Neuron
Review

The Neurobiology of Decision: Consensus and Controversy

Joseph W. Kable[1,*] and Paul W. Glimcher[2,*]
[1]Department of Psychology, University of Pennsylvania, Philadelphia, PA 19104, USA
[2]Center for Neuroeconomics, New York University, New York, NY 10003, USA
*Correspondence: kable@psych.upenn.edu (J.W.K.), glimcher@cns.nyu.edu (P.W.G.)
DOI 10.1016/j.neuron.2009.09.003

We review and synthesize recent neurophysiological studies of decision making in humans and nonhuman primates. From these studies, the basic outline of the neurobiological mechanism for primate choice is beginning to emerge. The identified mechanism is now known to include a multicomponent valuation stage, implemented in ventromedial prefrontal cortex and associated parts of striatum, and a choice stage, implemented in lateral prefrontal and parietal areas. Neurobiological studies of decision making are beginning to enhance our understanding of economic and social behavior as well as our understanding of significant health disorders where people's behavior plays a key role.

Introduction

Only seven years have passed since *Neuron* published a special issue entitled "Reward and Decision," an event that signaled a surge of interest in the neural mechanisms underlying decision making that continues to this day (Cohen and Blum, 2002). At the time, many scholars were excited that quantitative formal models of choice behavior—from economics, evolutionary biology, computer science, and mathematical psychology—were beginning to provide a fruitful framework for new and more detailed investigations of the neural mechanisms of choice. To borrow David Marr's (1982) famous typology for computational studies of the brain, decision scholars seemed for the first time poised to investigate decision making at the theoretical, algorithmic, and implementation levels simultaneously.

Since that time, hundreds of research papers have been published on the neural mechanisms of decision making, at least two new societies dedicated to the topic have been formed (the Society for Neuroeconomics and the Association for NeuroPsychoEconomics), and a basic textbook for the field has been introduced (Glimcher et al., 2009). In this review, we survey some of the scientific progress that has been made in these past seven years, focusing specifically on neurophysiological studies in primates and including closely related work in humans. In an effort to achieve brevity, we have been selective. Our aim is to provide one synthesis of the neurophysiology of decision making, as we understand it. While many issues remain to be resolved, our conviction is that the available data suggest the basic outlines of the neural systems that algorithmically produce choice. Although there are certainly vigorous controversies, we believe that most scientists in the field would exhibit consensus over (at least) the terrain of contemporary debate.

The Basic Mechanism

Any neural model of decision making needs to answer two key questions. First, how are the subjective values of the various options under consideration learned, stored, and represented? Second, how is a single highly valued action chosen from among the options under consideration to be implemented by the motor circuitry? Below, we review evidence that we interpret as suggesting that valuation involves the ventromedial sectors of the prefrontal cortex and associated parts of the striatum (likely as a final common path funneling information from many antecedent areas), while choice involves lateral prefrontal and parietal areas traditionally viewed as intermediate regions in the sensory-motor hierarchy. Based on these data, we argue that a "basic model" for primate decision making is emerging from recent investigations, which involves the coordinated action of these two circuits in a two-stage algorithm.

Before proceeding, we should be clear about the relationship between this neurophysiological model of choice and the very similar theoretical models in economics from which it is derived. Traditional economic models aim only to predict (or explain) an individual's observable choices. They do not seek to explain the (putatively unobservable) process by which those choices are generated. In the famous terminology of Milton Friedman, traditional economic models are conceived of as being "*as if*" models (Friedman, 1953). Classic proofs in utility theory (i.e., Samuelson, 1937), for example, demonstrate that any decision maker who chooses in a mathematically consistent fashion behaves *as if* they had first constructed and stored a single list of the all possible options ordered from best to worst, and then in a second step had selected the highest ordered of those available options. Friedman and nearly all of the neoclassical economists who followed him were explicit that the concept of utility was not meant to apply to anything about the algorithmic or implementation levels. For Friedman, and for many contemporary economists (i.e., Gul and Pesendorfer, 2008), whether or not there are "neurophysiological correlates of utility" is, by construction, irrelevant.

Neurophysiological models, of course, aim to explain the mechanisms by which choices are generated, as well as the choices themselves. These models seek to explain both behavior and its causes and employ constraints at the algorithmic level to validate the plausibility of behavioral predictions. One might call these models, which are concerned with algorithm and implementation as well as with behavior, *because* models. Although

a fierce debate rages in economic circles about the validity and usefulness of *because* models, we take as a given for the purposes of this review that describing the mechanism of primate (both human and nonhuman) decision making will yield new insights into behavior, just as studies of the primate visual system have revolutionized our understanding of perception.

Stage 1: Valuation

Ventromedial Prefrontal Cortex and Striatum: A Final Common Path

Most decision theories—from expected utility theory in economics (von Neumann and Morgenstern, 1944) to prospect theory in psychology (Kahneman and Tversky, 1979) to reinforcement learning theories in computer science (Sutton and Barto, 1998)—share a core conclusion. Decision makers integrate the various dimensions of an option into a single measure of its idiosyncratic subjective value and then choose the option that is most valuable. Comparisons between different kinds of options rely on this abstract measure of subjective value, a kind of "common currency" for choice. That humans can in fact compare apples to oranges when they buy fruit is evidence for this abstract common scale.

At first blush, the notion that all options can be represented on a single scale of desirability might strike some as a peculiar idea. Intuitively it might feel like complicated choices among objects with many different attributes would resist reduction to a single dimension of desirability. However, as Samuelson showed over half a century ago (Samuelson, 1937), any individual whose choices can be described as internally consistent can be perfectly modeled by algorithms that employ a single common scale of desirability. If someone selects an apple when they could have had an orange, and an orange when they could have had a pear, then (assuming they are in the same state) they should not select a pear when they could have had an apple instead. This is the core notion of consistency, and when people behave in this manner, we can model their choices as arising from a single, consistent "utility" ordering over all possible options.

For traditional economic theories, however, consistent decision makers only choose *as if* they employed a single hidden common currency for comparing options. There was no claim when these theories were first advanced that subjective representations of value were used at the algorithmic level during choice. However, there is now growing evidence that subjective value representations do in fact play a role at the neural algorithmic level and that these representations are encoded primarily in ventromedial prefrontal cortex and striatum (Figure 1).

One set of studies has documented responses in orbitofrontal cortex related to the subjective values of different rewards, or in the language of economics, "goods." Padoa-Schioppa and Assad (2006) recorded from area 13 of the orbitofrontal cortex while monkeys chose between pairs of juices. The amount of each type of juice offered to the animals varied from trial to trial, and the types of juices offered changed across sessions. Based on each monkey's actual choices, they calculated a subjective value for each juice reward, based on type and quantity of juice, which could explain these choices as resulting from a common value scale. They then searched for neurons that showed

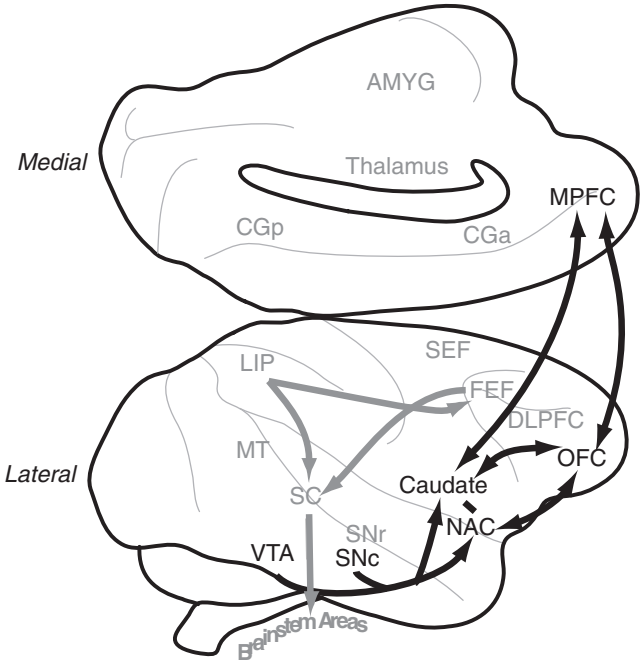

Figure 1. Valuation Circuitry
Diagram of a macaque brain, highlighting in black the regions discussed as playing role in valuation. Other regions are labeled in gray.

evidence of this hypothesized common scale for subjective value. They found three dominant patterns of responding, which accounted for 80% of the neuronal responses in this region. First and most importantly they identified *offer value* neurons, cells with firing rates that were linearly correlated with the subjective value of one of the offered rewards, as computed from behavior (Figure 2). Second, they observed *chosen value* neurons, which tracked the subjective value of the chosen reward in a single common currency that was independent of type of juice. Finally, they observed *taste* neurons, which showed a categorical response when a particular juice was chosen. All of these responses were independent of the spatial arrangement of the stimuli and of the motor response produced by the animal to make its choice. Perhaps unsurprisingly, offer value and chosen value responses were prominent right after the options were presented and again at the time of juice receipt. Taste responses, in contrast, occurred primarily after the juice was received.

