

Banburismus and the Brain: Decoding the Relationship between Sensory Stimuli, Decisions, and Reward

Review

Joshua I. Gold^{1,3} and Michael N. Shadlen²

¹Department of Neuroscience

University of Pennsylvania

Philadelphia, Pennsylvania 19104

²Howard Hughes Medical Institute

Department of Physiology & Biophysics and

Regional Primate Research Center University

of Washington

Seattle, Washington 98195

This article relates a theoretical framework developed by British codebreakers in World War II to the neural computations thought to be responsible for forming categorical decisions about sensory stimuli. In both, a weight of evidence is computed and accumulated to support or oppose the alternative interpretations. A decision is reached when the evidence reaches a threshold value. In the codebreaking scheme, the threshold determined the speed and accuracy of the decision process. Here we propose that in the brain, the threshold may be controlled by neural circuits that calculate the rate of reward.

In the early 1940's, Alan Turing and his colleagues at Bletchley Park broke the supposedly unbreakable Enigma code used by the German navy. They succeeded by finding in the encoded messages the barest hints of evidence to support or refute various hypotheses about the encoding scheme that they could exploit to determine the contents of the message. Their success rested, in part, on a mathematical framework with three critical components: a method of quantifying the weight of evidence provided by individual clues toward the alternative hypotheses under consideration, a method of updating this quantity given multiple pieces of evidence, and a decision rule to determine when the evidence was sufficient to render a judgment on the most likely hypothesis (Good, 1979). This framework has been shown to be of general use for making decisions about sequentially sampled data (Wald, 1947). Recent progress in neurobiology suggests that it may also offer an account of decisions about a different kind of encoded message: the representation of sensory stimuli in the brain. Our goal in this manuscript is to relate this theoretical framework to the operations that neurons perform to generate categorical decisions about sensory stimuli. A dividend of this framework is a prescription for how the representation of reward—more specifically, the rate of reward—is incorporated into neural computations that underlie decision formation.

In the first section, below, we review Turing's theoretical framework and its role in cracking the Enigma code. We describe how the codebreakers implemented this framework in a process that they called "Banburismus,"

named after the town of Banbury where they printed special sheets of paper that allowed them to carry out the computations that were central to this process. In the second section, we apply this framework to simple, two-alternative decisions on response-time perceptual tasks. We focus on how the critical components of the theory might be implemented in the brain. Central to this neural implementation is the relationship between the rate of reward and the decision rule that determines how and when to read out the sensory evidence. In the third section, we summarize recent experimental findings that support this model.

Background: Banburismus at Bletchley Park

To illustrate the framework developed by Turing, consider the following problem encountered by the codebreakers at Bletchley Park. The physical configuration of an Enigma machine, which included interchangeable rotors and a modifiable plugboard, determined which of over a billion billion encoding schemes was being used at any given time. Even more cunningly, the configuration of the machine changed for each character it encoded, so that knowing the identity of one encoded character would not necessarily reveal the identity of any other. Fortunately for the codebreakers, the Germans' procedures for using the machine greatly reduced the uncertainty associated with each message. For example, all of the machines in use on a given day were set to very similar configurations at the beginning of each message, differing only in the positions of the rotors (each of three rotors could be set in any of 26 positions, corresponding to 17,576 different schemes). Thus, it was reasonable to expect that some pairs of the thousands of messages intercepted in a single day would come from machines in exactly (or very nearly) the same state. This fact turned out to be essential for the British codebreakers, who realized that identifying these pairs could help them to infer other details about the state of the Enigma machines on that day (specifically, the identity of the first and sometimes second rotor; for more details, see <http://www.codesandciphers.org.uk> by Tony Sales; MacKay, 2002, currently available at <http://www.inference.phy.cam.ac.uk/mackay/itpmn>; Hodges, 2000). They therefore developed a set of methods for determining whether or not two different messages were encoded by machines in the same configuration (that is, whether they were encoded by the same scheme).

Formally, the problem encountered by the codebreakers came down to deciding which hypothesis— h_0 , that the encoding scheme for a given pair of messages was different, versus h_1 , that the encoding scheme was the same—was more probable, given the weight of evidence from the messages themselves. The key was in recognizing that if h_1 were true, the encoded messages would not be quite as random as they were otherwise. This lack of randomness would be evident in pairs of characters, one from each of the two messages. Specifically, they knew that the output of a given Enigma machine consisted of a roughly uniform distribution of charac-

³Correspondence: jgold@mail.med.upenn.edu

ters—a great virtue of the machine to its users. Moreover, two machines that were in different configurations (i.e., h_0) had equally random outputs. Such machines, encoding two different messages, thus produced pairs of characters that would be expected to match each other only as often as two letters picked independently from uniform distributions of all 26 letters, which is equal to a rate of 1/26. However, if two machines happened to be in the same configuration (h_1), the situation is quite different. In this case, even when encoding different messages, the two machines would produce pairs of characters that matched each other if (and only if) the letters they encoded also matched each other. Therefore, the expected “match rate” from the encoded messages would be the same as that from the unencoded messages, which is equal to $\sim 1/13$ in ordinary German text. In other words, the likelihood of getting a match (m) given h_1 , denoted $\Pr(m|h_1)$, was greater than the likelihood of getting a match given h_0 , denoted $\Pr(m|h_0)$. Accordingly, Turing defined the weight of evidence provided by a match (or, similarly, a non-match) in favor of h_1 over h_0 as the logarithm of the ratio of these likelihoods:

$$\text{Weight of evidence} = \log \left[\frac{\Pr(m|h_1)}{\Pr(m|h_0)} \right] \quad (1)$$

As a logarithm of probabilities, this quantity is measured in units that depend on the base of the logarithm (Good, 1979). In Banburismus, logarithms of base 10 were used and were called “bans” (the Banburismus process actually computed weights of evidence in units of $1/10^{\text{th}}$ of a ban—a deciban—that the codebreakers considered to be “about the smallest weight of evidence that is directly perceptible to human intuition”; Good, 1979). Likewise, logarithms of base e were called “natural bans.” It is worth noting that other, closely related quantities that are based on the logarithm of probabilities can be measured in the same units. Thus, “surprisal” and “mutual information” can be expressed in bans but are more familiar in units that use a logarithm of base 2: “bits.”

