

6

Proximate Mechanisms of Individual Decision- Making Behavior

Paul W. Glimcher

Abstract

In the early part of the twentieth century, neoclassical economic theorists began to explore mathematical models of maximization. The theories of human behavior that they produced explored how optimal human agents, who were subject to no internal computational resource constraints of any kind, should make choices. During the second half of the twentieth century, empirical work laid bare the limitations of this approach. Human decision makers were often observed to fail to achieve maximization in domains ranging from health to happiness to wealth. Psychologists working during this period responded to these failures by largely abandoning holistic theory in favor of large-scale multiparameter models that retained many of the key features of the earlier models. Over the course of the last two decades, scholars combining neurobiology, psychology, economics, and evolutionary approaches have begun to examine alternative theoretical approaches. Their data suggest explanations for some of the failures of neoclassical approaches and have revealed new theoretical avenues for exploration. While neurobiologists have largely validated the economic and psychological assumption that decision makers compute and represent a single decision variable for every option considered during choice, their data also make clear that the human brain faces severe computational resource constraints which force it to rely on very specific modular approaches to the processes of valuation and choice.

Introduction

In the first half of the twentieth century, economics experienced a revolution. Prior to that time, economics had been plagued by arbitrary models and a lack of focus. Responding to the challenge of Vilfredo Pareto (1971), Léon Walras (1954) and the Lausanne School as well as the American economist Paul Samuelson (1947) catalyzed a shift to mathematical rigor, theoretical clarity,

and a clear focus on the logical consistency of sets of decisions. Samuelson argued that simply assuming that a decision maker *should* be logically consistent in his or her choices could provide an economist with tremendous theoretical leverage in understanding human behavior. His argument was, in essence, that any decision maker who was logically consistent in his or her behavior (a property that unfortunately came to be known as *rationality*) could be described as if they computed and represented a decision variable for each option under consideration and selected the option from that set that had the highest idiosyncratic value (or utility) to that decision maker. Samuelson's insights showed such clarity and power that they came to define twentieth century economics along with the work of equally talented figures such as Kenneth Arrow and Gerard Debreu (1954).

Of course, the critical assumption upon which the neoclassical revolution relied was that when humans make decisions, they achieve an essentially perfect maximization of something (although there was no assumption that what they were maximizing was necessarily money or expected value). This assumes that the human brain, the product of sixty million years of primate evolution, is capable of perfectly representing any set of options and noiselessly computing the option that has the highest value to the chooser. This neglects the possibility that structural limitations to the human brain may make computations imperfect in highly systematic ways.¹ Put in evolutionary terms, it neglects the fact that accurate computation is extremely costly (in a metabolic sense) and that all organisms are under evolutionary pressure not to be abstractly perfect, but rather to be efficient; that is, to trade off the costs and benefits of accuracy in their environment. Of course, Samuelson and his cohort were well aware of this assumption, but they also recognized that their oversimplified approach was more powerful than any approach that had come before it.

The neoclassical program was famously challenged in the second half of the twentieth century, after it became clear that humans are, in fact, systematically inconsistent in their choice behavior under many circumstances. The first such observation of broad impact was made by the French economist Maurice Allais (1953) in the domain of probabilities. Humans, he noted, are inconsistent in their treatment of probabilities when those probabilities are very high or very low. Following Allais, others discovered a host of other examples of systematically inconsistent behavior. Daniel Kahneman and Amos Tversky (1979) attempted to classify and model these inconsistencies in a number of ways; most famously with their multiparameter *prospect theory*.

The great problem faced by prospect theory (and related approaches) however, was that it lacked any deep theoretical roots independent of the expected utility theory from which it was derived. Like expected utility theory, it postulated

¹ This point was made passionately by the decision scientist Herbert Simon (1955), who argued that humans do not choose optimally but rather approximate optimality as closely as possible given their cognitive limitations—a process he referred to as *satisficing*.

a utility-like function, though one with many more parameters than employed in the neoclassical model. Like expected utility theory, it also postulated a probability representation that combined multiplicatively with value, though a function with 1–2 free parameters rather than with no free parameters. Perhaps most importantly, and innovatively, the model also proposed the existence of a *reference point* against which all gains and losses were compared internally.

