



ELSEVIER

Available online at www.sciencedirect.com

SCIENCE @ DIRECT®

Games and Economic Behavior 52 (2005) 213–256

GAMES and
Economic
Behavior

www.elsevier.com/locate/geb

Physiological utility theory and the neuroeconomics of choice

Paul W. Glimcher*, Michael C. Dorris, Hannah M. Bayer

Center for Neural Science, New York University, New York, NY 10003, USA

Received 8 June 2004

Available online 5 October 2004

Abstract

Over the past half century economists have responded to the challenges of Allais [Econometrica (1953) 53], Ellsberg [Quart. J. Econ. (1961) 643] and others raised to neoclassicism either by bounding the reach of economic theory or by turning to descriptive approaches. While both of these strategies have been enormously fruitful, neither has provided a clear programmatic approach that aspires to a complete understanding of human decision making as did neoclassicism. There is, however, growing evidence that economists and neurobiologists are now beginning to reveal the physical mechanisms by which the human neuroarchitecture accomplishes decision making. Although in their infancy, these studies suggest both a single unified framework for understanding human decision making and a methodology for constraining the scope and structure of economic theory. Indeed, there is already evidence that these studies place mathematical constraints on existing economic models. This article reviews some of those constraints and suggests the outline of a neuroeconomic theory of decision.

© 2004 Elsevier Inc. All rights reserved.

1. Introduction

The history of economics has been marked by an iterative tension between prescriptive and descriptive advances. Prescriptive theories seek to define efficient or optimal deci-

* Corresponding author: Center for Neural Science, 4 Washington Place, 809 New York, NY 10003, USA.
E-mail address: glimcher@cns.nyu.edu (P.W. Glimcher).

sion making which descriptive advances then invariably suggest do not accurately describe human behavior. The neoclassical revolution, and the period that followed it, were no exception to this general paradigm. Working from the assumption that all of human behavior could be described as a rational effort to maximize utility, the neoclassical theorists largely succeeded in developing a coherent basic mathematical framework. What followed, beginning with the work of scholars like Allais (1953) and Ellsberg (1961), were a series of descriptive insights which indicated either that humans were poor utility maximizers or that the underlying assumptions of the neoclassical revolution were flawed.

Over the last two or three decades economists have responded to the descriptive challenge raised by these post-neoclassical studies by adopting one of two basic approaches. Either they argue that rational decisions based on utility theory occur only under some conditions and that defining those conditions is of paramount importance (cf. Simon, 1947, 1983), or they argue that standard utility theory requires modifications, additions, or novel approaches (cf. Savage, 1954; Kahneman and Tversky, 1979). The fundamental problem imposed by bounding rationality is that the resultant models have little or no predictive power outside of their bounded domains. The problem modified utility theories face is that these newer models often fail to be parsimonious and often appear ad hoc or under-constrained.

One recent trend in economic thought may reconcile this tension between prescriptive and descriptive approaches. There is some hope that it may yield an economic theory that is both highly constrained and parsimonious while still offering significant predictive power under a wide range of environmental conditions. That trend is the growing interest amongst both economists and neuroscientists in the physical mechanisms by which human decisions are made within the human brain. There is reason to believe, some of these *neuroeconomic* scholars argue, that the basic outlines of the human decision making architecture are already known and that studies of this architecture have already revealed some of the actual computations that the brain performs when making decisions. If this is true, then a combination of economic and neuroscientific approaches may succeed in providing a methodology for reconciling prescriptive and descriptive economics by producing a highly predictive and parsimonious model based on the actual economic computations performed by the human brain.

At this time there are, however, profound differences between the approaches taken by neuroscientists and economists interested in this problem. Neuroscientists tend to underestimate the complexity of actual human decision making and thus fail to take full advantage of the existing economic corpus, studying choice under conditions that economists often see as trivial. Indeed, to economists many of the recent neurobiological studies of decision making seem to be more about reflexes than about economic behavior. Economists, in a similar way, often employ overly simplistic or outdated notions of brain function that are only weakly related to the modern consensus views held by neuroscientists. As a result, many neurobiologists dismiss the work of economists as irrelevant to the study of the human mind and brain. This often leads members of the larger economic community to reject neuroeconomics as irrelevant to advancing economic knowledge and it leads members of the broader neuroscientific community to reject neuroeconomics as outdated or overly simplistic.

The primary goal of this paper is to attempt to resolve this discrepancy between economic and neuroscientific approaches to human decision making by demonstrating, for scholars with primary expertise in economics, how neuroscientific experiments can be used as tools for developing real economic theories and models. After a brief introduction, the paper reviews a series of empirical studies that many neuroscientists believe describe the basic architecture for decision making in both human and non-human primates. Working from an understanding of this architecture, we next describe an experiment aimed at deriving the neurobiological algorithm for calculating utilities in a two-alternative lottery task. The economic model derived from this physiological experiment is then used to predict the dynamic play-by-play behavior of real subjects engaged in the actual lottery task. We hope that presenting the material in this way will demonstrate that economic theory can be used to guide neurobiological experiments which can, in turn, yield new economic theories.

Along the way, we hope to explain three central points around which future developments in neuroeconomics will likely have to be organized. First, we hope to explain how very profoundly our current neuroscientific and current economic theories of brain function differ and how these differences can only be reconciled if economists become familiar with the highly quantitative models of brain function that are at the core of contemporary neuroscience. Second, we hope to stress the importance, to economics, of evolutionary biology. Humans are unique organisms, but there is growing evidence that we are far less unique in the production of economical behavior than most working economists suspect. For example, monkeys can play mixed strategy equilibrium games with the same efficiency as humans (Dorris and Glimcher, 2003) and birds can systematically alter the shape of their utility functions to adopt risk preferences appropriate for their environments (Caraco et al., 1980). There is now abundant evidence that our own economic behavior is evolved from, and very closely related to, the economic behaviors of our animal relatives. This may be the most critical point made in this paper because it calls into question the pervasive assumption amongst economists that our decision making process is both a uniquely human faculty and a broadly rational faculty. Third and finally, we hope to show that neuroscientific studies of economic behavior can be much more than efforts to locate a brain region associated with some hypothetical faculty like ‘justice’ or ‘cooperation.’ Such studies are valuable starting points, but have troubled some economists because they provide no predictive power with regard to economic behavior. We hope to demonstrate that neuroeconomic experiments can and will reveal the nature of the economic computations brains perform.

1.1. The gap between economic and neuroscientific conceptualizations of the brain

The neoclassical revolution had two profound effects during the second half of the twentieth century: it largely revealed how a rational utility maximizer would behave and essentially proved that humans could not be viewed as efficient utility maximizers under all conditions. This insight led a number of economists, perhaps most notably Herbert Simon (1997), to conclude that human decision makers could be viewed as rational utility maximizers in only a bounded sense. Conditions do occur under which humans behave rationally but there are also conditions under which humans behave in a clearly irrational

manner. One result of this insight has been a growing conviction in the economic community that human decision making can often be viewed as the product of two underlying processes, a bounded rational process well described by prescriptive economic theory and an irrational process which is best described empirically.

During the last decade a number of economists have begun to suggest that these two processes, the rational and irrational, may be instantiated within the human brain as two distinct mechanisms. Indeed, many have even suggested that irrational behavior can be uniquely attributed to limitations intrinsic to the neural architecture while rational behavior can be viewed as the product of a conscious faculty that somehow transcends this biological limitation. Vernon Smith put this in his 2003 Nobel prize lecture, when explaining the irrational effects of context on decision making, “[t]he brain, including the entire neurophysiological system, takes over gradually in the case of familiar mastered tasks and plays the equivalent of lightning chess ... all without conscious thinking by the mind.” Smith and others have argued that it is the mechanical processes of the brain itself which account for the irrationality that bounds the rational processes of the conscious mind. Arguing in more detail, Camerer et al. (2003) have suggested that human decision-making can be viewed as the product of one cognitive and one affective (or emotional) system and that these two systems co-exist as independent entities within the neural architecture because they have different evolutionary origins. These authors have even drawn on the existing neuroscientific literature to argue that each of these distinct modules for decision making can be localized to distinct anatomical regions within the human brain. For example they suggest that “regions that support cognitive automatic activity are concentrated in the back (occipital), top (parietal) and side (temporal) parts of the brain.”

At the same time that this revolution has been occurring in economic circles, neuroscientists interested in human decision making have begun to head in a surprisingly different direction. The revolution that gave birth to modern neuroscience in the early part of the twentieth century argued that all of human behavior could be conceived of as the product of two fundamentally distinct mechanisms: a sophisticated faculty that governed complex behavior and a simpler, cruder, mechanism that could produce reliable, but unavoidably simplistic, behaviors (Sherrington, 1906; Damasio, 1995; LeDoux, 1996; Glimcher, 2003a). This simpler mechanism, which came to be identified with the notion of a reflex, was widely believed to be tractable to neurophysiological analysis and formed the core of our understanding of brain function during the first half of that century.

During the last several decades, however, ongoing empirical work has begun to suggest to many neuroscientists that this view of the neural architecture is no longer tenable. More and more biological evidence now suggests to neuroscientists an essentially unitary view of the neural architecture that is much more deeply rooted in evolutionary theory than this original dualistic conception. What is emerging in neuroscientific circles is the view that a surprisingly holistic (though multi-component) decision making process governs behavior (Parker and Newsome, 1998; Schall and Thompson, 1999; Glimcher, 2003b). The varied inputs to this decision making process, it is argued, have all been shaped by evolution in order to yield a unified pattern of behavior that maximizes the *reproductive fitness* of organisms within the environments in which they operate (Maynard Smith, 1982; Stephens and Krebs, 1986; Krebs and Davies, 1991). Evolution makes animals, these scientists argue, fitness maximizers. But critically, evolution performs this role on all parts of

the organism simultaneously. Evolution yields a unitary organism, the global rationality of which is bounded by the requirements of the environment within which it evolved.

The economic capabilities of humans have, however, led many to conclude that we are fundamentally different from other animals in this regard, that we achieve rationality through a distinct and uniquely human mechanism than stands apart from the mechanisms possessed by other animals. The mechanisms that other animals possess may indeed still reside within our brains, but it is the irrational aspects of human behavior which can be uniquely attributed to this biological heritage. Quite compelling empirical data argue against this conclusion. First, it now seems clear that even animals with very small brains can behave in a surprisingly rational manner under a broad range of conditions. This seems to argue against the idea that in order to behave rationally humans would have needed to evolve some unique facility. Second, there is growing evidence that we share with our nearest relatives not just the ability to behave rationally, but we also share with them common boundaries to our rationality. If this is true then it is both the rational and irrational which we share with our nearest relatives, challenging the assumption that any of these aspects of behavior involve some uniquely human process. These data argue, in essence, that we differ more in degree than in nature from our nearest living relatives.

1.2. Evolutionary biology and economics: rational choice in simpler brains

In 1982, D.G.C. Harper published an influential experiment on the rationality with which mallard ducks forage for food (Harper, 1982). Mallard ducks were an interesting choice because their avian lineage evolved from dinosaurs about 200 million years ago and thus they are animals with an evolutionary heritage very different from our own. Further, they are animals with extremely small brains, typically less than 5 grams in weight. (In contrast, the human brain weighs about 1400 grams.) At an environmental level, these ducks live in small groups of about 10–50 individuals and normally obtain food by foraging together at waters edge. Finally, as with all animals who must maintain very low body weights in order to fly, they store little energy internally and thus their ability to survive and reproduce is well correlated with their ability to obtain food on a daily basis; at least amongst flighted birds, individuals who maximize the rate at which they obtain food each day maximize their long-term reproductive fitness (Krebs and Davies, 1991).

Harper's experiment focused on the behavior of a particular flock of 33 mallards that wintered on the main pond in the botanical gardens of Cambridge University in 1979. What specifically interested Harper was foraging strategies. To examine that possibility, Harper conducted a series of group decision-making experiments of a kind that will be familiar to most economists. At the beginning of each day two experimenters would approach the pond, each with a sack of bread-balls all having a particular size and weight. Standing at two separate locations the experimenters began throwing those bread-balls simultaneously but at different rates. The job of each duck was simply to decide in front of which experimenter to stand. On a typical day experimenter 1 would, for example, throw a 2-gram bread-ball once every 5 seconds while experimenter 2 would throw a 2-gram bread-ball once every 10 seconds. What Harper would measure was the moment-by-moment decisions of each duck, both while these conditions were held constant and when they changed, during a foraging period that lasted tens of minutes.

If we treat the situation as a 33-duck Nash-type game and assume that the ducks employ a standard concave utility curve for bread-balls, then a single Nash equilibrium emerges under these conditions. Taking a locally linear approximation of utility for the range of bread-ball sizes from 0 to 4 grams, then we can identify the precise Nash equilibrium and make quite specific predictions about what constitutes rational behavior for these animals. Under these conditions the expected utility of standing in front of experimenter 1 for any one duck should be equal to the probability of obtaining a bread-ball from that experimenter multiplied by the size of the bread-ball being thrown. If the likelihood, for any individual, of obtaining a bread-ball is a linear function of the number of other ducks standing before experimenter 1 (and Harper verified that it was) then the expected utility (EU) of standing in front of experimenter 1 is proportional to the bread per minute thrown by experimenter one divided by the number of ducks standing in front of experimenter one. At Nash equilibrium the EUs of standing before either experimenter must be equivalent, so if experimenter 1 were throwing a 2-gram bread-ball every 5 seconds and experimenter 2 were throwing a 2-gram bread-ball every 10 seconds then equilibrium would be reached when two-thirds of the ducks stood in front of experimenter 1 and the other one-third stood in front of experimenter 2. Equilibrium would occur when the ducks were probability matching.

Perhaps surprisingly, Harper found that this very accurately described the behavior of the ducks under a wide range of conditions. At all of the rates and bread-ball sizes Harper explored, within about 60 seconds of the start of bread-ball throwing, the population of ducks had assorted itself out at the Nash equilibrium solution. That means that they achieved this solution after as few as 6 bread-balls had been thrown by one of the experimenters. Further, when Harper and his assistants changed either their rate of throwing or the size of the bread-balls, the ducks re-assorted themselves, once again achieving a rational equilibrium within about 60 seconds. The ducks as a group behaved in a perfectly rational manner, in a manner that many economists would argue was evidence of a rational conscious process if this same behavior had been produced by humans operating under simple market conditions like these.