A quantity based on the logarithm of probabilities has another useful feature: additivity. Writing over 100 years ago about the degree of “belief” in a hypothesis that is provided by the evidence, Peirce noted that “two arguments which are entirely independent, neither weakening nor strengthening each other, ought, when they concur, to produce a belief equal to the sum of the intensities of belief which either would produce separately” (Peirce, 1878). Equation 1 fulfills this requirement, such that the weight of evidence provided by two independent sources is equal to the sum of the weights of evidence that they provide individually. Thus, for Turing, the overall weight of evidence for h_0 or h_1 that was provided by all the corresponding pairs of letters in two given messages could be computed by counting the number of matches and non-matches, then simply adding up the weight of evidence provided by each.

This tally of matches and non-matches was a form of sequential analysis, in which the overall weight of evidence bearing on the hypotheses under consideration is updated given each new piece of evidence (in this case, the next pair of letters). Consistent with this

kind of analysis, decisions in Turing’s Banburismus scheme were made by considering how much weight of evidence they needed to collect (in the form of matched and non-matched pairs) to be confident that one hypothesis has been sufficiently proven. Specifically, they considered two “barriers” to bookend the process of accumulating decibans; one was a positive number describing the weight of evidence needed to decide h_1 , the other a negative number describing the weight of evidence needed to decide h_0 . When the weight of evidence crossed either barrier, the decision was made. A similar formulation for statistical applications was developed independently by Wald shortly after the war (Wald, 1947).

As Turing was aware, this formulation had two distinct advantages. First, it had a definite stopping time that instructed the codebreaking team when to stop working on one pair of ciphers and turn their attention to another. Second, it predicted the accuracy of the decision, as follows. By Bayes’ theorem, the weight of evidence can be related to several other probabilities:

$$\log \left[\frac{\Pr(h_1|m)}{\Pr(h_0|m)} \right] = \log \left[\frac{\Pr(h_1)}{\Pr(h_0)} \right] + \text{weight of evidence}, \quad (2)$$

where $\Pr(h_0|m)$ and $\Pr(h_1|m)$ are called the posterior probabilities and describe the probability of each hypothesis after all the evidence has been sampled, and $\Pr(h_0)$ and $\Pr(h_1)$ are called the prior probabilities and describe the probability of each hypothesis before any evidence is sampled. If, for example, there is an equal prior probability of either hypothesis (which was the assumption made in Banburismus), then the weight of evidence is equal to the logarithm of the ratio of the posterior probabilities. In this case, for a given weight of evidence (e.g., the stopping point in Banburismus):

$$\log \left[\frac{\Pr(h_1|m)}{\Pr(h_0|m)} \right] = \text{weight of evidence} = B \quad (3)$$

where B is a constant that represents the barrier height in favor of h_1 . For two mutually exclusive hypotheses ($\Pr(h_0|m) = 1 - \Pr(h_1|m)$ or, equivalently, $\Pr(h_0) = 1 - \Pr(h_1)$), and assuming the weight of evidence is expressed in bans, this rearranges to:

$$\Pr(h_1|m) = \frac{1}{1 + 10^{-B}}. \quad (4)$$

Equation 4 indicates that the posterior probability of h_1 depends only on the value of the barrier, B , and not on the particular samples of evidence, m , encountered. In other words, as long as the weight of evidence reaches B bans, the probability that h_1 is correct is a fixed value.

Thus, in Banburismus, setting the height of the barrier that determined the weight of evidence to accumulate before committing to a decision was equivalent to setting an expected level of performance. For example, accumulating matches and non-matches until the weight of evidence reached 2 bans would provide approximately 100 to 1 odds that the decision was correct. Of course, the greater the weight of evidence required, the more time the codebreaker must spend on each message. It is this critical feature of the decision rule, explicitly linking the interpretation of the evidence to the

speed and accuracy of the decision process, that turns out to have important implications for how the brain might implement a kind of Banburismus to form decisions about sensory stimuli, as we discuss in the next section.

Implementation: Banburismus Performed by Neurons

In this section, we explore in detail the relationship between Banburismus and an application for which it appears to be particularly well suited: decisions about sensory stimuli. We focus on how the brain might implement the key components of this process, using operations that neurons perform. We address three specific questions: (1) How does the activity of sensory neurons relate to Turing’s “weight of evidence”? (2) How does the brain make use of all of the sensory evidence to which it has access? (3) What is the decision rule based on this evidence?

We consider these questions in the context of decisions between two alternatives (e.g., Is it present or not? Does it match or not? Which is larger?), because most psychophysical and electrophysiological studies to date address these kinds of decisions (Parker and Newsome, 1998; Schall, 2001). We place particular emphasis on decisions about the direction of dynamic, random dots, which have been studied extensively in behavioral experiments in humans (Blake and Hiris, 1993; Braddick, 1995; Morgan and Ward, 1980; Watamaniuk and Sekuler, 1992) and combined behavioral and electrophysiological experiments in monkeys (reviewed in Albright and Stoner, 1995; Britten, 2003; Orban, 1997). To formulate this decision as above, the problem is to decide between h_1 , that motion was (say) upward, versus h_0 , that motion was downward. It is known from recording, microstimulation, and inactivation studies that the neurons in extrastriate areas MT and MST of the macaque represent the motion information that the monkey uses to form its direction decision (Britten et al., 1992; Celebrini and Newsome, 1994, 1995; Newsome and Pare, 1988; Pasternak and Merigan, 1994; Salzman et al., 1992; Salzman and Newsome, 1994). Thus, neural activity in these areas represents the encoded, noisy evidence that the brain must interpret to decide which hypothesis was most probable.