Given the additional parameters that prospect theory added to the arsenal of choice modelers, it will come as no surprise that Kahneman and Tversky's model had greater predictive power than its antecedents. However, their model offered little in terms of new theory (the exception being the reference point, which was an enormous theoretical innovation) (for a modern analysis of the reference point, see Koszegi and Rabin 2006, 2007, 2009). By the beginning of the twenty-first century, it was commonplace for scholars to think of the computational objects within prospect theory as “really there” and “really true.” Still, there is no escaping the fact that many of the model's parameters are highly collinear and nonidentifiable (in the technical sense). This means, to take one example, that while it is possible to describe a class of behavior within the model as “risk seeking in losses,” the exact same behavior can be correctly described as a “distortion of probability.” That is an important point because it means that prospect theory does not provide clean links between empirical/experimental observations and underlying theoretical constructs. The theory, instead, provides predictive power via a rich set of fit-able parameters and model variants. Thus, while the “behavioral revolution” catalyzed by Kahneman and Tversky led to better prediction capabilities, it yielded surprisingly few fundamentally new approaches to understanding decision making.²

Understanding the Decision Maker at a Deeper Level

Over the last decade, in response to the empirical shortcomings of neoclassical economics and the theoretical boundaries of prospect theory (and related approaches), an intellectual revolution emerged in our understanding of the biological bases of human and animal decision making. It is hard to understate the scale of this revolution. In 1992, for example, at the annual meeting of

² An alternative strategy, suggested by Kahneman, Tversky and their colleague Eric Johnson, and later by Gerd Gigerenzer (2000), is that human decision making can best be modeled as the execution of a set of simple rules, or *heuristics*. The notion here is that at a psychological level, our decision-making process is organized around a set of task-specific semantic-level rules such as “try a dozen and pick the best.” These simple algorithms have also proven to be highly predictive of choice behavior under limited circumstances and when a sufficient number of them are incorporated into a larger model, these hierarchically organized multi-rule models can be quite broadly predictive. Biological studies of these processes conducted during the last decade, however, have generally failed to identify neural correlates of this class of psychological model. While this in no way reduces the usefulness of these models, biology remains the focus of this review and hence we move on without a review of the heuristics literature.

the Society for Neuroscience, more than 10,000 presentations were made, yet less than 10 addressed decision making in any form. At its 2014 meeting, by contrast, nearly a thousand presentations asserted conclusions related to the neural bases of decision making in species ranging from the 1000-cell nematode worm *Caenorhabditis elegans* to humans. Even more surprisingly, these studies point to a group of commonalities within and across species.

Nearly all current neurobiological evidence suggests that we can think of mammalian (including humans) decision making as being organized around three basic processes (Glimcher 2011a; Glimcher and Fehr 2013):

1. A suite of brain areas that construct an estimate of the value of each of the options being considered by the decision maker
2. A smaller overlapping suite of brain areas that actually compares these option values mechanistically and passes the resulting choice to movement control (or output) systems of the brain
3. A set of (again overlapping) learning-related areas that compares the quality of the obtained option with the expected quality of that option and updates the internal representation of option value to improve future decision making

This neurobiological work into the basis of decision making may prove critical for understanding how and why we make decisions even at a behavioral level, regardless of the implications of these studies for understanding health, disease, biology, and evolution. First, in terms of behavior, neurobiological studies may allow us to select among competing models using new classes of data. That selection will presumably prove particularly valuable when a simple neurological insight eliminates the need for detailed behavioral experiments in a host of domains. This would demonstrate that a particular class of computation (e.g., the representation of completely cardinal utilities) is physiologically impossible. Second, identifying the neural algorithms for choice might well inspire completely new classes of behavioral-level theories of choice which have greater predictive power at the behavioral level than do existing theories. Current economic models of “normalization” are an example of this type. Finally, neurobiological studies may lead to simplifications of existing models that improve the formal tractability of those models. Studies that seem to suggest that loss aversion and risk seeking in losses may employ a common mechanism provide an example of this kind. In any case, it seems likely that understanding the mechanism of choice is an endeavor that will proceed fastest and most accurately when engaged by a host of disparate communities, including neurobiologists, in this case.

