

c00028

# The Computation and Comparison of Value in Goal-directed Choice

Antonio Rangel

## OUTLINE

s0010	<b>Introduction</b>	423	<i>Neurobiological Evidence</i>	428	s0080		
s0020	<b>Preliminaries: What is Goal-directed Choice and What Needs to be Computed?</b>	424	<b>Simple Goal-directed Choice</b>	431	s0090		
s0100			<i>Theory</i>	432			
s0030			<i>Goal-directed Choice</i>	424	<i>Behavioral and Eye-tracking Evidence</i>	433	s0110
s0040	<i>What Needs to be Computed?</i>	425	<i>Neural Basis for the Model's Components</i>	433	s0120		
s0050	<b>Perceptual Decision-making</b>	425	<b>Final Remarks</b>	436	s0130		
s0060			<i>Model</i>	426	<i>Acknowledgments</i>	436	s0140
s0070			<i>Behavioral Evidence</i>	428	<i>References</i>	437	s0150

## INTRODUCTION

Neuroeconomics studies the neurobiological and computational basis of value-based decision-making. As the name suggests, these types of decisions involve the computation and comparison of values. This gives rise to one of the most basic questions in the field: how does the brain compute and compare values? This chapter provides an integrative review of computational and neurobiological studies addressing this question.

The chapter has several goals. It seeks to introduce economists and modelers to the neuroscience of the problem, and neuroscientists and psychologists to the relevant models (with the caveat that, in order to keep

the exposition accessible to both audiences, important simplifications have been made). This interplay between theory and neurobiology has been very productive in many areas of neuroeconomics, and our problem is no exception. The chapter also seeks to provide a framework that brings together computational and neurobiological findings by highlighting the connections between the relevant literatures in neuroscience, psychology, and economics. Finally, it seeks to motivate future work in this area by emphasizing some of the most important open questions.

Understanding how the brain computes and compares values is important for several reasons. First, in order to make good decisions, organisms need to assign

values to actions that are commensurate with the level of rewards that they generate, and they need to select the actions with the highest values. The extent to which they are able to do this depends on the properties of the algorithms that they use, and on the performance of the “wetware” implementing them.

p0040 Second, many models in economics (e.g., subjective experienced utility), psychology (e.g., prospect theory), and computer science (e.g., Q-learning) implicitly assume that the brain can flawlessly and instantaneously assign appropriate values to actions. Evidence from behavioral economics suggests that this is not always a correct assumption. A detailed understanding of how the brain computes and compares values should provide important insights on the strengths and limitations of existing models.

p0050 Third, evolution has endowed the brain with the capacity to make good choices in a large variety of environments, but not in all of them. In particular, organisms are susceptible to the influence of environmental variables that interfere with their ability to compute accurate values or select the best actions. Knowledge of the decision circuitry’s shortcomings has implications for public policy debates, such as the regulation of marketing, and for the design of economic institutions that can help people compensate for these shortcomings.

p0060 Finally, one of the most exciting areas of application for neuroeconomics is psychiatry. Damage to various neural systems may alter the way the brain assigns and compares values. In fact, psychiatric diseases such as obsessive-compulsive disorders, anorexia, psychopathy, and addiction can be characterized as impairments in decision-making. It follows that an improved understanding of the neurobiology of value-based choice should have immediate diagnostic and therapeutic applications.

p0070 The review is organized as follows. The first section defines goal-directed choice and places the problem on the context of the neuroeconomics research agenda, while the second reviews the literature on perceptual decision-making. The third section reviews the literature on goal-directed choice, before the chapter concludes with a final discussion.

## s0020 **PRELIMINARIES: WHAT IS GOAL-DIRECTED CHOICE AND WHAT NEEDS TO BE COMPUTED?**

### s0030 **Goal-directed Choice**

p0080 Value-based decision-making is a complex process that requires, at the very least, the deployment of five

basic types of computations (Rangel *et al.*, 2008). First, a representation of the decision problem needs to be constructed. This entails identifying potential courses of action (e.g., to pursue a prey item or not), as well internal states (e.g., the hunger level) and external states (e.g., the threat level). Second, a value needs to be assigned to the different actions under consideration. In order to make sound decisions, these values need to be aligned with the benefits associated with each action. Third, the valuations are compared in order to make a choice. Fourth, after implementing the decision, the brain needs to measure the desirability of the outcomes that follow. Finally, this feedback is used to update the other processes in order to improve the quality of future decisions. These are not rigid categories, but they provide a useful and reasonable decomposition of decision-making into its composite parts.

Another complication in understanding value-based choice is that a sizable body of evidence has shown that the brain utilizes multiple systems to evaluate the options under consideration (for reviews, see Balleine, 2005; Dayan, 2008; Rangel *et al.*, 2008; also Chapters 12 and 24 in this volume). At least three different types of systems have been identified so far: Pavlovian, habitual, and goal-directed. These systems are deployed in different circumstances, and have different relative advantages and disadvantages.

The Pavlovian systems assign values only to a small set of “hard-wired” behaviors that are evolutionarily appropriate responses to appetitive or aversive outcomes. Typical examples include preparatory behaviors such as approaching cues that predict the delivery of food, and consummatory responses to a reward such as pecking at a food magazine. An advantage of these systems is that they are fast, which can be invaluable in many circumstances. A wide range of behaviors with important economic consequences are likely to be controlled by the Pavlovian valuation systems. Examples include overeating in the presence of food, and perhaps even the harvesting of immediate present smaller rewards at the expense of delayed non-present larger rewards (Dayan *et al.*, 2006; Dayan and Seymour, 2008).

In contrast to the Pavlovian systems, which activate only a small set of hard-wired responses, habitual systems are capable of learning to assign values to a large number of actions through a process of trial and error. These systems exhibit the following characteristics. First, such systems learn to assign values to actions that are commensurate with the expected reward that they generate as long as sufficient practice is provided and the environment is stable enough (see Sutton and Barto, 1998; also Chapter 22 in this volume). Second, the systems learn and update values slowly – thus they forecast the value of actions incorrectly whenever there

is a sudden change in the rewards associated with the different actions. Third, the systems do not appear to assign values to novel actions, and thus cannot guide decision-making in new situations.

p0120 Goal-directed systems are capable of computing values in novel situations and in environments with rapidly changing action–outcome contingencies. These systems assign values to actions by computing action–outcome contingencies and then evaluating the rewards associated with the different outcomes. Under ideal conditions, the value assigned to an action equals the average expected reward obtained from that action. When the correct action–contingencies are used, and the brain evaluates outcomes correctly, such systems compute an accurate forecast of the expected reward generated by the action. A key disadvantage of goal-directed systems is that they are computationally intensive, especially in comparison to the habitual and Pavlovian systems, and require a large amount of accurate information to forecast values correctly (Daw *et al.*, 2005) Furthermore, in the presence of uncertainty about the action–outcome contingencies, their predictions are noisy forecasts of the true value of the action.

p0130 This chapter makes several simplifications to make our problem tractable. First, it focuses on the valuation and action-selection stages, and ignores the representation, evaluation, and learning processes (but see Chapters 22 and 24 for a more detailed discussion of these learning processes). Second, it considers the computation and comparison of values by goal-directed systems, but not by the habitual and Pavlovian systems. Third, it focuses on a simple type of goal-directed choice in which (1) there are only two options, (2) subjects are familiar with the outcomes generated by the two actions, and (3) consumption of the outcomes is immediate. The last simplification is useful, because it allows us to ignore the issue of how the value of outcomes is modulated by time and uncertainty (for a discussion of these issues, see Rangel *et al.*, 2008, and Chapter 23 of this volume).

### s0040 What Needs to be Computed?

p0140 A simple goal-directed choice involves a choice between a pair of actions – call them  $a$  and  $b$ . Subjects care about which action they take because the actions lead to different outcomes, denoted by  $o$ . The outcomes are associated with an amount of reward  $v(o)$ . The mapping of actions to outcomes may depend on the state of the world, denoted by  $s$ . For example, pulling the lever in a gambling machine (often called a bandit) leads to a desirable outcome if it is in the “payout” state, but to an undesirable one if it is in the “no payout” state.