Based on their timing and properties, these different responses likely play different roles during choice. Offer value signals could serve as subjective values, in a single common neuronal currency, for comparing and deciding between offers. They are exactly the kind of value representation posited by most decision theories, and they could be analogous to what economists call "utilities" (or, if they also responded to probabilistically delivered rewards to "expected utilities") and to what psychologists call "decision utilities." Chosen values, by contrast, can only be calculated *after* a choice has been made and, thus, could not be the basis for a decision. As discussed in the section below, however, learning the value of an action from experience depends on being able to compare two quantities—the

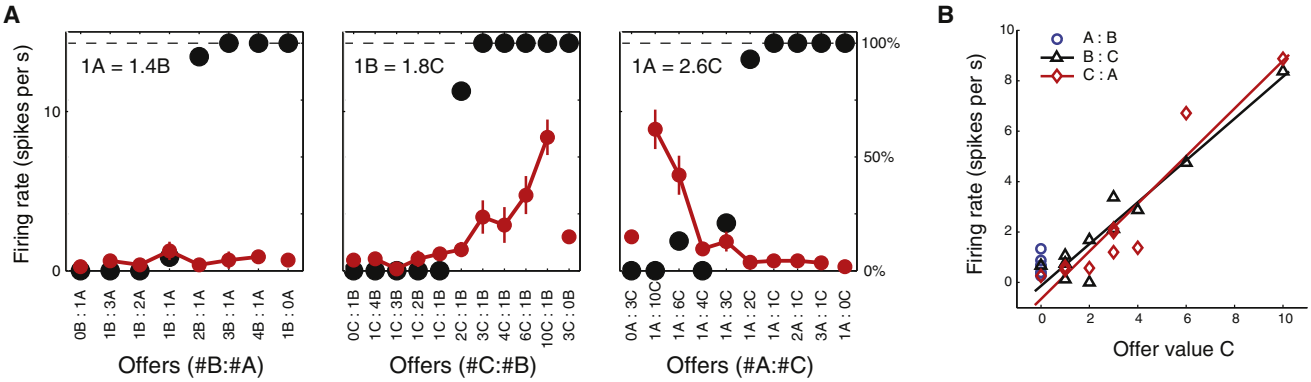

Figure 2. An Example Orbitofrontal Neuron that Encodes Offer Value, in a Menu-Invariant and Therefore Transitive Manner

(A) In red is the firing rate of the neuron (±SEM), as a function of the magnitude of the two juices offered, for three different choice pairs. In black is the percentage of time the monkey chose the first offer.

(B) Replots firing rates as a function of the offer value of juice C, demonstrating that this neuron encodes this value in a common currency in a manner that is independent of the other reward offered. The different symbols and colors refer to data from the three different juice pairs, and each symbol represents one trial type. Reprinted with permission from Padoa-Schioppa and Assad (2008).

forecasted value of taking an action and the actual value experienced when that action was taken. Chosen value responses in the orbitofrontal cortex may then signal the *forecast*, or for neurons active very late in the choice process the *experienced*, subjective value from that choice (see also Takahashi et al., 2009, for discussion of this potential function of orbitofrontal value representations).

In a follow-up study, Padoa-Schioppa and Assad (2008) extended their conclusion that these neurons provide utility-like representations. They demonstrated that orbitofrontal responses were "menu invariant"—that activity was internally consistent in the same way that the choices of the monkeys were internally consistent. In that study, choice pairs involving three different kinds of juice were interleaved from trial to trial. Behaviorally, the monkeys' choices obeyed *transitivity*: if the animal preferred apple juice over grape juice, and grape juice over tea, then he also preferred apple juice over tea. They observed the same three kinds of neuronal responses as in their previous study, and these responses did not depend on the other option offered on that trial (Figure 2). For example, a neuron that encoded the offer value of grape juice did so in the same manner whether the other option was apple juice or tea. This independence can be shown to be required of utility-like representations (Houthakker, 1950) and thus strengthens the conclusion that these neurons may encode a common currency for choice.

Importantly, Padoa-Schioppa and Assad (2008) distinguished the "menu invariance" that they observed, where neuronal responses do not change from trial to trial as the other juice offered changes, from a longer-term kind of stability they refer to as "condition invariance." Tremblay and Schultz's (1999) data suggest that orbitofrontal responses may not be "condition invariant," since these responses seem to adjust to the range of rewards when this range is stable over long blocks of trials. As Padoa-Schioppa and Assad (2008) argued, such longer-term rescaling would serve the adaptive function of allowing orbitofrontal neurons to adjust across conditions so as to encode value across their entire dynamic range. However, in discussing this study and the ones below, we focus primarily on the question of whether neuronal responses are "menu invariant," i.e., whether they adjust dynamically from trial to trial depending on what other options are offered.

Another set of studies from two labs has documented similar responses in the striatum, the second area that appears to represent the subjective values of choice options. Lau and Glimcher (2008) recorded from the caudate nucleus while monkeys performed an oculomotor choice task (Figure 3A). The task was based on the concurrent variable-interval schedules used to study Herrnstein's matching law (Herrnstein, 1961; Platt and Glimcher, 1999; Sugrue et al., 2004). Behaviorally, the monkeys dynamically adjusted the proportion of their responses to each target to match the relative magnitudes of the rewards earned for looking at those targets. Recording from phasically active striatal neurons (PANs), they found three kinds of task-related responses closely related to the orbitofrontal signals of Padoa-Schioppa and Assad (2006, 2008): *action value neurons*, which tracked the value of one of the actions, independent of whether it was chosen; *chosen value neurons*, which tracked the value of a chosen action; and *choice neurons*, which produced a categorical response when a particular action was taken. Action value responses occurred primarily early in the trial, at the time of the monkey's choice, while chosen value responses occurred later in the trial, near the time of reward receipt.

Samejima and colleagues (2005) provided important impetus for all of these studies when they gathered some of the first evidence that the subjective value of actions was encoded on a common scale (Figure 3B). In that study, monkeys performed a manual choice task, turning a lever leftward or rightward to obtain rewards. Across different blocks, the probability that each turn would be rewarded with a large (as opposed to a small) magnitude of juice was changed. Recording from the putamen, they found that one-third of all modulated neurons tracked *action value*. This was almost exactly the same percentage of action value neurons that Lau and Glimcher (2008) later found in the oculomotor caudate. Samejima and colleagues' design also allowed them to show that these responses did not depend

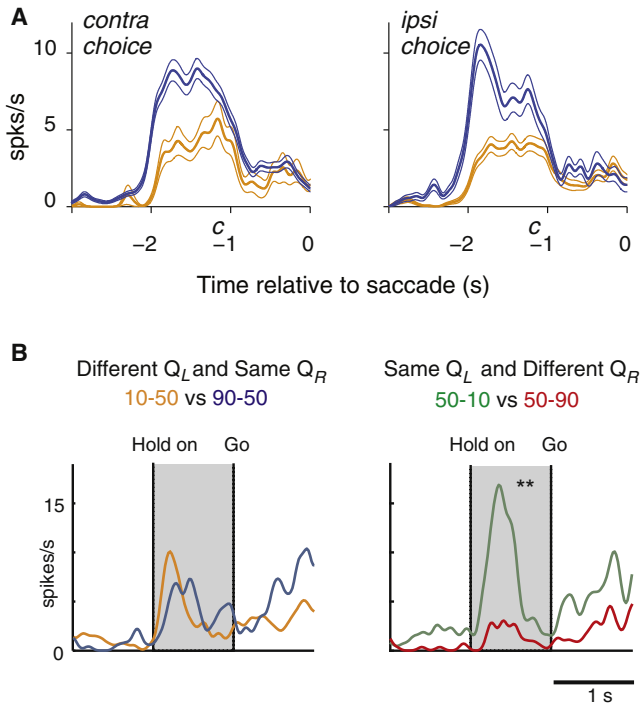

Figure 3. Two Example Striatal Neurons that Encode Action Value
(A) Caudate neuron that fires more when a contralateral saccade is more valuable (blue) compared to less valuable (yellow), independently of which saccade the animal eventually chooses. *c* denotes the average onset of the saccade cue, and the thin lines represent ±1 SEM. Reprinted with permission from Lau and Glimcher (2008).
(B) Putamen neuron that encodes the value of a rightward arm movement (Q_R), independent of the value of a leftward arm movement (Q_L). Reprinted with permission from Samejima et al. (2005).

on the value associated with the other action. For example, a neuron that tracked the value of a right turn would always exhibit an intermediate response when that action yielded a large reward with 50% probability, independent of whether the left turn was more (i.e., 90% probability) or less (i.e., 10% probability) valuable. This is critical because it means that the striatal signals, like the signals in orbitofrontal cortex, likely show the kind of consistent representation required for transitive behavior.

Thus, the responses in the caudate and putamen in these two studies mirror those found in orbitofrontal cortex, except anchored to the actions produced by the animals rather than to a more abstract goods-based framework as observed in orbitofrontal cortex. One key question raised by these findings is the relationship between the action-based value responses observed in the striatum and the goods-based value responses observed in the orbitofrontal cortex. The extent to which these representations are independent has received much attention recently. For example, Horwitz and colleagues (2004) have shown that asking monkeys to choose between "goods" that map arbitrarily to different actions from trial to trial leads almost instantaneously to activity in action-based choice circuits. Findings such as these suggest that action-based and goods-based representations of value are profoundly interconnected, although we acknowledge that this view remains controversial (Padoa-Schioppa and Assad, 2006, 2008).