Valuation Areas

When one searches for brain areas that encode the values of options under consideration, a small group of brain structures stand out. Like all brain structures,

these areas are made up of millions of neurons arranged in detailed, topographically organized aggregations that have many precise and restrictive properties. The neurons in these structures encode real-valued quantities in their *firing rates*. This is the rate at which each neuron generates electrochemical action potentials—the common currency of information representation and mathematical computation in the brain. For all neurons, these firing rates (the number of electrochemical potentials generated per second) are strictly bounded between zero and some upper limit (typically around 100 impulses per second). In addition, these firing rates are stochastic, showing well-understood Poisson-like variance (Tolhurst et al. 1983). One further restriction is that these firing rates typically present with a unique zero-point: a firing rate that they produce in the absence of any representational input. As a result, neurobiologists typically think of these firing rates as unique and fully cardinal, though stochastic, representations. Because these firing rates are both bounded and subject to significant stochasticity, neuroscientists tend to think of them as having a very finite information-carrying capacity.

The neurons that appear to play a privileged role in the representation of value are largely located in a discrete set of regions within the frontal cortex of the brain and in a second discrete set of regions within the striatum, a more evolutionarily ancient structure that lies beneath the cerebral cortex. What appears to be happening in this suite of areas, when one considers an option, is that these specialized areas work together to derive a single aggregate estimate of option value by the aggregation of signals (action potentials) from each of these many specialized subareas. Several very specific brain areas are key players in this aggregation process: the orbitofrontal cortex which appears to specialize in encoding a kind of immediate sensory-derived value for options under consideration; the dorsolateral prefrontal cortex which empirical data suggests plays a key role in both patience and self-control; the dorsal striatum which seems to play a key role in storing the values of frequently encountered habit-related options or actions; the amygdala which seems to be weakly specialized for encoding some classes of negative outcomes; the anterior cingulate cortex which seems to play a role in tracking opportunity costs—the value of staying with a current option versus switching to a new option (though this is less certain); the anterior insular cortex which seems to play a role in valuing aversive or disgusting options; and the hippocampus which appears to play a role in some kinds of explicit value recall processes (Bartra et al. 2013).

Our current evidence suggests that the aggregation of the signals from these areas physically occurs in a limited set of brain areas through a mechanical process of weighted impulse summation and that it is this single aggregated value signal that is the subject of the choice mechanism. This is a key point and one that has been a real constant in neurobiological studies. At least three meta-studies have now conclusively shown that activity in two aggregating brain areas robustly predicts human choice behavior, under almost any condition and when choices are made over almost any kind of good: the ventromedial

prefrontal cortex and the ventral striatum (Levy and Glimcher 2012; Bartra et al. 2013; Clithero and Rangel 2014). If, for example, brain activity is measured in either of these areas while human subjects view consumer goods, the relative levels of this activity predict, with fairly high accuracy, which consumer good an individual will choose (Tusche et al. 2010; Levy et al. 2011). The current record is about 94% accuracy (Gross et al. 2014). Importantly, this is true for many different kinds of goods and rewards, ranging from snack foods to social encounters. In fact, it is basically true for every class of choice object that has ever been studied.³

It is important to note, however, that from an economist's point of view, very few kinds of goods have actually been studied in this way. A wide range of consumer goods (from snack foods to drinks to consumer electronics) have been examined, as has music, visual art, social interactions, and social dominance. In addition, the neural representation of money, ranging from cents to hundreds of dollars, has also been examined. All of these categories of goods yield a virtually identical picture at the level of these final common aggregations (Bartra et al. 2013), but many critical categories of economic goods have not yet been studied. Life-changing wealth shocks, durable goods, and homes, to take three critical examples, have not yet been examined by neurobiologists. While there is as yet no reason to believe that these will be different from the many other kinds of goods that have been well studied, this is an important caveat to bear in mind.

What may be most interesting about this observation of a single common (or final) representation of choice value for all classes of goods so far studied is the similarity of this observation to the theoretical roots of neoclassical economic theory and prospect theory. Recall that both of those theories rested on the assumption that choice can be modeled as a representation of value; the theories assume that the desirability of a good (or more generally a reward or punishment) rests on a single numerical valuation for that good. Neoclassical theory had assumed that these valuations (or expected utilities) were stable and consistent. Prospect theory and the field of behavioral economics it spawned assumed that these values (or prospects) were highly context dependent. Both, however, posit some aggregate underlying value for comparing options, and this is exactly what has been observed. A small set of brain areas where expected values or prospects are represented for comparison.

