

Neuroscience and Economic Behavior

Paul W. Glimcher

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

Brains and Economics

Paul W. Glimcher
Center for Neural Science
4 Washington Place, 809
New York, NY 10003

glimcher@cns.nyu.edu
212-998-3904
Fax: 212-995-4011

Introduction

The history of economics has been marked by an iterative tension between prescriptive and descriptive advances. Prescriptive theories seek to mathematically define efficient or optimal decision making strategies, strategies which rational people are presumed follow. In contrast, descriptive theories use empirical data to describe how people, whether rational or irrational, actually behave. A tension between these two approaches invariably exists because descriptive studies often indicate that people do not reliably follow the rational predictions of prescriptive theories. One example of this tension can be found in the development of neoclassical economics at the turn of the twentieth century, an intellectual revolution which gave rise to modern economic theory. Working from the assumption that all of human behavior could be described as a rational effort to maximize wealth or happiness (or more formally the economic variable *utility*), the neoclassical theorists of the twentieth century largely succeeded in developing a coherent basic mathematical framework for economics. What followed, however, were a series of descriptive insights (for example Allais, 1953 and Ellsberg, 1961) which seemed to indicate either that humans were poor utility maximizers or that the underlying assumptions of the neoclassical revolution were in some way flawed.

One recent trend in economic thought that may reconcile this ongoing tension between prescriptive and descriptive approaches is the growing interest amongst both economists and neuroscientists in the physical mechanisms by which decisions are made within the human brain. These *neuroeconomic* scholars argue that study of the physical architecture for decision making will reveal the underlying computations that the brain performs during economic behavior. If this is true, then economic and neurobiological studies which seek to bridge the gap between these two fields may succeed in providing a methodology for reconciling prescriptive and descriptive economics. These studies may produce a highly predictive model of individual decision making that is based on the actual economic computations performed by the human brain.

Closing the Gap Between Economics and Neuroscience

By the second half of the twentieth century, theoretical economists had largely succeeded in developing prescriptive models that showed how an individual who wished to maximize his assets (or more precisely his *utility*) should behave. At the same time, experimental studies were beginning to demonstrate that much of human behavior was poorly predicted by these prescriptive models. This led a number of economists, most notably Herbert Simon (e.g., 1947; 1983; 1997), to conclude that human decision makers could be viewed as rational utility maximizers in only a limited, or *bounded* sense. Conditions under which humans behave rationally do occur, but there are also conditions under which human behavior violates the predictions of prescriptive theory and thus appears irrational. 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. Some have even begun to suggest that irrational behavior may be uniquely attributed to limitations intrinsic to the neural architecture, while rational behavior, when it occurs, may be viewed as the product of a conscious faculty that somehow transcends this biological limitation. Perhaps surprisingly, this conceptualization seems to correspond to the traditional framework that was employed in neuroscientific circles during the last several several centuries, but which modern neuroscientists are now widely beginning to challenge (e.g., Damasio, 1995; LeDoux, 1996; Glimcher, 2003a). Traditionally, brain scientists argued that all 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 (a point developed very clearly in Sherrington, 1906). This simpler mechanism proposed by the traditional view, which came to be identified with the physiological 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 the twentieth century.

During the last several decades, however, ongoing empirical work has begun to indicate to many neuroscientists that this view of the neural architecture is no longer tenable. Biological evidence now suggests to neuroscientists a more 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 clearly multi-component) decision making process governs behavior (Parker and Newsome, 1998; Schall and Thompson, 1999, Glimcher 2003b). The interdependent and 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* (rather than the utility) 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. It yields a single whole 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. Quite compelling empirical data now 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 argues 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 also 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.

In summary, this essay will make three main points: i) A successful effort to account for

economic decision making will only be possible if economists employ the highly quantitative, predictive, and surprisingly unitary models of brain function that are at the core of contemporary neuroscience. ii) While humans are unique organisms, 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 (Glimcher et al., 2004), 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 here, 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. iii) Finally, it is absolutely critical that the economists and other social scientists recognize that neuroeconomics and the kind of laboratory experiments that I and others undertake can be much more than efforts to locate a brain region associated with some hypothetical human faculty like 'cooperation'. Such studies are valuable starting points, but have troubled some economists because they provide no predictive power with regard to economic behavior. Really useful neuroeconomic studies, from the perspective of working economists, will have to fully describe the mechanisms by which economic computations, like the real-world calculations of costs and benefits, are achieved. It is an understanding of mechanism in that sense that will yield real predictive power in the marketplace.