Human imaging studies have provided strong converging evidence that ventromedial prefrontal cortex and the striatum encode the subjective value of goods and actions. While it is difficult to determine whether the single-unit neurophysiology and fMRI studies have identified *directly homologous* subregions of these larger anatomical structures in the two different species, there is surprising agreement across the two methods concerning the larger anatomical structures important in valuation. As reviewed elsewhere, dozens of studies have demonstrated reward responses in these regions that are consistent with tracking forecasted or experienced value, which might play a role in value learning (Delgado, 2007; Knutson and Cooper, 2005; O'Doherty, 2004). Here, we will focus on several recent studies that identified subjective value signals specific to the decision process. Two key design aspects that allow this identification in these particular studies are (1) no outcomes were experienced during the experiment, so that decision-related signals could be separated from learning-related signals as much as possible, and (2) there was a behavioral measure of the subject's preference, which allowed subjective value to be distinguished from the objective characteristics of the options.

Plassmann and colleagues (2007) scanned hungry subjects bidding on various snack foods (Figure 4). They used an auction procedure where subjects were strongly incentivized to report what each snack food was actually worth to them. They found that BOLD activity in medial orbitofrontal cortex was correlated with the subject's subjective valuation of that item (Figures 4B and 4C). Hare and colleagues have now replicated this finding twice, once in a different task where the subjective value of the good could be dissociated from other possible signals (Hare et al., 2008) and again in a set of dieting subjects where the subjective value of the snack foods was affected by both taste and health concerns (Hare et al., 2009).

In a related study, Kable and Glimcher (2007) examined participants choosing between immediate and delayed monetary rewards. The immediate reward was fixed, while both the magnitude and receipt time of the delayed reward varied across trials. From each subject's choices, an idiosyncratic *discount function* was estimated that described how the subjective value of money declined with delay for that individual. In medial prefrontal cortex and ventral striatum (among other regions), BOLD activity was correlated with the subjective value of the delayed reward as it varied across trials. Furthermore, across subjects, the neurometric discount functions describing how neural activity in these regions declined with delay matched the psychometric discount functions describing how subjective value declined with delay (Figure 5). In other words, for more impulsive subjects, neural activity in these regions decreased steeply as delay increased, while for more patient subjects this decline was less pronounced. These results suggest that neural activity in these regions encodes the subjective value of both immediate and delayed rewards in a common neural currency that takes into account the time at which a reward will occur.

Two recent studies have focused on decisions involving monetary gambles. These studies have demonstrated that modulation of a common value signal could also account for

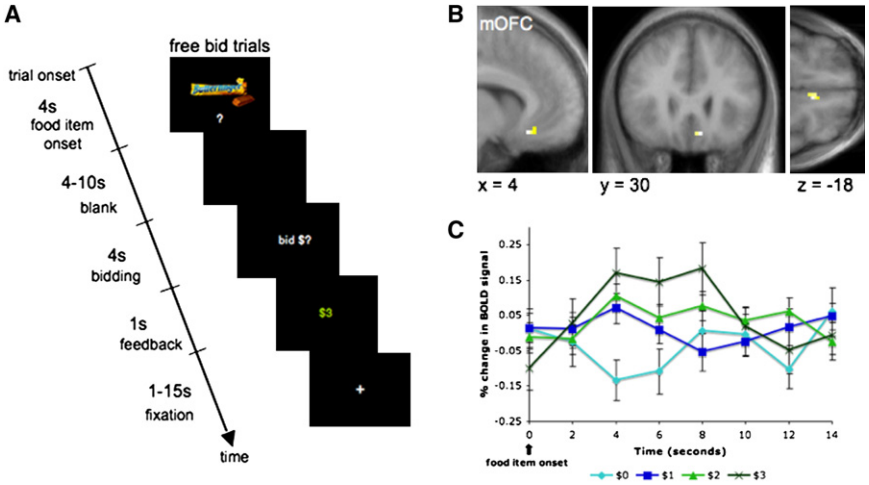

Figure 4. Orbitofrontal Cortex Encodes the Subjective Value of Food Rewards in Humans

(A) Hungry subjects bid on snack foods, which were the only items they could eat for 30 min after the experiment. At the time of the decision, medial orbitofrontal cortex (B) tracked the subjective value that subjects placed on each food item. Activity here increased as the subjects' willingness to pay for the item increased (C). Error bars denote standard errors. Reprinted with permission from Plassmann et al. (2007).

loss aversion and *ambiguity aversion*, two more recently identified choice-related behaviors that suggest important refinements to theoretical models of subjective value encoding (for a review of these issues see Fox and Poldrack, 2009). Tom and colleagues (2007) scanned subjects deciding whether to accept or reject monetary lotteries in which there was a 50/50 chance of gaining or losing money. They found that BOLD activity in ventromedial prefrontal cortex and striatum increased with the amount of the gain and decreased with the amount of the loss. Furthermore, the size of the loss effect relative to the gain effect was correlated with the degree to which the potential loss affected the person's choice more than the potential gain. Activity decreased faster in response to increasing losses for more *loss-averse* subjects. Levy and colleagues (2007) examined subjects choosing between a fixed certain amount and a gamble that was either risky (known probabilities) or ambiguous (unknown probabilities). They found that activity in ventromedial prefrontal cortex and striatum was correlated with the subjective value of both risky and ambiguous options.

Midbrain Dopamine: A Mechanism for Learning Subjective Value

The previous section reviewed evidence that ventromedial prefrontal cortex and striatum encode the subjective value of different goods or actions during decision making in a way that could guide choice. But how do these subjective value signals arise? One of the most critical sources of value information is undoubtedly past experience. Indeed, in physiological experiments, animal subjects always have to *learn* the value of different actions over the course of the experiment—for these subjects,

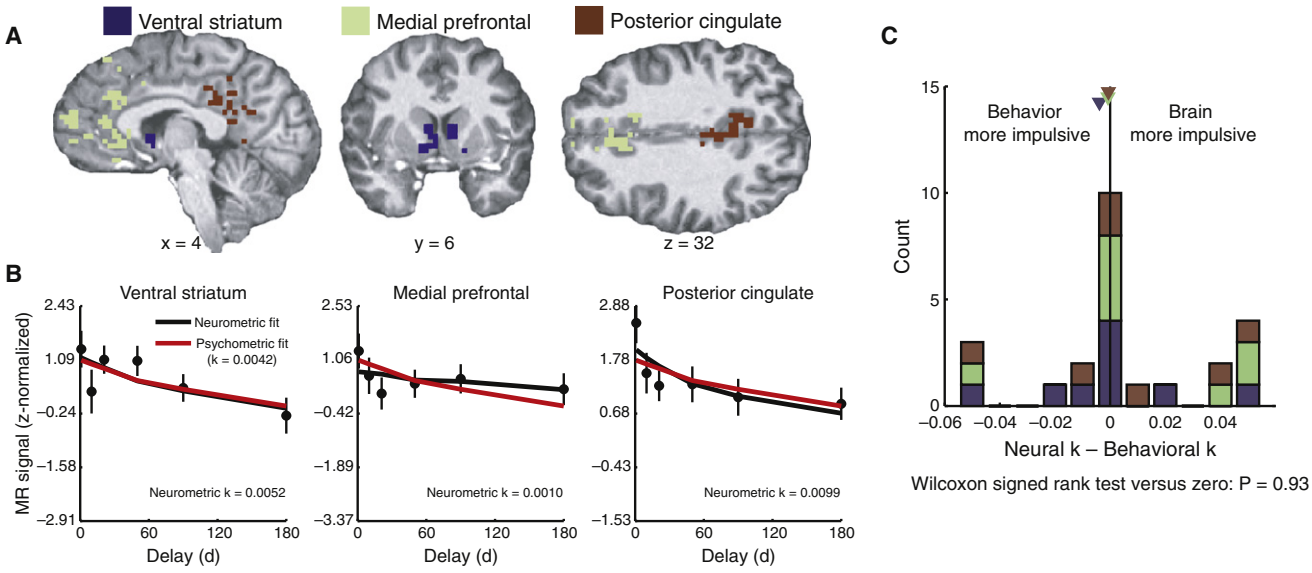

Figure 5. Match between Psychometric and Neurometric Estimates of Subjective Value during Intertemporal Choice

(A) Regions of interest are shown for one subject, in striatum, medial prefrontal cortex, and posterior cingulate cortex.

(B) Activity in these ROIs (black) decreases as the delay to a reward increases, in a similar manner to the way that subjective value estimated behaviorally (red) decreases as a function of delay. This decline in value can be captured by estimating a discount rate (k). Error bars represent standard errors.

(C) Comparison between discount rates estimated separately from the behavioral and neural data across all subjects, showing that on average there is a psychometric-neurometric match. Reprinted with permission from Kable and Glimcher (2007).