It is important to note that in many ways the data lean strongly toward a model that is more similar to prospect theory than to a strict version of neoclassical theory. If the final common valuation signal reflects the aggregation

³ It is worth noting that the activity of neurons is fundamentally cardinal (in the economic sense) and that these cardinal signals have been used by neurobiologists largely to predict choice in an ordinal utility-theoretic sense. Recently, there has been growing interest in employing more cardinal notions of utility derived from random utility models (McFadden 1974) for these cardinal signals. While details of those approaches are interesting and important, they lie outside the scope of this review.

of 5–15 inputs, there is every reason to believe that the aggregate values will be highly context dependent, and that is, of course, exactly what is generally observed in behavior. While many details are uncertain, we know, for example, that if the dorsolateral prefrontal cortex is artificially activated during choice, subjects become more transiently patient in their choices; the degree of this change in patience can be precisely predicted from the change in activity in the ventromedial prefrontal cortex induced indirectly by this activation. Thus while context dependencies of this kind are starting to be explored (Louie and De Martino 2013), interactions between these areas are also beginning to be understood in a crude way.

At the moment, we do not know exactly how these different areas interact in any real detail. Often, we do not even have credible theoretical frameworks for describing these empirically observed interactions. Basic questions, such as whether these areas compete for control of the final common representation (in a game theoretic sense) or smoothly work together in a weighted summation process, are currently under heated debate and experimental inquiry. Nonetheless, the pattern of connectivity and interaction is beginning to be mapped in an empirical way, and many interpret this pattern as being broadly supportive of a prospect theory-like mechanism.

Growing evidence indicates that the more evolutionarily ancient parts of the brain also seem to contribute (perhaps more indirectly) to the activity observed in these circuits, and these studies may shed fundamental insight into the role of motivations on preferences. As economists, we have known for centuries that changes in a chooser's motivational state can alter their preferences. In general, however, we have treated these motivational states as lying outside our economic models. Motivational changes may well modify preferences, but exactly how or why is something that economists have often been reticent to pursue. Neurobiologists, in contrast, have long been interested in how motivational states alter behavior and preferences, and several efforts to connect these motivational studies with economic models of preferences are now underway.

Changes in a human subject's hunger state, to take one example, has been known to alter activity in the evolutionarily ancient hypothalamus in a way that regulates hunger, thirst, and a host of other motivational variables. How does this happen, and could insights into this process be used to develop models of how preferences change in an economic sense? The answer seems to be yes. Changes in this ancient structure appear to be able to influence the activity of circuits that control human risk attitudes in a measurable and predictable way (Symmonds et al. 2010). Subjects who are hungry seem to behave differently, even in stock markets, due to changes in the hypothalamic state (Levy et al. 2013). There is a rich and growing literature on how values are influenced by these types of inputs.

Studies of how hormonal state influences choice and decision making also fall into this category, although their results are much more ambiguous than is reflected in the widespread public discussion of their findings. A classic

example stems from the work of Ernst Fehr and his coworkers (Kosfeld et al. 2005), which examined the effects of the hormone oxytocin (a hormone associated with pregnancy, care for the young, and pair bonding in females) on social decision making. They found that higher levels of oxytocin increase “an individual’s willingness to accept social risks arising through interpersonal interactions” (Kosfeld et al. 2005:673; for a more recent extension of our understanding of this phenomenon, cf. De Dreu et al. 2010). Related work on the hormone testosterone initially indicated that this masculinizing hormone led to greater interpersonal aggression and reduced social cooperation, but subsequent studies revealed this to be a placebo effect (Eisenegger et al. 2010). A rich literature is emerging which documents how a number of neurochemicals influence the activity of the brain areas that contribute to valuation, and thus alter valuation and choice (Crockett and Fehr 2013). These interactions identify another mechanism by which valuation and subsequent choice show context dependence and motivational sensitivity.