Relating Sensory Signals to Turing’s Weight of Evidence

It is common to consider categorical decisions about noisy sensory signals by formulating quantities like Turing’s weight of evidence that are monotonically related to the likelihood ratio (Green and Swets, 1966). Briefly, the response of a sensory neuron—in this case, a rate measured in spikes per second—can be thought of as a random variable: for a given set of stimulus conditions (e.g., upward motion), it has an expected value but can vary considerably from moment-to-moment or trial-to-trial. Thus, the neural response can take on numerous values, each with a certain probability. These probabilities can be estimated by taking repeated measurements of the response of a neuron (or group of neurons) to a stimulus and generating a curve describing the distribution of responses to that stimulus. The height of such a curve at a given spike rate for a given stimulus (e.g.,

h_1) is the likelihood of obtaining that response given the stimulus: $\Pr(\text{response}|h_1)$. Thus, with two such curves, one generated using h_1 and the other generated using h_0 , it is possible to compute the weight of evidence for any given neural response, which is $\log(\Pr(\text{response}|h_1)/\Pr(\text{response}|h_0))$. However, this cumbersome method seems unlikely to be implemented in the brain, which would need access to a representation of the complete distribution of neural responses to every stimulus encountered.

As we have shown previously, there is an easier way for the brain to compute the weight of evidence (Gold and Shadlen, 2001): it can use the difference in spike rates from two neurons (or two pools of neurons). For example, consider a pair of MT neurons that represent the evidence corresponding to the two alternative hypotheses on the direction-discrimination task. For a decision between upward and downward motion, this “neuron/antineuron” pair consists of an upward-sensitive MT neuron (its response, x , tends to be larger if h_1 is true) and a downward-sensitive neuron (its response, y , tends to be larger if h_0 is true). As shown in Figure 1, the weight of evidence (computed in natural bans) provided by these two responses together is proportional to their difference. Thus, a positive value of $(x - y)$ implies h_1 was more likely; a negative value implies h_0 was more likely.

This simple neural computation of the weight of evidence is robust with respect to several potentially complicating factors. The proportional relationship between $(x - y)$ and the weight of evidence does not depend on the assumption of normality (which is illustrated in Figure 1). For example, it also holds for the exponential and Poisson distributions (Gold and Shadlen, 2001). This relationship holds even if the responses x and y are not independent. For example, assume a bivariate normal distribution with equal variances for x and y under h_1 (and assume opponency, giving an analogous expression under h_0):

$$\Pr(x,y|h_1) = \frac{e^{-\frac{Q}{2}}}{2\pi\sigma^2\sqrt{1-\rho^2}}, \tag{5}$$

where

$$Q = \frac{1}{1-\rho^2} \left[\frac{(x-\mu_1)^2 + (y-\mu_0)^2 - 2\rho(x-\mu_1)(y-\mu_0)}{\sigma^2} \right],$$

σ^2 is the common variance, μ_1 and μ_0 are the means of x and y , respectively, under h_1 , and ρ is the covariance between x and y . Solving for the weight of evidence (in units of natural bans) for h_1 versus h_0 yields:

$$\begin{aligned} \text{weight of evidence} &= \log \frac{\Pr(x,y|h_1)}{\Pr(x,y|h_0)} \\ &= \frac{(\mu_1 - \mu_0)}{\sigma^2} \cdot \frac{1}{(1 - \rho)} \cdot (x - y) \end{aligned} \tag{6}$$

Note that numerous factors could cause nonzero values of ρ , including intrinsic factors like attention and arousal and extrinsic factors like the light level or other variations in the stimulus. For example, for the motion-discrimination task, motion in a given direction is typically presented at a variety of strengths (i.e., h_0 and h_1 each




Figure 1. Theoretical Relationship between Neural Responses and Turing's Weight of Evidence

The panels on the left depict hypothetical distributions of responses from a pair of sensory neurons when h_0 or h_1 is true, as indicated. For example, the upper curve could represent the responses of a direction-selective neuron that prefers upward motion. The lower curve would then represent the responses of the opposing “antineuron” that prefers downward motion. Here we make the simplifying assumption that the responses are normally distributed with mean values of μ_0 and μ_1 and equal variances (σ^2). Note that other forms of the distribution, including the exponential and Poisson, lead to similar results (Gold and Shadlen, 2001). The panel on the right shows Turing's weight of evidence computed from these responses plotted as a function of their difference; note the linear relationship.

encompass several different stimuli of different motion strengths). Because x and y both depend on motion strength, they are not independent: a strong stimulus that would tend to produce a relatively large value of x would also tend to produce a relatively small value of y . Equation 6 indicates that such a dependence between x and y does not alter the fact that the difference ($x - y$) is proportional to the weight of evidence.

Thus, under a variety of assumptions, the weight of evidence can be computed by taking a simple difference between the responses of two opposing neurons. The obvious advantage of this scheme is the simplicity of the computation. A potential pitfall is that it appears to require a constant of proportionality that depends on the parameters of the distributions describing the responses; e.g., μ_0 , μ_1 , σ , and ρ in Equation 6. However, as we discuss in the following sections, combining multiple pieces of evidence does not necessarily require the constant of proportionality to be known. Moreover, the choice of an appropriate decision rule can, in principle, bypass the need to scale the weight of evidence appropriately.

Making Use of All of the Evidence

In Banburismus, multiple pieces of evidence—the matched and non-matched pairs of characters from the two messages—were combined by adding together the weight of evidence computed from each pair individually. The resulting quantity, called a decision variable, was the sole factor used to make the decision. Recall that a critical advantage of this computation was that the value of the decision variable determined the level of performance. In this section, we describe how the brain might compute a similarly useful decision variable.

Consider a case in which the sensory evidence comes from the ongoing responses from the neuron/antineuron pair. A single “piece” of evidence can be thought of as representing the difference in spike rates at a given moment in time. As more time goes on, more evidence is available. If the responses are independent and identically distributed over time, then accumulating the evidence corresponds to a temporal integration of the difference in spike rates:

$$\text{total weight of evidence at time } t = k \int_0^t (\chi(\tau) - y(\tau)) d\tau \quad (7)$$

where k is a constant. Note that if k does not equal the constant of proportionality in Equation 6, then the value that is computed is merely proportional to the weight of evidence. Regardless of the value of k , however, the algorithm for updating the weight of evidence is the same: accumulate the difference in spike rates over time. As illustrated in Figure 2, this temporally accumulating evidence can be thought of as simply a single piece of evidence that lends more weight to the decision as more time passes.