In order to make good choices, the goal-directed valuation system needs to (1) assign values to actions that are commensurate with the amount of expected reward that they generate, and (2) choose the action that has been assigned the highest value. Thus, ideally the system ought to assign a value to action  $x$  equal to

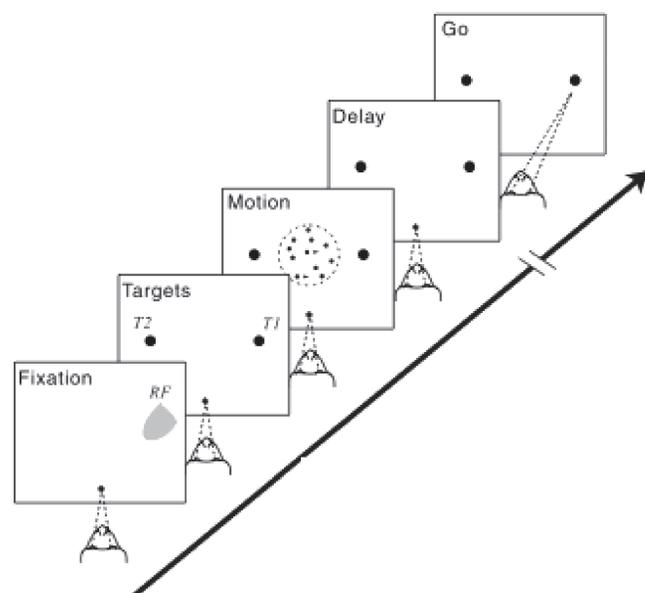
$$U_x = \sum_s p(s)v(o_x(s))$$

where  $p(s)$  is the probability of state  $s$ , and  $o_x(s)$  is the outcome generated by the action in state  $s$ .

Note that in order to estimate this value the brain needs to calculate in real time the probability of the potential states, the action–state–outcome associations, and the outcome–reward pairings. As we will show below, making these computations takes time.

## PERCEPTUAL DECISION-MAKING

We begin our discussion of the computation and comparison of goal-directed values with one of the best-studied problems in behavioral neuroscience: visual-saccadic decision-making. An example of this type of decision-making is the now classic experimental paradigm of Newsome *et al.* (1989), referred to here as Random Dot Motion (RDM) task. As illustrated in Figure 28.1a, monkeys are shown a patch of



**FIGURE 28.1** The Random Dot Motion task. The monkey is trained to perform a two-alternative forced choice discrimination task. He is shown a patch of randomly moving dots that contain a fraction of coherent movement to either T1 or T2, and is rewarded for correctly choosing the coherent direction. Choices are indicated by fixating, after a delay, on one of the two targets. Adapted from Shadlen and Newsome (2001), with permission.

moving dots. Most of the dots are moving randomly, but a fraction of them has a coherent direction of motion. The animal's task is to indicate which of two targets reflects the coherent direction of motion by executing an eye movement to the appropriate target. The difficulty of the task is varied across trials by changing the percentage of dots that are moving coherently. Subjects are rewarded a fixed amount for correct answers, and receive nothing for incorrect responses. (For more details on this task see Chapters 4 and 30 of this volume).

p0180 At first sight, the RDM task looks very different from value-driven choice. This is not the case. Subjects care about making the right choice because it affects the reward that they get. The value of making an eye movement equals the size of the reward for a correct response multiplied by the probability that that direction is the correct one. Subjects can estimate the action values using their senses to measure the net direction of motion in the stimulus display. This is easy when the percentage of coherent dots is high, and difficult when it is small. In either case, the task involves linking sensation to valuation and then using those values to select an action.

p0190 Note that the task is a special case of the model described above. There are two states of the world: the coherent motion goes in the direction of target 1 ( $s = 1$ ) or it goes in the direction of target 2 ( $s = 2$ ). The mapping of actions and states to outcome is given by:

$$o_1(s = 1) = o_2(s = ) = \text{one unit of reward}$$

and

$$o_1(s = ) = o_2(s = 1) = \text{zero units of reward.}$$

p0200 For simplicity, and without imposing any restrictions in the data, we can pick the units in which the reward is measured so that  $v(1) = 1$  and  $v(0) = 0$ . It then follows that

$$U_1 = p(s = 1)v(1) + p(s = 2)v(0) = p(s = 1)$$

and

$$U_2 = p(s = 2)v(1) + p(s = 1)v(0) = p(s = 2)$$

p0210 In other words, the value of choosing one of the targets is equal to the probability that it reflects the true coherent direction of motion. It follows that the problem of computing the value of the actions is closely related to the perceptual task of estimating the the direction of motion.

p0220 Although other perceptual discrimination tasks have been studied in the literature, the RDM task is

the one that is best understood from both computational and neurobiological perspectives. For this reason, the rest of this chapter focuses on the RDM task (see Romo and Salinas, 2003, and Gold and Shadlen, 2007, for reviews of alternative tasks).

## Model

A large number of computational models describing how the brain computes and compares values has been proposed. These include, among others, the independent accumulators model (Luce, 1986); the random-walk and diffusion models (Stone, 1960; Ratcliff, 1978; Laming, 1979; Link, 1992; and Ratcliff *et al.*, 2003); and several hybrid models (Smith, 1995; Usher and McClelland, 2001). For detailed discussions of the differences and similarities among all of them, see Smith (2000), Ratcliff and Smith (2004), Ditterich (2006), and Bogacz (2007).

All of these models share some features that are illustrated in Figure 28.2a. First, they contain processes that encode the moment-to-moment sensory evidence in support of the two potential states of the world. For example, in this figure there is a population of neurons encoding evidence in favor of the first target, and another population encoding evidence in favor of the second. Importantly, the models assume that these perceptual measurement processes are noisy. Second, there is an accumulator process associated with each of the two states that integrates the cumulative sensory evidence which it has received in its favor. Note how, as the trial progresses, the accumulators average out some of the noise. Finally, the models contain a criterion that needs to be satisfied for a choice to be made. For example, in the accumulators model, the first accumulator to reach a threshold is selected. The models listed in the previous paragraph differ in the details of the integration and comparison processes they employ to accomplish this sequence of events – see Figure 2 in Smith and Ratcliff (2004) and Figure 1 in Bogacz (2007) for useful summaries.

Here, we focus on the random-walk model (which is a version of the diffusion model in which computations are made in discrete units of time). This version of the model is attractive because it is mathematically simpler than its counterparts, and it makes predictions that can account for important aspects of the psychometric and reaction-time functions. Furthermore, a growing body of research has found signals in the brain that resemble some of the variables computed by the random-walk model (more on this below).

The random-walk model is illustrated in Figures 28.2b–e. The key variable is a measure of relative evidence: positive values of this variable denote that the

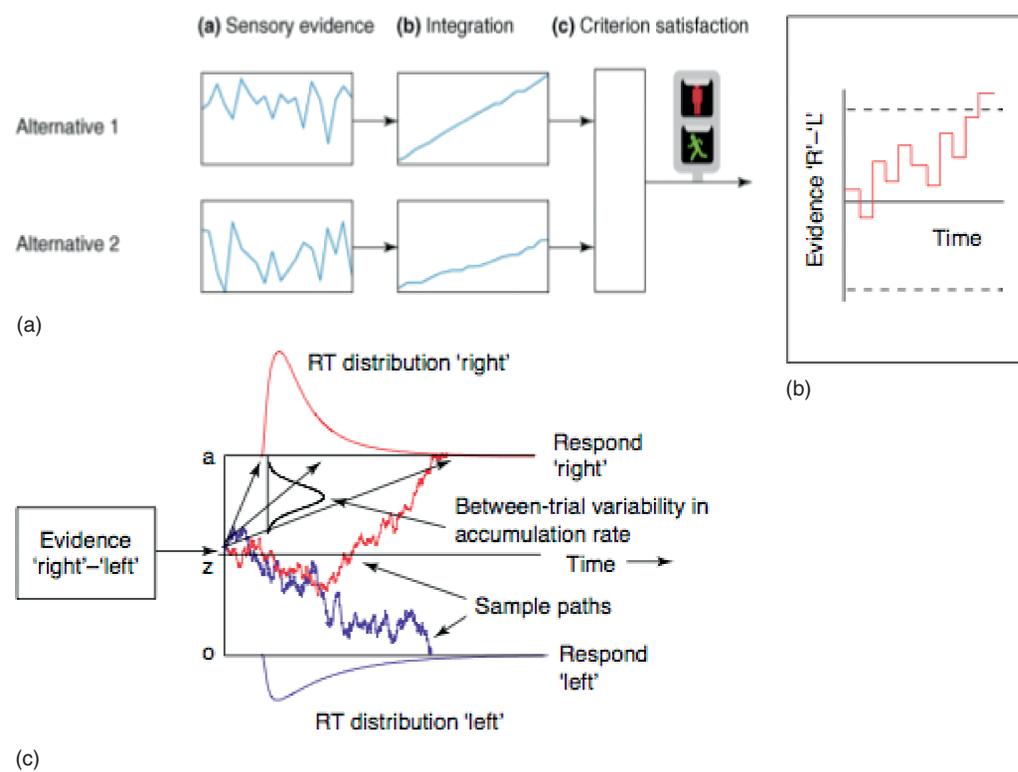
s0060

p0230

p0240

p0250

p0260



**FIGURE 28.2** Models of the RDM-task. (a) Schematic illustrating the main components of the models. Reproduced from Bogacz (2007), with permission. (b) A typical run of the random-walk model. The red line represents the amount of relative evidence accumulated for the “right” target. The process starts at a middle point and stops the first time this variable crosses one of the thresholds (depicted by the bracketed horizontal lines). “Right” is chosen when it crosses the upper threshold; “left” is chosen when it crosses the lower one. Time advances in discrete steps. The size of every step is given by a Gaussian distribution with a mean that is proportional to the true direction of motion. This noise is meant to capture the variability in the perceptual processes. (c) A typical run of the continuous version of the model. Note that the red and blue lines represent two different sample paths and that, due to the variability in the model, it is possible for “left” to be chosen even if the net coherent motion goes right. Adapted from Smith and Ratcliff (2004), with permission. f0020

estimation process favors the “right” decision, whereas negative values denote the opposite. The process starts at a middle point and stops the first time that this variable crosses one of the thresholds. Time advances in discrete steps. The size of every step is given by a Gaussian distribution with a mean that is proportional to the true direction of motion. This noise is meant to capture the variability in the sensory processes.