Choice Areas

Much more is known now about the choice process in monkeys than in any other species. Our understanding of how humans actually choose is currently somewhat speculative (Glimcher 2013; Gold and Heekeren 2013). We do know that a class of brain mechanisms, well studied in monkeys, seems designed to implement the choice process by selecting one of several options based on the relative value of those options. Broadly speaking, two research threads have explored the choice mechanism under two different conditions: during *perceptual* decision making and during *economic* decision making. Perceptual decision making refers to choices between two options, when the information required for making that choice is provided slowly and the decision maker integrates that information over time: subjects must trade off speed-of-decision against accuracy; thus, models describe not only choices but also reaction times. A host of models of this process have been applied to neurobiological data. All of them essentially describe neural circuits that integrate imperfect incoming signals to some choice threshold. Activity above that threshold then mechanically triggers the output machinery of the brain that implements the actual external-world manifestation of the choice (Usher and McClelland 2001; Gold and Shadlen 2007; Cisek and Kalaska 2010). The second category of choice, economic decision making, describes non-reaction-time processes and has many similarities. Like the perceptual models, they instantiate a winner-takes-all competition between the option values, with the higher option value ultimately taking control of the network and activating output mechanisms (Padoa-Schioppa 2011; Glimcher 2013).

While an enormous literature attempts to adjudicate between (or to reconcile) these two classes of models, such debate falls outside the scope of this overview. What is important is that the basic biophysical mechanism for

choosing is an area of very active research and steady progress. This work indicates that these choice mechanisms likely operate in parallel in a number of brain areas, including the posterior parietal cortex, dorsal striatum, and orbitofrontal cortex as well as several other areas. The most compelling hypothesis is that the mechanism for choice is composed of a set of distributed mechanisms that aggregate value signals for each of the options being analyzed and that these aggregated values are then compared using a neural mechanism which forces the activity representing value for each of the available options to compete for control of the representational network. The “winner” of that competition, by “winning,” exceeds a static biophysical threshold which is the physical instantiation of choice. Choice seems likely to be something that can be accomplished through this competitive mechanism by any one of several parallel simultaneous processes and may involve levels of specialization not yet understood. Thus the act of making a choice may be less of a singular process than a suite of interacting processes which all operate on the shared common representation of value mentioned above.

Learning-Related Processes

The third area in which significant progress has been made is in our understanding of the mechanisms by which the brain learns from experience. These are the mechanisms that allow us to alter the values we place on goods or options based on our experience. This kind of feedback-driven learning has been the subject of an enormous amount of work in computer science, psychology, and neuroscience. Parallel developments in these three fields have now made clear the outlines of the learning system.

Much of the trial-and-error learning system of the brain is organized around the neurotransmitter dopamine (Schultz 2006; Glimcher 2011b), a chemical released from a specific group of neurons located in the midbrain. We now know that the instantaneous concentration of this chemical encodes discrepancies between experienced and expected rewards (more precisely, utilities). If one expects an action to yield a dollar but it instead yields two, then immediately after one receives the two dollars, dopamine concentrations are briefly elevated. Had one received three dollars against this same expectation, the concentration would have been higher. Had one received only one dollar, as expected, the concentration would have been unchanged. Had one received only fifty cents, the concentration in the brain would have been transiently depressed. Dopamine thus encodes what is referred to as a reward prediction error signal.

The dopaminergic neurons of the midbrain are now known to compute this term from two independent inputs: one encodes the amount of reward (or utility) predicted; the other encodes the amount of reward (or utility) experienced. The neurons themselves compute this difference using a biophysical form of subtraction and then broadcast the numerical result of their computation—by releasing dopamine—throughout the front half of the brain.

The dopamine concentration thus offers a kind of teaching signal that can be used to update representations of how much actions, goods, and situations are worth. Like the value system mentioned above, of which it is a crucial part, dopamine neurons encode these reward prediction errors for all types of rewards (and punishments) on a single common scale. Thus, again, in a way not entirely unlike the models of neoclassical and behavioral economics, reward prediction errors closely resemble the economic notion of a utility shock (Rutledge et al. 2010).

These reward prediction error signals are then used in a variety of ways to update the internal representations of value that guide choice prospectively (Doll et al. 2014). The simplest of these mechanisms is often referred to as reinforcement learning, and it is simply the process of gradually incrementing or decrementing the strength of synapses in the brain with dopamine, until the strength of those synapses matches the experienced value of the good or action being encoded by that synapse. More complicated mechanisms use this same signal to test and improve models of the causal structure of the environment, which are also stored as synaptic strengths.