Thus, as in Turing's scheme, an accumulated difference of opposing sensory responses in the brain corresponds to the weight of evidence that can distinguish between the alternative hypotheses. For equal prior probabilities, the value of this decision variable is particularly easy to interpret: zero implies that both hypotheses are equally probable, an increasingly positive value implies a growing probability that one of the hypotheses is correct, and an increasingly negative value implies the alternative (Equation 2). It should be noted that the decision variable can, in principle, take into account many other factors and still distinguish between the hypotheses in such a straightforward manner. These factors include sensory evidence from multiple sources, unequal prior probabilities, and anticipated costs and benefits associated with the potential outcomes (Graham, 1989; Green and Swets, 1966). How these complicated and in some cases subjective factors are computed in the brain is beyond the scope of this article. Instead, we will simply assume that the decision variable reflects the posterior probabilities as if it were computed as an accumulation of evidence from a single source and equal prior probabilities. As we describe in the next section, a particularly useful consequence of this assumption is that it implies a close relationship between the decision rule and the rate of reward.

The Decision Rule

Here we consider a decision rule that places barriers on the value of the decision variable. In Banburismus, the height of the barriers represented a trade-off be-

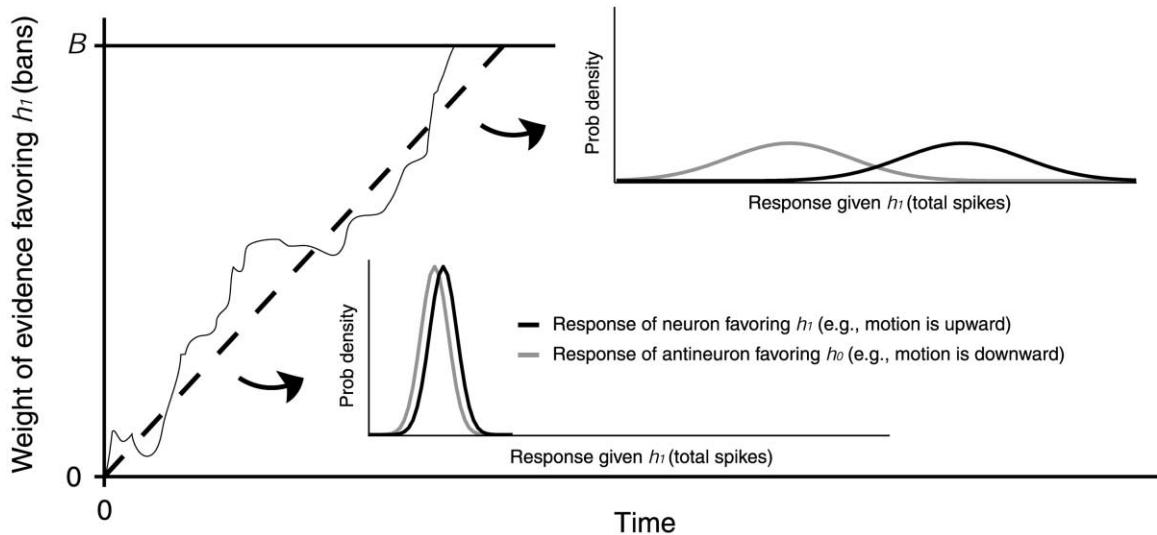


Figure 2. Decision Model

The weight evidence in favor of h_1 over h_0 is the accumulated difference between responses of a neuron that prefers h_1 and an antineuron that prefers h_0 (see Equation 7). The curves depicted are computed under the condition that h_1 is true. The thin, wavy line depicts a simulated trajectory that represents how the weight of evidence might grow on a single trial as a function of time. The dashed line depicts the expectation (mean value) of this trajectory at each time point. Note that changing the constant of proportionality used to relate the accumulated difference to the weight of evidence (see text for details) simply scales the ordinate. The two insets illustrate the correspondence between the weight of evidence and the underlying neuron and antineuron responses at two time points, viewed as hypothetical (normal, equivariate) probability density functions. These functions describe the distributions of the total number of spikes generated by the neuron and antineuron up to that time point, given that h_1 is true. If the weight of evidence reaches the barrier at B , the process is stopped and a decision is rendered for h_1 . This is the expected outcome when h_1 is true. If the weight of evidence reaches $-B$ (not shown), a decision for h_0 is issued in error.

tween accuracy and speed. Higher barriers meant that more evidence was accumulated, which, in turn, implied a greater probability of getting the correct answer. However, accumulating more evidence took more of the codebreakers' valuable time. In the brain, a similar trade-off occurs. We quantify this trade-off for a particular set of assumptions and show that the barrier can be adjusted to maximize a quantity like the rate of reward that depends on both accuracy and speed.

To illustrate this idea, we will focus on a reaction-time version of the direction-discrimination task, in which the subject is allowed to control the amount of time to make a response (Roitman and Shadlen, 2002). The sensory evidence is a difference signal from visual cortex. We assume that the difference is a random variable that is drawn from a normal distribution with a variance of one and a mean value that depends on the strength of motion (measured as a percentage of coherently moving dots, which typically takes values of 0, 3.2, 6.4, 12.8, 25.6, and 51.2). Upward and downward motion strengths are assigned positive and negative means, respectively. We also assume that there are equal prior probabilities of upward and downward motion (that is, both are presented with equal frequency). The decision process is simply to accumulate the difference (in 1 ms steps) until it reaches a threshold of evidence, $\pm B$. If the evidence reaches $+B$, the decision is upward (see Figure 2); if it reaches $-B$, the decision is downward. A class of models known as sequential sampling models share many of these features and have been particularly successful at explaining accuracy and response times on this kind of task (Audley and Pike, 1965; LaBerge, 1962; Laming,

1968; Link and Heath, 1975; Luce, 1986; Ratcliff and Rouder, 1998; Stone, 1960; Usher and McClelland, 2001; Vickers, 1979).