The model has several interesting properties. First, it predicts a logistic relationship between choice and the extent to which the dots are moving in a coherent direction (Figure 28.2d). Second, it predicts that the distribution of reaction times is right-skewed, and that reaction times are longer for error trials (Figure 28.2c). Third, it also predicts that reaction times decrease with stimulus coherence and that accuracy improves (Figure 28.2e). Finally, it can be shown that it implements a sequential probability ratio test (Wald and Wolfowitz, 1948), which is the optimal statistical test for the type of inference problem faced by the brain in the RDM task (see Gold and Shadlen, 2001, 2002, for details).

Although significant progress has been made in modeling this problem, many important theoretical questions remain unsolved. Perhaps the most important limitations of existing models are the fact that they only apply to the case of two alternatives, and that it is not obvious how to adapt them to the general case (for some initial attempts to resolve this issue, see McMillen and Holmes, 2005, and Bogacz and Gurney, 2007). Another problem is the proliferation of competing models in the literature. For example, in addition to those listed above, Brown and Heathcote (2005) have recently proposed a new class of “ballistic models” where the initial state for the accumulators is random but the path of integration is deterministic. An important task for the upcoming years is to use novel behavioral tests, as well neural data, to compare these models in depth. Most likely, the model that will best explain the data is yet to be developed. Finally, in order to improve the connection between the theory and the neurobiology, neuroeconomists need to develop a new generation of models that describe

more closely the role of the relevant brain nuclei and neurotransmitter systems. For some important initial efforts along these lines, see Wang (2002), Brown and Heathcote (2005), Brown *et al.* (2005), Lo and Wang (2006), and Wong and Wang (2006).

p0290 Note that the models of perceptual decision-making described here are ambitious in their scope. They provide a unified theory of value computation, action selection, and reaction times. They make stark behavioral predictions about how the psychometric choice function and the reaction time distribution should be affected by the coherence of the stimulus. Most importantly, they also make predictions about which decision variables should be encoded in the brain, and about how they are computed. As we will see in the subsequent two sections, the random-walk model has been quite successful in all of these respects.

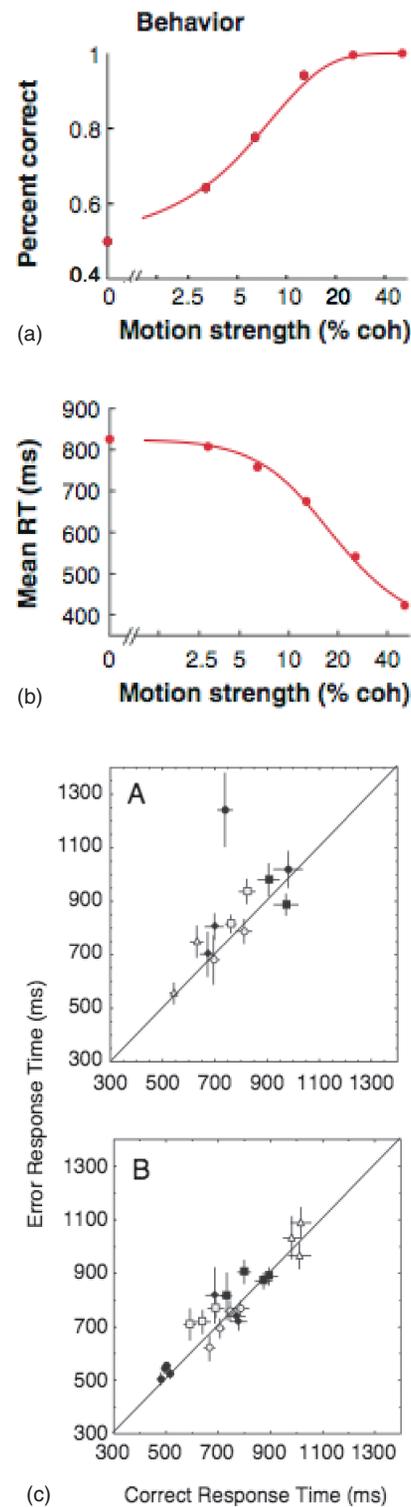
### s0070 Behavioral Evidence

p0300 The behavioral predictions of the random-walk model and its variants are largely consistent with the data from the RDM task. Figure 28.3 illustrates this using data from a human version of the task (Palmer *et al.*, 2005). As predicted by the model, choices are a logistic function of stimulus strength (Figure 28.3a), reaction times decrease in stimulus strength (Figure 28.3b), and response times are longer in error than in correct trials (Figure 28.3c). See Smith and Ratcliff (2004) for a further discussion of the behavioral predictions made by the different models.

### s0080 Neurobiological Evidence

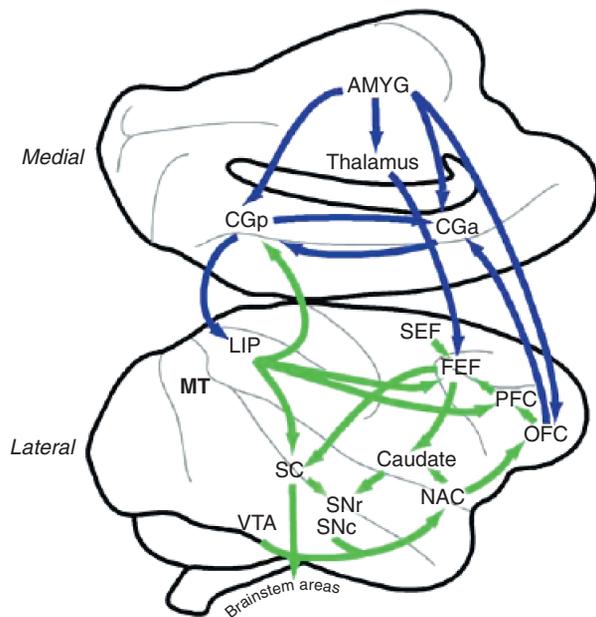
p0310 A central goal of neuroeconomics is to construct theories of value-based decision-making that are neurobiologically grounded. Therefore, a key test of any neuroeconomic model is whether there is evidence that the computations it specifies are being carried out by the brain. Over the past 15 years, a sizable number of electrophysiology studies have probed the neural basis of the random-walk model. This section reviews the most salient findings and describes some of the open questions. For more comprehensive reviews of the neurobiology of perceptual decision-making, see Platt (2002), Glimcher (2003), Romo and Salinas (2003), Schall (2005), and Gold and Shadlen (2007).

p0320 Since some readers of this chapter will be unfamiliar with the brain areas discussed below, a mini-primer of their location and known general function is provided (see Kandell *et al.*, 2000, for more details). Figure 28.4 depicts two networks: the green one involves structures and connections that have been associated with



**FIGURE 28.3** Behavioral predictions of the random-walk model in a human version of the RDM-task. (a) The random-walk model leads to a logistic choice function. (b) Reaction times and choice accuracy increase with stimulus coherence. Adapted from Gold and Shadlen (2007), with permission. (c) Mean response times for errors were slower than those for correct responses. Each point corresponds to a condition-observer pair. Adapted with permission from Palmer *et al.* (2005).