But perhaps just as interesting as these observations on the behavior of the group were Harper's observations on how individual ducks behaved. Within the flock ducks have an established pecking order and conflict between the ducks continually challenges and renews this order. Harper observed that this pecking order was evident within the flock as they foraged. Not all ducks obtained the same amount of bread (the likelihood of obtaining any given bread-ball was proportional to rank) and ducks conflicted with each other for access to the bread. Mechanisms of conflict, aggression and competition were operating while this rational solution was being achieved.

The ducks behaved rationally. Does the fact that it was ducks who behaved in this way make decisions of this type uninteresting to economists or irrelevant to studies of human choice? Or do these results suggest that the classical models of rationality based in utility theory could in principle be used by biologists to study brain function in non-human animals? If such a study were undertaken could it tell us anything of interest to economics? To begin to answer those critical questions we need to re-examine two final issues before turning to the body of this paper, one issue in economics and one issue in neuroscience. First, we need to review the development of classical utility theory. Second, we need to ask

how classical utility theory may be related to the neural architecture for decision making that exists in animals ranging from ducks to humans.

1.3. Using neuroscience as an economic tool

Modern utility theory has its origins in the theory of expected value first proposed by Pascal. He argued that the value of any course of action could be determined by multiplying the gain that could be realized from that action by the likelihood of receiving that gain. This product, which we now call expected value, was presumed to represent a rational decision variable. While Pascal and his colleagues recognized that not all human decision making could be accurately described with expected value theory, they argued that all rational decision making should follow this prescriptive theory (cf. Arnauld and Nicole, 1996; Pascal, 1966).

By the mid-1700s, however, it was clear that the Pascalian approach did an extremely poor job of predicting human choice behavior under conditions of significant risk. Daniel Bernoulli made this point in 1738 (see Bernoulli, 1954). “To make this clear it is perhaps advisable to consider the following example: somehow a very poor fellow obtains a lottery ticket that will yield with equal probability either nothing or twenty thousand ducats. Will this man evaluate his chance of winning at ten thousand ducats? Would he not be ill-advised to sell his lottery ticket for nine thousand ducats? To me it seems that the answer is in the negative.” Bernoulli argued for a model of rational decision making in which the likelihood of a gain was multiplied by the utility, rather than the value, of that gain. His notion was that gains were represented in the decision process by a roughly logarithmic function of value that also incorporated a representation of the chooser’s wealth.

Modern work has, however, made it abundantly clear that this theory also falls far short of the descriptive goal of predicting actual human behavior. For example, Allais (1953) demonstrated that human choice can be non-transitive, Kahneman and Tversky (cf. Kahneman et al., 1982) demonstrated that human choice behavior deviates widely from monotonicity, and most recently game theorists have even shown that under some conditions (cf. Guth et al., 1982) humans knowingly make choices that will result in losses rather than gains. All of these experiments point out the limits of classical utility theory as a tool for understanding human choice behavior. As a result many economists have proposed that decision making is best viewed as involving the interaction of a utility-based mechanism and a second, perhaps less rational, process. The central argument that we will make in this paper is that current neurobiological data contradicts this view and instead supports a model of human and animal decision making more closely tied to the core insight that Pascal and Bernoulli provided.

Utility theory proposes that decision makers must represent the desirability of each possible course of action using a common scale and that choosing is the process of selecting the most desirable of these possible courses of action. Pascal had argued that desirability should be computed as the product of value and likelihood of gain. Bernoulli had taken an important step forward by arguing that desirability involved a more complex computation influenced by properties intrinsic to the chooser, like current wealth. Although Bernoulli clearly meant utility theory to be an objective and prescriptive model for decision making, in this regard he came very close to introducing a subjective model (cf. Savage, 1954)

for the decision making process. In Bernoulli's model, two variables from the external world were modified by processes internal to the chooser and the product of these internal computations, expected utility, was then represented and used to make choices. Although there is still significant uncertainty about the precise form of that internal computation, current neurobiological evidence seems to strongly support this early claim. The brains of primates, almost certainly including humans, appear to represent a complex variable which under many circumstances closely parallels classical expected utility. In the final stages of decision making, the neural architecture seems to select the most desirable action from amongst representations of the desirability of all available actions by a winner-take-all process.

This model, which will be developed at length below, does however depart from neo-classical economic theory in an important way. Neoclassical theory has always made the famous *as if* argument: it is *as if* expected utility was computed by the brain. Modern neuroscience suggests an alternative, and more literal, interpretation. The available data suggest that the neural architecture actually does compute a desirability for each available course of action. This is a real physical computation, accomplished by neurons, that derives and encodes a real variable. The process of choice that operates on this variable then seems to be quite simple; it is the process of executing the action encoded as having the greatest desirability. Of course the challenge that this emerging view poses is thus to determine exactly how this desirability is computed. It is this process which combines elements of Bernoulli's utility theory and other operators in an evolutionary context to achieve efficient decision making in the environments for which each species evolved.

While neuroscientists are only just now beginning to describe the computation that transduces objective measures from the outside world into this representation of desirability, several factors are already becoming clear. First, under many conditions, conditions under which choice appears rational, this desirability encoded by the neurons of the brain very closely approximates expected utility. Second, under conditions in which choice behavior is poorly predicted by rational choice models, these neural representations still encode the desirability of each course of action, although under these conditions desirability and expected utility are of necessity not identical. The available data suggest that the neural decision-making process is always rational with regard to these internal representations of desirability. When choosers deviate from rationality it is this physiological encoding of desirability, which we refer to as *physiological expected utility*, that departs from neoclassical theory.

Together, these observations raise an intriguing possibility which forms a central subject of this paper: the neural architecture may indeed compute and represent the physiological expected utility of many possible courses of action, much as neoclassical utility theory proposes. Evolution may have shaped the neural architecture to perform efficiently under many, but not all, environmental circumstances. When choosers are efficient in the economic sense, that architecture accurately represents the expected utility of available choices. When physiological and objective utility differ, it reflects inefficiency not in the mechanism that chooses, but in the ability of the neural architecture antecedent to the choice mechanism to compute physiological expected utilities efficiently. In some cases inefficiencies of these types will arise when the most complicated cortical mechanisms for estimating likelihoods encounter problems that they did not evolve to solve. In other cases,

inefficiencies will occur because simpler brainstem systems encounter problems that they did not evolve to solve efficiently. All of these biologically generated inefficiencies would therefore bound rational behavior. The available evidence thus suggests a synthesis of modern economic and neuroscientific approaches. By biologically defining the mechanisms which compute physiological expected utility we should be able to derive a mechanistically accurate economic theory which is by necessity predictive.

In the following sections we hope to present a case study of this approach. Beginning with a physiological investigation of choice mechanisms, we will derive a mathematical description of the process by which a particular class of dynamic decision making is accomplished. Having derived an algorithm for choice under these conditions we will use that surprisingly parsimonious model to predict the dynamic play-by-play choice behavior of individuals under novel conditions. The accuracy of those predictions will then be tested to assess this neurobiologically derived model. In essence, what we hope to do is to use neurobiological techniques to develop a simple economic theory that is both testable and parsimonious.

2. The neuroscience of connecting sensation and action

2.1. Overview of sensory and motor neuroscience

During the second half of the twentieth century neuroscience made huge advances, particularly towards understanding both the structure and function of the sensory systems that gather data about the outside world and the movement control systems through which all behavioral responses are generated. For the most part, these studies provided the insights upon which our current understanding of the human brain rests. These studies provide, essentially, a core theory of brain function which, like the neoclassical approach in economics, organizes the ways scholars address almost all questions of neural function. In order to understand how neurobiologists attempt to understand decision making, it is therefore necessary to know a little about the organizing principles of these input and output systems.

2.1.1. Sensory systems

Tremendous progress has been made towards understanding all of our senses but the brain system we understand best is the visual system. (For an introductory overview of the visual system see the vision chapter in the excellent textbook by Rosenzweig et al. (2002). For a more detailed overview see the textbook by Squire et al. (2002)). Insights from the study of this system organize neurobiological approaches not just to sensory systems but to brain function in general. The work of this system begins in the retina, a five layer thick sheet of cells lining the inner surface of the eyeball like a sheet of photographic film. At each location on this sheet lies a single photoreceptor, a cell which transduces individual photons of light into electrochemical signals that can be passed to the brain. These electrochemical signals are, in turn, passed by a class of retinal neurons called *retinal ganglion cells*, through the optic nerve which leaves the eyeball and connects to the neurons




Fig. 1. The basic flow of information in the primate visual and eye movement systems shown superimposed on a monkey brain. Vision: photons entering the eyeball activate neurons in the retina. That activity is relayed through the optic nerve to the lateral geniculate nucleus (LGN) of the thalamus. From there information passes to the primary visual cortex (VI) and on to higher level visual cortices (V2, V3, V4, MT, etc.). These signals gain access to movement control systems via a number of pathways, one of which involves the parietal cortex. One subregion of the parietal cortex, the lateral intraparietal area (LIP) is known to be of particular importance. Movement: movements of the eyes are controlled by many areas acting in concert. Of particular importance is the lateral intraparietal area (LIP) of parietal cortex. Activity in this area influences both the cortical frontal eye field (FEF) and the subcortical superior colliculus (SC). These areas in turn regulate the brainstem regions (BS) that govern the muscles that surround the eye.

of the *lateral geniculate nucleus of the thalamus* which lies inside the mammalian brain (Fig. 1).

The lateral geniculate nucleus in humans and monkeys is a laminar structure, composed of six pancake-like sheets of neurons stacked one on top of each other. Each sheet receives a topographically organized set of projections from one of the two retinæ. This topographic organization means that at a particular location in, for example, the second layer of the lateral geniculate, all the neurons receive inputs from a single fixed location in one of the two retinæ. Because individual locations in a retina monitor a single location in visual space (like a single location on a photographic negative) each location in the geniculate is thus specialized to monitor a particular position in space.

It has also been shown that adjacent positions within any given geniculate layer receive projections from adjacent positions within the referring retina. This adjacent topographic mapping means that each layer in the geniculate forms a complete and topographically organized map of the images that fall on the retinæ. In this map, as in almost all structures made of nerve cells, information is encoded by the level of electrochemical activity of the individual cells that make up the map. Geniculate neurons highly activated by light falling on the portion of the retina they monitor respond by producing pulses of electrical activity, called *action potentials*, at a high rate. In these neurons it is essentially the rate of action potential generation that is used to encode properties like the contrast or brightness of a location in the visual world. This set of organizing principles means that each layer in the geniculate forms a complete and topographically organized screen on which are projected, as a pattern of action potentials, the visual images that fall on one of the retinæ. Thus activation of a particular neuron at a particular location within the geniculate map indicates that a visual stimulus has appeared at that location in the visual world.

These geniculate maps project, in turn, to the *primary visual cortex*. Lying against the back of the skull, the primary visual cortex, also called *area VI*, is composed of several million neurons. These neurons form their own complex topographic map of the visual

world organized into roughly 1 square millimeter patches. Each square millimeter of tissue is specialized to perform a common basic set of analyses on the light that falls on a specific region of the retina. Within these 1 mm by 1 mm chunks of cortex, individual neurons have been shown to be highly specialized in ways that allow many different analyses to proceed simultaneously. For example, some neurons in each patch become active, producing action potentials, whenever a vertically oriented boundary between light and dark falls on the region of the retina they monitor. Others are specialized for light-dark edges tilted to the right or to the left. Some respond to input exclusively from one retina, others respond equally well to inputs from either retina. Yet others respond preferentially to colored stimuli. Amongst neurophysiologists this complex pattern of sensitivities in area V1, or of *receptive field properties*, is of tremendous conceptual importance. It suggests that information coming from the retina is sorted, analyzed and recoded before being passed on to visual areas that lay farther along in the processing stream.

The topographic, or *retinotopic*, map in area V1 projects to a host of other areas which also contain topographically mapped representations of the of the visual world. Areas with names like *V2*, *V3*, *V4* and *MT* construct a maze of ascending and descending projections amongst what may be more than 30 mapped representations of the visual environment. Each of these maps appears to be specialized, extracting specific types of information from the visual image. One of these areas, for example, forms a topographic map that encodes the speed and direction at which visual images move across the retina. Others encode information about the presence of faces or objects. This network of maps is the neural hardware with which we perceive the visual world around us.

The most critical features of this brain organization are first, that incoming sensory information is organized in a massively parallel topographic fashion and second that there seems to be an orderly progression of information from the peripheral receptors to the cerebral cortex where some of the most complex analyses are performed. Fortunately, most of the other sensory systems follow a very similar organizational plan. The sense of touch, for example, involves the passage of signals originating in the skin to a topographically mapped nucleus in the thalamus. From there these signals pass to the topographically organized *somatosensory area I* of the cortex and from there to higher order somatosensory areas in the cortex. Our understanding of the visual system therefore serves as a guide for understanding how essentially all information about the outside world is gathered by the brain.

2.1.2. Motor systems

Within neurobiology, studies of movement control areas, usually referred to collectively as the components of the *motor system*, are segregated into two main divisions: those that control systems that regulate movements of the body, hands, feet and mouth (the skeleto-muscular system) and those that move the eyes (the oculomotor system). As with the sensory systems, there seem to be strong parallels between the multiple motor systems of the brain and as in studies of the sensory systems, our core framework largely derives from studies of one system, in this case the oculomotor system. The oculomotor system has provided especially fertile ground for study because of the simplicity of the mechanics of the eyeball. While movements of the arm, for example, involve dozens of muscles and complex inertial moments, movements of each eye involve only 6 muscles and no detectable inertia.

(For an introductory overview of the motor system see the motor chapters in Rosenzweig et al., 2002. For more detail see Squire et al., 2002.)

When an eye movement is produced, for example an orienting eye movement or *saccade* that rapidly shifts the point-of-gaze from one location to another in the outside world, the six muscles that control the position of each eye are activated by six groups of neurons that lie deep in the brainstem. These *motor neurons* are, in turn, controlled by two systems in the brainstem. One that regulates the horizontal position of the eye and one that regulates the vertical position of the eye. These two control centers receive inputs from the *superior colliculus* which lies just beneath the thalamus and the colliculus, in turn, receives its principal input from the *frontal eye field* of the cerebral cortex. Like the visual areas described above, the superior colliculus and the frontal eye field are also constructed in topographic fashion. In this case, their constituent neurons form topographic maps of all possible eye movements. Imagine a photograph of a landscape. Now lay a transparent coordinate grid that shows the horizontal and vertical rotation of the eye that would be required to look directly at any point on the underlying photograph. Both the superior colliculus and the frontal eye fields contain maps very like these transparent coordinate grids. Activation of neurons at a particular location in the frontal eye field produces activation in a corresponding position in the superior colliculus which in turn activates the brainstem areas that cause a saccade of a particular amplitude and direction to be executed. If this point of activation were to move across the cortical map of the frontal eye field, the amplitude and direction of the elicited movement would change in a lawful manner specified by the horizontal and vertical lines of the coordinate grid around which it is organized. The neurons of the superior colliculus and the frontal eye field can thus be viewed as topographically organized command arrays in which every neuron sits at a location in the map dictated by the direction and length of the saccade it produces.