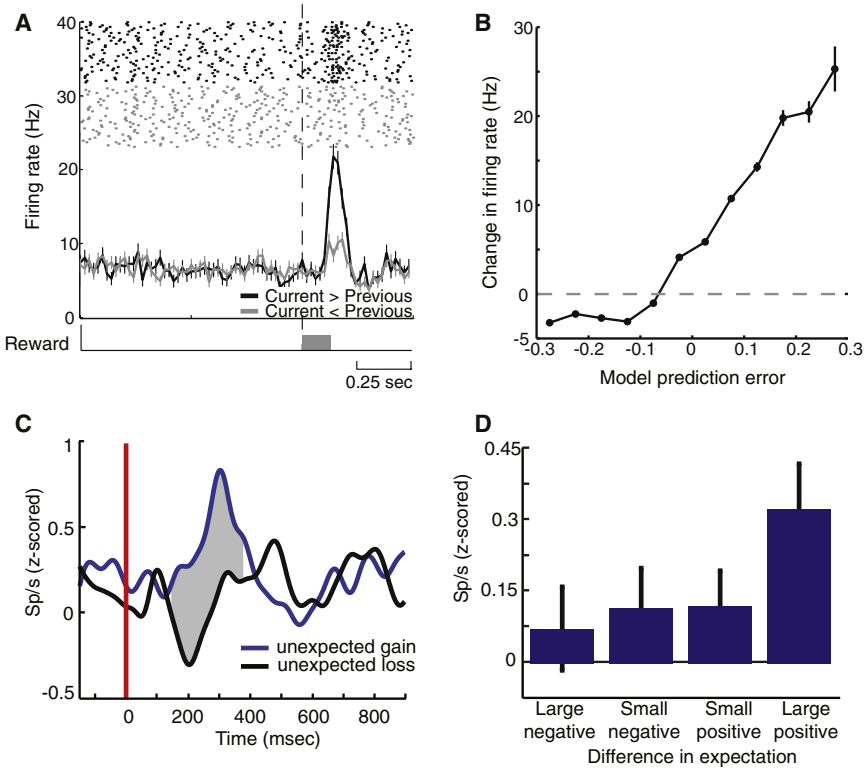

Figure 6. Dopaminergic Responses in Monkeys and Humans

(A) An example dopamine neuron recorded in a monkey, which responds more when the reward received was better than expected.

(B) Firing rates of dopaminergic neurons track positive reward prediction errors.

(C) Population average of dopaminergic responses (n = 15) recorded in humans during deep-brain stimulation (DBS) surgery for Parkinson's disease, showing increased firing in response to unexpected gains. The red line indicates feedback onset.

(D) Firing rates of dopaminergic neurons depend on the size and valence of the difference between the received and expected reward. All error bars represent standard errors.

Panels (A) and (B) reprinted with permission from Bayer and Glimcher (2005), and panels (C) and (D) reprinted with permission from Zaghloul et al. (2009).

the consequences of each action cannot be communicated linguistically. Although there are alternative viewpoints (Dommett et al., 2005; Redgrave and Gurney, 2006), unusually solid evidence now indicates that dopaminergic neurons in the midbrain encode a teaching signal that can be used to learn the subjective value of actions (for a detailed review, including a discussion of how nearly all of the findings often presented as discrepant with early versions of the dopaminergic teaching signal hypothesis have been reconciled with contemporary versions of the theory, see Niv and Montague, 2009). Indeed, these kinds of signals can be shown to be sufficient for learning the values of different actions from experience. Since these same dopaminergic neurons project primarily to prefrontal and striatal regions (Haber, 2003), it seems likely that these neurons play a critical role in subjective value learning.

The computational framework for these investigations of dopamine and learning comes from reinforcement learning theories developed in computer science and psychology over the past two decades (Niv and Montague, 2009). While several variants of these theories exist, in all of these models, subjective values are learned through iterative updating based on experience. The theories rest on the idea that each time a subject experiences the outcome of her choice, an updated value estimate is calculated from the old value estimate and a *reward prediction error*—the difference between the experienced outcome of an action and the outcome that was forecast. This reward prediction error is scaled by a *learning rate*, which determines the weight given to recent versus remote experience.

Pioneering studies of Schultz and colleagues (1997) provided the initial evidence that dopaminergic neurons encode a reward prediction error signal of the kind proposed by a class of theories called *temporal-difference learning* (TD-models; Sutton and Barto, 1998). These studies demonstrated that, during conditioning tasks, dopaminergic neurons (1) responded to the receipt of unexpected rewards, (2) responded to the first reliable predictor of reward after conditioning, (3) did not respond to the receipt of fully predicted rewards, and (4) showed a decrease in firing when a predicted reward was omitted. Montague et al. (1996) were the first to propose that this pattern of results could be completely explained if the firing of dopamine neurons encoded a reward prediction error of the type required by TD-class models. Subsequent studies, examining different Pavlovian conditioning paradigms, demonstrated that the qualitative responses of dopaminergic neurons were entirely consistent with this hypothesis (Tobler et al., 2003; Waelti et al., 2001).

Recent studies have provided more quantitative tests of the reward prediction error hypothesis. Bayer and Glimcher (2005) recorded from dopaminergic neurons during an oculomotor task, in which the reward received for the same movement varied in a continuous manner from trial to trial. As is required by theory, the response on the current trial was a function of an exponentially weighted sum of previous rewards obtained by the monkey. Thus, dopaminergic firing rates were linearly related to a model-derived reward prediction error (Figures 6A and 6B). Interestingly, though, this relationship broke down for the most negative prediction error signals, although the implications of this last finding have been controversial.

Additional studies have demonstrated that, when conditioned cues predict rewards with different magnitudes or probabilities, the cue-elicited dopaminergic response scales with magnitude or probability, as expected if it represents a cue-elicited prediction error (Fiorillo et al., 2003; Tobler et al., 2005). In a similar manner, if different cues predict rewards after different delays, the cue-elicited response decreases as the delay to reward increases, consistent with a prediction that incorporates

discounting of future rewards (Fiorillo et al., 2008; Kobayashi and Schultz, 2008; Roesch et al., 2007).

Until recently, direct evidence regarding the activity of dopaminergic neurons in humans has been scant. Imaging the midbrain with fMRI is difficult for several technical reasons, and the reward prediction error signals initially identified with fMRI were located in the presumed striatal targets of the dopaminergic neurons (McClure et al., 2003; O'Doherty et al., 2003). However, D'Ardenne and colleagues (2008) recently reported BOLD prediction error signals in the ventral tegmental area using fMRI. They used a combination of small voxel sizes, cardiac gating, and a specialized normalization procedure to detect these signals. Across two paradigms using primary and secondary reinforcers, they found that BOLD activity in the VTA was significantly correlated with positive, but not negative, reward prediction errors.

Zaghloul and colleagues (2009) reported the first electrophysiological recordings in human substantia nigra during learning. These investigators recorded neuronal activity while individuals with Parkinson's disease underwent surgery to place electrodes for deep-brain stimulation therapy. Subjects had to learn which of two options provided a greater probability of a hypothetical monetary reward, and their choices were fit with a reward prediction model. In the subset of neurons that were putatively dopaminergic, they found an increase in firing rate for unexpected positive outcomes, relative to unexpected negative outcomes, while the firing rates for expected outcomes did not differ (Figures 6C and 6D). Such an encoding of unexpected rewards is again consistent with the reward prediction error hypothesis.

Pessiglione and colleagues (2006) demonstrated a causal role for dopaminergic signaling in both learning and striatal BOLD prediction error signals. During an instrumental learning paradigm, they tested subjects who had received L-DOPA (a dopamine precursor), haloperidol (a dopamine receptor antagonist), or placebo. Consistent with other findings from Parkinson's patients (Frank et al., 2004), L-DOPA (compared to haloperidol) improved learning to select a more rewarding option but did not affect learning to avoid a more punishing option. In addition, the BOLD reward prediction error in the striatum was larger for the L-DOPA group than for the haloperidol group, and differences in this response, when incorporated into a reinforcement learning model, could account for differences in the speed of learning across groups.

Stage 2: Choice

Lateral Prefrontal and Parietal Cortex: Choosing Based on Value

Learning and encoding subjective value in a common currency is not sufficient for decision making—one action still needs to be chosen from among the set of alternatives and passed to the motor system for implementation. What is the process by which a highly valued option in a choice set is selected and implemented?

While we acknowledge that other proposals have been made regarding this process, we believe that the bulk of the available evidence implicates (at a minimum) the lateral prefrontal and parietal cortex in the process of selecting and implementing choices from among any set of available options. Some of the best evidence has come from studies of a well-understood model decision making system: the visuo-saccadic system of the monkey. For largely technical reasons, the saccadic-control system has been intensively studied over the past three decades as a model for understanding sensory-motor control in general (Andersen and Buneo, 2002; Colby and Goldberg, 1999). The same has been true for studies of choice. The lateral intraparietal area (LIP), the frontal eye fields (FEF), and the superior colliculus (SC) comprise the core of a heavily interconnected network that plays a critical role in visuo-saccadic decision making (Figure 7) (Glimcher, 2003; Gold and Shadlen, 2007). The available data suggest a parallel role in nonsaccadic decision making for the motor cortex, the premotor cortex, the supplementary motor area, and the areas in the parietal cortex adjacent to LIP.

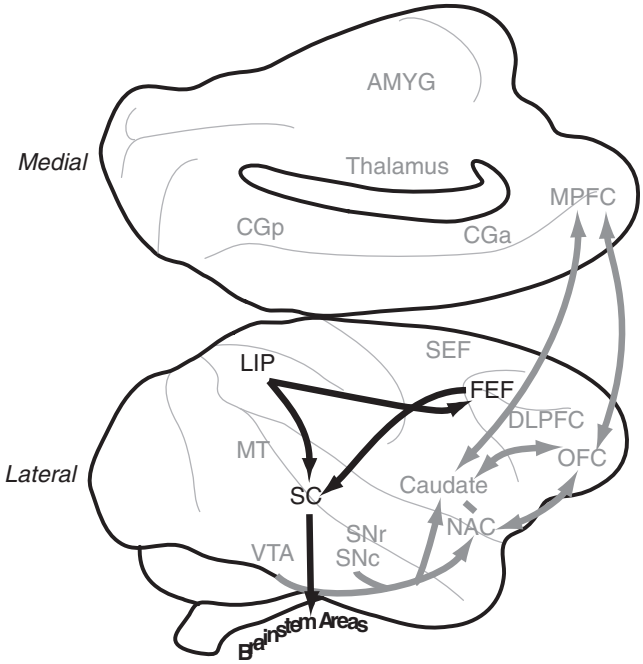

Figure 7. Choice Circuitry for Saccadic Decision Making
Diagram of a macaque brain, highlighting in black the regions discussed as playing a role in choice. Other regions are labeled in gray.