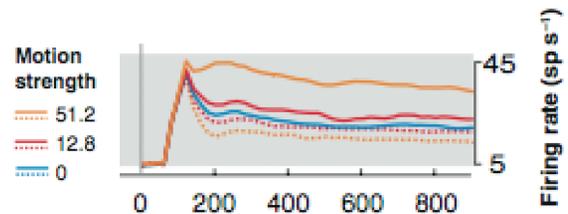
f0030



f0040

**FIGURE 28.4** Network of brain areas that have been implicated in saccadic decision-making. The figure depicts medial and lateral views of a monkey brain. The blue arrows denote signals that have been implicated by previous studies in the planning and execution of eye-movements. The green arrows denote signals that have been implicated by previous studies in the computation and comparison of goal-values, as well as in other reward-related computations. AMYG, amygdala; CGp, posterior cingulated cortex; CGa, anterior cingulated cortex; LIP, lateral intraparietal area; MT, middle temporal area; SEF, supplementary eye fields; FEF, frontal eye fields; PFC, prefrontal cortex; OFC, orbitofrontal cortex; SC, superior colliculus; NAC, nucleus accumbens; SNr, substantia nigra pars reticulata; SNc, substantia nigra pars compacta; VTA, ventral tegmental area. Adapted from McCoy and Platt (2004), with permission.

visual decision-making, while the blue one involves areas that have been associated with the computation and comparison of value. It is important to emphasize that the exact roles played by the different parts of these networks are still not fully understood, and in some cases are hotly debated. Here, we focus the discussion only on the subset of these areas that have been closely associated with the computations of the random-walk model and its near relatives. Neurons in the middle temporal area (MT) encode the motion properties of objects in the visual field. The role of the lateral intraparietal area (LIP) in saccadic decision-making has been the subject of past controversy: some have claimed that it is involved in spatial attention, others that it is involved primarily in the planning of saccadic eye movements. The superior colliculus (SC) is widely believed to be involved in controlling orienting behaviors such as eye movements. This area receives visual inputs from LIP, and sends motor projects to the brainstem and the spinal cord. The frontal eye-fields (FEF)



**FIGURE 28.5** Activity in MT during the RDM task. The average activity in MT neurons is plotted as a function of time within the trial aligned to stimulus presentation. The panel describes the motion coherence of the stimulus. Solid and dashed curves represent trials in which the direction of motion coherence fell in and out of the neuron's response field, respectively. Note that the time-courses are fairly constant. Adapted from Gold and Shadlen (2007), with permission.

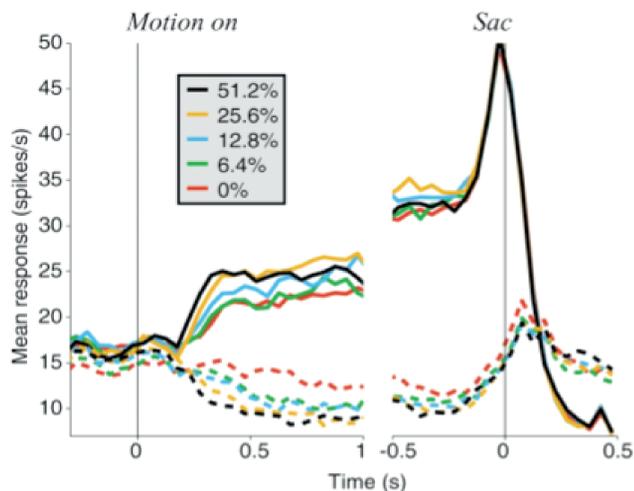
f0050

are also believed to form part of the network that controls the generation of eye movements.

As we saw in Figure 28.2a, the random-walk model has three key components: a set of processes capturing the incoming sensory evidence, an accumulator measuring the amount of net evidence in favor of the alternatives, and a set of thresholds implementing the decision criterion. This provides a useful way to organize the neurobiological literature.

Consider first the sensory processes. Direction-sensitive neurons in area MT have been shown to encode information about the direction of motion in the stimulus display (Dubner and Zeki, 1971; Zeki, 1974; Britten *et al.*, 1993). MT is organized around groups of neurons that respond preferentially to stimuli moving in a certain direction (Born and Bradley, 2005). Neuroscientists refer to the preferred parameters to which the group of neurons responds as the response field. Figure 28.5 illustrates the activity of a population of MT neurons with similar response fields. Activity increases with stimulus coherence when the dot movement is in the direction of the neuron's response field, and decreases with stimulus coherence otherwise. Note, as assumed by the model, that activity in the MT neurons is sustained throughout the trial. Similar results have been found for area MST (Celebrini and Newsome, 1994).

Although the previous findings are very suggestive, they do not establish a causal role for MT in the computation of goal values. After all, MT neurons could be measuring movement for purposes other than making a decision. To establish a causal role, it is necessary to use either lesion or stimulation methods. Using the later technique, Salzman *et al.* (1990, 1992) showed that microstimulation of the MT cells encoding for a particular direction of motion biased perceptual judgments in favor of that direction. Similar results were subsequently obtained for area MST (Celebrini and Newsome, 1994).



**FIGURE 28.6** Activity in LIP during the RDM task. Average activity in LIP neurons is plotted as a function of time within the trial aligned to either the stimulus onset (left panel) or the time of the saccadic response (right panel). Solid and dashed curves represent trials in which the monkey judged the direction to be towards and away from the LIP response field, respectively. Note two important features of the time-courses: first, for trials in which a choice is made into the response field, firing rates are increasing on stimulus coherence. The opposite is true when the choice is out of the response field. Second, when a choice is made into the response field, all of the time-courses collapse into a single threshold about 100ms before the initiation of the saccade. Adapted from Shadlen and Newsome (2001), with permission.

Next, consider the accumulator signal. Shadlen and Newsome (1999, 2001) found that activity in the lateral intraparietal area (LIP) resembles the accumulation of sensory signals predicted by the models. This is illustrated in Figure 28.6. Note a few things about these neurons. First, LIP neurons are also organized around response fields – in this case, LIP neurons respond to eye movements in a particular direction. Second, for trials in which a choice is made into the response field, firing rates are increasing on stimulus coherence. The opposite is true when the choice lies outside of the response field. Third, when a choice is made into the response field, all of the time-courses rise to a single threshold about 100ms before the initiation of the saccade. A limitation of these studies is that they impose fixed response times plus a delay between stimulus presentation and response, which is inconsistent with the assumptions of the random-walk model. Follow-up experiments by Roitman and Shadlen (2002) have shown that the findings do, however, extend to the case of free and non-delayed responses.

To address the issue of causality, Hanks *et al.* (2006) microstimulated LIP neurons and found that they could bias the proportion of choices that are made towards the stimulated response field. Not only that: just as predicted by the random-walk model,

stimulation generated faster reaction times when movements were made into the neurons' response field, and slower reaction times when they were made out of it. Based on this evidence, Shadlen and others have suggested that LIP neurons encode a form of the accumulator variable and that this information is used in making perceptual choices (Roitman and Shadlen, 2002; Mazurek *et al.*, 2003; Huk and Shadlen, 2005; Hanks *et al.*, 2006). Horwitz and Newsome (1999, 2001a, 2001b) have shown that there are cells in the superior colliculus (SC) that exhibit properties similar to those in LIP.

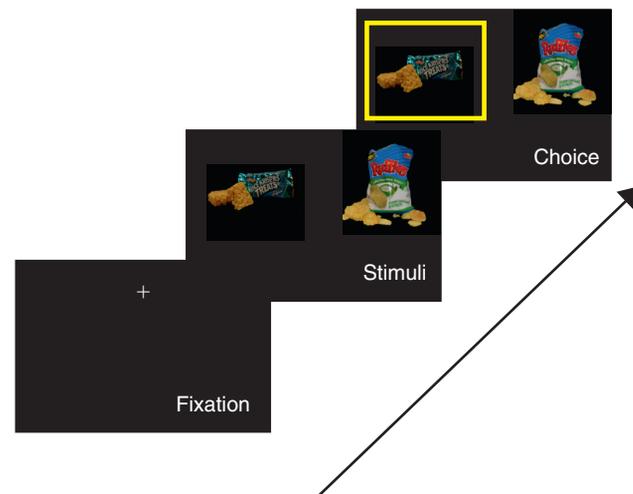
Finally, consider the neural implementation of the threshold. Although LIP neurons exhibit a threshold-like behavior (see Figure 28.4; Shadlen and Newsome, 2001; Roitman and Shadlen, 2002), existing experiments cannot determine whether the threshold is implemented within LIP or in other parts of the decision-making network. Lo and Wang (2006) have argued that a natural candidate for the implementation of the threshold is the SC. This midbrain area receives strong input from cortical areas, such as LIP, and contains "burst neurons" that fire just before a saccade is initiated. Lo and Wang have proposed that this cluster of neurons might be the *de facto* threshold that needs to be crossed to initiate a saccade. (For more details on this model, see Chapter 31 in this volume).

Many behavioral experiments have shown that subjects adjust their performance in response to time pressure in a way that is consistent with a decrease in the size of the thresholds (Smith and Ratcliff, 2004). This raises the question of how the threshold is modulated. Lo and Wang (2006) have proposed a model that involves a cortico-basal ganglia circuit. The SC receives inhibitory input from the substantia nigra pars reticulata (SNr), which in turn receives inhibitory input from the caudate nucleus. In the model, cortical activity increases firing rates in the caudate, which, by increasing the inhibition of the SNr, diminish the amount of activity required for the burst neurons to fire, thus decreasing the threshold. Further experiments are needed to determine the validity of this hypothesis.