Studies of the arm movement or verbal movement systems are at a comparatively early stage, but it seems fair to say that the basic features of these systems appear quite similar at this level of analysis. Again, a few critical features of the nervous system seem to emerge from this knowledge. First, outgoing signals are organized in a massively parallel topographic fashion. Second, there seems to be an orderly progression of information from higher areas like the cortex down to lower areas that ultimately control the movements of the muscles themselves.

2.2. Early studies of decision making

Over the course of the last 15 or 20 years a number of very influential studies, initially conducted in monkeys, have begun to examine the simplest possible connections between these sensory and motor architectures. As such, these studies constituted the first serious neural examination of decision making, albeit decision making of a very simple kind. Jeffrey Schall and his colleagues at Vanderbilt University conducted some of the first of these studies (cf. Hanes and Schall, 1996; Schall and Thompson, 1999), training thirsty rhesus monkeys to stare straight ahead at a centrally located spot of light presented on a video display (Fig. 2). Shortly after the monkey began staring straight ahead eight secondary targets appeared, arranged radially around the central fixation stimulus. Seven of those targets appeared in a common color and one appeared in a different color, an *oddball*. If the animal




Fig. 2. The Oddball Task. In the original Schall experiments thirsty monkeys were seated staring at a cross in the center of a blank display. Eight spots of light then illuminated, seven in one color and an eighth, the oddball, in a different color. The monkey had to decide where to look, and only if he looked at the oddball did he earn a fluid reward. While monkeys made these decisions Schall and his colleagues monitored the activity of single neurons in the frontal eye fields.

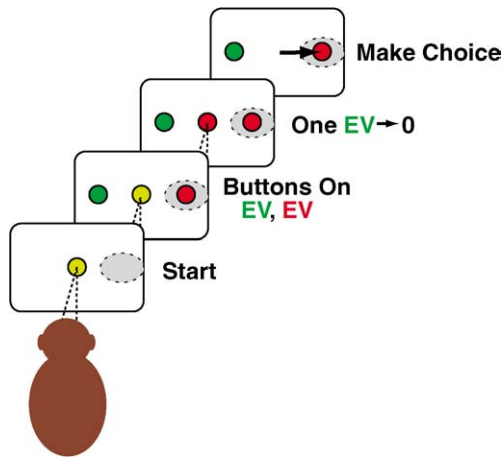


Fig. 3. Platt and Glimcher's Two-Choice Cued Lottery. Once again thirsty monkeys are seated staring at the center of a display. Red and green response targets, acting as buttons, then appear. At this point in the play the red and green responses have widely divergent expected values which the monkey is expected to have learned. After an imposed delay the central spot changes color identifying one of the targets as the rational choice and the other target as valueless on this trial. If the monkey looks at the target that is the same color as the central target he earns the appropriate reward.

looked at any of the 7 common color targets the play, or trial, ended immediately. If he looked at the oddball, he received a drop of fruit juice as a reward.

Under conditions like these, we know quite a lot about both the sensory and motor processes that must become active in the monkey's brain. When the targets illuminate,

we know that eight locations in the retinal, lateral geniculate and visual cortical maps become active. One for each of the eight visual targets. These signals propagate through the visual system towards saccadic eye movement control centers like the frontal eye fields and the superior colliculus. Only one of the 8 locations, however, represents the oddball and ultimately leads to activation of the eye movement control circuitry in those areas. So how is the translation, from 8 visual signals to one motor command, actually accomplished? To answer that question Schall and his colleagues studied the activity of single nerve cells in the saccadic movement maps of the frontal eye fields while monkeys performed this oddball detection task.¹

Schall found that rate of action potential generation by neurons at each of the eight locations in the frontal eye field map rose to an early peak shortly after the 8 targets were illuminated, but only after about 0.08 seconds was there evidence, in these neurons, of an underlying decision process in operation. At that point, neuronal action potential firing rates continued to grow only in neurons at the one location encoding the oddball. After the level of activity at that location crossed an apparently fixed threshold value, the movement was produced. This led Schall to suggest the existence of a decisional threshold in each neuron in the topographic map of the frontal eye field and raised the possibility that the topographic map was constructed in such a way that only one local cluster of neurons within the map could reach the decisional threshold at a time. The topographic map seemed to serve as an organizational framework for imposing something like a winner-take-all decision making strategy. Of course a winner-take-all strategy is critical because one cannot usefully look in two directions at once.

Where do the sensory signals that trigger the threshold activations of these neurons originate and how, if at all, are these signals related to the more complex decisions that are the subject of economic study? William Newsome and his colleagues at Stanford University provided a critical set of data for answering that question (cf. Parker and Newsome, 1998). They were initially interested in understanding how the brain generates the perception of motion so they began by training rhesus monkeys to watch a visual display that humans see as moving in an ambiguous fashion. They then asked their monkey subjects to report the direction in which the display appeared most likely to be moving. In these experiments the monkeys looked through a circular window at a cloud of white dots that appeared to move against a black background for two seconds. Critically, whenever the dots appeared, not all of them moved in the same direction. During any given two second display, many of the individual dots were moving in different, randomly selected, directions. Only a small fraction of the dots actually moved in a coordinated direction and it was this coordinated direction of movement that the monkeys were trained to detect.

¹ The technology for this kind of study is now in wide use with animals. Although quite safe, it is invasive and thus moral considerations preclude its use in humans. Single neuron studies of this type provide unequalled spatial and temporal resolution, a level of resolution critical for quantitative economic analysis and modeling. Using much cruder techniques like brain scanning, essentially all of these insights have been shown to be relevant to understanding the human brain, but our quantitative understanding of these processes rests almost entirely on studies in animals. To understand even these simple decisions animal studies are thus a necessary starting point to which Schall and his colleagues turned.

Newsome hypothesized that activity in one of the cortical visual areas might be both necessary and sufficient for the perceptual experience we have when we see an object move, activity in that area might be the physical instantiation of the subjective experience of seeing motion. Quite a bit was also known about the activity of individual neurons in this topographic map, the map in cortical *area MT*. Each MT neuron was known to become active whenever a visual stimulus moved in a *particular direction* across the portion of the visual world scrutinized by that cell's location in the MT topographic map. Each neuron thus had an idiosyncratic *preferred direction* and because each neuron prefers motion in a different direction and because many neurons work together to encode motion at each location in the visual world, the population of neurons in area MT could, in principle, discriminate motion in all possible directions at all visible locations.

In a series of experiments Newsome and his colleagues (Newsome et al., 1989; Salzman et al., 1990) were able to demonstrate that area MT forms a topographic map of the visual world in which the strength of motion in the visual world is encoded by the rate at which each neuron in the map fires action potentials. In principle this map thus provides, in Newsome's task, instantaneous and independent estimates of the strength and direction of motion at all locations in the visual world. If, for example, the animals were rewarded for correctly determining whether the spots they were watching tended to drift rightwards or leftwards, the activity in area MT encodes the information used by the animals to perform the task.

In a series of subsequent experiments and simulations Newsome and his colleague Michael Shadlen (cf. Shadlen et al., 1996; Shadlen and Newsome, 2001) sought to extend these observations by trying to determine how the signal originating in area MT was actually analyzed and used by the animal to produce eye movements that would consistently yield juice rewards in this environment, presumably by triggering the appropriate eye movement in the motor map of the frontal eye fields. It was known that neurons in area MT are functionally connected to maps in the posterior parietal cortex (Fig. 1) which are themselves connected to the maps in the frontal eye fields. This led Shadlen to propose that while the monkeys stared at the moving dot display, neurons in the posterior parietal cortex mathematically integrated the output of neurons in area MT with respect to time, yielding a new topographical map that encoded a time averaged estimate of motion direction at each location in the visual world. And critically, it was this time averaged estimate that should have served as the critical decision variable in the task that their monkeys had been taught. The precise conditions of the task they employed defined this as the optimal strategy for analyzing the visual motion. In other words, they proposed that the map in area MT made topographic connections to the posterior parietal cortex which extracted a decision variable from the MT activity and passed this decision variable, presumably topographically, to the frontal eye fields. In their model, which was developed formally (Shadlen et al., 1996), the neurons of the posterior parietal cortex thus served as topographically organized accumulators that could be used to trigger topographically aligned neurons in the frontal eye field map, thus generating the saccade most likely to be reinforced.

In a series of elegant experiments Shadlen and his colleagues went on to test this hypothesis and were able to verify the accuracy of many of their predictions. They were even able to demonstrate that the activity of neurons at each location in one of the areas of the poste-

rior parietal cortex, *area LIP*, was tightly correlated with the log of the likelihood that the eye movements encoded at that location would yield a reward (Gold and Shadlen, 2001).

2.3. Summary

These results suggest that the simplest kind of connection between sensation and action can be described as a process by which topographic parallel representations of signals from the outside world are used to trigger behavioral responses in topographically organized output maps, perhaps through the intermediate representation of a simple decision variable (see Glimcher (2003a) for a more in depth survey of this work). This much is uncontroversial. Also uncontroversial is that these sensory-motor connections do not constitute decision making in the economic sense. Neoclassical variables like value and expected utility, which are central to formal rational decision making, do not occur in a very clear fashion during these experiments. One possibility that this raises is that these are precisely the kinds of crude and primitive processes that are responsible for economically irrational behavior. Rational choice models may break down, be bounded, because mechanisms like these “take over.” But there is an alternative hypothesis. These mechanisms may be much more complicated than they appear from the experiments that have already been presented. Indeed, these experiments may reveal only the tip of the neurobiological iceberg. Distinguishing between these two hypotheses is, fortunately, an empirical problem. We can begin to ask whether these circuits and this general model can account for more complicated classes of decision making by examining these same neurons under conditions that more closely approximate the kinds of rational choice that are of interest to economists.

3. Economic studies of decision making in the brain

Ever since Pascal, economic analysis has focused on two variables which play an important role in rational choice: the likelihood of realizing a gain or loss and the magnitude of that gain or loss. Certainly other variables are important determinants of behavior, for example the time at which the gains will be realized, but magnitude and likelihood almost always influence rational choice. Do these variables influence the neural circuits we have already described? One of the first experiments to address that question systematically was conducted in 1999 by Platt and Glimcher. In that series of studies we asked whether the neurons in *area LIP* which Shadlen and his colleagues had examined might form a topographical map of either the magnitude or likelihood of gain associated with particular eye movements. Shadlen’s results had strongly suggested that the neurons in *area LIP* connected stimulus and response during simple perceptual decision making, and that the rate of action potential generation in these neurons might encode some kind of decision variable. We hoped to verify this hypothesis by constructing an experiment in which monkeys could make one of two possible movements while we systematically varied either the likelihood or magnitude of gain associated with each movement. This would allow us to determine whether the neurons in *area LIP* carried signals that might be useful for real economic decision making.

In that experiment, thirsty monkeys were trained to stare straight ahead at a central visual stimulus (Fig. 3) while two eye movement targets, which served the same role as response buttons serve in a typical economic task, were illuminated. At the end of each trial, or play, monkeys would have to choose whether to look at the left target, the right target, or to do nothing. Immediately before they had to make that decision, however, the color of the fixation light would change, identifying one of the two targets as valueless on that particular trial. The critical manipulation was that on sequential blocks of 100 trials the amount of juice that the monkeys would earn for each of the leftwards and rightwards movements was systematically manipulated. Finally, while the monkeys made decisions under these varying conditions, the activity of single neurons in area LIP was recorded.² Each neuron was examined while 5–7 different reward magnitude conditions were presented.

At a theoretical level, these subjects faced an exceedingly simple task. At the end of each play the color of the fixation light indicated what movement had both the highest expected value and expected utility and a rational chooser would be expected to produce that movement. We recognized, however, that expected utilities could be computed for each response early in each play, before the color of the fixation light simplified the task. Consider a block of 100 plays during which a leftward movement would yield 0.1 ml of juice and a rightward movement would yield 0.3 ml of juice. At the beginning of each play there was a 50% chance that a leftward movement would be identified as reinforced and, similarly, a 50% chance that a rightward movement would be reinforced. At that time the expected value of the two movements would be 0.05 ml and 0.15 ml of juice respectively. Then, after the fixation light changes color identifying, for example, the leftward movement as rewarded, the expected values change. After that point in the play the expected value of the leftward movement is 0.1 ml and the expected value of the rightward movement is 0 ml. We hoped to determine whether these early estimates of expected value were related to the activity of neurons in area LIP.

What we found was that the activity of LIP neurons was a surprisingly, though not precisely, linear function of these values (Fig. 4). More precisely, we found that both early and late in a play the firing rate for neurons associated with the leftward movement encoded:

$$\frac{LeftReward}{LeftReward + RightReward} = FiringRate. \quad (1)$$

Early in each play the firing rates of all LIP neurons were correlated with the relative expected value of their movements with regard to other possible movements. Late in the play, after the relative expected values of all but one movement had been reduced to 0, the neurons encoding the reinforced movement rose to a fixed firing rate near the maximum for these neurons.

² To economists it may seem very odd that the color of the fixation light was changed at the very end of the trial. This was essential, neurobiologically, because it allowed us to retain control over the actual movement that the animal produced. The color change allowed us to disassociate the movement produced at the end of the trial from the expected utility of the targets during the early part of the trial, allowing us to rule out the possibility that LIP activity was directly controlling *movement production* rather than abstractly encoding the likelihood or magnitude of gains that influence movement choice. In subsequent experiments, some of which are discussed below, animals were required to make decisions without the aid of this color change (see Platt and Glimcher (1999) for details). Those results confirmed the conclusions described here.




Fig. 4. The firing rate of a neuron in area LIP is correlated with relative expected value. Early in the play, when the expected value of each of the two possible movements is non-zero, firing rate is a roughly linear function of $[\text{Reward}_1 \div (\text{Reward}_1 + \text{Reward}_2)]$. After the cue indicates that the value of $\text{Reward}_2 = 0$, firing rate elevates. These results, and other described in the text, suggested that firing rates in area LIP may encode either the relative expected value or the relative expected utility of movements.

