. Non-cooperative games. *Ann. Math.* 54:286–95
- Newsome WT, Britten KH, Movshon JA. 1989. Neuronal correlates of a perceptual decision. *Nature* 341:52–54
- O'Doherty J, Rolls ET, Francis S, Bowtell R, McGlone F, et al. 2000. Sensory-specific satiety-related olfactory activation of the human orbitofrontal cortex. *Neuro. Rep.* 11:893–97
- Pascal B. 1670. *Pensees*. Transl. AJ Krailsheimer, 1966. New York: Penguin Books (From French)
- Pettit DL, Helms MC, Lee P, Augustine GJ, Hall WC. 1999. Local excitatory circuits in the intermediate gray layer of the superior colliculus. *J. Neurophysiol.* 81:1424–27
- Platt ML, Glimcher PW. 1999. Neural correlates of decision variables in parietal cortex. *Nature* 400:233–38
- Ratcliff R. 1978. A theory of memory retrieval. *Psychol. Rev.* 85:59–104
- Reddi BA, Carpenter RH. 2000. The influence of urgency on decision time. *Nat. Neurosci.* 3:827–30
- Rinvik E, Grofova I, Ottersen OP. 1976. Demonstration of nigroretectal and nigroreticular projections in the cat by axonal transport of proteins. *Brain Res.* 112:388–94
- Robinson DA. 1972. Eye movements evoked by collicular stimulation in the alert monkey. *Vision Res.* 12:1795–1808
- Robinson DA, Fuchs AF. 1969. Eye movements evoked by stimulation of frontal eye fields. *J. Neurophysiol.* 32:637–48
- Rolls ET. 2000. The orbitofrontal cortex and reward. *Cereb. Cortex* 10:284–94

- Salzman CD, Britten KH, Newsome WT. 1990. Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346:174–77
- Schiller PH, Sandell JH. 1983. Interactions between visually and electrically elicited saccades before and after superior colliculus and frontal eye field ablations in the rhesus monkey. *Exp. Brain. Res.* 49:381–92
- Schultz W. 1998. Predictive reward signal of dopamine neurons. *J. Neurophysiol.* 80:1–27
- Schultz W, Apicella P, Ljungberg T. 1993. Responses of monkey dopamine neurons to reward and conditioned stimuli during successive steps of learning a delayed response task. *J. Neurosci.* 13:900–13
- Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science* 275:1593–99
- Schultz W, Romo R. 1990a. Dopamine neurons of the monkey midbrain: contingencies of responses to active touch during self-initiated arm movements. *J. Neurophysiol.* 63:592–606
- Schultz W, Romo R. 1990b. Dopamine neurons of the monkey midbrain: contingencies of responses to stimuli eliciting immediate behavioral reactions. *J. Neurophysiol.* 63:607–24
- Salzman CD, Britten KH, Newsome WT. 1990. Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346:174–177
- Seltzer B, Pandya DN. 1984. Further observations of parieto-temporal connections in the rhesus monkey. *Exp. Brain Res.* 55:301–12
- Shadlen MN, Britten KH, Newsome WT, Movshon JA. 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* 16:1486–510
- Shadlen MN, Newsome WT. 1996. Motion perception: seeing and deciding. *Proc. Natl. Acad. Sci. USA* 93:628–33
- Shadlen MN, Newsome WT. 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86:1916–36
- Sherrington CS. 1906. *The Integrative Action of the Nervous System*. New York: Scribner
- Sherrington CS. 1947. *The Integrative Action of the Nervous System*. New Haven, CT: Yale Univ. Press. 2nd ed.
- Sparks DL, Holland R, Guthrie BL. 1976. Size and distribution of movement fields in the monkey superior colliculus. *Brain Res.* 113:21–34
- Sparks DL, Mays LE. 1990. Signal transformations required for the generation of saccadic eye movements. *Annu. Rev. Neurosci.* 13:309–36
- Sternberg S. 1969a. The discovery of processing stages: extensions of Donders' method. In *Attention and Performance II*, ed. WG Koster. Amsterdam: North-Holland
- Sternberg S. 1969b. Memory scanning: mental processes revealed by reaction-time experiments. *Am. Sci.* 57:421–57
- Takikawa Y, Kawagoe R, Hikosaka O. 2002. Reward-dependent spatial selectivity of anticipatory activity in monkey caudate neurons. *J. Neurophysiol.* 87:508–15
- Thompson KG, Bichot NP, Schall JD. 1997. Dissociation of target selection from saccade planning in macaque frontal eye field. *J. Neurophysiol.* 77:1046–50
- Thompson KG, Hanes DP, Bichot NP, Schall JD. 1996. Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *J. Neurophysiol.* 76:4040–55
- Von Neumann JV, Morgenstern O. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton Univ. Press
- Wurtz RH, Goldberg ME. 1972. Activity of superior colliculus in the behaving monkey. III. Cells discharging before eye movements. *J. Neurophysiol.* 35:575–86
- Zeki SM. 1974. Functional organization of a visual area in the superior temporal sulcus of the rhesus monkey. *J. Physiol. (Lond.)* 236:549–73




Figure 5 Average firing rate in area LIP during the moving dot task. The colored lines plot the mean firing rates, averaged across trials, of an LIP neuron in response to dot motion in the preferred and nonpreferred directions. The *left* set of curves plot the data synchronized to the onset of the moving dot stimulus. The *right* set of curves plot data from a later period in the trials, synchronized to saccade onset. Note that as the fraction of dots moving in the preferred direction is increased the rate at which firing rate grows is increased.




Figure 7 (A) Gold & Shadlen (2000) first determined the movement, shown by the black arrow, produced when stimulation was delivered to the frontal eye fields during passive fixation. They then presented a moving dot stimulus, in this case an upward stimulus, while the animals performed a reaction time version of the moving dot task. On some of these trials electrical stimulation was delivered to the frontal eye fields during the reaction time interval. The resulting movement, shown schematically as a grey arrow, was hypothesized to be the product of the stimulation-induced movement (the black arrow) and a hypothetical signal (shown in cyan) related to the likelihood that an upward movement would be reinforced. (B) Plots of the deviation in stimulation-induced saccade endpoint, the purple arrow in panel A, as a function of dot correlation and reaction time, show that the hypothesized signal grows over time and encodes a value accurately described by a likelihood calculation. [From Gold & Shadlen 2000. Reproduced with permission from *Nature*.]