Electrophysiological studies have also provided useful evidence to distinguish between the various models. Consider, for example, the experiments by Ditterich *et al.* (2003), which showed that microstimulation of MT speeds up decisions in the preferred direction of the neurons but slows down saccades in the opposite direction. These findings are consistent with the existence of a common accumulator, as in the random-walk model, but not with the existence of independent accumulators, as in the race-to-barrier model. A similar implication follows from the microstimulation experiments in LIP by Hanks *et al.* (2006).

p0410

Even though saccadic decision-making is one of the better understood problems in neuroeconomics and behavioral neuroscience, much work remains to be done. First, how do neural circuits integrate the sensory information to generate the accumulator signal that has been found in LIP? Do neural ensembles in this area compute it, or does it arise from the interaction of a larger circuit that remains to be characterized? Furthermore, what is the functional difference between the accumulator signals found in LIP and those found in the SC and the frontal eye-fields? Second, it has been shown that activity in LIP is modulated by prior beliefs and expected rewards (Platt and Glimcher, 1999). These variables change the equation for the goal values described in the model above. An important open problem is to understand how these types of variables are computed and how they become incorporated into the accumulator signal found in LIP. Third, another interesting puzzle has to do with the existence of multiple accumulator signals in LIP, one for each favored direction. How do all of these signals become combined into a single decision? Existing models are silent on this, because they assume that there is a single accumulator signal. Fourth, what is the selection criterion used by the brain, and how is it implemented? The findings reviewed above are consistent with the existence of thresholds. However, it is also possible that the observed thresholds are epiphenomena and that the underlying selection criterion is substantially different. In addition, how do executive control processes modulate the selection criteria? Finally, does the random-walk model capture all of the computations required to make perceptual decisions, or are some essential processes missing?



**FIGURE 28.7** An example of simple goal-directed choice. After an initial fixation cross, pictures of two snacks are displayed. Subjects are then free to look at the items until they are ready to make a choice by pressing one of two keys. The pattern of fixations during the deliberation process is recorded with an eye-tracker. Adapted from Krajbich *et al.* (2007), with permission.

f0070

notation below). Let  $o_{\text{left}}$  denote the snack associated with the “left” action and  $o_{\text{right}}$  that associated with the “right” choice. The value of the two actions is given by

$$U_{\text{left}} = v(o_{\text{left}}) \text{ and } U_{\text{right}} = v(o_{\text{right}}).$$

Note some of the key differences from the RDM task. The value of an action is simply the reward obtained by eating the food to which it leads. Unlike the case of perceptual choice, the values do not depend on experimenter-controlled states that need to be estimated using sensory information. However, since the snacks associated with each action change from trial to trial, the brain needs to (1) identify the outcome associated with each action and (2) estimate its reward value. This is in contrast to the case of perceptual choice, in which the outcomes associated with the actions were constant but there was uncertainty about their probability. Despite the differences, in both cases the choice problem boils down to a computation and comparison of value signals.

p0440

Note that since no two consumption situations are identical, every time an action is evaluated the brain needs to characterize anew the outcome associated with it. For example, are the current snacks safe? What exact caloric value do they have? Since not all of the characteristics are observable before consuming the good (e.g., taste), there is an inescapable amount of uncertainty associated with every choice – for example, the trustworthiness of the experimenter and recent news about the safety of the food supply might affect evaluations of the outcome. Understanding how the

p0450

s0090

## SIMPLE GOAL-DIRECTED CHOICE

p0420

Now consider the case of simple goal-directed choice. A typical example is shown in Figure 28.7. Subjects see a pair of pictures of familiar snacks and have to decide which one to eat at the end of the experiment. They may signal their choice by making an eye movement to one of the items, through a voice command, or by pressing one of two buttons. As in the previous section, the discussion is simplified by focusing on this particular type of goal-directed choice.

p0430

The simple model of valuation derived above highlights some of the key computations that need to be made. In this case there are two actions, “left” and “right”. There is only one state of the world given by the nutritional demands of the subject at the time of the decision (for simplification, we omit the state from the

brain deals with this type of uncertainty is an important open question beyond the scope of this review. Here, the model and discussion are simplified by assuming that this type of uncertainty is not present, which is a reasonable approximation for the experiments discussed below.

p0460 A more important difference from the RDM-task has to do with the role of attention. Although animals can make goal-directed choices among items while maintaining a central fixation (Padoa-Schioppa and Assad, 2006), subjects normally pay sequential attention to two stimuli during the process of choice – for example, they might fixate first on the left item, observe it for a few hundred milliseconds, and then switch to the other one. This introduces a role for attention, and makes the problem of simple goal-directed choice significantly more complicated.

### s0100 Theory

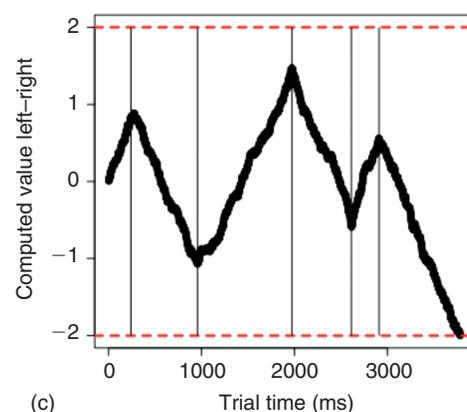
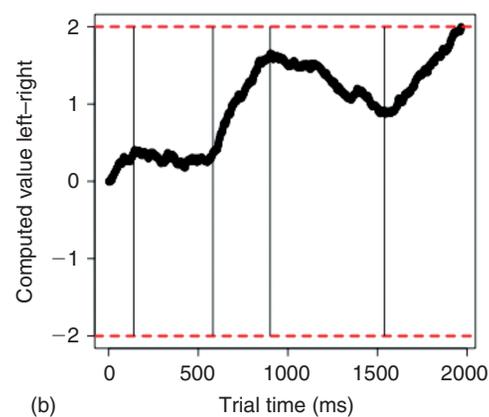
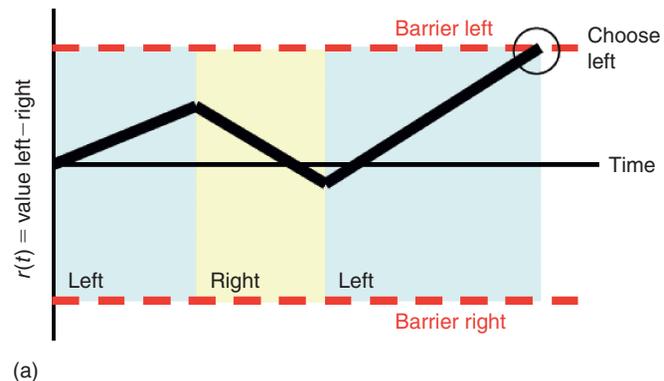
p0470 This section reviews the work of Krajbich and colleagues (2007), which has extended the random-walk model of perceptual decision-making to the case of attention-driven choice. The idea behind the model is depicted in Figure 28.8a. The brain computes a relative value variable that is analogous to the accumulator of the random-walk model. This variable measures the relative value of “right” action. A decision is made when the decision value crosses one of the two thresholds (upper for “left” and lower for “right”).

p0480 The key innovation of the model has to do with the role of attention. The relative value variable follows a random-walk process with Gaussian noise in which the mean rate of change is determined by the item that the subject is looking at. In particular, the variable changes from one instant to another according to the formula

$$r(t) = (t - 1) + \lambda(v(o_{\text{target}}) - \theta v(o_{\text{non-target}})) + \varepsilon(t)$$

where  $\varepsilon(t)$  is white Gaussian noise,  $v(o_{\text{target}})$  is the value of the item currently being attended,  $v(o_{\text{non-target}})$  is the value of the other item,  $\theta$  is a positive constant between 0 and 1 measuring the advantage given to the item being attended in the value construction process, and  $\lambda$  is a positive constant that scales the value functions (see Krajbich *et al.*, 2007, for further details).

p0490 The model of visual attention is extremely simple. The identity of the first fixation is random, and independent of the value of the two items. Fixation lengths are drawn from an equal distribution, and are also assumed to be unaffected by the value of the items. Note that the subject fixates on an item and then alternates fixations between both of them until a barrier is reached.