What this suggested was that something very close to an economic choice variable was indeed being carried by the firing rates of these neurons. We next had the monkeys perform the same task in a way that would let us determine whether the likelihood of a gain influenced LIP firing rates. To do that, we held the magnitude of the juice reward constant for both leftwards and rightwards movements across a series of blocks and varied the likelihood that at the end of each play the left or right movements would be reinforced. Thus on a block of trials in which both movements yielded 0.15 ml of juice but there was an 0.8 probability that the right movement would be reinforced and an 0.2 probability that the left movement would be reinforced, the expected values of the two movements early in each play would be 0.12 and 0.03 ml, respectively. Under these conditions the early firing rates of the neurons were again a roughly linear function of these values. Specifically, the firing rate of a neuron associated with the leftward movement was a linear function of the probability that the leftward movement would yield the juice reward. Together, these results suggested an interesting possibility, that the topographic map in area LIP encodes something like the relative expected value, or perhaps even the relative expected utility, of each possible eye movement under the conditions we had been studying.

If these neurons encoded relative expected utility under these simple conditions, what happens when behavior deviates from prescriptive economic theory, what happens when choice is only weakly related to expected value? Does some other less rational system gain control of decision-making while these neurons continue to encode prescriptive economic variables? The neurons in this area had been originally identified as a link in a very simple sensory-motor behavior, the kind of brain system that might have been expected to account for the bounds of rationality. Instead, in this experiment we had gathered evidence that these neurons might carry a signal that could be predicted by prescriptive theory. What exactly do neurons in area LIP encode and how are they related to rational and irrational choice? To answer this question we next turned to an experiment more like those employed by experimental economists.

3.1. Game theory and parietal maps

Our first goal in this next set of experiments (Dorris and Glimcher, 2004) was to be-
 havioral task which engaged humans in voluntary decision making and which could also
 be employed in a neurophysiological setting with monkeys. To this end, we had both hu-
 man and animal subjects play the role of the employee in the classic *inspection game* (cf.
 Kreps, 1990). The general form of the 2×2 payoff matrix for this game is shown in Fig. 5.
 We selected the inspection game because the payoff matrix can be easily adjusted to yield
 any mixed strategy equilibrium. We accomplished this exclusively, in our version of the
 game, by varying the cost of inspection to the employer (Fig. 5, left panel, variable I) such
 that at equilibrium the probability of shirking for the employee ranged from 10 to 90% in
 randomly ordered sequentially presented blocks of trials.

Rational decision-makers should choose the option with the highest expected utility on
 each play. If both subjects act rationally, then a mixed strategy equilibrium will be reached
 when the expected utility for each choice is equal for both players.

Thus at Nash equilibrium for the employee:

$$EU(S) = EU(W) \tag{2}$$

where $EU(S)$ is the expected utility for choosing to shirk, $EU(W)$ is the expected utility
 for choosing to work. If $p(I)$ is the probability of the employer inspecting and $1 - p(I)$
 is the probability of the employer not inspecting when at equilibrium (and we assume, for an
 initial analysis, that utility can be approximated as a linear function), W is the wage paid
 by the employer to the employee, and C is the cost of work to the employee then the payoff
 matrix (Fig. 5) expands to

$$p(I) * 0 + (1 - p(I)) * W = p(I) * (W - C) + (1 - p(I)) * (W - C); \tag{3}$$

solving for $p(I)$:

$$p(I) = C / W. \tag{4}$$

Similarly for the subject acting as the employer, at Nash equilibrium (again assuming
 for the initial analysis a linear utility function) the expected utility for inspecting is equal
 to the expected utility for not inspecting. Solving for $p(S)$:

$$p(S) = I / W \tag{5}$$




Fig. 5. The Inspection Game. Left panel shows the game in normal form. W = wage earned by the employee. C = cost of working for employee, V = value of work to employer, I = cost of inspection to employer. Right panels show representative payoffs yielding mixed strategy Nash equilibria when a linear utility function is assumed.

where $p(S)$ is the probability of the employee shirking when at equilibrium.

Because the employee payoffs remained the same for all blocks of trials, $p(I)$ for the employer should remain constant at 50% at all equilibria. Between blocks, $p(S)$ for the employee varied from 10 to 90% in 20% steps and was manipulated by varying the employer's cost of inspection from 0.1 to 0.9 in steps of 0.2 (see Eq. (5)).

3.1.1. *Human vs. human*

In the first set of experiments, pairs of human subjects were placed in separate rooms and played a repeated version of the inspection game. They were not aware of the nature of their opponent. It was another human in this case but could also have been a dynamic computer algorithm (see below). All subjects were naive to the nature of the payoff matrix and the game and were simply instructed to “make as much money as possible.” From the point of view of the employee, on each trial, it was necessary to use a mouse to chose one of two unlabeled buttons on a computer screen that corresponded to either working or shirking. After each play the payoff was presented on the screen along with a cumulative total of earnings over the last 10 trials. The first block of 50 trials was a practice session at a 50% shirking rate Nash equilibrium. Afterwards, 5 separate unsignaled Nash equilibrium blocks of 150 trials each were played in a randomized order over the course of two 1.5-hour sessions.³ A total of 8 subjects were tested. At the end of each session, subjects were paid their cumulative earnings which were typically about \$35 US.

The equilibrium equations presented above gave us a crude prescriptive theory of what subjects would do, and their actual behavior provided descriptive data. Our task as neurophysiologists would be to determine which, if either, of these two approaches predicted the activity of neurons in area LIP. To begin to do that we compared the observed rates of “working” and “shirking” with the rates predicted at equilibrium, once again assuming linear utility over the range of values we examined.

Figure 6 shows a 20-trial running average of the behavior of a human employee playing a human employer during two sequentially presented blocks. Although both players were free to choose either of two actions on every trial, we found that the overall behavior of our human subjects was surprisingly well predicted by game theory given our simple utility assumption. The gray lines show the unique Nash equilibrium solution for each block. Note that the employee quickly reached, and then fluctuated unpredictably around, these prescriptively defined equilibrium shirk rates.

To examine equilibrium behavior in greater detail we quantified, for each subject, the probability of shirking during the last half of each block. We then plotted this shirk rate against the predicted equilibrium rate (Fig. 7, Eq. (4)). We found that the responses of humans closely tracked these prescriptively defined shirk rates at behavioral equilibrium for rates above 40%. When, however, our prescriptive theory predicted rates below about 40%, we observed that subjects shirked more than predicted. There may be a temptation to conclude that this deviation reflects a non-linearity in the true underlying utility function. Were the underlying utility function for money to be significantly concave over the range of our measurements, as might be expected, then shirk rates would have decreased rather than

³ Block switches were unsignaled in order to keep the human task as similar as possible to the monkey task described below.

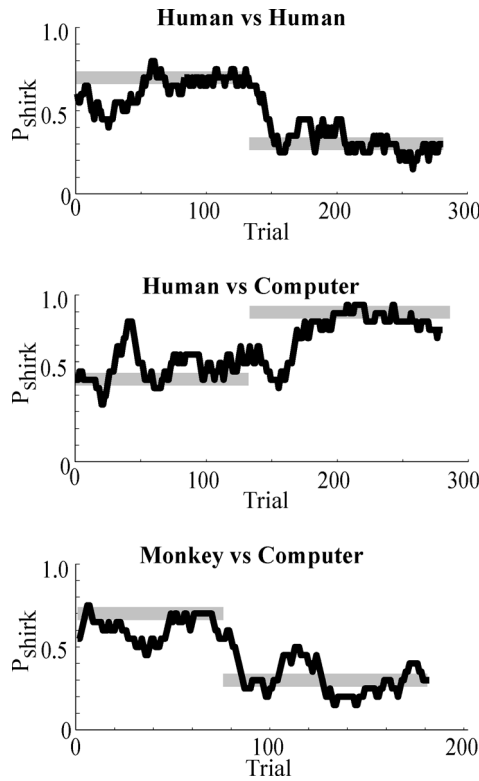


Fig. 6. Dynamic behavior of humans and monkeys playing the inspection game. Black lines plot a 20-play running average of the employee’s behavior during two equilibrium blocks. Grey lines plot Nash equilibrium solutions during both blocks assuming a linear utility curve.

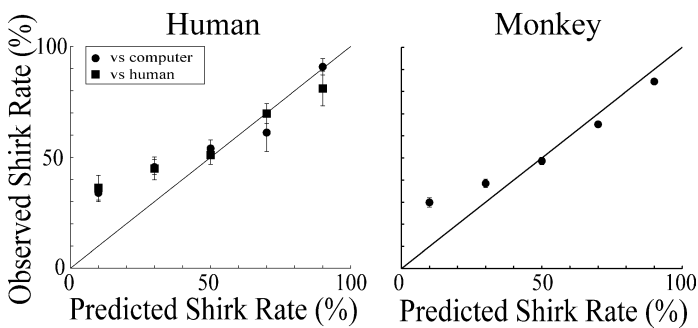


Fig. 7. Plot of equilibrium behavior for humans and monkeys. During the last half of each block of plays we computed the average rates of shirking (\pm S.E.M.) for our three groups during a total of 5 equilibrium conditions.

increased at these lower points. A more plausible explanation is either that this deviation reflects some sort of sampling strategy that maximizes the ability of players to detect block switches, or reflects an irrationality. In either case, what we had obtained were behavioral

measures during voluntary choice and at least some of these measures were well predicted by our prescriptive theory.

The implications of the measurements poorly predicted by our prescriptive theory were, however, uncertain. As with almost all experimental economic data, this either indicated an inadequacy in our prescriptive theory or a frank irrationality in our subjects. As physiologists, however, our goal would be to engage this deviation from our prescriptive theory using a neurobiological approach.

3.1.2. Human vs. computer

A second set of experiments were then conducted which were identical to the human versus human experiment except that the role of the employer was played by a standardized computer algorithm (see http://www.cns.nyu.edu/glimcher/inspection_game for MATLAB code of complete algorithm). That was critical because methodological constraints imposed by single neuron recording experiments essentially precluded our use of a real opponent against the monkeys. Briefly, the computer algorithm worked by tracking two variables of the employee's behavior:

- (1) the history of employee's choices to give an estimate of the overall probability that the employee would shirk ($p(S)$),
- (2) the employee's repetition rate (rep_{actual}), that is, how often a subject repeated the response of the previous play.

We calculated the expected repetition rate (rep_{expected}) for a given proportion of shirking assuming the probability of a response on each trial was controlled by a random process independent of previous choices:

$$rep_{\text{expected}} = (p(S) * p(S)) + ((1 - p(S)) * (1 - p(S))); \quad (6)$$

the difference in the rep_{actual} from rep_{expected} was used to bias the computer's estimate of $p(S)$ for the upcoming trial

$$p(S)_{\text{corrected}} = p(S) + \lambda(rep_{\text{expected}} - rep_{\text{actual}}) \quad (7)$$

in which λ was set to 0.1.

The variable $p(S)_{\text{corrected}}$ represents an estimate of the probability of the employee shirking on the current play given his past proportion of shirking and allows the algorithm to exploit any dependency of upcoming behavior on actions taken during the previous play. The variable $p(S)_{\text{corrected}}$ was substituted for $p(S)$ in calculating the relative expected utilities of inspect and no inspecting (for the computer employer) on the upcoming play which, in turn, was used to guide the computer's choice. In addition, an *exploration bonus* was added which gradually increased as the algorithm continued to produce a single response. This was necessary so the computer employer did not maintain a fixed $p(I)$ calculation (and thus fixed expected utility calculation) after every work trial, but rather continued to search for a maximally efficient behavior throughout the changing conditions of the experiment.

Of course, the computer employer would be deterministic if it always chose the option with the highest expected utility on every trial. If a human or monkey subject had sufficient

precision in a play-by-play estimate of $p(S)$, they could then accurately predict the actions of the algorithm. In order to incorporate stochasticity into the actions of the algorithm we employed a decision-rule which converted relative expected utility into a response probability. When inspecting and not inspecting had equal values, the decision rule randomly selected the inspect and no inspect options with equal likelihood. As the expected utility of one response increased, the probability that the more valuable response would be selected increased gradually.

Eight additional human subjects again played this version of the inspection game, playing five 150 trial blocks over 2 sessions. As in the first experiment they were not aware of the nature of their opponent and were simply instructed to “make as much money as possible.” Blocks of plays were presented exactly as before and subjects were paid their cumulative earnings which were again about \$35 US.

Figure 6 shows a 20-trial running average of the dynamic behavior of a human employee playing the inspection game during a typical session. The gray lines show the unique Nash equilibrium solutions. Just as when playing a human employer, our human subjects quickly reached and then fluctuated around, the shirk rates associated with the two sequentially presented Nash equilibrium states studied in this session.

During the last half of each block, once subjects had reached a stable state, we again determined the average shirk rate and plotted this against the equilibrium shirk rate prescribed by a linear utility function (Fig. 7, Eq. (4)). As was the case in the preceding experiment, these human subjects tended to deviate from the prescribed solution by over-shirking when shirk rates of 10 or 30% were predicted at equilibrium ($p < 0.05$, t -test against zero assuming unequal variance). Our standardized computer opponent thus elicited behavior from our employees that was statistically indistinguishable from their behavior when playing against human employers (two-way ANOVA, $F = 0.22$, $p > 0.05$, d.f. = 1).

3.1.3. Monkey vs. computer

We then trained monkeys to play a version of the inspection game against our computer employer and assessed whether their behavior was comparable to that of humans. In monkey experiments, thirsty animals competed for a water reward delivered after each play and indicated their choices on each play with a saccadic eye movement directed to one of two eccentric visual targets. Plays began with the illumination of a centrally located yellow fixation target. Once subjects were looking at this target two eccentric targets were illuminated, a red *shirk* target that was positioned so that the neuron under study was active when the monkey picked that target and a green *work* target that appeared opposite to the red target. Halfway through each play, the fixation point blinked and when it reappeared yellow the subject had 0.75 seconds to select and execute a response.

As Fig. 6 indicates, the dynamic behavior of monkeys playing this game appeared remarkably similar to the behavior of humans. At the beginning of each block monkeys, just like humans, quickly reached and then fluctuated unpredictably around the shirk rates associated with the Nash equilibrium states. When we examined the equilibrium behavior of the monkeys we found that it too appeared to be very similar to the equilibrium behavior produced by humans, even deviating from the prescriptive predictions in a similar manner. As Fig. 7 indicates, just like humans, during the inspection game the monkeys tracked the Nash equilibrium solutions (again assuming linear utility) and deviated from those solu-

tions when shirking rates of 30% or less were prescribed ($p < 0.01$). Two monkeys were studied while they played 8 sets of 100–200 trial blocks. Although the behavior of monkeys and humans was statistically differentiable during the 70 and 90% Nash equilibrium blocks, monkeys appeared to provide a surprisingly accurate model of humans overall.