At a theoretical level, the process of choice must involve a mechanism for comparing two or more options and identifying the most valuable of those options. Both behavioral evidence and theoretical models from economics make it clear that this process is also somewhat stochastic (McFadden, 1974). If two options have very similar subjective values, the less-desirable option may be occasionally selected. Indeed, the probability that these "errors" will occur is a smooth function of the similarity in subjective value of the options under consideration. How then is this implemented in the brain?

Any system that performed such a comparison must be able to represent the values of each option before a choice is made and then must effectively pass information about the selected option, but not the unselected options, to downstream circuits. In the saccadic system, among the first evidence for such a circuit came from the work of Glimcher and Sparks (1992), who

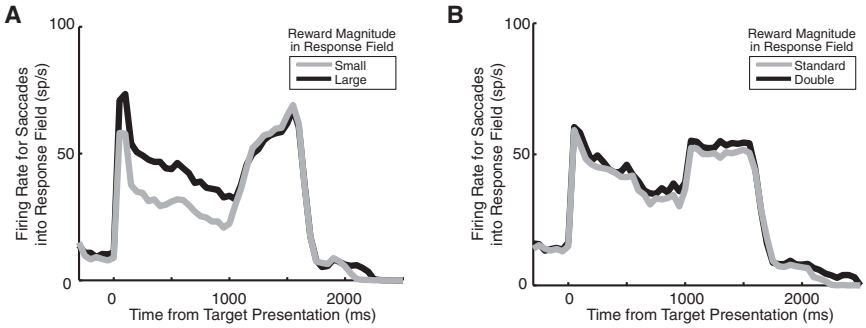

Figure 8. Relative Value Responses in LIP
LIP firing rates are greater when the larger magnitude reward is in the response field (n = 30) (A) but are not affected when the magnitude of all rewards are doubled (n = 22) (B). Adapted with permission from Dorris and Glimcher (2004).

essentially replicated in the superior colliculus Tanji and Evarts' classic studies of motor area M1 (Tanji and Evarts, 1976). The laminar structure of the superior colliculus employs a topographic map to represent the amplitude and direction of all possible saccades. They showed that if two saccadic targets of roughly equal subjective value were presented to a monkey, then the two locations on this map corresponding to the two saccades became weakly active. If one of these targets was suddenly identified as having higher value, this led almost immediately to a high-frequency burst of activity at the site associated with that movement and a concomitant suppression of activity at the other site. In light of preceding work (Van Gisbergen et al., 1981), this led to the suggestion that a winner-take-all computation occurred in the colliculus that effectively selected one movement from the two options for execution.

Subsequent studies (Basso and Wurtz, 1998; Dorris and Munoz, 1998) established that activity at the two candidate movement sites, during the period before the burst, was graded. If the probability that a saccade would yield a reward was increased, firing rates associated with that saccade increased, and if the probability that a saccade would yield a reward was decreased, then the firing rate was decreased. These observations led Platt and Glimcher (1999) to test the hypothesis, just upstream of the colliculus in area LIP, that these premovement signals encoded the subjective values of movements. To test that hypothesis, they systematically manipulated either the probability that a given saccadic target would yield a reward or the magnitude of reward yielded by that target. They found that firing rates in area LIP before the collicular burst occurred were a nearly linear function of both magnitude and probability of reward.

This naturally led to the suggestion that the fronto-parietal network of saccade control areas formed, in essence, a set of topographic maps of saccade value. Each location on these maps encodes a saccade of a particular amplitude and direction, and it was suggested that firing rates on these maps encoded the desirability of each of those saccades. The process of choice, then, could be reduced to a competitive neuronal mechanism that identified the saccade associated with the highest level of neuronal activity. (In fact, studies in brain slices have largely confirmed the existence of such a mechanism in the colliculus—see for example Isa et al., 2004, or Lee and Hall, 2006.) Many subsequent studies have bolstered this conclusion, demonstrating that various manipulations that increase (or decrease) the subjective value of a given saccade also increase (or decrease) the firing rate of neurons within the frontal-parietal maps associated with that saccade (Dorris and Glimcher, 2004; Janssen and Shadlen, 2005; Kim et al., 2008; Leon and Shadlen, 1999, 2003; Sugrue et al., 2004; Wallis and Miller, 2003; Yang and Shadlen, 2007). Some of these studies have discovered one notable caveat to this conclusion, though. Firing rates in these areas encode the subjective value of a particular saccade *relative* to the values of all other saccades under consideration (Figure 8) (Dorris and Glimcher, 2004; Sugrue et al., 2004). Thus, unlike firing rates in orbitofrontal cortex and striatum, firing rates in LIP (and presumably other frontal-parietal regions involved in choice rather than valuation) are *not* "menu invariant." This suggests an important distinction between activity in the parietal cortex and activity in the orbitofrontal cortex and striatum. Orbitofrontal and striatal neurons appear to encode absolute (and hence transitive) subjective values. Parietal neurons, presumably using a normalization mechanism like the one studied in visual cortex (Heeger, 1992), rescale these absolute values so as to maximize the differences between the available options before choice is attempted.

At the same time that these studies were underway, a second line of evidence also suggested that the fronto-parietal networks participate in decision making, but in this case decision making of a slightly different kind. In these studies of perceptual decision making, an ambiguous visual stimulus was used to indicate which of two saccades would yield a reward, and the monkey was reinforced if he made the indicated saccade. Shadlen and Newsome (Shadlen et al., 1996; Shadlen and Newsome, 2001) found that the activity of LIP neurons early in this decision-making process carried stochastic information about the likelihood that a given movement would yield a reward. Subsequent studies have revealed the dynamics of this process. During these kinds of perceptual decision-making tasks, the firing rates of LIP neurons increase as the evidence that a saccade into the response field will be rewarded accrues. However, this increase is bounded; once firing rates cross a maximal threshold, a saccade is initiated (Churchland et al., 2008; Roitman and Shadlen, 2002). Closely related studies in the frontal eye fields lead to similar conclusions (Gold and Shadlen, 2000; Kim and Shadlen, 1999). Thus, both firing rates in LIP and FEF and behavioral responses in this kind of task can be captured by a race-to-barrier diffusion model (Ratcliff et al., 1999).

Several lines of evidence now suggest that this threshold represents a value (or evidence based) threshold for movement selection (Kiani et al., 2008). When the value of any saccade crosses that preset minimum, the saccade is immediately initiated. Importantly, in the models derived from these data, the intrinsic stochasticity of the circuit gives rise to the stochasticity observed in actual choice behavior.

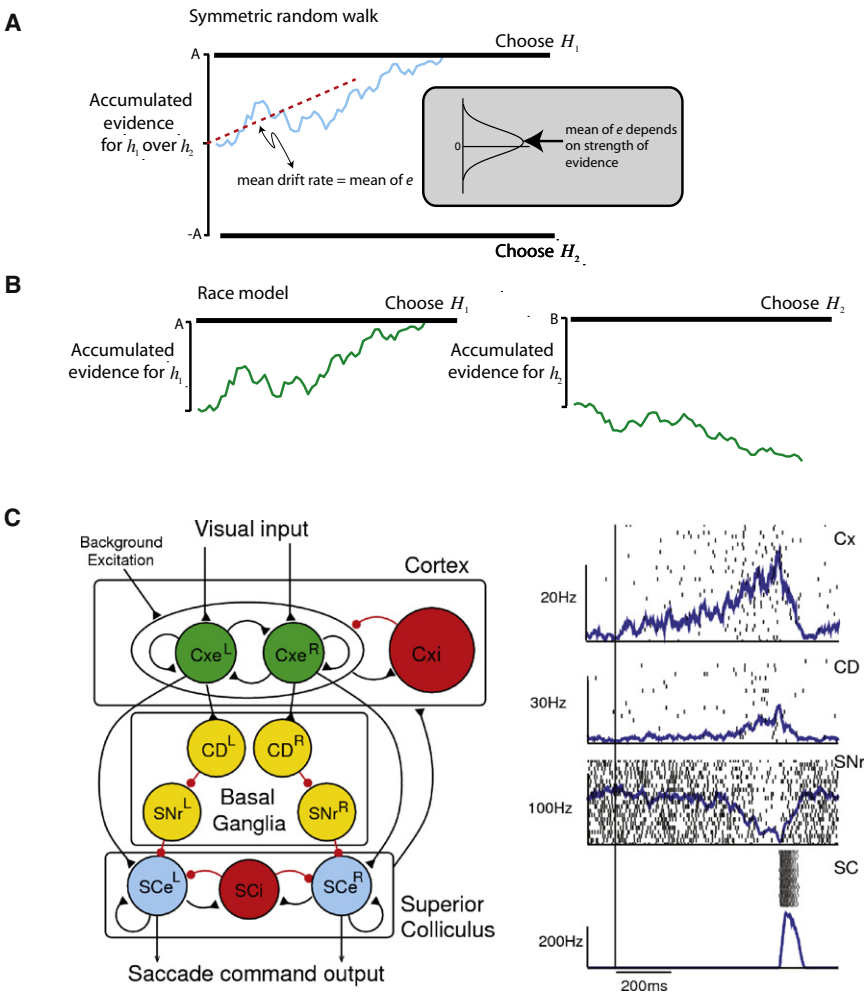

Figure 9. Theoretical and Computational Models of the Choice Process
Schematic of the symmetric random walk (A) and race models (B) of choice and reaction time. (C) Schematic neural architecture and simulations of the computational model that replicates activation dynamics in LIP and superior colliculus during choice. Panels (A) and (B) reprinted with permission from Gold and Shadlen (2007). Panel (C) adapted with permission from Lo and Wang (2006).