**FIGURE 28.8** Model of simple goal-directed choice. (a) A typical run of the model. The thick black line represents the relative value of the “right” minus the “left” actions. The white and yellow vertical bands represent the time of the fixation. In the example depicted here, the fixations are left, right, left. The process starts at a relative value of zero and accumulates following a random-walk model that depends on the direction of the fixation. In particular, subject to the qualifications described in the text, the process approaches the left barrier while the subject is looking at the left item, and *vice versa*. A decision is made when the relative decision value signal crosses a barrier. Not shown are the random fluctuations associated with the random walk. (b) A simulated run of the model using parameters estimated from real data for a trial in which the “left” items is more valuable. (c) Another simulated run of for a trial in which both items have equal value. See Krajbich *et al.* (2007) for details.

f0080

p0500 Figures 28.8b and 28.8c depict simulated runs with parameters estimated from real experimental data. The left panel depicts a case in which  $v(o_{\text{left}}) > v(o_{\text{right}})$ . Here, the signal climbs on average toward the left barrier and that choice is made with higher probability. The right panel depicts a case in which  $v(o_{\text{left}}) = v(o_{\text{right}})$ . In this case, the signal can fluctuate widely between the two barriers before a choice is made, and both options are selected with equal probability.

p0510 The model's predictions regarding the psychometric choice function and reaction-time profiles are similar to those of the random-walk model: choice probabilities are a logistic function of the underlying difficulty of the choices (as measured by  $v(o_{\text{left}}) - v(o_{\text{right}})$ ), reaction times decrease with the difficulty of the choice, and reaction times are larger for error trials (in which the item with the lowest value is chosen) than in correct trials.

p0520 The model also makes predictions about the connection between visual attention and choice. For example, it predicts that last fixations are shorter (because they are interrupted when a barrier is reached), that the last fixation is to the chosen item as long as it is sufficiently better than the item seen in the next-to-last trial, and that there is a choice bias in favor of items seen first and last.

p0530 Although this model is able to account for several important features of the data, as we will see below, it leaves unanswered several important questions about goal-directed choice. First, and most importantly, is the issue of why it takes time to construct the relative value signal. In the case of perceptual decision-making, there is a perceptual state that needs to be measured. However, what are the analogous processes here? One possibility is that the value signal might be computed by a process that involves observing the characteristics of the outcomes, retrieving from memory reward values associated with those characteristics, and finally integrating them. Each of these more fundamental steps might play the analogous role to sensory evidence. Second, the model assumes that the fixations are random, but this is probably not the case – for example, the first fixation might be driven by low-level visual features of the stimuli. Given the central role that visual attention plays in this class of models, it is essential to improve our understanding of the processes guiding visual attention. Third, since most choice situations involve more than two alternatives, the theory needs to be extended in this direction. For a preliminary attempt, see Pulst-Korenberg and Rangel (2008).

p0540 The theory of attention-driven goal-directed choice presented here is not the only one that has been proposed. In particular, the model described here is closely related to decision-field theory (Roe *et al.*,

2001; Busemeyer and Townsend, 1993; Busemeyer and Johnson, 2004). The main difference between this model and decision-field theory is that whereas the former is an extension of the random-walk model, the latter is more closely related to the hybrid connectionist models of McClelland and Rumelhart (1981) and Usher and McClelland (2001).

### Behavioral and Eye-tracking Evidence

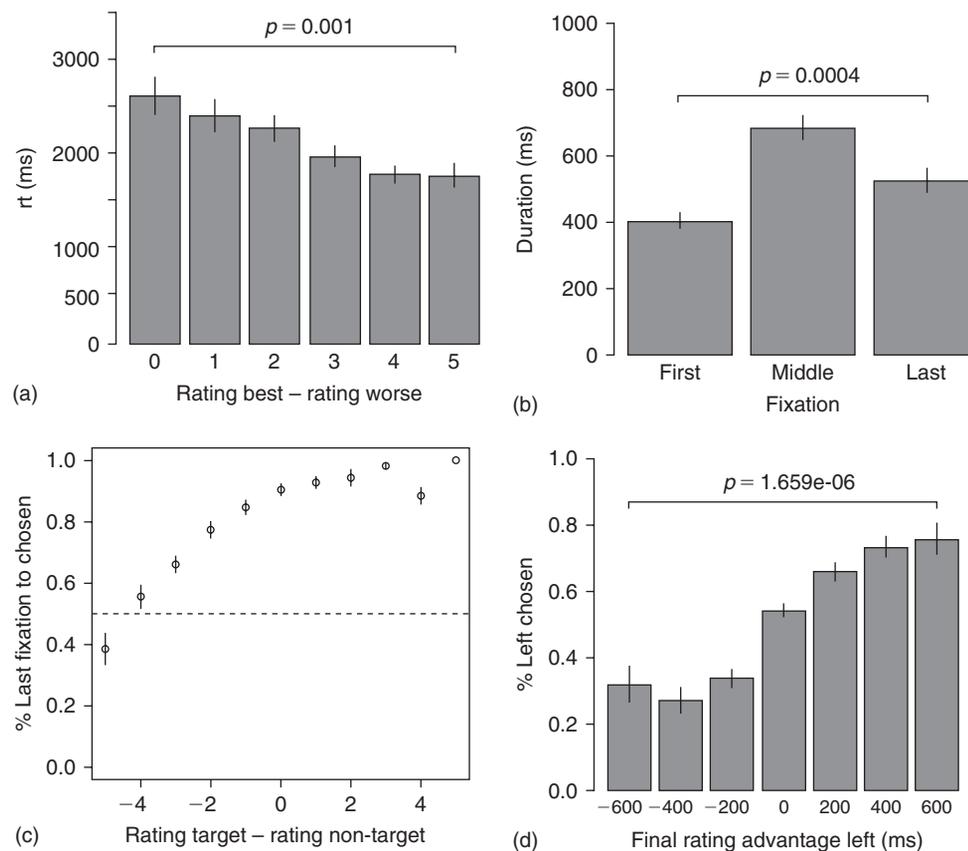
Krajbich and colleagues (2007) tested the behavioral implications of the model using eye-tracking as well as behavioral responses and reaction times. Figure 28.9 illustrates some of their results. The first panel shows that reaction times increase with the difficulty of the choice, as measured by the difference in an independently obtained measure of the value of the items. The second panel shows that last fixations are shorter. The third panel shows that the last fixation tends to be to the chosen item as long as the last-seen item is sufficiently better than the next-to-last item. Finally, the last panel shows that there is a duration fixation bias: the probability of choosing the left item increases with the relative amount of time for which it was seen during the choice process.

An interesting implication of the model is that it should be possible to manipulate choices by *exogenously* manipulating visual attention. Armel and colleagues (2008) have tested this prediction using a version of the experiment in which the fixation durations are controlled by the experimenter. Consistent with the predictions of the model, they show that longer exposure times increase the probability that appetitive items be chosen, but decrease the probability of choice for aversive items.

### Neural Basis for the Model's Components

As in the case of perceptual decision-making, a key test of the model of is the presence of neurobiological signals that resemble the computations that it makes. Unfortunately, much is less is known about the neural basis of this type of choice.

Several studies have found that activity in the orbitofrontal cortex (OFC) is correlated with behavioral measures of the goal-value assigned to objects. For example, in an fMRI study depicted in Figure 28.10, Plassmann *et al.* (2007) showed pictures of snacks to hungry subjects, who then placed bids for the right to eat them at the end of the experiment. The size of the bids was positively correlated with BOLD activity in the OFC. For related fMRI findings, see Erk *et al.* (2002), Arana *et al.* (2003), Paulus and Frank (2003),



**FIGURE 28.9** Eye-tracking and behavioral predictions of the goal-directed choice model. (a) Reaction times increase with the difficulty of the choice, as measured by the difference in an independently obtained measure of value. (b) Last fixations are shorter. (c) The last fixation tends to be to the chosen item, unless the last seen item is sufficiently worse than that seen next-to-last. (d) There is a duration fixation bias: the probability of choosing the left item increases with the relative amount of time for which it was seen during the choice process. See Krajbich *et al.* (2007) for more details, predictions, and analyses.

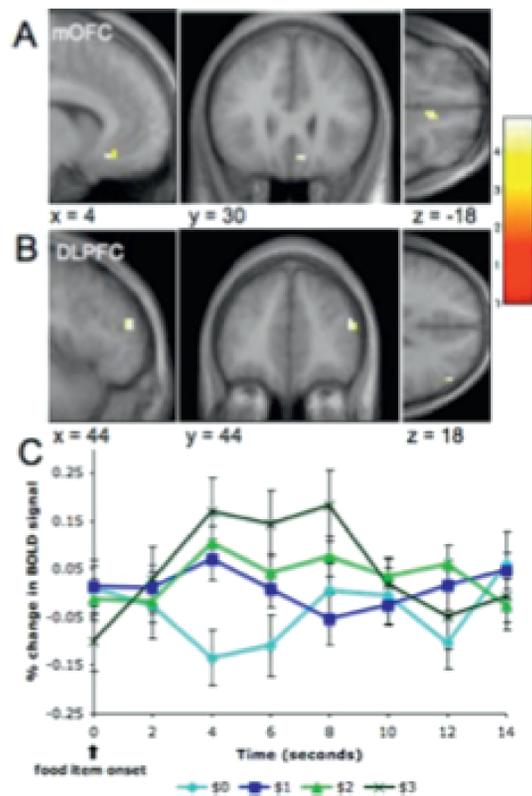
Valentin *et al.* (2007), and Hare *et al.* (2008). In a closely related monkey electrophysiology study of binary choice, Padoa-Schioppa and Assad (2006) showed that a population of neurons in OFC encoded for the subjective value of the two outcomes independently of the action that had to be taken to obtain them. Also related, Fellows and Farah (2007) showed that patients with lesions to the OFC were incapable of making consistent hypothetical food choices. Together, these findings suggest that activity in the medial OFC encodes for the value of the outcomes associated with the actions and, based on the study of Fellows and Farah, that these computations are essential for making consistent goal-directed choices.