The Neuroscience of Economic Behavior

Modern *utility theory*, arguably the foundation of modern economics, has its origins in the theory of expected value first proposed by Pascal as early as 1670. 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, a product we now call *expected value*. 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. 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. As a result, Daniel 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 should be represented in the decision making process by a roughly logarithmic function of value that also incorporated a representation of the chooser's wealth, a variable we now refer to as *expected utility*. Building on these ideas, modern utility theory proposes that decision makers represent the desirability of each possible course of action using a common scale and that making a decision involves selecting the available course of action with the highest expected utility.

Neurobiological studies conducted over the past decade have revealed that the brains of both human and nonhuman primates do indeed represent a complex variable which, under many circumstances, closely parallels this classical notion of expected utility (e.g., Platt and Glimcher, 1999; Gold and Shadlen, 2000; Brieter et al., 2001, Knutson, et al., 2001; Paulus et al., 2001). For example, the rate at which nerve cells in the posterior parietal cortex generate electrochemical action potentials is very precisely correlated with prescriptively computed estimates of expected utility under many conditions (Glimcher et al., 2004). Further, some of these studies even suggest that in the final stages of the decision making process, the neural architecture selects a course of action by mechanistically generating the response associated with the greatest activity in the posterior parietal cortex.

Identifying the Neurobiological Representation of Expected Utility

One of the first studies to make this suggestion that something like expected utility is actually instantiated within the nervous system was Platt and Glimcher (1999). In those experiments, trained rhesus monkeys were allowed to participate in repeated rounds of a simple lottery while the activity of nerve cells in the posterior parietal cortex was monitored. At the beginning of each round a red and a green spot were illuminated on a screen directly in front of the monkey. This began the lottery phase of the round, a period during which the monkey did not know whether the red or green light would be offered to him as a prize at the end of that round. At the end of this phase, a third light changed color to red or green indicating which of the two initial lights had been randomly selected to yield a fruit juice reward on that particular round. The monkey then received the fruit juice if he simply made contact with the selected light at the end of the round. While monkeys played hundreds of rounds of this game, Platt and Glimcher systematically varied the prize associated with the red and green lights. They did this by varying either the size of the reward associated with each light, the value of that light, or the relative probabilities that the red or green lights would be selected at the end of the round, the likelihood that each light would yield a reward.

These two variables, *value* and *likelihood*, were selected for manipulation because essentially all utility theories since Pascal are based on the assumption that rational decision makers should assess the desirability of any course of action by combining the value and likelihood of gain. Even though in this experiment the monkeys did not need to monitor these values in order to behave efficiently, Platt and Glimcher hoped to determine whether these economic variables were encoded in the nervous system while monkeys participated in the repeated lotteries, just as we might expect them to be in human players.

Platt and Glimcher found that a discrete group of nerve cells in the posterior parietal cortex encoded, separately for each light, a combination of the value and likelihood of reinforcement associated with each light during the lottery phase of each round. It appeared from this result that under these conditions the brains of their monkeys explicitly encoded something very much like the economically defined expected value or expected utility of each light in a simple lottery task.

Game Playing Monkeys

Dorris and Glimcher (Glimcher et al., 2004) extended this finding when they examined the activity of this same brain region while a new group of rhesus monkeys were engaged in a strategic conflict known as the *inspection game*. In the human version of that game, two opponents face each other, an *employer* and an *employee*. On each round of the game the employee must decide whether to *go-to-work*, in which case he earns a fixed wage, or whether to *shirk*, in hopes of earning his wage plus a bonus. (In the human version of the game, the free time gained by shirking in itself is conceived of as a sort of bonus.) The goal of the employee is simply to maximize his gains in terms of salary and bonus. The employer, on the other hand, must decide between trusting his employee to arrive for work or spending money to hire an inspector who can actually check up on the employee. The goal of the employer is to spend as little as possible on inspections while maximizing the employee's incentive to work.

The inspection game is of particular interest to economists because when players employ prescriptively rational strategies, the outcome of the game can be predicted using a theory developed by John Nash in the 1950s. Nash equilibrium theory describes how, when the cost of inspection to the employer is set high, the efficient strategy for both players converges on a solution in which the employee manages to shirk fairly often. Conversely, a low inspection cost to the employer defines a theoretical equilibrium solution in which shirk rates are low. (For a more detailed description of the inspection game and an introduction to game theory see chapter 12 of Glimcher, 2003a.)