This problem is analogous to one-dimensional Brownian motion to a pair of barriers (Link, 1992; Ratcliff and Rouder, 1998). The means of the normal distributions represent the drift rates, μ_i (where the subscript i reflects the particular motion strength). The psychometric function describing the probability of correctly reaching either the "up" or "down" barrier on a given trial is:

$$p_i = \frac{1}{1 + e^{-2B|\mu_i|}}, \quad (8)$$

which, like Equation 4, is the logistic function (here using base e). Unlike in Equation 4, however, the probability correct in this function depends on not just the barrier height B , but also a coherence-dependent term, μ_i . This extra term is to account for the coherence-dependent scale factor that relates the weight of evidence to the difference signal that reaches B (Equation 6). Specifically, a fixed B corresponds to a smaller weight of evidence for lower coherences (i.e., weaker signals) than for higher coherences (stronger signals). The mean time to reach the barrier is:

$$t_i = \frac{B}{|\mu_i|} \tanh(B|\mu_i|), \quad (9)$$

where

$$\lim_{\mu \rightarrow 0} (t) = B^2.$$




Figure 3. Predicted Effect of Barrier Height on Speed, Accuracy, and Reward Rate in a Direction-Discrimination Task. The calculations assume that an experiment uses six levels of difficulty (see text for details). (A) Accuracy. The probability of a correct choice averaged across all six motion strengths is plotted as a function of barrier height. A higher barrier ensures that decisions are based on a larger weight of evidence and thus are more accurate on average. The dashed vertical line is the barrier height associated with the maximum rate of reward. (B) Reward rate. The reward rate depends on the decision time plus fixed times in and between trials plus a “time out” for an error. A maximum reward rate occurs when the decision process is sufficiently accurate but not overly time consuming. We suggest that the brain adjusts the barrier height to achieve the maximum rate of reward. (C) Psychometric function using the optimal barrier height. The probability of a correct response is plotted for each motion strength. By definition, the performance on the 0% coherent motion stimulus is at chance (not shown). The line is a best fitting cumulative Weibull, which is a reasonable approximation to the logistic function (Equation 8). (D) Response times using the optimal barrier height. The time from stimulus onset until behavioral response is plotted at each of the motion strengths. The response time is decision time (from Equation 9) plus 250 ms.

Note that the average decision time also depends on B and the particular list of motion strengths chosen by the experimenter. For details of these derivations, refer to Link (1992) and Ross (1983).

The critical question is how the brain sets the barrier height. Equations 8 and 9 quantify for this decision process the trade-off evident in Banburismus. A fixed barrier height corresponds to a fixed level of performance. Raising (or lowering) the barrier raises (lowers) performance and increases (decreases) the mean time to reach a decision. This trade-off between accuracy and time is inherent to this kind of decision rule and therefore is likely to be central to the neural mechanisms responsible for controlling the barrier. We suggest that these mechanisms involve maximizing the overall rate of reward in a given experiment. These equations allow us to quantify this rate, by calculating both the overall performance and the amount of time per trial.

Overall performance depends on the probability of crossing the correct barrier, p_i , at each motion strength, c_i . To calculate each p_i , we assume that the difference from upward and downward motion sensors—and thus the drift rate μ_i —is a linear function of motion strength (Britten et al., 1993). For illustration purposes (Figure 3), we have chosen $\mu_i = 0.4c_i$, which leads to a psychometric function that is qualitatively similar to data from monkeys (Roitman and Shadlen, 2002) and humans. For simplicity, we assume that accuracy is the same for upward or downward motion, and that all six motion strengths are shown with equal frequency. In this case, the overall performance is $\bar{p} = \text{Mean}[p_i]$, which is depicted in Figure 3A as a monotonically increasing function of the barrier height B .

Given the overall accuracy, determining the rate of reward requires calculating the amount of time per trial. This is the decision time, t_i , plus several other intervals. For simplicity in the following calculations, we assume

that there are three other intervals, each of a fixed duration: the non-decision time t_{fix} on each presentation (e.g., visual latencies and motor preparation), equal to ~ 250 ms; the time between trials $t_{\text{ITI}} = \sim 1$ s; and a penalty time added for errors $t_{\text{penalty}} = \sim 2$ s. Of course, the values are open to debate and even experimental manipulation. Exploiting a well-known theoretical result that the decision time on error trials is the same as on correct trials in this “diffusion-to-barrier” model (Luce, 1986), the average rate of reward is given by:

$$r = \bar{p} \left\{ \frac{1}{N} \sum_{i=1}^N p_i (t_i + t_{\text{fix}} + t_{\text{ITI}}) + (1 - \bar{p})(t_i + t_{\text{fix}} + t_{\text{ITI}} + t_{\text{penalty}}) \right\}^{-1} \quad (10)$$

Thus, the rate of reward r depends on the probability of a correct response and the average time per trial, which, in turn, depend on the particular list of motion strengths as well as the barrier height.

Unlike overall performance, the rate of reward is not a monotonic function of barrier height (Figure 3B). If the barrier height is too low, then there are too many errors, which do not yield a reward and add additional time. If the barrier height is too high, then time is wasted. The key insight is that the subject can maximize the rate of reward by setting B to the value shown by the vertical dashed line. This level determines the psychometric and response-time functions shown in Figures 3C and 3D. Our goal here is not to compare these functions to data (although they appear to be qualitatively similar; see, for example, Roitman and Shadlen, 2002) but to expose the relationship between the representation of reward and decision making.

How the brain might go about adjusting the level of the barrier, B , is an open question. In principle, this adjustment could be accomplished by a process of trial-

and-error, in which the barrier is raised and lowered until the maximum rate of reward is achieved. Interestingly, in this case, the constant of proportionality that relates the accumulated difference ($x - y$) to the weight of evidence is not needed to find the barrier height that leads to the maximum rate of reward. That information is not lost, however: once the barrier height is fixed, it corresponds to a particular level of overall performance and thus can be expressed in units of the weight of evidence, such as natural bans. Note that this quantity is not the weight of evidence that would be calculated based on knowledge of the stimulus motion strength and the associated sensory response distributions because that information would lead to perfect performance at all motion strengths. Rather, it is the weight of evidence that corresponds to a fixed level of uncertainty across all stimulus strengths in an experiment (this quantity will tend to overestimate the evidence from weak stimuli and underestimate the evidence from strong stimuli). Accordingly, the evidence that accumulates during a trial can be interpreted as a fraction of this quantity and thus in units of natural bans—even when the scaling between the decision variable and the weight of evidence is not known (e.g., if the brain does not know the shapes of the sensory response distributions).