Activity correlated with the value of outcomes has also been recorded in the dorsolateral prefrontal cortex (Leon and Shadlen, 1999; Wallis and Miller, 2003; Hare *et al.*, 2007; Plassmann *et al.*, 2007). Camus and colleagues (2008) have recently shown that disruption of activity in the right dorsolateral prefrontal

cortex using rTMS interferes with subjects' ability to compute the goal-values of items. The relationship between the value computations in DLPFC and in OFC is an important open question.

A crucial question for the models described above is whether we can find a neural signal resembling the relative value variable. Pioneering studies in this area include those by Platt and Glimcher (1999), Dorris and Glimcher (2004), and Sugrue *et al.* (2004), which have recorded neurons in monkey LIP during a goal-directed choice task. Figure 28.11a describes the experimental set-up of Sugrue *et al.* (2004), which is a little different from the simple goal-directed choices that we have discussed so far. Monkeys make choices between two saccade targets (red and green), and the location of the targets is randomized in every trial. The likelihood of getting a juice reward after selecting a target depends on the history of previous actions. Choices are indicated by making a saccade to either the left or right targets. Note that this task closely resembles

p0600



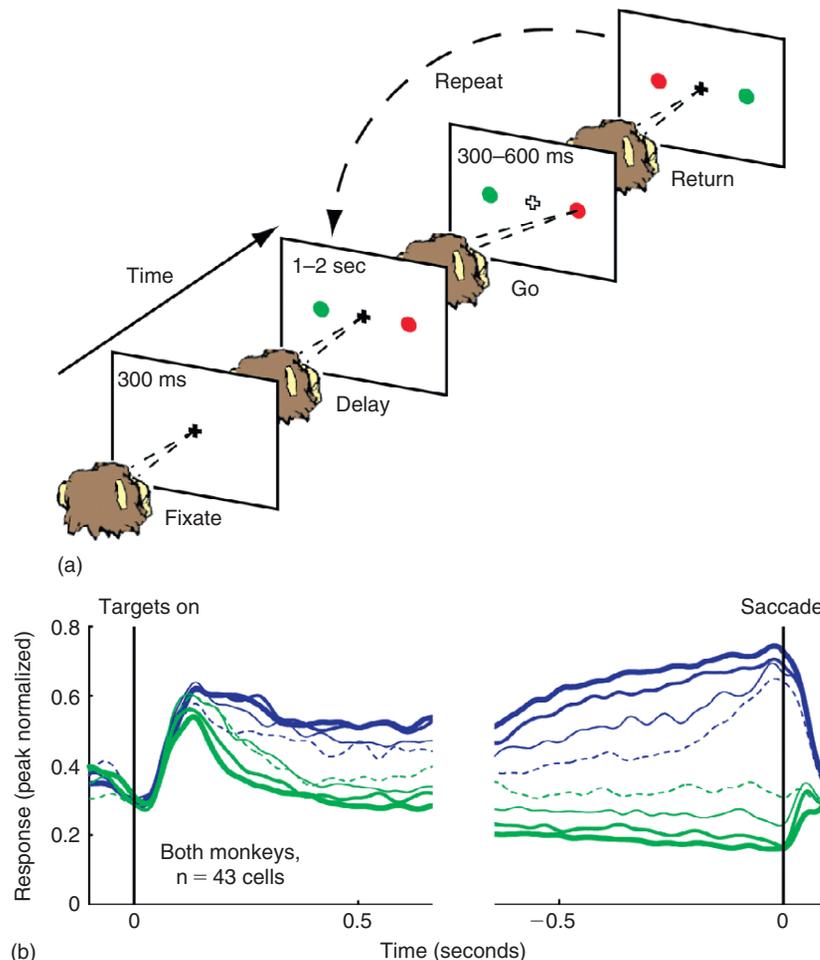
**FIGURE 28.10** fMRI study of goal-directed valuation in orbitofrontal cortex. Hungry subjects were shown pictures of snacks and placed bids for the right to eat them at the end of the experiment. (a, b) The size of the bids, which are a measure of the goal values computed by the brain, were positively correlated with BOLD activity in the orbitofrontal cortex and right dorsolateral prefrontal cortex. (c) The average time-course of the BOLD signal is higher in response to items with larger goal values. Adapted from Plassmann *et al.* (2007), with permission.

the DRM task, except that the value of the targets now depends on the history of actions, and not on the state of a visual stimulus. Figure 28.11b depicts activity in LIP during the course of a trial, and several aspects of the result are worth highlighting. First, there are strong parallels with the case of perceptual decision-making, which is perhaps not surprising, given that both are instances of goal-directed choice. Second, activity in LIP is increasing with the relative value of the value of the target in the response field, and it follows a time-course that is remarkably similar to the one in Figure 28.6 – the key difference being that activity is now sorted by the value of the target instead of the strength of the stimulus. Third, activity in LIP seems to reach a threshold just before a movement into the neuron's response field, also just as before. Although so far the evidence is only correlational, the parallel with the RDM task suggests that LIP is encoding something akin to the relative value signal (Dorris

and Glimcher, 2004). For more details on these experiments, see Chapter 31 of this volume.

Note that since the monkey maintains a central fixation during the evaluation of the targets, it is unclear whether attention plays a role in the task. One possibility is that the monkey is using covert attention in a way that is not captured by the aggregated data analysis; another is that information about the two values is being received and processed in parallel.

The excitement of these findings notwithstanding, there are many open questions regarding the neurobiology of goal-directed choice. First, what is the role of the signal encoded in LIP? Second, as discussed above, an integrator-like signal was found in the SC and FEF in the RDM task. Are similar signals present in goal-directed choice? Since these areas seem to be involved in the preparation and execution of eye movements, it is natural to speculate that this will be the case when the goal-directed choices are made through a saccade, but not when they require other types of motor output (e.g., a vocal command or a hand movement). Third, which brain regions perform computations that are specific to each type of motor output, and which perform computations that apply regardless of how the action is executed? Although much research remains to be done, the existing evidence suggests the following conjecture. Areas that are primarily involved in sensory-specific processing, such as MT in the case of vision, should play a central role in perceptual choice, but not in general goal-directed choice. Areas that play a central role in the preparation and execution of particular motor outputs, such as LIP, FEF, and SC in the case of eye movements, should play a central role in both types of choices as long as the decision is expressed through that motor system. Finally, areas of the prefrontal cortex such as the DLPFC and the OFC, as well as parts of the basal ganglia not discussed here, should play a role in both types of decisions regardless of the system used to signal the choice. (For a detailed review of the role of the basal ganglia, see Chapters 25 and 32 of this volume). Fourth, the models presented in this section assume that choices are made over actions, not over outcomes. Is this a justified assumption? In the absence of further evidence, an equally plausible possibility is for the choices to be made over outcomes, and then for the necessary motor commands to be executed. Another possibility is for the brain to make choices simultaneously over outcomes and actions, and then somehow to coordinate the result. Finally, how are considerations such as risk, uncertainty, time, and social effects incorporated into the machinery of goal-directed choice? (For more on this point, see Chapter 29 of this volume.)



**FIGURE 28.11** Activity in LIP during a goal-directed choice task. (a) The task. To begin a run of trials, the animal must fixate on the central cross. Saccade targets (red and green) then appear at randomized locations while the animal maintains fixation. Dimming of the fixation cross then signals that a saccadic response can be made. After a period of maintained fixation, the targets appear again at a random location and the trial starts again. The reward from choosing a target depends on the history of previous actions (see Sugrue *et al.*, 2004, for details). (b) Average activity in LIP neurons is plotted as a function of time and is aligned to either the appearance of the targets at the beginning of the trial (left panel) or to the time of the saccadic response (right panel). Blue curves denote trials in which the monkey made a choice into the neuron's response field. Green curves denote trials in which the choice was outside of the response field. In both cases, thicker lines indicate that the chosen target has a higher associated expected reward. Note that these time-courses closely resemble those in Figure 28.3; the only difference is that now LIP is discriminating based on the expected value of the chosen targets instead of stimulus coherence.

s0130

## FINAL REMARKS

p0630

A central goal of neuroeconomics is to construct a theory of value-based choice that is grounded in neurobiology. The problem of the construction and comparison of values, especially for the special case of perceptual decision-making, is perhaps the most advanced example of this research agenda. As has been shown in this review, we now have models describing the computations that the brain makes in order to make simple perceptual and goal-directed choices. The models make adequate behavioral predictions about choices, reaction times, and fixation

patterns in a wide class of experimental settings. Furthermore, electrophysiology, fMRI, rTMS, lesion, and microstimulation studies are beginning to demonstrate that some of the key computations made in the models are instantiated in the brain. Thus, although much remains to be done, there is reason to hope that rapid progress will be made on these issues over the next decade.