One of Nash's fundamental insights was that at a *mixed strategy equilibrium*, a situation in which a prescriptively rational player would distribute his actions amongst two or more alternatives at a fixed frequency, the expected utility of the two or more actions amongst which his responses are distributed must be equivalent. This means that during the inspection game, the expected utilities of working and shirking must be equal during rational play, regardless of how frequently the Nash solution requires that the player work. The Nash approach argues, essentially, that a behavioral equilibrium occurs when the desirability of working and shirking are rendered equal by the behavior of one's opponent irrespective of how often that equilibrium requires that one work. The

Nash equations themselves go a step further, defining the precise rates of working and shirking that are prescriptively rational for a given computation of expected utilities.

Dorris and Glimcher found that the probability a human playing the inspection game for money would chose to shirk was well predicted by the prescriptive Nash equilibrium computations whenever those computations predicted shirking rates of 40% or more. When, however, this prescriptive theory predicted shirking rates below approximately 40%, human subjects were observed to shirk more frequently than was predicted. Importantly, this descriptive assessment of humans seemed to differ from the prescriptive assessment provided by the Nash equilibrium equations. When the expected utility calculations indicated that the humans should have shirked 10% of the time, the humans behaved instead as if working and shirking were equally desirable when a shirk rate of about 30% was produced.

When Dorris and Glimcher analyzed the behavior of their monkeys, they found that the behavior of the monkeys was surprisingly similar, even essentially identical, to the behavior of the human employees. Just like humans, the monkeys seemed to precisely track the Nash equilibrium solutions and deviated from those solutions only when shirking rates of less than 40% were prescribed during the inspection game (Figure 2). This was a critical advance because it allowed Dorris and Glimcher to examine the role of the posterior parietal cortex during a voluntary strategic game in which monkeys and humans seemed to employ similar, or identical, strategies.

Dorris and Glimcher noted from the Nash approach that the desirabilities of working and shirking must be equivalent whenever strategic competition yields a mixed strategy, and thus that mixed strategies when they occurred must be associated with the equal desirability of working and shirking as represented in the nervous system. Prescriptive theory proposes that this point of equal desirability should be computable from Nash equations. These descriptive studies, however, indicated that the Nash equations failed to accurately identify this point of equal desirability when shirk rates of less than 40% were predicted. This permitted Dorris and Glimcher to examine the activity of neurons in the posterior parietal cortex while monkeys played the inspection game in order to determine exactly what these neurons encoded. They found that when the monkey's

behavior was well predicted by the Nash equations neural activity was equivalent to the expected utility of economic theory. Interestingly, when the monkeys deviated from the prescriptive predictions of economic theory, the neurons continued to behave as if they were still representing two courses of action that had equal desirabilities. It was as if the neurons were unaware that the computations they were performing under those conditions violated prescriptive theory; the activity of the neurons continued to predict behavior, whether or not it was formally rational in the way described by the Nash equations. The neurons thus seemed to encode a *physiological expected utility* that was, in essence, a hybrid of the expected utility employed in prescriptive theories and the subjective preferences measured in descriptive studies. This observation is critical because it is one of the first suggestions that neurobiological studies may be able to resolve the tension between prescriptive and descriptive approaches in economics. Neurobiological measurements may identify variables that are related to prescriptive theory and highly predictive of behavior. The challenge that these observations raise is whether or not the neurobiological computations that yield these variables can be fully determined. This is a critical point we will return to in a moment.

Humans Playing Lotteries

Studies of monkeys and other animals continue to demonstrate the surprising similarities between the economic behavior of humans and our nearest relatives. They also employ highly precise brain measurement technologies which cannot be used in humans. Recently, however, the less precise brain scanning technologies that can be employed in humans have also begun to yield significant insights into the neural basis of economic behavior, which also hints at a reconciliation between prescriptive and descriptive approaches to economics (McCabe et al., 2001; Montague and Berns, 2001). One of the first and most compelling of these studies examined the behavior of humans during a lottery similar to the one employed by Dorris and Glimcher for the study of monkeys (Breiter et al., 2001). In that experiment, human subjects were presented, on sequential rounds, with one of three possible lotteries (Figure 3). In lottery one (the good lottery), they faced equal chances of winning \$10, \$2.50 or \$0. In lottery two (the intermediate lottery) they faced an

equal chance of winning \$2.50, winning \$0, or losing \$1.50. In lottery three (the bad lottery) they faced an equal chance of winning \$0, losing \$1.50 or losing \$6.