To summarize, we have described a decision variable based on equal prior probabilities and a scaled version of the weight of evidence. A decision rule that places positive and negative boundaries on this variable leads to a fixed percentage of correct responses and predictable decision times. Raising the barriers improves performance but slows down the decision process. Lowering the barriers does the opposite, degrading performance as it speeds up. A measurement like the overall rate of reward is sensitive to this trade-off between accuracy and speed and thus could, in principle, be used to adjust the barrier height in the brain.

Neurophysiology: Banburismus in the Brain

In this section, we briefly review experimental evidence supporting the idea that a neural implementation of a form of Banburismus is responsible for generating categorical decisions about sensory stimuli. We focus first on a recent study of neural activity in the lateral intraparietal area (LIP) measured during a reaction-time version of the direction-discrimination task (Roitman and Shadlen, 2002). Results from that study provide the clearest example to date of neural mechanisms that appear to correspond to the key features of Banburismus, including the accumulation of sensory evidence to a threshold value. We then consider the general properties of neurons that may implement a process like Banburismus to form perceptual decisions.

The study by Roitman and Shadlen (2002) adds to a growing body of work that addresses how the brain transforms motion information into a categorical direction decision (Gold and Shadlen, 2000; Horwitz and Newsome, 1999; Kim and Shadlen, 1999; Shadlen et al., 1996; Shadlen and Newsome, 2001). A distinguishing feature of their study was the use of a response-time version of the direction-discrimination task, in which the monkey was trained to indicate its judgment of direction by making an eye-movement response as soon as it formed its




Figure 4. Banburismus in the Brain?

Responses from neurons in the lateral intraparietal area (LIP) of the monkey during a reaction-time direction-discrimination task. The left side shows responses averaged from a population of LIP neurons aligned to the onset of the motion stimulus. After an initial dip, the responses recover and soon after increase in a roughly linear fashion with a slope that is approximately proportional to the motion strength. These coherence-dependent responses are thought to represent the accumulated weight of evidence in favor of one direction of motion. The right side shows the LIP responses aligned to the beginning of the monkey's eye movement. When the response reaches a value of ~ 65 spikes/s, the monkey is committed to a decision, which is communicated by the eye movement ~ 80 ms later. Graph is modified from Roitman and Shadlen (2002), with permission.

decision. Performance accuracy on this task was similar to that found on fixed-duration versions of the task. Response times were, as expected, inversely related to accuracy: more difficult (lower motion strength) trials took longer than less difficult trials. Correlates of this motion strength- and time-dependent decision process were found in the activity of single neurons in area LIP. In a given experiment, they recorded from an LIP neuron and arranged the task so that one of the two choice targets was located in that neuron's response field. They then sorted trials in terms of whether the monkey's eye-movement response was into or out of the neuron's response field. They found that when the monkey correctly made an eye movement into the neuron's response field, its activity (measured in spikes/s) gradually built up while the monkey was viewing the dots and forming its decision (Figure 4). This build-up depended on motion strength: more difficult trials corresponded to a slower build-up than less difficult trials. In all cases, however, the build-up appeared to progress until reaching a fixed value, at which point the monkey made its eye-movement response.

Several aspects of these results are closely related to Banburismus. First, LIP responses during motion viewing, like those measured in other brain regions like the frontal eye field (FEF; Kim and Shadlen, 1999) and the superior colliculus (Horwitz and Newsome, 1999), appear to represent the “read-out” of motion activity represented in areas MT and MST. Computational models suggest that this read-out involves taking the difference from opposing pools of motion sensors (M.E. Mazurek and M.N. Shadlen, submitted; Shadlen et al., 1996); i.e., computing a form of Turing’s weight of evidence. Second, LIP activity appears to represent the integration over time of these motion signals. This integration can be thought of as updating the overall weight of evidence provided by the dynamic sensory signals (Equation 7). A transformed version of this accumulating weight of evidence also appears to be represented in signals related to the preparation of the oculomotor response itself (Gold and Shadlen, 2000). Third, LIP activity during motion viewing appears to reach a barrier height that, in principle, may be related to the weight of evidence required to make the decision and execute the appropriate eye-movement response.

These results exemplify what appears to be a general property of neurons involved in the formation of perceptual decisions: the ability to link sensory evidence to motor intention. That is, there seems to be a close relationship in the brain between forming a categorical decision and planning behavior. Neurons that form categorical decisions accumulate sensory (and other) evidence that bears on hypotheses. If, as in Banburismus, the evidence is accumulated to a threshold value, then reaching that value implies that the decision process is over and the brain is committed to a proposition. Typically, the proposition is linked to some behavior that is then permitted to ensue. Thus, it is convenient to think about the threshold for commitment as a shift from sensory analysis to behavioral intention. For this reason, it is natural to consider the process of accumulating evidence to threshold in terms similar to sensorimotor integration. For example, the brain regions that appear to represent the developing direction decision on the motion-discrimination task—including area LIP, the FEF, and the superior colliculus—are known also to play a role in preparing the associated eye-movement response (Schall, 1991).

It should be emphasized that this idea about the role of action-oriented circuits in decision formation does not imply that decisions must be coupled with overt behavioral responses. At least some of these circuits may be able to interpret sensory information in a more general, “pragmatic” framework that includes representations of potential or deferred actions (Jeannerod, 1994; Rizzolatti et al., 1994). For example, neural signals related to sensory-motor transformations that involve delayed actions, multiple potential actions, or complex behavioral rules are found in certain areas of the prefrontal cortex (Murray et al., 2000; Wallis et al., 2001). These kinds of neurons may prove central to the ability to form categorical decisions about sensory stimuli across a host of behavioral contexts.