These two lines of evidence, one associated with the work of Shadlen, Newsome, and their colleagues and the other associated with our research groups, describe two classes of models for understanding the choice mechanism. In reaction-time tasks, the race-to-barrier model describes a situation in which a choice is made as soon as the value of any action exceeds a preset threshold (Figures 9A and 9B). In non-reaction time economics-style tasks, a winner-take-all model describes the process of selecting the option having the highest value from a set of candidates. Wang and colleagues (Lo and Wang, 2006; Wang, 2008; Wong and Wang, 2006) have recently shown that a single collicular or parietal circuit can be designed that performs both winner-take-all and thresholding operations in a stochastic fashion that depends on the inhibitory tone of the network (Figure 9C). Their models suggest that the same mechanism can perform two kinds of choice—a slow competitive winner-take-all process that can identify the best of the available options and a rapid thresholding process that selects a single movement once some preset threshold of value is crossed. LIP, the superior colliculus, and the frontal eye fields therefore seem to be part of a circuit that receives as input the subjective value of different saccades and then, representing these as relative values, stochastically selects from all possible saccades a single one for implementation.

Open Questions and Current Controversies

Above, we have outlined our conclusion that, based on the data available today, a minimal model of primate decision making includes valuation circuits in ventromedial prefrontal cortex and striatum and choice circuits in lateral prefrontal and parietal cortex. However, there are obviously many open questions about the details of this mechanism, as well as many vigorous debates that go beyond the general outline just presented. With regard to valuation, some of the important open questions concern what all of the many inputs to the final common path are (i.e., Hare et al., 2009), how the function of ventromedial prefrontal cortex and striatum might differ (i.e., Hare et al., 2008), and how to best define and delineate more specific roles for subcomponents of these large multipart structures. In terms of value learning, current work focuses on what precise algorithmic model of reinforcement learning best describes the dopaminergic signal (i.e., Morris et al., 2006), how sophisticated the expectation of the future rewards that is intrinsic to this signal is, whether these signals adapt to volatility in the environment as necessary for optimal learning (i.e., Behrens et al., 2007), and how outcomes that are worse than expected are encoded (i.e., Bayer et al., 2007; Daw et al., 2002). In terms of choice, some of the important open questions concern what modulates the state of the choice network between the thresholding and winner-take-all mechanisms, what determines which particular options are passed to the choice circuitry, whether there are mechanisms for constructing and editing choice sets, how the time accorded to a decision is controlled, and whether this allocation adjusts in response to changes in the cost of "errors."

While space does not permit us to review all the current work that addresses these questions, we do want to elaborate on one of these questions, which is perhaps the most hotly debated in the field at present. This is whether there are multiple valuation subsystems, and if so, how these systems are defined, how

independent or interactive they are, and how their valuations are combined into a final valuation that determines choice. Critically, this debate is *not* about whether different regions encode subjective value in a different manner—for example, the menu-invariant responses in orbitofrontal cortex compared to the relative value responses in LIP. Rather, the question is whether different systems encode *different and inconsistent* values for the same actions, such that these different valuations would lead to diverging conclusions about the best action to take. Many proposals along these lines have been made (Balleine et al., 2009; Balleine and Dickinson, 1998; Bossaerts et al., 2009; Daw et al., 2005; Dayan and Balleine, 2002; Rangel et al., 2008). One set builds upon a distinction made in the psychological literature between: *Pavlovian systems*, which learn a relationship between stimuli and outcomes and activate simple approach and withdrawal responses; *habitual systems*, which learn a relationship between stimuli and responses and therefore do not adjust quickly to changes in contingency or devaluation of rewards; and *goal-directed systems*, which learn a relationship between responses and outcomes and therefore do adjust quickly to changes in contingency or devaluation of rewards. Another related set builds upon a distinction between *model-free* reinforcement learning algorithms, which make minimal assumptions and work on "cached" action values, and more sophisticated *model-based* algorithms, which use more detailed information about the structure of the environment and can therefore adjust more quickly to changes in the environment. These proposals usually associate the different systems with different regions of the frontal cortex and striatum (Balleine et al., 2009) and raise the additional question of how these multiple valuations interact or combine to control behavior (Daw et al., 2005). It is important to note that most of the evidence we have reviewed here concerns decision making by what would be characterized in much of this literature as the "goal-directed" system. This highlights the fact that our understanding of valuation circuitry is in its infancy. A critical question going forward is how multiple valuation circuits are integrated and how we can best account for the functional role of different subregions of ventromedial prefrontal cortex and striatum in valuation. While no one today knows how this debate will finally be resolved, we can identify the resolution of these issues as critical to the forward progress of decision studies.

Another significant area of research that we have neglected in this review concerns the function of dorsomedial prefrontal and medial parietal circuits in decision making. Several recent reviews have focused specifically on the role of these structures in valuation and choice (Lee, 2008; Platt and Huettel, 2008; Rushworth and Behrens, 2008). Some of the most recently identified and interesting electrophysiological signals have been found in dorsal anterior cingulate (Hayden et al., 2009; Matsumoto et al., 2007; Quilodran et al., 2008; Seo and Lee, 2007, 2009) and the posterior cingulate (Hayden et al., 2008; McCoy and Platt, 2005). Decision-related signals in these areas have been found to occur *after* a choice has been made, in response to feedback about the result of that choice. One key function of these regions may therefore be in the monitoring of choice outcomes and the subsequent adjustment of both choice behavior and sensory acuity in response to this monitoring. However, there is also evidence suggesting that parts of the anterior cingulate may encode action-based subjective values, in an analogous manner to the orbitofrontal encoding of goods-based subjective values (Rudebeck et al., 2008). This reiterates the need for work delineating the specific functional roles of different parts of ventromedial prefrontal cortex in valuation and decision making.

Finally, a major topic area that we have not explicitly discussed in detail concerns decisions where multiple agents are involved or affected. These are situations that can be modeled using the formal frameworks of game theory and behavioral game theory. Many, and perhaps most, of the human neuroimaging studies of decision making have involved social interactions and games, and several reviews have been dedicated to these studies (Fehr and Camerer, 2007; Lee, 2008; Montague and Lohrenz, 2007; Singer and Fehr, 2005). For our purposes, it is important to note that the same neural mechanisms we have described above are now known to operate during social decisions (Barraclough et al., 2004; de Quervain et al., 2004; Dorris and Glimcher, 2004; Hampton et al., 2008; Harbaugh et al., 2007; King-Casas et al., 2005; Moll et al., 2006). Of course, these decisions also require additional processes, such as the ability to model other people's minds and make inferences about their beliefs (Adolphs, 2003; Saxe et al., 2004), and many ongoing investigations are aimed at understanding the role of particular brain regions in these social functions during decision making (Hampton et al., 2008; Tankersley et al., 2007; Tomlin et al., 2006).

Conclusion

Neurophysiological investigations over the last seven years have begun to solidify the basic outlines of a neural mechanism for choice. This breakneck pace of discovery makes us optimistic that the field will soon be able to resolve many of the current controversies and that it will also expand to address some of the questions that are now completely open.

Future neurophysiological models of decision making should prove relevant beyond the domain of basic neuroscience. Since neurophysiological models share with economic ones the goal of explaining choices, ultimately there should prove to be links between concepts in the two kinds of models. For example, the neural noise in different brain circuits might correspond to the different kinds of stochasticity that are posited in different classes of economic choice models (McFadden, 1974; Selten, 1975). Similarly, rewards are experienced through sensory systems, with transducers that have both a shifting reference point and a finite dynamic range. These psychophysically characterized properties of sensory systems might contribute both to the decreasing sensitivity and reference dependence of valuations, which are both key aspects of recent economic models (Koszegi and Rabin, 2006). Moving forward, we think the greatest promise lies in building models of choice that incorporate constraints from both the theoretical and mechanistic levels of analysis.

Ultimately, such models should prove useful to questions of human health and disease. There are already several elegant examples of how an understanding of the neurobiological mechanisms of decision making has provided a foundation for

understanding aberrant decision making in addiction, psychiatric disorders, autism, and Parkinson's disease (Bernheim and Rangel, 2004; Chiu et al., 2008a, 2008b; Frank et al., 2004; King-Casas et al., 2008; Redish, 2004). In the future, we feel confident that understanding the neurobiology of decision making also points the way toward improved treatments of these diseases and others, where people's choices play a key role.

ACKNOWLEDGMENTS

We thank N. Senecal and J. Stump for their comments on a previous version of this manuscript. P.W.G. is supported by grants from the National Institutes of Health (R01-EY010536 and R01-NS054775).

REFERENCES

Adolphs, R. (2003). Cognitive neuroscience of human social behaviour. Nat. Rev. Neurosci. *4*, 165–178.

Andersen, R.A., and Buneo, C.A. (2002). Intentional maps in posterior parietal cortex. Annu. Rev. Neurosci. *25*, 189–220.