At the beginning of each round the subjects were told which lottery they would be playing, and the average activity in many brain areas was simultaneously measured. After that measurement was complete, the lottery was actually played and the humans were then told how much real money they had earned on that round. This design was particularly interesting because of an important and well described deviation of human behavior from prescriptive expected utility theory (Kahneman and Tversky, 1979; Kahneman, et al., 1982). All three of these particular lotteries present a one-third possibility of winning \$0, but they do so under different conditions. In the good lottery winning \$0 is the worst possible outcome while in the bad lottery it is the best. Kahneman, Tversky and their colleagues have noted that while humans should, according to prescriptive expected utility theory, always see winning \$0 as a neutral event this is not the case. When a human participates in the good lottery, they find winning \$0 to be an intensely negative outcome while when a human participates in the bad lottery, they find winning \$0 to be a positive outcome. What Breiter and colleagues hoped to determine was whether the activity of some brain area might track this human response to the context within which the 0\$ outcome was embedded.

What they found was that the activity of a brain region called the sub-lenticular extended amygdala did seem to encode both the expected utility of each lottery and the context within which gains and losses are realized. When humans were first presented with the lottery they would face on that round, activity in this brain area was closely related to the expected utility of the overall lottery in which they were placed. This was an outcome well described by prescriptive theory. After the lottery ran, however, they found that the activity of this area was a rough function of the response of the human to the particular context. Activity in this area was higher when the subjects won \$0 in the bad lottery than when they won \$0 in the good lottery.

Once again the neural results lead to an interesting and perhaps unexpected conclusion. When human behavior is rational, as defined by our prescriptive economic theories, we can find evidence that some brain areas encode expected utility. When, however, human behavior deviates

from a prescriptive theory, the brain seems to encode something that predicts behavior more accurately than the prescriptive theory.

Together, these observations raise an intriguing possibility: The neural architecture may indeed compute and represent something like the expected utility of many possible courses of action, much as prescriptive theory proposes. When choosers are efficient in the economic sense, that architecture accurately represents the expected utility of available choices. When prescriptive and descriptive conclusions differ, however, the neural architecture seems to be closely related to the descriptive process. While it may be counterintuitive to believe that sensitivity to things like context may reflect the principled output of highly developed neural circuits, this may simply reflect the fact that evolution shaped our neural architecture to perform efficiently under many, but not all, environmental circumstances. In some cases our irrational sensitivities to things like context may arise when the most complicated cortical mechanisms inside our skulls encounter problems that they did not evolve to solve. It is these biologically-based inefficiencies which might place boundaries on the circumstance in which we might be expected to produce economically rational behavior. The available evidence thus suggests a synthesis of modern economic and neuroscientific approaches. By biologically defining the mechanisms which make decisions we should be able to derive a mechanistically accurate economic theory which is, by necessity, predictive.

Neuroscience and Economic Theory

Can the neurobiological computations that yield real-world decisions be fully determined given these observations? The answer to that question is not yet entirely clear but my colleagues and I (Glimcher et al., 2004) have been attempting to extend this approach by studying how the brain computes, or learns, the physiological expected utilities that guide choice behavior under some more complicated circumstances. Bayer and Glimcher have attempted to do this by studying the activity of a group of nerve cells in the substantia nigra pars compacta which use the neurochemical dopamine to communicate with other nerve cells. These cells are widely believed to compute the difference between the gains that a human or animal expects to receive and the gains that they

actually receive (cf. Schultz 2002), and a growing body of evidence now suggests that this is the substrate from which physiological expected utilities may be calculated. What Bayer and Glimcher attempted to do was to examine these neurons during a simple choice task in an effort to derive a precise mathematical equation that could describe exactly what they encode. We then used this equation to predict the behavior of monkeys during a simple economic task. In that task, which was studied by Lau and Glimcher, monkeys were faced with a dynamic lottery. On each round the monkeys could select either a red or green light placed in front of them. Before each round the lottery was designed so that there was a fixed probability that each of the two lights would be *armed* with a reward. For example, there might be a 10% chance that the red light would be armed before each round and a 20% chance that the green light would be armed before each round, and the lights were always armed with the same amount of fruit juice. Critically, once a light was armed it remained armed until chosen by the monkey in a subsequent round.