3.1.4. *The physiological basis of strategic decision making*

Having thus established that humans and monkeys play this strategic game in an very similar fashion both when their behavior is predicted by the equilibrium equations we were using and when it was not, we were able to move on to our neurophysiological question: How is LIP activity related to choice during strategic decision making? One of Nash's (1951) fundamental insights was that at a mixed strategy equilibrium the desirability of the actions in equilibrium must be equivalent. This means that during the inspection game the expected utilities of working and shirking must be equal at equilibrium.

For the purposes of the foregoing analysis we had assumed a linear utility function and were able to thus prescriptively define a rational equilibrium. Our descriptive data indicated that this rational equilibrium did a fair job of predicting the behavior of our subjects under some conditions but failed under others. One could, of course, extend this particular prescriptive approach by incorporating a more realistic utility function and adding to that a learning algorithm that might even predict the over-shirking observed at low shirk rates. Our enhanced prescriptive theory would then better account for our descriptive observations.

If, however, the neurons in area LIP are the substrate upon which actual choice is generated and that choice is generated at equilibrium by a process similar to the one Nash envisioned, then we might be able to take an alternative approach. The Nash approach argues, essentially, that equilibrium occurs when the desirability of working and shirking are equal. Economists define those desirabilities as rational when they are well predicted by the expected utility of prescriptive theory. But rational or not, those desirabilities might well be represented at some point in the neural architecture. What we were trying to determine was whether the quantitative desirability of an action is encoded by the activity of neurons in area LIP not just for some categories of behavior, rational or irrational, but for behavior in general. If our economic approach was sound, then at behavioral equilibrium the desirability of working and shirking should have been equivalent. If our neurobiological approach was sound, then at behavioral equilibrium the level of neuronal activity associated with working and shirking should also have been equivalent. This should be true regardless of whether choice is rational or not. Put another way, if LIP encodes a physiological form of expected utility and this *physiological expected utility* is the actual substrate from which choice is produced, then rational behavior could be defined as occurring when prescriptive theory accurately predicts this physiological expected utility.

To begin to test the validity of this schema we began by repeating the Platt and Glimcher experiment described in the preceding section on our game playing monkeys. During the first block of plays highlighted in grey (Fig. 8), when a movement to the red target was instructed it yielded 0.25 ml of juice and when a movement to the green target was instructed it yielded 0.5 ml of juice. In the second block of trials, the payoffs associated with each target were reversed. Before the change in fixation point color indicated which movement would be rewarded, the neuron responded more strongly if the red target (which would

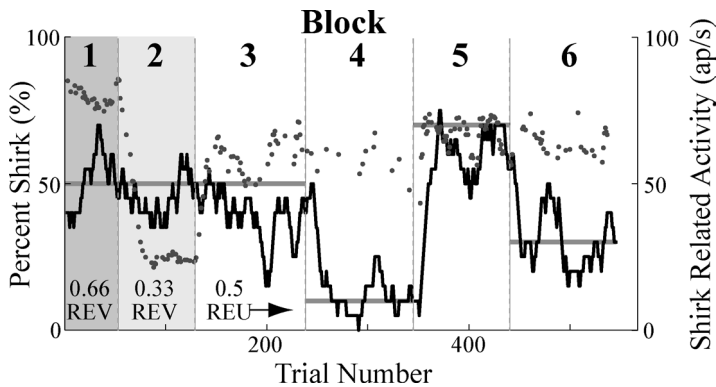


Fig. 8. Behavior of a monkey and an LIP neuron during the inspection game. The black line plots a 20-trial running average of the behavior of the monkey across six blocks of plays. In the first two blocks, highlighted in grey, no game was played but rather fixed expected values were presented. (Relative expected values of 0.66 and 0.33, respectively.) Blocks 3–6 presented 4 sequential payoff matrices corresponding to Nash mixed equilibria of 50, 10, 70, and 30% shirking. Grey dots plot the firing rate of the neuron on each play in which the monkey shirked.

serve as the shirk target later in the session) yielded a larger reward. We refer to these as *instructed trials*, and this difference in firing rate was typical of our population in this experiment as it was in the Platt and Glimcher experiments (Fig. 8; $p < 0.01$, paired t -test for visual and delay epochs, $N = 20$).

Figure 8 examines the relationship between physiological expected utility, behavior, and firing rate for this neuron. The lower axis plots the play numbers during which 6 sequential blocks were presented. In the first block, instructed trials were presented which reinforced a movement to the shirk target with twice as much juice as a movement to the work target, a relative expected value⁴ of 0.66. The second block presented a relative expected value of 0.33. Blocks 3–6 presented inspection trials in which the dynamic interactions of the two players should have maintained a relative desirability for the two responses of near 0.5. The solid grey lines initially plot the probability of the red target being rewarded during the first two instructed blocks followed by the predicted equilibrium rate of picking the red (shirk) target during the four inspection trial blocks. At a purely behavioral level, the animal seemed to closely approximate the response strategies predicted by our simple prescriptive model. Initially the probability of picking the red target was roughly 50% during the instructed blocks and then shifted dynamically to each of the equilibrium strategies in the subsequent 4 inspection trial blocks. The gray dots plot the neuronal firing rate after target onset. Note that when the relative expected value of a movement to the red target is high, firing rate is high. When the relative expected value of the red target is low, firing rate is low, and when the animal is engaged in a strategic conflict the firing rate associated with this same movement is fairly constant at an intermediate level. This is the specific result that would be expected if LIP neurons encode the desirability, that is they

⁴ [Value of Red Target \div (Value of Red Target + Value of Green Target)].

instantiate the physiological expected utility, of movements into their neuronal response fields.

3.2. *Encoding shirk targets versus work targets*

For a subset of 20 neurons we also examined the effects of reversing the locations of the work and shirk targets during 50% shirking rate Nash equilibrium blocks of the inspection game. This effectively changed both the probability of being rewarded and the magnitude of that reward associated with the target monitored by our neuron, while the relative physiological expected utility of working and shirking should have remained at equilibrium. Firing rates should differ across blocks if they reflect either of these individual decision variables but they should remain constant if they reflect physiological expected utility. The firing rates were indistinguishable, a finding consistent with the hypothesis that LIP firing rates encode the physiological expected utility of choices ($p > 0.05$, paired t -test, $N = 20$, for all 6 epochs).

3.3. *Encoding relative versus absolute desirability*

In order to test the hypothesis that LIP neurons encode the *relative* desirability of movements rather than the absolute expected value (or utility) of movements, we examined 18 neurons while monkeys completed a block of trials in which the magnitude of both working and shirking rewards was doubled. If LIP activity is sensitive to the absolute desirability of the response encoded by the neuron, cells should fire more for blocks in which the rewards are doubled. If, however, LIP activity is sensitive to the relative desirability of choices the firing rate should be the same for both blocks of trials. We found that there was no change ($p > 0.05$, paired t -test, $N = 18$, for all 6 epochs) in the firing rate of LIP neurons when absolute reward magnitude was doubled. This, and the results from our other experiments, further support the possibility that LIP neurons encode the relative physiological expected utility of a movement.

3.4. *Summary*

These data make a relatively simple suggestion. Map-like structures in the brain may encode, quantitatively, the relative desirabilities of all possible courses of action. These maps may form something like a final common path for decision making. If that is true, then it is these maps, some of which have already been identified physiologically, which are actually the subject of economic theory.

If this is true, what leverage does this knowledge give economists? Is neurobiology relevant to economics or does it simply push back the problem of studying behavioral responses to a biological measurement that adds no particular insight? To answer that question one needs to develop a slightly more complete model of the neural decision making process that this hypothesis implies. One can then ask whether the physiological process provides opportunities for economic insight.

4. A utility-theory based neurobiological model of decision making

The concept that emerges from the preceding data is that the brain employs a map of eye movement desirabilities to generate the most utile eye movement via well understood brainstem circuitry. In our model of this process (Fig. 9), area LIP forms a map of all possible eye movements. Each neuron in the map encodes the relative desirability of a particular eye movement. The instantaneous action potential generation rate by any neuron thus encodes the relative physiological expected utility of that movement. The nature of the inputs to these neurons are, at this point in time, uncertain. In the brain, inputs might be in the form of independent maps of expected reward magnitude and reward likelihood that project into area LIP from the frontal cortices, and relative physiological expected utilities might be calculated within area LIP from these inputs. Alternatively, expected utility could be calculated elsewhere and passed to area LIP where a local normalization simply converts these values into *relative* physiological expected utilities. Any number of input schemes are possible at this point, but for the purposes of this initial model the output of LIP is the focus and is presumed to encode the product of an evaluative process that yields relative physiological expected utilities.

The output of the LIP map is passed to the frontal eye fields. We know from the work of Schall and his colleagues (see introduction for details) that once a single region in the frontal eye field map is driven over a threshold level of activity, the map of the superior colliculus generates an eye movement having the appropriate length and direction. In the model, the physiological expected utility map in LIP thus drives the serially arranged maps of the frontal eye fields and colliculus towards a winner-take-all state. For biophysical reasons, only one of the regions in each of these aligned and interconnected topographic maps can be active above threshold, which effectively constrains the decision making system to produce one movement at a time. (An important constraint since the eyes can only move

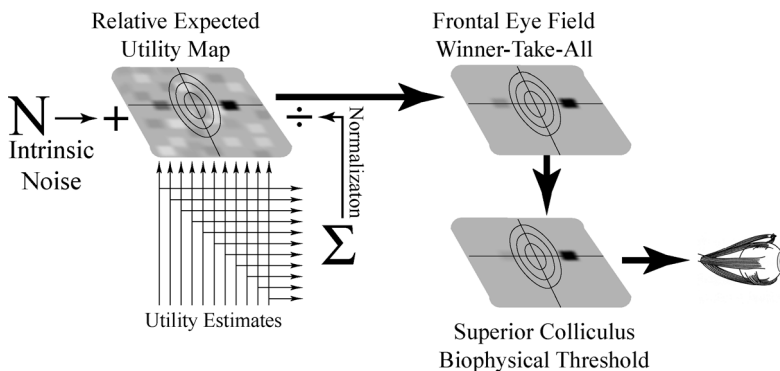


Fig. 9. Model of the Decision-Making Process. Area LIP is presumed to form a map of relative desirabilities of all possible saccadic eye movements, a map of relative physiological expected utility. The LIP map received as inputs estimates of the desirabilities of all possible saccades. This input is contaminated by intrinsic noise which can impose stochastic patterns on behavior under some conditions. The input utility estimates are summed and used to normalize the map so that it represents relative physiological expected utility. The output of the map passes serially to the frontal eye fields and superior colliculus which filter the output imposing a winner-take-all outcome. This winner-take-all outcome is finalized in the colliculus by a biophysically imposed threshold.

in one direction at a time.) The interaction between the utility representation in area LIP and the movement execution circuitry in the frontal eye fields and superior colliculus thus forces a convergence to a single output, the single output associated with the highest expected utility.

Such a system could account for much of the decision making we have observed in humans and animals, and it is consistent with the available neurobiological data, but it is unclear how such a system could account for the kinds of stochastic behavior that are often observed during strategic games. In the inspection game, for example, we studied mixed strategy equilibria and found that both humans and monkeys were highly stochastic under those conditions. A Markov chain analysis that searched for patterns of working and shirking in our data suggested that our subjects were being highly stochastic. In a closely related study, Barraclough et al. (2002) made a similar observation. In their experiment, monkeys played matching pennies against an intelligent opponent thousands of times and their data indicated that the monkeys adopted almost perfectly stochastic strategies. It therefore seems reasonable to ask whether our basic model can be extended to account for behavioral stochasticity during mixed strategy equilibria.

One known source of stochasticity at the neuronal level is the mechanism by which synaptic inputs give rise to action potentials in cortical neurons. Abundant evidence indicates that when cortical neurons are repeatedly activated by precisely the same stimulus, the neurons do not deterministically generate action potentials in precisely the same pattern. Instead, the pattern of stimulation delivered to cortical neurons appears to determine only the average firing rates of those neurons, the instant-by-instant dynamics of action potential generation are highly variable and appear to defy precise prediction (Tolhurst et al., 1981; Dean, 1981). The available data suggests that this moment-by-moment variation, the overall variance in cortical firing rate, is related to mean firing rate by a roughly fixed constant of proportionality that has a value near 1.07 over a very broad range of mean rates (Tolhurst et al., 1981; Dean, 1981; Zohary et al., 1994; Lee et al., 1998), and this seems to be true of essentially all cortical areas that have been examined including parietal cortex (Lee et al., 1998). This has led to the suggestion that action potential production can be described as something like a Poisson process, a probabilistic operation at the root of neuronal computation.

More recently, there have been several efforts to identify the biophysical source of this Poisson-like stochasticity. Mainen and Sejnowski (1995) sought to determine whether the process of action potential generation, the biophysical mechanism that converts the analogue input voltages that cells receive into action potentials, was the source of this variability. Their work led to the conclusion that action potential generation is quite deterministically tied to membrane voltage, and thus that this process was not a source of intrinsic action potential variability. Subsequent studies have begun to suggest that it may instead be the process of synaptic transmission which imposes a stochastic pattern on cortical action potential production (for review, see Stevens, 1994). The actual dynamic pattern of membrane voltage which controls action potential generation deterministically is induced by synaptic transmission, a process that now appears to be irreducibly stochastic. All of these data suggest that the precise pattern of activity in cortical neurons is stochastic. The dynamics of exactly when an action potential is generated seems to depend on truly random physical processes.

Shadlen et al.'s (1996) study demonstrated that relating neuronal firing rates to behavior required a knowledge of two critical parameters; the intrinsic variance in instantaneous firing rate evidenced by each cortical neuron (the Poisson-like variability of the action potential generation process) and the correlation in action potential patterns (imposed by the cortical microcircuitry) between the many neurons that participate in any neural computation (the inter-neuronal cross-correlation). Shadlen and colleagues demonstrated that both of these properties contribute to the unpredictability evidenced by behavior. The variability in the firing rate of each neuron contributes to the unpredictability of behavior by producing an initial stochasticity in the neuronal architecture and the degree to which that stochasticity influences behavior depends on how tightly correlated are the firing patterns of the many neurons in a population.