Balleine, B.W., and Dickinson, A. (1998). Goal-directed instrumental action: Contingency and incentive learning and their cortical substrates. Neuropharmacology *37*, 407–419.

Balleine, B.W., Daw, N.D., and O'Doherty, J.P. (2009). Multiple forms of value learning and the function of dopamine. In Neuroeconomics: Decision Making and the Brain, P.W. Glimcher, C.F. Camerer, E. Fehr, and R.A. Poldrack, eds. (New York, NY: Academic Press), pp. 367–387.

Barraclough, D.J., Conroy, M.L., and Lee, D. (2004). Prefrontal cortex and decision making in a mixed-strategy game. Nat. Neurosci. *7*, 404–410.

Basso, M.A., and Wurtz, R.H. (1998). Modulation of neuronal activity in superior colliculus by changes in target probability. J. Neurosci. *18*, 7519–7534.

Bayer, H.M., and Glimcher, P.W. (2005). Midbrain dopamine neurons encode a quantitative reward prediction error signal. Neuron *47*, 129–141.

Bayer, H.M., Lau, B., and Glimcher, P.W. (2007). Statistics of midbrain dopamine neuron spike trains in the awake primate. J. Neurophysiol. *98*, 1428–1439.

Behrens, T.E.J., Woolrich, M.W., Walton, M.E., and Rushworth, M.F.S. (2007). Learning the value of information in an uncertain world. Nat. Neurosci. *10*, 1214–1221.

Bernheim, B.D., and Rangel, A. (2004). Addiction and cue-triggered decision processes. Am. Econ. Rev. *94*, 1558–1590.

Bossaerts, P., Preuschoff, K., and Hsu, M. (2009). The neurobiological foundations of valuation in human decision-making under uncertainty. In Neuroeconomics: Decision Making and the Brain, P.W. Glimcher, C.F. Camerer, E. Fehr, and R.A. Poldrack, eds. (New York, NY: Academic Press), pp. 353–365.

Chiu, P.H., Kayali, M.A., Kishida, K.T., Tomlin, D., Klinger, L.G., Klinger, M.R., and Montague, P.R. (2008a). Self responses along cingulate cortex reveal quantitative neural phenotype for high-functioning autism. Neuron *57*, 463–473.

Chiu, P.H., Lohrenz, T.M., and Montague, P.R. (2008b). Smokers' brains compute, but ignore, a fictive error signal in a sequential investment task. Nat. Neurosci. *11*, 514–520.

Churchland, A.K., Kiani, R., and Shadlen, M.N. (2008). Decision-making with multiple alternatives. Nat. Neurosci. *11*, 693–702.

Cohen, J.D., and Blum, K.I. (2002). Reward and decision. Neuron *36*, 193–198.

Colby, C.L., and Goldberg, M.E. (1999). Space and attention in parietal cortex. Annu. Rev. Neurosci. *22*, 319–349.

D'Ardenne, K., McClure, S.M., Nystrom, L.E., and Cohen, J.D. (2008). BOLD responses reflecting dopaminergic signals in the human Ventral Tegmental Area. Science *319*, 1264–1267.

Daw, N.D., Kakade, S., and Dayan, P. (2002). Opponent interactions between serotonin and dopamine. Neural Netw. *15*, 617–634.

Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nat. Neurosci. *8*, 1704–1711.

Dayan, P., and Balleine, B.W. (2002). Reward, motivation, and reinforcement learning. Neuron *36*, 285–298.

de Quervain, D.J., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., and Fehr, E. (2004). The neural basis of altruistic punishment. Science *305*, 1254–1258.

Delgado, M.R. (2007). Reward-related responses in the human striatum. Ann. N Y Acad. Sci. *1104*, 70–88.

Dommett, E., Coizet, V., Blaha, C.D., Martindale, J., Lefebvre, V., Walton, N., Mayhew, J.E., Overton, P.G., and Redgrave, P. (2005). How visual stimuli activate dopaminergic neurons at short latency. Science *307*, 1476–1479.

Dorris, M.C., and Munoz, D.P. (1998). Saccadic probability influences motor preparation signals and time to saccadic initiation. J. Neurosci. *18*, 7015–7026.

Dorris, M.C., and Glimcher, P.W. (2004). Activity in posterior parietal cortex is correlated with the relative subjective desirability of action. Neuron *44*, 365–378.

Fehr, E., and Camerer, C.F. (2007). Social neuroeconomics: The neural circuitry of social preferences. Trends Cogn. Sci. *11*, 419–427.

Fiorillo, C.D., Tobler, P.N., and Schultz, W. (2003). Discrete coding of reward probability and uncertainty by dopamine neurons. Science *299*, 1898–1902.

Fiorillo, C.D., Newsome, W.T., and Schultz, W. (2008). The temporal precision of reward prediction in dopamine neurons. Nat. Neurosci. *11*, 966–973.

Fox, C.R., and Poldrack, R.A. (2009). Prospect theory and the brain. In Neuroeconomics: Decision Making and the Brain, P.W. Glimcher, C.F. Camerer, E. Fehr, and R.A. Poldrack, eds. (New York, NY: Academic Press), pp. 145–173.

Frank, M.J., Seeberger, L.C., and O'Reilly, R.C. (2004). By carrot or by stick: Cognitive reinforcement learning in parkinsonism. Science *306*, 1940–1943.

Friedman, M. (1953). Essays in Positive Economics (Chicago, IL: University of Chicago Press).

Glimcher, P.W. (2003). The neurobiology of visual-saccadic decision making. Annu. Rev. Neurosci. *26*, 133–179.

Glimcher, P.W., and Sparks, D.L. (1992). Movement selection in advance of action in the superior colliculus. Nature *355*, 542–545.

Glimcher P.W., Camerer C.F., Fehr E., and Poldrack R.A., eds. (2009). Neuroeconomics: Decision Making and the Brain (New York, NY: Academic Press).

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

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Gul, F., and Pesendorfer, W. (2008). The case for mindless economics. In The Foundations of Positive and Normative Economics, A. Handbook, A. Caplin, and A. Schotter, eds. (New York, NY: Oxford University Press), pp. 3–41.

Haber, S.N. (2003). The primate basal ganglia: Parallel and integrative networks. J. Chem. Neuroanat. *26*, 317–330.

Hampton, A.N., Bossaerts, P., and O'Doherty, J.P. (2008). Neural correlates of mentalizing-related computations during strategic interactions in humans. Proc. Natl. Acad. Sci. USA *105*, 6741–6746.

Harbaugh, W.T., Mayr, U., and Burghart, D.R. (2007). Neural responses to taxation and voluntary giving reveal motives for charitable donations. Science *316*, 1622–1625.

Hare, T.A., O'Doherty, J., Camerer, C.F., Schultz, W., and Rangel, A. (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. J. Neurosci. *28*, 5623–5630.

Hare, T.A., Camerer, C.F., and Rangel, A. (2009). Self-control in decision-making involves modulation of the vmPFC valuation system. Science *324*, 646–648.

Hayden, B.Y., Nair, A.C., McCoy, A.N., and Platt, M.L. (2008). Posterior cingulate cortex mediates outcome-contingent allocation of behavior. Neuron *60*, 19–25.

Hayden, B.Y., Pearson, J.M., and Platt, M.L. (2009). Fictive reward signals in the anterior cingulate cortex. Science *324*, 948–950.

Heeger, D.J. (1992). Normalization of cell responses in cat striate cortex. Vis. Neurosci. *9*, 181–197.

Herrnstein, R.J. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. J. Exp. Anal. Behav. *4*, 267–272.

Horwitz, G.D., Batista, A.P., and Newsome, W.T. (2004). Representation of an abstract perceptual decision in macaque superior colliculus. J. Neurophysiol. *91*, 2281–2296.

Houthakker, H.S. (1950). Revealed preference and the utility function. Economica *17*, 159–174.

Isa, T., Kobayashi, Y., and Saito, Y. (2004). Dynamic modulation of signal transmission through local circuits. In The Superior Colliculus: New Approaches for Studying Sensorimotor Integration, W.C. Hall and A. Moschovakis, eds. (New York: CRC Press), pp. 159–171.

Janssen, P., and Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. Nat. Neurosci. *8*, 234–241.

Kable, J.W., and Glimcher, P.W. (2007). The neural correlates of subjective value during intertemporal choice. Nat. Neurosci. *10*, 1625–1633.

Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica *47*, 263–291.

Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. J. Neurosci. *28*, 3017–3029.

Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. *2*, 176–185.

Kim, S., Hwang, J., and Lee, D. (2008). Prefrontal coding of temporally discounted values during intertemporal choice. Neuron *59*, 161–172.

King-Casas, B., Tomlin, D., Anen, C., Camerer, C.F., Quartz, S.R., and Montague, P.R. (2005). Getting to know you: Reputation and trust in a two-person economic exchange. Science *308*, 78–83.

King-Casas, B., Sharp, C., Lomax-Bream, L., Lohrenz, T., Fonagy, P., and Montague, P.R. (2008). The rupture and repair of cooperation in borderline personality disorder. Science *321*, 806–810.

Knutson, B., and Cooper, J.C. (2005). Functional magnetic resonance imaging of reward prediction. Curr. Opin. Neurol. *18*, 411–417.

Kobayashi, S., and Schultz, W. (2008). Influence of reward delays on responses of dopamine neurons. J. Neurosci. *28*, 7837–7846.

Koszegi, B., and Rabin, M. (2006). A model of reference-dependent preferences. Q. J. Econ. *121*, 1133–1166.