The accompanying figure shows the free choices made by a monkey while playing this game as a thick grey line. One can see that the behavior of the animal is apparently chaotic, fluctuating from red to green. The thin black line shows the prediction of a neuroeconomically derived model developed from a study of these dopamine neurons. What is critical is that the model does a remarkably good job of predicting the behavior of the animal on a step-by-step basis. The model, which is neither truly prescriptive nor descriptive, is highly constrained by neurobiological observations and makes clear behavioral predictions. Of course the model is making predictions about a very simple economic behavior, but it seems likely that more sophisticated models of this type will soon be developed. And it is these forthcoming models which will either validate or invalidate the promise of the developing neuroeconomic approach.

Implications

One of the critical and persistent issues in economics has been our inability to reconcile the rational choice model at the core of modern theory with the fact that humans are the product of a 600 million year evolutionary lineage. We all recognize that non-human animals have limited me-

chanical 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 until recently been a central premise of mainstream economic thought that humans have transcended these limitations. That humans rely on a more fundamentally rational process, which many presume is subjectively experienced as consciousness and which is often assumed to be mechanistically located within the cerebral cortex, and that this uniquely human process 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 biological-evolutionary heritage influences our actions. 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 neoclassical 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 the simple and 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 unitary 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; evolutionary theory defines the field within which mechanism is optimized by neoclassical constraints; 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 basically 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 and neural maps of these physiological expected utilities seem to be the substrate upon which decisions are actually made (Glimcher et al., 2004).

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 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 economic model. Studies like these seem to be elucidating the mechanisms by which economic behavior is accomplished, and a critical advantage of this approach to irrational behaviors is that once the mechanism is understood, all economic behavior 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. More recently the biologist E. O. Wilson (1998) 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.

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.

Breiter, H. C., Aharon, I., Kahneman, D., Dale, A. and Shizgal, P. (2001) Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron* 30:619-639

Caraco, T., Martindale, S. and Whittam, T. S. (1980) An empirical demonstration of risk-sensitive foraging preferences. *Animal Behavior*. 28: 820-830.

Damasio, A. R. (1995) *Descartes Error: Emotion Reason and the Human Brain*. London: Pan Macmillan.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics*. 75: 643-669.

Gold, J. I. and Shadlen, M. N. (2000) Representation of a perceptual decision in developing oculomotor commands. *Nature* 404: 390 - 394

Glimcher, P. W. (2003a) *Decisions, Uncertainty and the Brain: The Science of Neuroeconomics*. Cambridge: MIT Press.

Glimcher, P. W. (2003b) Neural Correlates of Primate Decision-Making. *Annual Review Neurosci.* 25: 133-179.

Glimcher, P. W., Dorris, M. C. and Bayer, H. M. (2004) Physiological expected utility theory and the neuroeconomics of choice. *Games and Economic Behavior*. *In Press*.

Kahneman, D. and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econo-*

metrica. 47: 263-291.

Kahneman, D., Slovic, P. and Tversky, A. (1982) *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.

Knutson, B., Adams, C. M., Fong, G. W. and Hommer, D. (2001) Anticipation of Increasing Monetary Reward Selectively Recruits Nucleus Accumbens. *J. Neurosci* 21: RC159

Krebs, J.R. and Davies, N.B. (eds.) (1991) *Behavioural Ecology*, 3rd edition. Oxford: Blackwell Scientific Publications.

LeDoux, J. (1996) *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. New York: Simon and Schuster.

Maynard Smith, J. (1982) *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.

McCabe, K., Houser, D., Ryan, L., Smith, V. and Trouard, T. (2001) A functional imaging study of cooperation in two-person reciprocal exchange. *Proceedings of the National Academy of Science (USA)* 98: 11832-11835.

Montague, P. R. and Berns, G. S. (2002) Neural economics and the biological substrates of valuation. *Neuron*. 36:265-284.

Parker, A. J. and Newsome, W.T. (1998) Sense and the single neuron: probing the physiology of perception. *Annual Review Neurosci.* 21:227-77.

Platt, M. L. and Glimcher, P. W. (1999) Neural Correlates of Decision Variables in Parietal Cortex. *Nature*. 400: 233-238

Paulus, M. P., Hozack, N., Zauscher, B., McDowell, J. E., Frank, L., Brown, G. G. and Braff, D. L.