In addition to representing the transformation of sensory information into intention, many of these action-oriented circuits possess several other properties that

may help to implement a process like Banburismus. The first property is persistent activity, which may be a hallmark of an ability to accumulate over time the salient information needed to form the decision variable. Among the many locations in the brain known to carry persistent activity are frontal and parietal association areas of cortex (Goldman-Rakic, 1995; Snyder et al., 2000). In these regions, the persistent activity appears to be closely related to the formation of categorical decisions. For example, neurons in both LIP and the prefrontal cortex that exhibit sustained activity through a delay period on a visually guided saccade task are exactly those neurons that appear to be involved in forming the decision on the direction-discrimination task (Kim and Shadlen, 1999; Roitman and Shadlen, 2002). Likewise, the activity of some neurons in the FEF persists throughout a delay period between the appearance of a visual target and the initiation of an eye movement to that target. This activity can represent the selection of the target among distractors and the preparation of the eye-movement response (Bichot and Schall, 1999). Additionally, in some cases the initiation of this kind of visually guided eye movement appears to be tied to FEF activity reaching a threshold level (Hanes and Schall, 1996).

The second property of these circuits is that they appear to be sensitive to psychological factors that can affect the underlying decision variable. A particularly compelling example comes from area LIP, where both the prior probability of getting a reward and the expected magnitude of the reward have been shown to modulate neural activity related to a visually guided saccade task (Platt and Glimcher, 1999). Prior probabilities can similarly affect the activity of neurons in the superior colliculus (Basso and Wurtz, 1998; Dorris and Munoz, 1998). In addition, regions of the lateral prefrontal cortex and the orbitofrontal cortex have been shown to represent information related to the value and probability of an anticipated reward in the context of sensory-motor tasks (Hikosaka and Watanabe, 2000; Leon and Shadlen, 1999; Tremblay and Schultz, 1999, 2000; Watanabe, 1996). Mechanisms responsible for shifting selective attention may also help to shape the decision variable, possibly by controlling the scale factor that determines the weight of evidence provided by a given set of sensory neurons. Indeed, attentional modulation is a central feature of sensory-motor transformation signals in association areas of cortex (Colby and Goldberg, 1999; Desimone and Duncan, 1995).

The third property suggested by the framework we have presented is that decision-making circuits in the brain are likely to access information about accuracy and the rate of reward. Consistent with this idea, psychophysical studies in both humans (Johnstone and Alsop, 2000) and pigeons (McCarthy and Davidson, 1984) suggest that reward schedules can affect the decision rule on simple perceptual tasks (Sternberg, 2001). Where this occurs in the brain is unknown, although recent studies suggest several possible candidates (for a review, see Schultz, 2000). For example, the supplementary eye field appears to represent an evaluation of performance on sensory-motor tasks, including whether a task was completed successfully and whether a reward was obtained (Stuphorn et al., 2000). Neurons in a region

of the basal ganglia called the caudate reflect an anticipation of reward in a manner that may be linked to the decision to make a particular eye movement (Lauwereyns et al., 2002). Additionally, a quantity resembling the anticipated rate of reward has been reported in neurons in the anterior cingulate cortex of monkeys performing a visual discrimination task (Shidara and Richmond, 2002). Of course, how these or other reward-related signals might serve to adjust the barrier height in the decision process (see Figures 2 and 3) is far from clear.

Thus, an implementation of Banburismus in the brain to form decisions about sensory stimuli is likely to depend on neurons with several characteristics. These characteristics include involvement in the control of behavior (the preparation for action, the control of intention and attention, the ability to invoke more complex behavioral rules); a capacity for persistent activity or temporal integration; sensitivity to a host of factors like sensory evidence, prior probabilities, and anticipated costs and benefits that may be incorporated into a decision variable; and access to information related to an evaluation of performance. As we have discussed, neurons with some of these characteristics appear to be capable of forming decisions in simple, two-alternative forced-choice tasks. However, their role in a wider variety of tasks, as well as their sensitivity to changes in the animal's reward schedule, as we have proposed, have yet to be established. Indeed, whether real neurons in the brain implement a form of Banburismus to form decisions about sensory stimuli remains an open question. Hopefully, enough bans of evidence will be obtained in this active area of research to provide a rewarding answer.

Acknowledgments

We thank John Palmer, Ruchi Kapoor, and Jochen Ditterich for helpful discussions on the ideas in this paper, and Brian Salzberg for a critical reading of the manuscript. This work was supported by HHMI, EY11378, and the Burroughs-Wellcome Fund.

References

- Albright, T.D., and Stoner, G.R. (1995). Visual motion perception. *Proc. Natl. Acad. Sci. USA* 92, 2433–2440.
- Audley, R.J., and Pike, A.R. (1965). Some stochastic models of choice. *Br. J. Math. Stat. Psychol.* 18, 207–225.
- 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.
- Bichot, N.P., and Schall, J.D. (1999). Effects of similarity and history on neural mechanisms of visual selection. *Nat. Neurosci.* 2, 549–554.
- Blake, R., and Hiris, E. (1993). Another means for measuring the motion aftereffect. *Vision Res.* 33, 1589–1592.
- Braddick, O. (1995). Visual perception. Seeing motion signals in noise. *Curr. Biol.* 5, 7–9.
- Britten, K.H. (2003). The middle temporal area: motion processing and the link to perception. In *The Visual Neurosciences*, L.M. Chalupa and J.S. Werner, eds. (MIT Press), in press.
- Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J. Neurosci.* 12, 4745–4765.
- 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.
- Celebrini, S., and Newsome, W.T. (1994). Neuronal and psychophys-