## Acknowledgments

Financial support from the NSF (SES-0134618) and the Gordon and Betty Moore Foundation is gratefully

s0140

p0640

acknowledged. I thank Paul Glimcher, Colin Camerer, and Todd Hare for their useful comments.

## References

- Arana, F.S., Parkinson, J.A., Hinton, E. *et al.* (2003). Dissociable contributions of the human amygdala and orbitofrontal cortex to incentive motivation and goal selection. *J. Neurosci.* 23, 9632–9638.
- Armel, C., Beuamel, A., and Rangel, A. (2008). The impact of computation time on decision values. *Am. Econ. Rev.*, (forthcoming).
- Balleine, B.W. (2005). Neural bases of food-seeking: affect, arousal and reward in corticostriatal limbic circuits. *Physiol. Behav.* 86, 717–730.
- Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends Cogn. Sci.* 11, 118–125.
- Bogacz, R. and Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput.* 19, 442–477.
- Born, R.T. and Bradley, D.C. (2005). Structure and function of visual area MT. *Annu. Rev. Neurosci.* 28, 157–189.
- Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Vis. Neurosci.* 10, 1157–1169.
- Brown, E., Gao, J., Holmes, P. *et al.* (2005). Simple neural networks that optimize decisions. *Intl J. Bifurc. Chaos* 15, 803–826.
- Brown, S. and Heathcote, A. (2005). A ballistic model of choice response time. *Psychological Rev.* 112, 117–128.
- Busemeyer, J.R. and Johnson, J.G. (2004). Computational models of decision making. In: D. Koehler and N. Narvey (eds), *Handbook of Judgment and Decision Making*. New York, NY: Blackwell Publishing Co., pp. 133–154.
- Busemeyer, J.R. and Townsend, J.T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Rev.* 100, 432–459.
- Camus, M., Shimojo, S., Camerer, C. *et al.* (2008). rTMS of right dorsolateral prefrontal cortex disrupts the computation of goal-directed values. Working Paper, Caltech.
- Celebrini, S. and Newsome, W.T. (1994). Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. *J. Neurosci.* 14, 4109–4124.
- 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. (2008). The role of value systems in decision making. In: C. Engel and W. Singer (eds), *Better Than Conscious? Implications for Performance and Institutional Analysis*. Cambridge, MA: MIT Press, (in press).
- Dayan, P., Niv, Y., Seymour, B., and Dayan, P. (2006). The misbehavior of value and the discipline of the will. *Neural Networks* 19, 1153–1160.
- Ditterich, J. (2006). Stochastic models of decisions about motion direction: behavior and physiology. *Neural Networks* 19, 981–1012.
- Ditterich, J., Mazurek, M.E., and Shadlen, M.N. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nat. Neurosci.* 6, 891–898.
- 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.
- Dubner, R. and Zeki, S.M. (1971). Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus in the monkey. *Brain Res.* 35, 528–532.
- Erk, S., Spitzer, M., Wunderlich, A. *et al.* (2002). Cultural objects modulate reward circuitry. *NeuroReport* 13, 2499–2503.
- Fellows, L.K. and Farah, M.J. (2007). The role of ventromedial prefrontal cortex in decision making: judgment under uncertainty or judgment *per se*? *Cerebral Cortex* 17(11), 2669–2674.
- Glimcher, P.W. (2003). The neurobiology of visual-saccadic decision making. *Annu. Rev. Neurosci.* 26, 133–179.
- Gold, J.I. and Shadlen, M.N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16.
- Gold, J.I. and Shadlen, M.N. (2002). Banburisms and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* 36, 299–308.
- Gold, J.I. and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 374–535.
- Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nat. Neurosci.* 9, 682–689.
- Hare, T., O’Doherty, J., Camerer, C. *et al.* (2008). Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *J. Neurosci.*, (forthcoming).
- Horwitz, G.D. and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science* 284, 1158–1161.
- Horwitz, G.D. and Newsome, W.T. (2001a). Target selection for saccadic eye movements: direction-selective visual responses in the superior colliculus. *J. Neurophysiol.* 86, 2527–2542.
- Horwitz, G.D. and Newsome, W.T. (2001b). Target selection for saccadic eye movements: prelude activity in the superior colliculus during a direction-discrimination task. *J. Neurophysiol.* 86, 2543–2558.
- Huk, A.C. and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *J. Neurosci.* 25, 10420–10436.
- Kandell, E.R., Schwartz, J., and Jessell, T. (2002). *Principles of Neural Science*. New York, NY: McGraw-Hill.
- Krajbich, I., Armel, C., and Rangel, A. (2007). Visual attention drives the computation of value in goal-directed choice. Working Paper, Caltech.
- Laming, D. (1979). A critical comparison of two random-walk models for choice reaction time. *Acta Psychologica* 43, 431–453.
- 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.
- Link, S.W. (1992). *The Wave Theory of Difference and Similarity*. Hillsdale, NJ: Lawrence Erlbaum.
- 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.
- Luce, R.D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford: Oxford University Press.
- Mazurek, M.E., Roitman, J.D., Ditterich, J., and Shadlen, M.N. (2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex* 13, 1257–1269.
- McClelland, J. and Rumelhart, D. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Rev.* 88, 375–407.
- McCoy, A. and Platt, M.L. (2005). Expectations and outcomes: decision-making in the primate brain. *J. Comp. Physiol. A. Neuroethol. Sens. Neural Behav. Physiol.* 191, 201–211.
- McMillen, T. and Holmes, P. (2005). The dynamics of choice among multiple alternatives. *J. Math. Psychol.* 50, 30–57.
- Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. *Nature* 341, 52–54.
- Padoa-Schioppa, C. and Assad, J.A. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441, 223–226.

- Palmer, J., Huk, A.C., and Shadlen, M.N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vision* 5, 376–404.
- Paulus, M.P. and Frank, L.R. (2003). Ventromedial prefrontal cortex activation is critical for preference judgments. *NeuroReport* 14, 1311–1315.
- 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. (2002). Neural correlates of decisions. *Curr. Opin. Neurobiol.* 12, 141–148.
- Platt, M.L. and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Pulst-Korenberg, J. and Rangel, A. (2008). The role of visual attention in trinary economic choice. Working Paper, Caltech.
- Rangel, A., Camerer, C., and Montague, R. (2008). Neuroeconomics: the neurobiology of value-based decision making. Working Paper, Caltech.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Rev.* 85, 59–108.
- Ratcliff, R., Cherian, A., and Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *J. Neurophysiol.* 90, 1392–1407.
- Ratcliff, R. and Smith, P. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Rev.* 111, 333–367.
- Roe, R.M., Busemeyer, J., and Townsend, J.T. (2001). Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychological Rev.* 108, 370–392.
- 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.
- Romo, R. and Salinas, E. (2003). Flutter discrimination: neural codes, perception, memory and decision making. *Nat. Rev. Neurosci.* 4, 203–218.
- Salzman, C.D., Britten, K.H., and Newsome, W.T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature* 346, 174–177.
- Salzman, C.D., Murasugi, C.M., Britten, K.H., and Newsome, W.T. (1992). Microstimulation in visual area MT: effects on direction discrimination performance. *J. Neurosci.* 12, 2331–2355.
- Schall, J.D. (2005). Decision making. *Curr. Biol.* 15, R9–11.
- Shadlen, M.N. and Newsome, W.T. (1996). Motion perception: seeing and deciding. *Proc. Natl Acad. Sci. USA.* 93, 628–633.
- 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.
- Smith, P. (1995). Psychophysically-principled models of visual simple reaction time. *Psychological Rev.* 102, 567–593.
- Smith, P. (2000). Stochastic dynamic models of response time and accuracy: a foundational primer. *J. Math. Psychol.* 44, 408–463.
- Smith, P.L. and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika* 25, 251–260.
- Sugrue, L.P., Corrado, G., 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: MIT Press.
- Usher, M. and McClelland, J. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Rev.* 108, 550–592.
- Valentin, V.V., Dickinson, A., and O'Doherty, J. (2007). Determining the neural substrates of goal-directed learning in the human brain. *J. Neurosci.* 27, 4019–4026.
- Wald, A. and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Ann. Math. Stats* 19, 326–339.
- 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. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36, 955–968.
- Wong, K.F. and Wang, X.J. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314–1328.
- Zeki, S.M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *J. Physiology*, 236, 549–573.