- (2001) Prefrontal, parietal, and temporal cortex networks underlie decision-making in the presence of uncertainty. *Neuroimage*. 13:91-100.
- Schall, J. D. and Thompson, K. G. (1999) Neural selection and control of visually guided eye movements. *Annual Review Neurosci*. 22: 241-59.
- Sherrington, C.S. (1906) *The Integrative Action of the Nervous System*. New York: Charles Scribner's Sons.
- Schultz, W. (2002) Getting formal with dopamine and reward. *Neuron*. 36:241-263
- Simon, H. A. (1947). *Administrative Behavior*. New York: Free Press.
- Simon, H.A. (1983) *Reason in Human Affairs*. Stanford: Stanford University Press.
- Simon, H.A. (1997) *Models of Bounded Rationality: Empirically Grounded Economic Reason*. Cambridge: MIT Press.
- Stephens, D.W. and Krebs, J.R. (1986) *Foraging Theory*. Princeton: Princeton University Press.
- Veblen, T. (1898) Why is economics not an evolutionary science? *Quarterly Journal of Economics*. 12: 373-397.
- Wilson, E. O. (1998) *Consilience*. New York: Knopf.

Inspection Game Earnings

Computer Employer

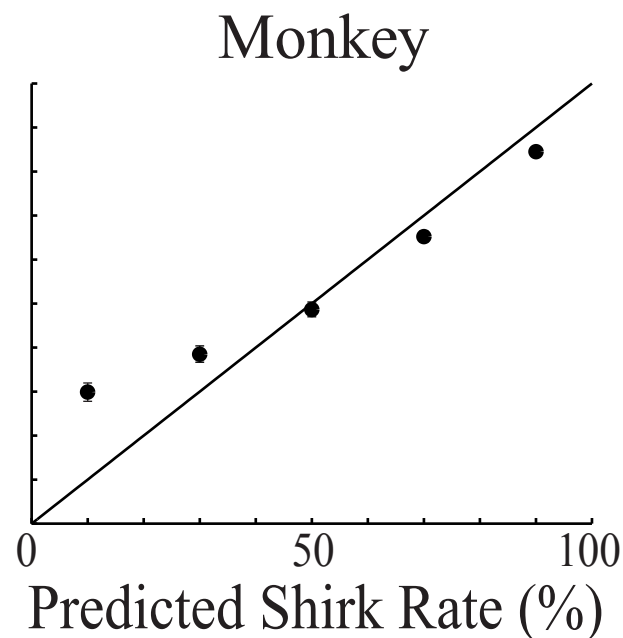
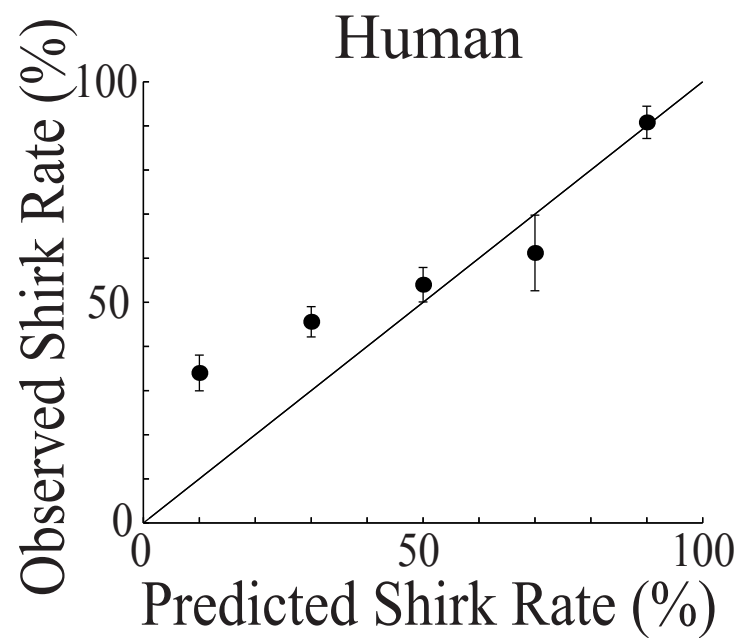
		Employer Action	
		Inspect	No Inspect
Employee Action	Work	\$1.70	\$2.00
	Shirk	-\$0.30	-\$2.00

Human Employee

		Employer Action	
		Inspect	No Inspect
Employee Action	Work	\$1.00	\$1.00
	Shirk	\$0.00	\$2.00


Human Employee

		Employer Action	
		Inspect	No Inspect
Employee Action	Work	1 oz	1 oz
	Shirk	0 oz	2 oz



Lotteries

Good



Intermediate



Bad