To make this insight clear consider a population of 1000 neurons, all of which fire action potentials with the same mean rate and which have the same level of intrinsic variability but are generating action potentials independently of each other. The members of such a population would be generating moment-by-moment patterns of action potentials that were completely uncorrelated; the only thing that they would share is a common underlying mean firing rate. Because of this independence, globally averaging the activity of all of these independent neurons would allow one to recover the underlying mean rate at any instant. Thus if, for example, 1000 neurons in the real LIP map represented each movement (which is not an unreasonable number) and the firing rates of those neurons were sufficiently uncorrelated, decisions produced by the outputs of the map could deterministically reflect physiological expected utility encoded in the average firing rate. Consider, as an alternative, a circuit in which a population of 1000 LIP neurons encoding a single movement all fire with the same mean rate, and have the same level of intrinsic variability, but in which each of the 1000 source neurons were tightly correlated in their activity patterns. Under these conditions, it is the stochastic and synchronous pattern of activity shared by all of the neurons in the population encoding a particular movement that is available to areas like the frontal eye fields, rather than the underlying mean rate. In a highly correlated system of this type, the output at any moment is irreducibly stochastic. Of course these are just two extreme examples. Many levels of correlation between neurons are possible and each would provide areas like the frontal eye fields with a slightly different level of access to the underlying mean rate, and a different level of intrinsic randomness.

As a result of these speculations and data that appear to support them (Zohary et al., 1994; Britten et al., 1996; Parker et al., 2002), if the LIP map does encode the relative physiological expected utility of all available eye movements, the instant-by-instant representation of those utilities would be necessarily stochastic. Consider the activity of area LIP during the inspection game. Under conditions in which the relative physiological expected utility of working is much higher than shirking, the map, although stochastic, provides an output that always leads to working. As working and shirking approach the same relative physiological expected utility, however, the behavior becomes more and more stochastic until, at equilibrium, the behavior is completely dominated by the neuronal stochasticity. Under those conditions it is the neuronal variance which generates the decision. A model of this type thus generates a smooth transition from determinate to stochastic responding as a function of relative expected utility of exactly the type called for by game theory.

These data and this model seem to suggest a stochastic decision making process that can account for a broad range of choice behavior. We began this neurobiological investigation, however, hoping to account for simple, crude, reflexive kinds of decision making. The kinds of processes that we might expect to occur when, as Vernon Smith put it, “the brain takes over.” Instead what we seem to have found is a system that can at least partially account for rational decision under a broad range of conditions. If that is so, then two critical questions remain. First, what brain areas are responsible for computing the physiological expected utilities that appear to be represented in area LIP? Second, what happens when subjects behave irrationally?

Knowing what we do about the neural basis of decision making, and presuming that a model of this general type explains that process, then we can begin to answer these questions empirically: The output of the evaluative systems of the brain are represented in areas like LIP. These representations are used by the stochastic decision making architecture to produce behavior. It is these representations, and their generation, which form the subject of most of economic theory. With this knowledge in hand it seems that we really can begin to perform experiments that might yield mechanistic explanations of economic behavior that tie classes of decisions to the underlying hardware.

5. Extending the model

Over the past ten years significant progress has been made towards understanding the neurobiology of primate learning mechanisms. Of particular value have been studies of two groups of neurons in the primate brain; neurons in the *ventral tegmental area* and neurons in the *substantia nigra pars compacta*. Both of these clusters of nerve cells are chemically distinct from the cells around them; they all employ the chemical *dopamine* as a neurotransmitter for communicating with their targets, the neurons whose firing rates they influence (Fig. 10). Both of these groups of neurons are relatively small, consisting of only a few thousand neurons. Both groups send their output fibers, or axons, long distances and these axons, in a manner unusual in the primate brain, terminate throughout the cerebral cortex and other structures including the basal ganglia. (For a review of this literature, see Schultz, 2002.)

For decades it has been known that these neurons and the dopamine they release play a critical role in brain mechanisms of reinforcement. Many of the drugs currently abused in our society mimic the actions of dopamine in the brain. This led many researchers to believe that dopamine neurons directly encoded the rewarding value of events in the outside world. Wolfram Schultz and his colleagues, however, made a critical observation that widely altered how these neurons were viewed. When an animal sits passively in a quiet environment these neurons produce action potentials at a fixed rate of about 3 per second (3 Hz). This is the resting state of these neurons. Schultz and his colleagues measured the activity of these neurons while quite thirsty monkeys sat passively and listened to a tone which was followed by a tiny squirt of fruit juice into their mouths. When this paired presentation was repeated dozens of times, they found that the dopamine neurons of both areas continued to produce action potentials at their resting rate following the delivery of juice, a curious result because we can assume that fruit juice is reinforcing to thirsty monkeys.

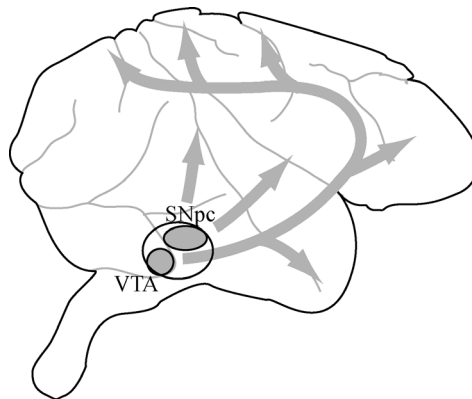


Fig. 10. Dopamine neurons of the substantia nigra pars compacta (SNpc) and Ventral Tegmental Area (VTA). These brainstem neurons project throughout the frontal, temporal and parietal cortices as shown here on a monkey brain. When an action potential is generated in the SNpc or VTA it propagates outward to the synapses at the end of these nerve cells. There, the electrochemical impulses cause the neurochemical dopamine to be released into a synapse. Dopamine then alters the electrochemical activity of the neuron on the far side of that synapse.

Next, without warning, Schultz doubled the magnitude of reward delivered to the monkeys for a series of trials. In response to the first of these unexpectedly large reinforcements the dopamine neurons responded with a dramatic increase in the rate at which they were generating action potentials. The firing rate of these neurons immediately after the juice was delivered increased from 3 per second to about 80 per second for a duration of about a tenth of a second. On the second trial of this group, the firing rate of the neurons was still elevated, but not as strongly. And as the new magnitude of reward was repeated, the magnitude of the sudden increase in firing rate was reduced until, after 10–30 trials, it had returned entirely to the 3 Hz resting state, or baseline. Next, without warning Schultz presented a series of trials in which the magnitude of juice reward was reduced to the original level. On the first of these trials the dopamine neurons responded with a transient decrease in firing rate which, after many repetitions, also eventually returned to baseline.

Based on these observations Schultz argued that the dopamine neurons seemed to encode the difference between the reward that an animal expected to receive and the reward that an animal actually received (Schultz et al., 1997). In the syntax of the reinforcement learning literature, the neurons appeared to encode a *reward prediction error* (Sutton and Barto, 1998). At a neurocomputational level what Schultz suggested was that the dopamine neurons probably received an input that encoded, as a firing rate, the magnitude of the reward that the animal expected to receive after the tone. The dopamine neurons were also presumed to receive an input that encoded the magnitude of reward that the animal actually received. The surfaces of these neurons were then presumed to employ well understood biophysical mechanisms to compute the mathematical difference between these two inputs, the reward prediction error, and to transmit this value to their target neurons in the cortex and basal ganglia via the neurotransmitter dopamine.

In an elegant series of subsequent experiments Schultz has essentially validated this initial proposal and has even demonstrated that several interesting properties of human and

animal learning can be predicted by the activity patterns of these neurons. More formally, these available data point strongly towards the hypothesis that the dopamine neurons carry a reward prediction error, at least in the conditioning tasks that Schultz and his colleagues have explored.

5.1. *Mechanisms for computing physiological expected utilities*

Given what we already know about economic decision making, the implications of the dopamine data for economics may be significant. If these neurons do encode the reward prediction error of learning theory, then the exact computational properties of these neurons may explain how some physiological expected utilities are computed, or learned. The mechanism of satisficing, under some conditions, may be the mathematical computations that these neurons evolved to perform.

In order to test that hypothesis, we (HMB and PWG) recently began to characterize the computations that the dopamine neurons perform while animals engage in a simple rewarded task. In that experiment, thirsty monkeys stared straight ahead at a yellow spot of light and waited for a second yellow spot of light to appear directly above it. Once that spot of light appeared it remained illuminated for 4 seconds. The monkey was free, at any time during that 4 seconds period, to look towards the second light, and when the monkey chose to look determined how much juice he earned.

The 4 seconds interval was divided into 16 logarithmically scaled subintervals, the first of which was 0.01 seconds long and the last of which was 0.8 seconds long. Before the first play of each day we would randomly select 4 sequential intervals from these 16 and assign rewards to them without indicating this assignment in any way to the monkey. If the monkey made his movement during the first of the reinforced intervals he received 0.22 milliliters of juice, if during the second 0.24 ml of juice, if during the third 0.26 ml of juice and if during the fourth 0.28 ml. If he made his movement at any other time he received no reward. Finally, irrespective of when he looked and what juice he earned, the entire 4 seconds interval elapsed before a new play began. These reward contingencies persisted for a minimum of 90 trials after which there was a fixed 5% chance, on each subsequent trial, that the reward contingencies would be secretly changed. When the reward contingencies did change, a new group of four sequential intervals were randomly selected for reinforcement.

Under these conditions, the task of an efficient monkey is simply to learn when to move in order to maximize juice intake. The monkey begins (Fig. 11) the session with a fairly random set of movements and then gradually learns when to move. At an unpredictable time the rewarded intervals are switched and the monkey responds by learning the new time which yields a maximal reward.

The dopamine data gathered by Schultz suggests that the monkeys may accomplish this goal, at least in part, with the dopamine neurons of the ventral tegmental area and the substantia nigra pars compacta. This raises the interesting possibility that a behavioral experiment like this one could be used to extract the algorithm that the dopamine neurons use to compute the reward prediction error.

We therefore began by recording the activity of dopamine neurons during this experimental task. On each play we quantified the firing rate of the dopamine neuron we were

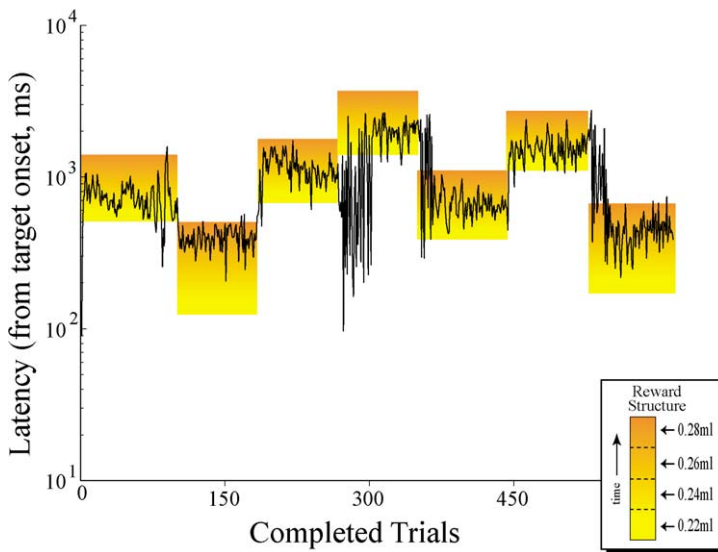


Fig. 11. Behavior of a monkey during the dopamine experiment. The black line plots the play-by-play delay selected by the subject. The colored bars indicate the unsignaled intervals during which a juice reward could be earned.

studying immediately after the reward was (or would have been) delivered. On a typical day a monkey would be presented with a series of about 600 plays and would encounter about 5 reward contingencies. At the end of the day we would then have a complete record of the rewards that the subject had received and the firing rates of the dopamine neuron on each of those plays. We began by making one simplifying assumption, that the algorithm the monkey employed in this task was linear. Though perhaps unjustified, this was an essential first step because it allowed us to use linear regression to extract the computation that the neurons were performing. We therefore computed a linear regression on the firing rates, asking what combination of the preceding 20 rewards best accounted for the firing rate of the neuron on each trial. What we found was remarkably simple and remarkably like the hypotheses of reinforcement learning theorists. Given our assumptions, the neurons could best be described as subtracting from the reward magnitude received on the current play an exponentially weighted average of the rewards received on the preceding 8–10 plays. Expressed as an iterative computation:

$$\begin{aligned} \text{Firing Rate} &= \text{Reward Prediction Error}_T \\ &= \alpha(\text{Current Reward} - \text{Reward Prediction Error}_{T-1}). \end{aligned} \quad (8)$$

From this value, the reward prediction at time $T - 1$ could then be used to derive a new reward prediction. In sum, what we found was a descriptive algorithm for learning reward contingencies that was measured mechanistically. We derived an equation for the dopamine neurons that we hoped could be used to predict economic behavior under other circumstances.

5.2. *Using neural data to build economic models*

The real test of whether these neurobiological data are of any use to economists is whether they can be used to develop or refine economic models. As a result, with Brain Lau we undertook to model a simple and well described behavior that both humans and animals show using the neurally derived algorithms that the preceding experiments had defined. The behavior that we sought to model was a simple choice between two alternatives, *a* or *b*, which had different values that could change unpredictably.

We modeled the chooser as a two stage process based on the neural data we had already acquired (Fig. 12). The first stage of our model learned a physiological value for each of the two choices using an exponentially weighted average of exactly the type computed iteratively by the dopamine neurons (Eq. (8)) and the second stage employed a decision process modeled after our understanding of area LIP and its targets.

The first stage thus consisted of two estimators, *A* and *B*. Whenever the animal selected option *A*, the reward earned by the monkey was used to iteratively compute a weighted average of the value of that action. Formally, we used the equation derived from the dopamine neurons to do this, leaving the exponential weighting parameter free. An identical module computed the physiological expected utility of the *B* response and was required to use the same exponential weighting parameter employed for *A*. The decision process then compared these two physiological expected utilities and stochastically selected the response having the higher physiological expected utility. Formally,




Fig. 12. Neuroeconomic model of probability matching. The model computes an exponentially weighted average of the gains obtained on prior executions of each possible response. This weighted average is presumed to be computed iteratively based on the reward prediction error signal carried by dopamine neurons. This weighted average yields a physiological expected utility (PEU) for all movements which are then fed to the physiological expected utility map in Area LIP. In area LIP physiological expected utilities are normalized to relative physiological utilities. Finally, the intrinsic variance of cortical neurons in LIP ensures that the relative PEUs are encoded with weak stochastic variance. (See Fig. 9 for additional details.) Subsequent levels of processing force a winner-take-all extraction of the most desirable movement from the LIP map.

the decision rule used a sigmoidal function to convert the difference in value of the two movements to a probability of choosing each movement. The slope of the sigmoid, which we refer to as the *stochastic transfer function*, was left as the second free parameter in this model and can be loosely thought of as encapsulating the level of stochasticity in the LIP decision circuit. We left this as a free parameter because there is growing evidence that the stochasticity of a neuronal population like LIP may be variable and may be set dynamically to maximize task performance (Barraclough et al., 2002; Dorris and Glimcher, 2004).