Lau, B., and Glimcher, P.W. (2008). Value representations in the primate striatum during matching behavior. Neuron *58*, 451–463.

Lee, D. (2008). Game theory and neural basis of social decision making. Nat. Neurosci. *11*, 404–409.

Lee, P., and Hall, W.C. (2006). An in vitro study of horizontal connections in the intermediate layer of the superior colliculus. J. Neurosci. *26*, 4763–4768.

Leon, M.I., and Shadlen, M.N. (1999). Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. Neuron *24*, 415–425.

Leon, M.I., and Shadlen, M.N. (2003). Representation of time by neurons in the posterior parietal cortex of the macaque. Neuron *38*, 317–327.

Levy, I., Rustichini, A., and Glimcher, P.W. (2007). A single system represents subjective value under both risky and ambiguous decision-making in humans. In 37th Annual Society for Neuroscience Meeting (San Diego, CA).

Lo, C.C., and Wang, X.J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. Nat. Neurosci. *9*, 956–963.

Marr, D. (1982). Vision: A Computational Investigation into the Human Representation and Processing of Visual Information (New York, NY: Henry Holt and Co., Inc.).

Matsumoto, M., Matsumoto, K., Abe, H., and Tanaka, K. (2007). Medial prefrontal cell activity signaling prediction errors of action values. Nat. Neurosci. *10*, 647–656.

McClure, S.M., Berns, G.S., and Montague, P.R. (2003). Temporal prediction errors in a passive learning task activate human striatum. Neuron *38*, 339–346.

McCoy, A.N., and Platt, M.L. (2005). Risk-sensitive neurons in macaque posterior cingulate cortex. Nat. Neurosci. *8*, 1220–1227.

McFadden, D. (1974). Conditional logit analysis of quantitative choice behavior. In Frontiers in Econometrics, P. Zarembka, ed. (New York, NY: Academic Press), pp. 105–142.

Moll, J., Krueger, F., Zahn, R., Pardini, M., de Oliveira-Souza, R., and Grafman, J. (2006). Human fronto-mesolimbic networks guide decisions about charitable donation. Proc. Natl. Acad. Sci. USA *103*, 15623–15628.

Montague, P.R., and Lohrenz, T. (2007). To detect and correct: Norm violations and their enforcement. Neuron *56*, 14–18.

Montague, P.R., Dayan, P., and Sejnowski, T.J. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. J. Neurosci. *16*, 1936–1947.

Morris, G., Nevet, A., Arkadir, D., Vaadia, E., and Bergman, H. (2006). Midbrain dopamine neurons encode decisions for future action. Nat. Neurosci. *9*, 1057–1063.

Niv, Y., and Montague, P.R. (2009). Theoretical and empirical studies of learning. In Neuroeconomics: Decision Making and the Brain, P.W. Glimcher, C.F. Camerer, E. Fehr, and R.A. Poldrack, eds. (New York, NY: Academic Press), pp. 331–351.

O'Doherty, J.P. (2004). Reward representations and reward-related learning in the human brain: insights from neuroimaging. Curr. Opin. Neurobiol. *14*, 769–776.

O'Doherty, J.P., Dayan, P., Friston, K., Critchley, H., and Dolan, R.J. (2003). Temporal difference models and reward-related learning in the human brain. Neuron *38*, 329–337.

Padoa-Schioppa, C., and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. Nature *441*, 223–226.

Padoa-Schioppa, C., and Assad, J.A. (2008). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. Nat. Neurosci. *11*, 95–102.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R.J., and Frith, C.D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. Nature *442*, 1042–1045.

Plassmann, H., O'Doherty, J., and Rangel, A. (2007). Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. J. Neurosci. *27*, 9984–9988.

Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. Nature *400*, 233–238.

Platt, M.L., and Huettel, S.A. (2008). Risky business: The neuroeconomics of decision making under uncertainty. Nat. Neurosci. *11*, 398–403.

Quilodran, R., Rothe, M., and Procyk, E. (2008). Behavioral shifts and action valuation in the anterior cingulate cortex. Neuron *57*, 314–325.

Rangel, A., Camerer, C., and Montague, P.R. (2008). A framework for studying the neurobiology of value-based decision making. Nat. Rev. Neurosci. *9*, 545–556.

Ratcliff, R., Van Zandt, T., and McKoon, G. (1999). Connectionist and diffusion models of reaction time. Psychol. Rev. *106*, 261–300.

Redgrave, P., and Gurney, K. (2006). The short-latency dopamine signal: a role in discovering novel actions? Nat. Rev. Neurosci. *7*, 967–975.

Redish, A.D. (2004). Addiction as a computational process gone awry. Science *306*, 1944–1947.

Roesch, M.R., Calu, D.J., and Schoenbaum, G. (2007). Dopamine neurons encode the better option in rats deciding between differently delayed or sized rewards. Nat. Neurosci. *10*, 1615–1624.

Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J. Neurosci. *22*, 9475–9489.

Rudebeck, P.H., Behrens, T.E., Kennerley, S.W., Baxter, M.G., Buckley, M.J., Walton, M.E., and Rushworth, M.F. (2008). Frontal cortex subregions play distinct roles in choices between actions and stimuli. J. Neurosci. *28*, 13775–13785.

Rushworth, M.F., and Behrens, T.E. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. Nat. Neurosci. *11*, 389–397.

Samejima, K., Ueda, Y., Doya, K., and Kimura, M. (2005). Representation of action-specific reward values in the striatum. Science *310*, 1337–1340.

Samuelson, P. (1937). A note on the measurement of utility. Rev. Econ. Stud. *4*, 155–161.

Saxe, R., Carey, S., and Kanwisher, N. (2004). Understanding other minds: Linking developmental psychology and functional neuroimaging. Annu. Rev. Psychol. *55*, 87–124.

Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. Science *275*, 1593–1599.

Selten, R. (1975). A reexamination of the perfectness concept for equilibrium points in extensive games. Int. J. Game Theory *4*, 819–825.

Seo, H., and Lee, D. (2007). Temporal filtering of reward signals in the dorsal anterior cingulate cortex during a mixed-strategy game. J. Neurosci. *27*, 8366–8377.

Seo, H., and Lee, D. (2009). Behavioral and neural changes after gains and losses of conditioned reinforcers. J. Neurosci. *29*, 3627–3641.

Shadlen, M.N., and 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.

Shadlen, M.N., Britten, K.H., Newsome, W.T., and Movshon, J.A. (1996). A computational analysis of the relationship between neuronal and behavioral responses to visual motion. J. Neurosci. *16*, 1486–1510.

Singer, T., and Fehr, E. (2005). The neuroeconomics of mind reading and empathy. Am. Econ. Rev. *95*, 340–345.

Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. Science *304*, 1782–1787.

Sutton, R.S., and Barto, A.G. (1998). Reinforcement Learning: An Introduction (Cambridge, MA: The MIT Press).

Takahashi, Y.K., Roesch, M.R., Stalnaker, T.A., Haney, R.Z., Calu, D.J., Taylor, A.R., Burke, K.A., and Schoenbaum, G. (2009). The orbitofrontal cortex and ventral tegmental area are necessary for learning from unexpected outcomes. Neuron *62*, 269–280.

Tanji, J., and Evarts, E.V. (1976). Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. J. Neurophysiol. *39*, 1062–1068.

Tankersley, D., Stowe, C.J., and Huettel, S.A. (2007). Altruism is associated with an increased neural response to agency. Nat. Neurosci. *10*, 150–151.

Tobler, P.N., Dickinson, A., and Schultz, W. (2003). Coding of predicted reward omission by dopamine neurons in a conditioned inhibition paradigm. J. Neurosci. *23*, 10402–10410.

Tobler, P.N., Fiorillo, C.D., and Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. Science *307*, 1642–1645.

Tom, S.M., Fox, C.R., Trepel, C., and Poldrack, R.A. (2007). The neural basis of loss aversion in decision-making under risk. Science *315*, 515–518.

Tomlin, D., Kayali, M.A., King-Casas, B., Anen, C., Camerer, C.F., Quartz, S.R., and Montague, P.R. (2006). Agent-specific responses in the cingulate cortex during economic exchanges. Science *312*, 1047–1050.

Tremblay, L., and Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. Nature *398*, 704–708.

Van Gisbergen, J.A., Robinson, D.A., and Gielen, S. (1981). A quantitative analysis of generation of saccadic eye movements by burst neurons. J. Neurophysiol. *45*, 417–442.

von Neumann, J., and Morgenstern, O. (1944). The Theory of Games and Economic Behavior (Princeton, NJ: Princeton University Press).

Waelti, P., Dickinson, A., and Schultz, W. (2001). Dopamine responses comply with basic assumptions of formal learning theory. Nature *412*, 43–48.

Wallis, J.D., and Miller, E.K. (2003). Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. Eur. J. Neurosci. *18*, 2069–2081.

Wang, X.J. (2008). Decision making in recurrent neuronal circuits. Neuron *60*, 215–234.

Wong, K.F., and Wang, X.J. (2006). A recurrent network mechanism of time integration in perceptual decisions. J. Neurosci. *26*, 1314–1328.

Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. Nature *447*, 1075–1080.

Zaghloul, K.A., Blanco, J.A., Weidemann, C.T., McGill, K., Jaggi, J.L., Baltuch, G.H., and Kahana, M.J. (2009). Human substantia nigra neurons encode unexpected financial rewards. Science *323*, 1496–1499.