- ical sensitivity to motion signals in extrastriate area MST of the macaque monkey. *J. Neurosci.* 14, 4109–4124.
- Celebrini, S., and Newsome, W.T. (1995). Microstimulation of extrastriate area MST influences performance on a direction discrimination task. *J. Neurophysiol.* 73, 437–448.
- Colby, C.L., and Goldberg, M.E. (1999). Space and attention in parietal cortex. *Annu. Rev. Neurosci.* 22, 319–349.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* 18, 193–222.
- Dorris, M.C., and Munoz, D.P. (1998). Saccadic probability influences motor preparation signals and time to saccadic initiation. *J. Neurosci.* 18, 7015–7026.
- 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. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends Cogn. Sci.* 5, 10–16.
- Goldman-Rakic, P.S. (1995). Cellular basis of working memory. *Neuron* 14, 477–485.
- Good, I.J. (1979). Studies in the history of probability and statistics. XXXVII. A.M. Turing's statistical work in World War II. *Biometrika* 66, 393–396.
- Graham, N.V.S. (1989). *Visual Pattern Analyzers* (Oxford: Oxford University Press).
- Green, D.M., and Swets, J.A. (1966). *Signal Detection Theory and Psychophysics* (New York: Wiley).
- Hanes, D.P., and Schall, J.D. (1996). Neural control of voluntary movement initiation. *Science* 274, 427–430.
- Hikosaka, K., and Watanabe, M. (2000). Delay activity of orbital and lateral prefrontal neurons of the monkey varying with different rewards. *Cereb. Cortex* 10, 263–271.
- Hodges, A. (2000). *Alan Turing: The Enigma* (New York: Walker Publishing Company, Inc.).
- Horwitz, G.D., and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. *Science* 284, 1158–1161.
- Jeannerod, M. (1994). The representing brain: Neural correlates of motor intention and imagery. *Behav. Brain Sci.* 17, 187–245.
- Johnstone, V., and Alsop, B. (2000). Reinforcer control and human signal-detection performance. *J. Exp. Anal. Behav.* 73, 275–290.
- 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.
- LaBerge, D.A. (1962). A recruitment theory of simple behavior. *Psychometrika* 27, 375–396.
- Laming, D.R.J. (1968). *Information Theory of Choice-Reaction Times* (London: Academic Press).
- Lauwereyns, J., Watanabe, K., Coe, B., and Hikosaka, O. (2002). A neural correlate of response bias in monkey caudate nucleus. *Nature* 418, 413–417.
- 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: Erlbaum).
- Link, S.W., and Heath, R.A. (1975). A sequential theory of psychological discrimination. *Psychometrika* 40, 77–105.
- Luce, R.D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*, Oxford Psychology Series no. 8 (New York: Oxford University Press).
- MacKay, D.J.C. (2002). *Information Theory, Inference and Learning Algorithms* (Cambridge: Cambridge University Press), in press.
- McCarthy, D., and Davidson, M. (1984). Isobias and alloibias functions in animal psychophysics. *J. Exp. Psychol. Anim. Behav. Process.* 10, 390–409.
- Morgan, M.J., and Ward, R. (1980). Conditions for motion flow in dynamic visual noise. *Vision Res.* 20, 431–435.

- Murray, E.A., Bussey, T.J., and Wise, S.P. (2000). Role of prefrontal cortex in a network for arbitrary visuomotor mapping. *Exp. Brain Res.* 133, 114–129.
- Newsome, W.T., and Pare, E.B. (1988). A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J. Neurosci.* 8, 2201–2211.
- Orban, G.A. (1997). Visual processing in macaque area MT/V5 and its satellites (MSTd and MSTv). In *Cerebral Cortex: Extrastriate Cortex in Primates*, K.S. Rockland, J.H. Kaas, and A. Peters, eds. (New York: Plenum Press), pp. 359–434.
- Parker, A.J., and Newsome, W.T. (1998). Sense and the single neuron: probing the physiology of perception. *Annu. Rev. Neurosci.* 21, 227–277.
- Pasternak, T., and Merigan, W.H. (1994). Motion perception following lesions of the superior temporal sulcus in the monkey. *Cereb. Cortex* 4, 247–259.
- Peirce, C.S. (1878). The probability of induction. *Popular Science Monthly* 12, 705–718.
- Platt, M.L., and Glimcher, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature* 400, 233–238.
- Ratcliff, R., and Rouder, J.N. (1998). Modeling response times for two-choice decisions. *Psychol. Sci.* 9, 347–356.
- Rizzolatti, G., Riggio, L., and Sheliga, B.M. (1994). Space and selective attention. In *Attention and Performance, Vol. 15: Conscious and Nonconscious Information Processing*, C. Umiltà, ed. (Cambridge, MA: The MIT Press), pp. 232–265.
- 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.*, in press.
- Ross, S.M. (1983). *Stochastic Processes* (New York: Wiley).
- Salzman, C.D., and Newsome, W.T. (1994). Neural mechanisms for forming a perceptual decision. *Science* 264, 231–237.
- 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. (1991). Neural basis of saccadic eye movements in primates. In *The Neural Basis of Visual Function*, A.G. Leventhal, ed. (London: Macmillan Press), pp. 388–442.
- Schall, J.D. (2001). Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.* 2, 33–42.
- Schultz, W. (2000). Multiple reward signals in the brain. *Nat. Rev. Neurosci.* 1, 199–207.
- 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.
- Shidara, M., and Richmond, B.J. (2002). Anterior cingulate: single neuronal signals related to degree of reward expectancy. *Science* 296, 1709–1711.
- Snyder, L.H., Batista, A.P., and Andersen, R.A. (2000). Intention-related activity in the posterior parietal cortex: a review. *Vision Res.* 40, 1433–1441.
- Sternberg, S. (2001). Separate modifiability, mental modules, and the use of pure and composite measures to reveal them. *Acta Psychol. (Amst.)* 106, 147–246.
- Stone, M. (1960). Models for choice reaction time. *Psychometrika* 25, 251–260.
- Stuphorn, V., Taylor, T.L., and Schall, J.D. (2000). Performance monitoring by the supplementary eye field. *Nature* 408, 857–860.
- Tremblay, L., and Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature* 398, 704–708.
- Tremblay, L., and Schultz, W. (2000). Reward-related neuronal activity during go-nogo task performance in primate orbitofrontal cortex. *J. Neurophysiol.* 83, 1864–1876.
- Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592.
- Vickers, D. (1979). *Decision Processes in Visual Perception* (New York: Academic Press).
- Wald, A. (1947). *Sequential Analysis* (New York: Wiley).
- Wallis, J.D., Anderson, K.C., and Miller, E.K. (2001). Single neurons in prefrontal cortex encode abstract rules. *Nature* 411, 953–956.
- Watamaniuk, S.N.I., and Sekuler, R. (1992). Temporal and spatial integration in dynamic random-dot stimuli. *Vision Res.* 32, 2341–2347.
- Watanabe, M. (1996). Reward expectancy in primate prefrontal neurons. *Nature* 382, 629–632.