In order to generate a behavioral dataset for comparison with the physiologically derived model, we had monkeys perform a simple 300–700 play two-alternative lottery based on the matching law experiments developed by Herrnstein (1961, 1997). On each play two visual targets appeared, one to the left and one to the right of fixation. Animals had about 1 second to select one of these two targets. Before each play there was a fixed probability that each of the two targets would be *armed* with a reward. For example, there might be a 0.1 probability that the right target would be armed before each play and a 0.2 probability that the left target would be armed before each play. Targets were always armed with 0.25 ml of juice and data were gathered during a period when we had reason to believe that the monkeys' thirst was at a fairly constant level. Critically, as in the original Herrnstein experiments, the arming was cumulative. Once a target was armed it remained armed until it was chosen in a subsequent play. Staddon and Motheral (1978) and Staddon (1980), amongst others, have shown that under these conditions matching the probability of each response to the probability of each reward is a nearly optimal solution (at a molar level) and a number of researchers have shown that a wide variety of primate species, including humans, show probability matching behavior under these conditions (Schrier, 1965; De Villiers and Herrnstein, 1976; De Villiers, 1977; Bradshaw and Szabadi, 1988).

As in the earlier neurophysiological experiment, we also instituted unpredictable block shifts during which the arming probabilities changed. Blocks were a minimum of 90 plays in length after which there was a 0.05 probability that the targets would be changed to a new randomly selected probability of arming. On a typical day we were able to examine 3–7 blocks of 90–125 plays.

Figure 13 shows the average behavior of the monkey subjects who demonstrate approximate probability matching. More precisely, the probability that they will pick the left target is a linear function of the relative likelihood that the left target will be armed. The slope of the line is, in fact, expected to be slightly less than 1 to allow for over sampling of low probability reinforcements in order to maximize detection of block shifts (cf. Sutton and Barto, 1998) and this is exactly what we observed.

What we hoped to determine was not, however, anything about the global average behavior of our subjects during probability matching. What we hoped to model was the actual dynamics of the choice behavior by using our neurobiologically derived model. To do that, we ran the model described above against exactly the same behavioral contingencies that a real monkey had encountered on a particular day and asked how closely the behavior of the model predicted the behavior of the monkeys. In essence, we asked the model to make a one-step-ahead prediction of what the monkey would do under these conditions. Figure 14 shows a typical example of that prediction. The exponential averaging rate and the steepness of the stochastic transfer function were determined with a maximum likeli-

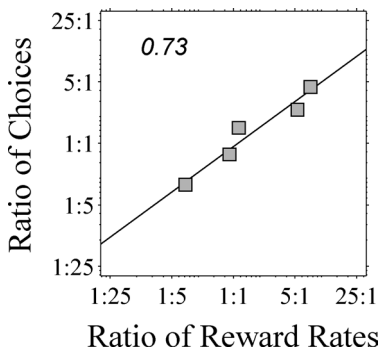


Fig. 13. Monkey behavior on the two-alternative lottery. Subjects must select, on each play, between two alternatives with different, although cumulative, independent probabilities of yielding 0.25 ml of juice. The subjects probability of making either response is a linear function of the ratio of reward rates. This has been shown to be an efficient strategy under these conditions. The experiment replicates earlier matching law studies of humans and monkeys.




Fig. 14. Dynamic one-step-ahead prediction of the model shown in Fig. 12. Black line plots a 20 trial moving average of a subject’s choice behavior during three consecutive unsignaled blocks of plays. Red line plots the dynamic behavior of the model.

hood nonlinear iterative fit. The vertical dashed lines show the 3 sequential blocks of plays examined in this particular experiment. Notice that the model does a remarkably good job of capturing the dynamics of the real monkey.

The model we developed here would, of course, behave irrationally under some conditions. The computational process it employs does not directly compute expected utilities but approximates them with a weighted average. As a result, if one were to observe that the monkeys rationally probability matched under some circumstances and irrationally persisted in probability matching under other conditions where this was inefficient, there might be no need to postulate different mechanisms for these two classes of behavior. Instead, one might be able to propose a single identified mechanism for both classes of behavior. This, at heart, is the central goal of the neuroeconomic approach.

5.3. Summary

Prescriptively, one can certainly define a rational strategy for a chooser who encounters these environmental conditions. Indeed, a number of psychologists and economists have presented analyses that specify rational behavioral responses to these conditions (Fretwell, 1972; Staddon, 1980). At a descriptive level, one can also gather data demonstrating both how behavior under some conditions is well predicted by prescriptive models and data demonstrating ways in which observed behavior differs from the rational choice model. The advantage of the prescriptive model is its parsimony and efficiency, the advantage of the descriptive model is its predictive power. The model presented in the preceding sections is in some ways a hybrid of these approaches. It begins by identifying the mechanisms for rational choice and then builds from there to empirically identify the algorithms actually used to achieve choice behavior. While the model itself may be fairly unremarkable, the intermediate representations that it employs are not artificial constructs that may lack parsimony, but instead the intermediate constructs aspire to represent real neural systems.

It is also, of course, important to recognize that this is one of the simplest possible models of this type. Neuroeconomics is in its infancy and the sophistication of the models we employ reflects this fact. For example, we know that this model does a poor job of predicting the behavior of either our monkeys or our human subjects during the inspection game. Clearly, this algorithm does not describe how physiological expected utilities are computed during strategic games. We do not yet know how the dopamine neurons respond during strategic games, an observation that suggests that one way to begin to understand the dynamics of strategic decision-making would be to study the dopamine neurons under those conditions. Indeed a number of important experiments like this one suggest themselves, all experiments that can be done and which will constrain mechanistic explanations of choice behavior.

6. Generalizing neuroeconomic results to complex behaviors

Perhaps one of the most critical questions that these experiments raise is whether these findings about decisions expressed by eye movements will generalize to more complex behavioral responses. Will future neurobiological studies reveal similar algorithms for the control of hand movements? How would hand movement and eye movement systems interact and what can this tell us about how verbally expressed decisions are generated? And even more important, how would tasks that require several sequential and inter-related decisions be represented in this architecture? These are absolutely critical questions, about which we only have the first hints of understanding, but they are so critical to the long-term significance of these results that we turn next to engage these issues.

6.1. Generalizing to humans

Over the last four years several research groups have begun to employ brain scanning technologies to search for area LIP in the human brain. Fortunately, an area has been

identified in the intraparietal sulcus of humans that is active before saccadic eye movements. Work by Sereno et al. (2001) suggests that this area is the human homologue of area LIP. Perhaps even more encouraging is that it appears to be located amongst a cluster of human movement-related areas that seem very similar to the areas in the monkey intraparietal sulcus. A number of groups are now beginning to test the hypothesis that these areas participate in human decision making by replicating some of the monkey experiments we have described here using functional magnetic resonance imaging, a brain scanning technique, in humans. Their results are very promising and very similar to the available monkey literature (Sereno et al., 2001; Paulus et al., 2001; Connolly et al., 2002). While it is much too soon to be certain, it does seem very likely that the results we have described in monkeys will generalize to humans when it comes to the control of saccadic eye movements. What about the control of arm movements? This seems particularly important because most economic experiments ask human decision makers to respond by pressing buttons.

6.2. Generalizing to arm movements

We now know that in both monkeys and humans there are one or more brain areas in parietal cortex adjacent to area LIP (Colby and Goldberg, 1999; Sereno et al., 2001; Connolly et al., 2002) that seem to play a critical role in the control of several types of arm and hand movements. The best studied of these areas is probably the parietal reach region (Snyder et al., 2000), which has been described in both monkeys (Snyder et al., 1997) and humans (Connolly et al., 2002) and which lies very near to area LIP. In the monkey, the parietal reach region is active before monkeys make movements with their hands and the patterns of activity in this area seems to very closely parallel the patterns of activity already described in area LIP. Each neuron in the parietal reach region seems to become maximally active before a reach to a particular location in the extrapersonal space immediately in front of a subject. Snyder and colleagues, for example, have shown that when a monkey sits in front of a set of illuminated buttons and must decide which of those buttons to press in order to receive a reward, neurons in the parietal reach region seem to encode the movement that the monkey is about to make just as neurons in LIP encode upcoming saccadic eye movements. Together these observations suggest that decision making, whether expressed by eye movement or hand movement, and whether in monkeys or humans, may well employ the same basic neural architecture. That is encouraging, but to many economists it still seems unlikely that these kinds of systems which make decisions about such simple responses can be related to decisions about pension funds or market investments. To begin to determine whether the kinds of insights we have already gained about decision making can be generalized to explain much more complicated behavior one needs next to look at the neural architecture by which complicated sequences of responses are generated over much longer periods of time.

6.3. Generalizing to behavioral sequences

In the case of movements produced with the skeletal muscular system, movements for walking, bidding at auction, or even to some extent for speaking, three areas play a critical

role in movement generation and these areas seem roughly analogous to the frontal eye field and superior colliculus: the primary motor cortex, the premotor cortex and the supplementary motor area. The primary motor cortex is the final common pathway for output to the spinal circuits that control the arm and hand musculature directly. This area, in turn, receives projections from the premotor cortex and the supplementary motor area. Both the premotor cortex and the supplementary motor areas receive projections from the parietal areas we have examined here, but appear to perform different functions. The premotor area seems to initiate simple single response movements of the type we have been discussing and it seems likely that when the parietal reach region triggers a movement it does so via the premotor cortex. In contrast, when complex sequences of movements are produced over an extended period the supplementary motor cortex seems to play a privileged role. Under these conditions neurons in the premotor cortex are surprisingly silent and highly specialized neurons in the supplementary motor area seem to generate the complex behaviors that are composed of many independent movements (Tanji, 2001). While it is not yet clear what form economic decision making will take in these areas, it seems plausible to assume that the same organizational principles will obtain; the parietal cortex may play a very similar role in the control of these movements. While we have no idea how humans make decisions about when to invest in their 401 k plans, the act of making that investment involves an activation of the supplementary motor area, amongst other things.

6.4. Coordinating multiple brain areas in human decision making

Perhaps the most critical point that can be made when considering how these many brain areas interact is to note that these multiple areas appear always to be both highly interconnected at an anatomical level and well coordinated at a physiological level. To make the level of that coordination clear consider a simple eye movement task in which a monkey subject will decide to look at the single red target presented amongst 7 green targets. The presentation of these 8 visual stimuli leads to the transient activation of 8 groups of neurons in the parietal cortex, in the frontal eye fields and in the superior colliculus. Over the course of a few tenths of a second all three of these areas quickly suppress activity associated with the 7 green targets while activity associated with the red target is enhanced. These three distinct but heavily interconnected areas converge together towards a single solution. There may be strong interactions between these areas, and there may even be circumstances in which the areas initially provide conflicting signals (Curtis and D'Esposito, 2003) but the areas largely converge together in order to ultimately activate the motor system at the level of the brainstem and spinal cord. There is no doubt that one of the principle challenges facing neuroeconomists will be to learn how this complex coordination is achieved and how these areas differ. Unfortunately, we know almost nothing about how to answer these questions today so any schema would be purely speculative.

7. General summary

One of the most critical points that we hope to have made to an economic audience is that there is no evidence that hidden inside the brain are two fully independent systems, one

rational and one irrational. There is, for example, no evidence that there is an emotional system, per se, and a rational system, per se, for decision making at the neurobiological level. There is no evidence that some sort of primitive system exists which can wrest control of the motor system from a more recently evolved (or uniquely human) system. For neurobiologists this is a central point. This is such an important point, and one in which so much acrimony has arisen between neurobiologists and economists that it requires further elaboration.

Recently, a number of economists have begun to suggest, at a psychological level, that human decision making can be broken down into two categories; typically rational and irrational. At a formal economic level there really can be no challenge to this argument. At least since Herbert Simon argued for satisficing this has been widely accepted in economic circles, and since rational and irrational decision making are defined at the economic level this is an irrefutable point. These economists go on to argue that there is subjective personal experience that rational decision making can, at least under some conditions, be viewed as the product of conscious introspection. Because of a widespread belief amongst the lay public that conscious experience resides uniquely in the cerebral cortex (a belief that may have its roots in the relatively recent evolutionary origin of that organ and a lay conviction that conscious experience is unique to humans) these economists speculate that rational decision making must therefore be the product of the cerebral cortex. These economists go on to further speculate that irrational decision making must therefore be the product of some other brain system. Typically, these economists argue that this other brain system must be evolutionarily more ancient because, it is presumed, animals with less complicated nervous systems cannot possibly engage in the kinds of efficient rational decision making studied by economists. Based on recent neurobiological work which suggests that brain systems associated with emotion have a relatively ancient evolutionary origin, these economists suggest that irrational behavior is the product of these ancient emotional systems within the brain.

What we cannot stress strongly enough is that the vast majority of evolutionary biologists and neurobiologists reject this view. There are probably two principle reasons that biologists reject this dualist view of the nervous system; one neurobiological and one behavioral. First there is no neurobiological evidence that emotional and non-emotional systems are fully distinct in the architecture of the primate brain. Second, there is no evidence that rational and irrational behavior are the product of two distinct brain systems, one of which is uniquely rational and one of which is uniquely irrational.

7.1. Conclusion

One of the critical persistent issues in economics has been an inability to reconcile the rational choice model at the core of modern economics with the fact that humans are the product of a 600 million year evolutionary lineage. We all recognize that non-human animals have limited mechanical and neural capacity. Fish that live in total darkness have neither eyes nor the neural architecture for vision. We all accept that even our closest living relatives, the great apes, face fundamental conceptual limitations that are probably not apparent to them. But it has long been the central premise of economic thought that humans are different from all of these other organisms. That humans rely on a more funda-

mentally rational neural machinery and that this machinery, which economists presume is subjectively experienced as consciousness and which they often assume is mechanistically located within the cerebral cortex, endows us with nearly perfect rationality.