While we do not know for certain, it seems highly likely that there are dozens of algorithmic mechanisms for building estimates of the values of choice objects in different ways. Just as we know that there are many specialized brain areas that contribute to the aggregate value signal, we know that each of these mechanisms uses slightly different algorithms to extract value estimates from experience. Most, if not all, of these mechanisms seem to rely on dopamine but do so in different ways. This enforces both the great commonality and the great heterogeneity of the human decision-making system. The great commonality is the use of a common currency that is the product of (among other things) the universality of dopamine. The heterogeneity of mechanisms for extracting meaning from the dopamine signal is likely what accounts for the many different neurobiological modules that contribute to aggregate valuation.

The existence of a single common currency does not, however, imply that this common currency yields a perfectly stable valuation for all choice objects. Indeed the existing data argues against this conclusion. The bulk of the existing evidence suggests that this common value is quite reference dependent (Louie and De Martino 2013); just as in choice behavior, the dopaminergic value attributed to a given outcome depends on the context in which it is presented (Tobler et al. 2005). This is a critical point. Neural representations of value are highly context dependent.

Decisions about Which We Are Largely Ignorant

While enormous progress has been made in laying many of the basic features of the human architecture for decision making, we know almost nothing about decision making that is guided by *symbolic computation*. To understand this

concept, consider a situation in which a human subject is asked to affix a value to a stock in a stock market. Our decision maker has two basic ways to do this. In the simplest possible mechanism, the decision maker simply experiences the rises and falls of the stock value each day and derives, from trial-and-error experience, an estimate of what the stock “feels like” it will be worth tomorrow. Such a decision maker can be shown to rely on the kinds of dopaminergic valuation mechanisms discussed above, and we know much about how she sets her estimate of value and how she will decide whether or not to buy that stock.

Alternatively, our decision maker could resort to a very symbolic and trained mechanism for valuing the stock. She might gather large quantities of numerical data about the company, aggregate that data by computing derivatives and integrals using tools from calculus, and then produce, using these symbolic mathematical methods, a valuation and a subsequent choice. It is very important to note that once our decision maker starts using symbolic mathematical tools to derive a value, we as neurobiologists know almost nothing about how her brain is operating. While we do know something about how numerical value is represented in the brain, we are only just beginning to understand how to apply the models of decision making we have developed to these kinds of situations (Dehaene and Sigman 2012). Fortunately, however, circumstances in which symbolic calculation guides decision making are rare in simple consumer choice, the kind of decision making that has been most broadly studied by neurobiologists.

Summary

During the first half of the twentieth century, neoclassical economists put forward mathematically beautiful theories of human decision. These theories rested on the notion that human decision makers were highly consistent in their behavior, and they modeled our decisions by proposing that we represent some internal variable that encodes just how valuable a good or service is to us. These internal values, or expected utilities, were then assumed to be compared numerically so as to identify the option or action having the greatest desirability. By the end of the twentieth century, however, it was clear that humans are not reliably consistent in their choice behavior. Although behavioral economists retained many of the features of neoclassical theory, they accounted for these inconsistencies by proposing a number of alternative algorithms.

Over the last two decades, the actual algorithms by which the human brain makes decisions have finally begun to be elucidated and, perhaps unexpectedly, there are clear parallels in the structural algorithms employed by the brain to many of the ideas put forth by the neoclassical and behavioral economists of the twentieth century. The values of options in a choice set do seem to be represented in a single common currency and compared during the choice process; however, that single common currency representation reflects the input

of many weakly independent brain modules interacting in a highly context-dependent way. These modules learn and store values based on our experience, but extract meaning from those experiences in very different ways. The result is a heterogeneity of valuation systems that interact in a largely cooperative and context-dependent way to yield the often fluid estimates of value that guide our choices. Fortunately, the precise mechanisms employed by each of these valuation areas are the subject of very rapidly accumulating research, and many of these systems are beginning to be deeply understood.

The next couple of decades should see a dramatic refinement in the predictive power of neurobiological models of choice behavior. Even today, some of these models are beginning to dramatically outperform more traditional, purely behavioral models. While a detailed summary of these models and mechanisms lies beyond the scope of this chapter, it is hoped that this brief overview will offer some insight into the mechanisms and algorithms employed by the human brain to produce our choices. One cannot help but note that understanding these mechanisms is an endeavor that lies outside the scope of traditional neoclassical economics. Indeed, many economists would argue that inquiry into the causal mechanism of choice is anathema to economics. That is, of course, a reasonable position but not one likely to be taken by neuroeconomists seeking to understand the proximate mechanisms by which we choose.