In the last half century a number of influential economists have begun to challenge this assumption. These economists have argued that we must begin to recognize that our evolutionary heritage influences the actions that we take. Many of the decisions that we make, they argue, may be inefficient because of that evolutionary history. Surprisingly, however, many of these same economists argue that an efficient model of human behavior will have to be two-tiered. There is, these economists accept from classical economic theory, a fundamentally rational conscious decision maker within our skulls. This is, they presume, an evolutionary development unique to our species which has arisen within the very recent past. But there is also a second more ancient and mechanistic system, and when inefficient decision making occurs it can be attributed to the activity of this evolutionarily ancient mechanism.

For many neurobiologists studying the mechanisms by which choice is accomplished, this seems to be an oddly dualist approach to the physiology of mind. In the seventeenth century Descartes proposed that all of human behavior could be divided into two principle classes and that each of these categories of behavior could be viewed as the product of distinct processes. The first of those classes Descartes defined as those simple predictable behaviors which both humans and animals could express. Behaviors which predictably linked sensory stimuli with motor responses. Their simple deterministic nature suggested to him that for these behaviors the sensory to motor connection lay within the material body, making those simple connections amenable to physiological study. For the second class, behaviors in which no deterministic connection between sensation and action was obvious, he followed Aristotle's lead, identifying the source of these actions as the rational, but nonmaterial, soul.

Over the last several decades neurobiologists have begun to broadly reject this dualistic formulation for several reasons. First because there seems to be no physiological evidence that such a view can be supported and second because it seems to fly in the face of evolutionary theory which forms the basis of modern biology. Instead, what seems to be emerging is a much more synthetic view in which economic theory can serve as the core for a monist approach to understanding the behavior not just of simple organisms that survive in narrowly defined environments but also for understanding the most complex and generalist of extant species, *homo sapiens*.

In sum, neuroeconomics seeks to unify the prescriptive and descriptive approaches by relating evolutionary efficiencies to underlying mechanisms. Neoclassical economics and the utility theory on which it is based provide the ultimate set of tools for describing these efficient solutions; and evolutionary theory defines the field within which mechanism is optimized by neoclassical constraints; and neurobiology provides the tools for elucidating those mechanisms.

Over the past decade a number of researchers in both neuroscience and economics have begun to apply this approach to the study of decision making by humans and animals. What seems to be emerging from these early studies is a remarkably economic view of the primate brain. The final stages of decision making seem to reflect something very much like a utility calculation. The desirability, or physiological expected utility, of all available

courses of action seem to be represented in parallel. Topographic maps of the physiological expected utilities of movements or actions seem to be the substrate upon which decisions are actually made.

These representations, in turn, seem to be the product of many highly coordinated brain circuits. Some of these brain circuits, like the dopamine neurons of the ventral tegmental area and the substantia nigra pars compacta, are already beginning to be described. The algorithms by which these circuits compute the economic variables from which physiological expected utilities are derived are now under intensive study. Indeed, several of these mechanistic studies are even now being used to make economic predictions about the behavior of human and non-human primates both when that behavior follows, and when it deviates from, the prescriptive neoclassical model. Studies like these seem to be elucidating the mechanisms by which satisficing is accomplished and a critical advantage of this approach to irrational behaviors is that once mechanism is understood, satisficing should become broadly predictable. In essence, neuroeconomics argues that it is mechanism which can serve as the bridge between the prescriptive and descriptive approaches that dominate economics.

As early as 1898 the economist Thorstein Veblen made this point in an essay entitled “Why is economics not an evolutionary science?” He suggested that in order to understand the economic behavior of humans one would have to understand the mechanisms by which those behaviors were produced (Veblen, 1898). More recently, the biologist Wilson (1998) and the economists Zak and Denzau (2001) have made a similar point. Arguing that a fusion of the social and natural sciences is both inevitable and desirable, Wilson has suggested that this fusion will begin with a widespread recognition that economics and biology are two disciplines addressing a single subject matter. Ultimately, economics is a biological science. It is the study of how humans choose. That choice is inescapably a biological process. Truly understanding how and why humans make the choices that they do will undoubtedly require a neuroeconomic science.

Acknowledgments

The authors acknowledge the support of the National Eye Institute of the US National Institutes of Health, The McKnight Foundation, The James S. McDonnell Foundation, The Klingenstein Foundation, and the Human Frontier Science Program. The authors express their gratitude to Paul Zak for helpful discussion and to Andrew Schotter and Maggie Grantner.

References

- Allais, M., 1953. Le comportement de l’homme rationnel devant le risque, critique des postulats et axiomes de l’école américaine. *Econometrica* 21, 53–526.
- Arnauld, A., Nicole, P., 1996. *Logic or the Art of Thinking*. Cambridge Univ. Press, Cambridge. Reprinted edition edited by Biroker, J.V., original 1662.
- Barracough, D.J., Conroy, M.L., et al., 2002. Stochastic decision-making in a two-player competitive game. *Society for Neuroscience Abstracts* 285.16.

- Bernoulli, D., 1954. Exposition of a new theory on the measurement of risk. *Econometrica* 22, 23–36. Reprinted from 1738.
- Bradshaw, C.M., Szabadi, E., 1988. Quantitative analysis of human operant behavior. In: Davey, G., Cullen, C. (Eds.), *Human Operant Conditioning and Behavior Modification*. Wiley, New York.
- Britten, K.H., Newsome, W.T., Shadlen, M.N., Celebrini, S., Movshon, J.A., 1996. A relationship between behavioral choice and the visual responses of neurons in macaque area MT. *Visual Neurosci.* 13, 87–100.
- Camerer, C., Loewenstein, G., Prelec, D., 2003. *Neuroeconomics: How neuroscience can inform economics*. In press.
- Caraco, T., Martindale, S., Whittam, T.S., 1980. An empirical demonstration of risk-sensitive foraging preferences. *Animal Behav.* 28, 820–830.
- Colby, C.L., Goldberg, M.E., 1999. Space and attention in parietal cortex. *Annual Rev. Neurosci.* 22, 319–349.
- Connolly, J.D., Goodale, M.A., Menon, R.S., Munoz, D.P., 2002. Human fMRI evidence for the neural correlates of preparatory set. *Nature Neurosci.* 5, 1345–1352.
- Curtis, C.E., D'Esposito, M., 2003. Success and failure suppressing reflexive behavior. *J. Cognitive Neurosci.* 15, 409–418.
- Damasio, A.R., 1995. *Descartes Error: Emotion Reason and the Human Brain*. Pan Macmillan, London.
- Dean, A.F., 1981. The variability of discharge of simple cells in the cat striate cortex. *Exper. Brain Res.* 44, 437–440.
- De Villiers, P.A., 1977. Choice in congruent schedules and a quantitative formulation of the law of effect. In: Honig, W.K., Staddon, J.E.B. (Eds.), *Handbook of Operant Behavior*. Prentice Hall, Englewood Cliffs, NJ.
- De Villiers, P.A., Herrnstein, R.J., 1976. Toward a law of response strength. *Psychological Bull.* 83, 1131–1153.
- Dorris, M.C., Glimcher, P.W., 2003. Monkeys as an animal model of human decision making during strategic interactions. Submitted for publication.
- Dorris, M.C., Glimcher, P.W., 2004. Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. *Neuron*. In press.
- Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. *Quart. J. Econ.* 75, 643–669.
- Fretwell, S.D., 1972. *Populations in a Seasonal Environment*. Princeton Univ. Press, Princeton, NJ.
- Glimcher, P.W., 2003a. *Decisions, Uncertainty and the Brain: The Science of Neuroeconomics*. MIT Press, Cambridge.
- Glimcher, P.W., 2003b. Neural correlates of primate decision-making. *Annual Rev. Neurosci.* 25, 133–179.
- Gold, J.I., Shadlen, M.N., 2001. Neural computations that underlie decisions about sensory stimuli. *Trends Cognitive Sci.* 5, 10–16.
- Güth, W., Schmittberger, R., Schwarze, B., 1982. An experimental analysis of ultimatum bargaining. *J. Econ. Behav. Organ.* 3, 367–388.
- Hanes, D.P., Schall, J.D., 1996. Neural control of voluntary movement initiation. *Science* 271, 427–430.
- Harper, D.G.C., 1982. Competitive foraging in mallards: “Ideal free” ducks. *Animal Behav.* 30, 575–584.
- Herrnstein, R.J., 1961. Relative and absolute strength of response as a function of frequency of reinforcement. *J. Exper. Anal. Behav.* 4, 267–272.
- Herrnstein, R.J., 1997. In: Rachlin, H., Laibson, D.I. (Eds.), *The Matching Law*. Harvard Univ. Press, Cambridge.
- Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–291.
- Kahneman, D., Slovic, P., Tversky, A., 1982. *Judgment under Uncertainty: Heuristics and Biases*. Cambridge Univ. Press, Cambridge.
- Krebs, J.R., Davies, N.B. (Eds.), 1991. *Behavioural Ecology*. Third ed. Blackwell Scientific Publications, Oxford.
- Kreps, D.M., 1990. *A Course in Microeconomic Theory*. Princeton Univ. Press, Princeton, NJ.
- LeDoux, J., 1996. *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. Simon and Schuster, New York.
- Lee, D., Port, N.L., Kruse, W., Georgopoulos, A.P., 1998. Variability and correlated noise in the discharge of neurons in motor and parietal areas of the primate cortex. *J. Neurosci.* 18, 1161–1170.
- Mainen, Z.F., Sejnowski, T.J., 1995. Reliability of spike timing in neocortical neurons. *Science* 268, 1503–1506.
- Maynard Smith, J., 1982. *Evolution and the Theory of Games*. Cambridge Univ. Press, Cambridge.
- Nash, J.F., 1951. Non-cooperative games. *Ann. of Math.* 54, 286–295.
- Newsome, W.T., Britten, K.H., Movshon, J.A., 1989. Neuronal correlates of a perceptual decision. *Nature* 341, 52–54.

- Parker, A.J., Newsome, W.T., 1998. Sense and the single neuron: Probing the physiology of perception. *Annual Rev. Neurosci.* 21, 227–277.
- Parker, A.J., Krug, K., Cumming, B.G., 2002. Neuronal activity and its links with the perception of multi-stable figures. *Philos. Trans. Roy. Soc. London Ser. B* 357, 1053–1062.
- Pascal, B., 1966. *Pensees*. Penguin Books, London. Translated by Krailsheimer, A.J.
- Paulus, M.P., Hozack, N., Zauscher, B., McDowell, J.E., Frank, L., Brown, G.G., Braff, D.L., 2001. Prefrontal, parietal, and temporal cortex networks underlie decision-making in the presence of uncertainty. *Neuroimage* 13, 91–100.
- Platt, M.L., Glimcher, P.W., 1999. Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Rosenzweig, M.R., Breedlove, S.M., Leiman, A.L., 2002. *Biological Psychology*. Sinauer, Sunderland, MA.
- Salzman, C.D., Britten, K.H., Newsome, W.T., 1990. Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346, 174–177.
- Savage, L.J., 1954. *The Foundations of Statistics*. Wiley, New York.
- Schall, J.D., Thompson, K.G., 1999. Neural selection and control of visually guided eye movements. *Annual Rev. Neurosci.* 22, 241–259.
- Shrier, A.M., 1965. Response rates of monkeys under varying conditions of sucrose reinforcement. *J. Compar. Physiological Psych.* 59, 378–384.
- Schultz, W., 2002. Getting formal with dopamine and reward. *Neuron* 36, 241–263.
- Schultz, W., Dayan, P., Montague, P.R., 1997. A neural substrate of prediction and reward. *Science* 275, 1593–1599.
- Sereno, M.I., Pitzalis, S., Martinez, A., 2001. Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* 294, 1350–1354.
- Shadlen, M.N., Britten, K.H., Newsome, W.T., Movshon, J.A., 1996. A computational analysis of the relationship between neuronal and behavioral responses to visual motion. *J. Neurosci.* 16, 1486–1510.
- Shadlen, M.N., Newsome, W.T., 2001. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916–1936.
- Sherrington, C.S., 1906. *The Integrative Action of the Nervous System*. Scribner's, New York.
- Simon, H.A., 1947. *Administrative Behavior*. Free Press, New York.
- Simon, H.A., 1983. *Reason in Human Affairs*. Stanford Univ. Press, Stanford.
- Simon, H.A., 1997. *Models of Bounded Rationality: Empirically Grounded Economic Reason*. MIT Press, Cambridge.
- Snyder, L.H., Batista, A.P., Andersen, R.A., 1997. Coding of intention in the posterior parietal cortex. *Nature* 386, 167–170.
- Snyder, L.H., Batista, A.P., Andersen, R.A., 2000. Intention-related activity in the posterior parietal cortex: a review. *Vision Res.* 40, 1433–1441.
- Squire, L.R., Bloom, F.E., McConnell, S.K., Roberts, J.R., Spitzer, N.C., Zigmond, M.J., 2002. *Fundamental Neuroscience*. Academic Press, New York.
- Staddon, J.E.R., 1980. *Limits to Action: The Allocation of Individual Behavior*. Academic Press, New York.
- Staddon, J.E.R., Motheral, S., 1978. On matching and maximizing in operant choice experiments. *Psychological Rev.* 85, 436–444.
- Stephens, D.W., Krebs, J.R., 1986. *Foraging Theory*. Princeton Univ. Press, Princeton.
- Stevens, C.F., 1994. Neuronal communication. Cooperativity of unreliable neurons. *Current Biol.* 4, 268–269.
- Sutton, R.S., Barto, A.G., 1998. *Reinforcement Learning*. MIT Press, Cambridge.
- Tanji, J., 2001. Sequential organization of multiple movements: Involvement of cortical motor areas. *Annual Rev. Neurosci.* 24, 631–651.
- Tolhurst, D.J., Movshon, J.A., Dean, A.F., 1981. The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vision. Res.* 23, 775–785.
- Veblen, T., 1898. Why is economics not an evolutionary science? *Quart. J. Econ.* 12, 373–397.
- Wilson, E.O., 1998. *Consilience*. Knopf, New York.
- Zak, P., Denzau, A.T., 2001. Economics is an evolutionary science. In: Somit, A., Peterson, S. (Eds.), *Evolutionary Approaches in the Social Sciences: Toward a Better Understanding of Human Nature*. JAI Press.
- Zohary, E., Shadlen, M.N., Newsome, W.T., 1994. Correlated neuronal discharge and its implications for psychophysical performance. *Nature* 370, 140–143.