10 *The Speed and Accuracy of a Simple Perceptual Decision: A Mathematical Primer*

Michael N. Shadlen, Timothy D. Hanks, Anne K. Churchland, Roozbeh Kiani, and Tianming Yang

10.1 Introduction

A decision is a commitment to a proposition among multiple options. Often this commitment leads to a particular course of action. It might be said that the life of an organism consists of a series of decisions made in time [29]. Leaving aside the moral dimension of choices, it is reasonable to view even pedestrian decisions as a window on higher cognitive function, for they offer a glimpse of how the brain connects information to behavior in a contingent manner. Indeed, even simple visual decisions are based on a sophisticated confluence of available sensory information, prior knowledge, and the potential costs and benefits associated with the possible courses of action. Thus, understanding the neural basis of these decisions provides a model for the principles that govern higher brain function. An important step in making progress toward this goal is to develop frameworks for understanding simple decisions.

In this chapter, we describe one such framework, sequential sampling, which has been used to explain a variety of reaction-time decision tasks. This framework has a rich history in statistical decision theory and mathematical psychology. Our review is at best cursory, but we provide some key citations to the literature. Our main goal in this chapter is to provide an introductory tutorial on the mathematics that explains the psychometric and chronometric functions — that is, accuracy and speed of decisions as a function of difficulty. Then, to underscore our enthusiasm for this framework, we briefly summarize recent data from our laboratory that suggest a possible neural implementation of the computational principles in the parietal cortex of the monkey.

10.2 The Diffusion-to-Bound Framework

In this chapter, we consider simple binary decisions made upon the sequential analysis of evidence. In principle this evidence could arrive continuously in time or in discrete time steps. At each time step, the decision-maker either stops the decision process by committing to one of the alternatives or continues the process by waiting for the next piece of evidence. Thus, in a very general sense, the process involves tradeoffs between speed and accuracy and costs associated with obtaining more information. The framework is sufficiently rich to show up in a variety of problems, from quality control decisions (ship or reject a lot), to pricing bonds, and perception [15, 19]. Our interest is in perception and higher brain function. Perhaps it is worth saying at the outset that not all problems in perception involve the sequential arrival of information, and even when information is provided this way, there is no guarantee that it will be accumulated across time. However, in the particular case we study, it clearly is. Furthermore, we believe this gives us some insight into how the brain can solve more complex decisions that involve the accumulation of evidence obtained at different points in time.

For students of neuroscience and perception, the sequential analysis of information to some termination point lies at the heart of understanding measurements of reaction time and perceptual accuracy. Reaction-time tasks are important because they highlight the tradeoff in speed and accuracy of perception and because they permit identification of the time epoch in which a decision is forming but has not yet been completed. Simultaneous measurement of choices and reaction times provides multiple quantitative constraints that must be satisfied by any theory that aims to explain the neural mechanisms underlying the decision process.

For these reasons, we study decision-making in the context of a reaction-time motion discrimination task. In this task, fixating subjects are presented with a patch of moving dots. Some of these dots move together, or "coherently" in a given direction, while the remaining dots move randomly. At any time after motion onset, subjects can indicate their choice about the direction of motion of the stimulus by making a saccade to one of two targets. To make the task easier or more difficult, the percentage of coherently moving dots is increased or decreased. This task can be performed by both human and monkey subjects, allowing for both psychophysical and physiological investigation. Multiple studies have shown that accuracy depends on the coherence of the moving dot stimulus (figure 10.1 A) [6, 27]. Furthermore, both humans and monkeys reach their decisions faster when the motion coherence is larger, that is, when the task is easier (figure 10.1 B). Ultimately, we would like to understand the neural mechanisms that underlie both the choices and reaction times measured in this task.

The diffusion-to-bound framework, illustrated in figure 10.2, explains the pattern of behavior shown in figure 10.1. In this framework, noisy momentary evidence for one alternative or the other displaces a decision variable in the positive or negative direction. The expected size of the momentary evidence

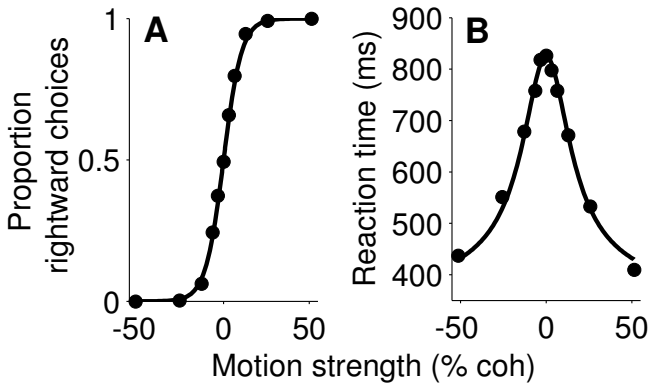

Figure 10.1 Behavioral data from one monkey performing reaction time (RT) version of the direction discrimination task. A. Psychometric function. The probability of a rightward direction judgment is plotted as a function of motion strength. Positive coherence refers to rightward motion and negative coherence to leftward motion. B. Effect of motion strength on RT. Mean RT for correct trials is plotted as a function of motion strength as in A. Error bars are smaller than the symbols. The solid lines show a combined diffusion model fit to the choice and RT data.

is related to the direction and strength of the motion stimulus, but in any one moment, the evidence is a random number. Over time, these random momentary evidence values are accumulated, giving rise to a random trajectory. The decision process terminates when the trajectory encounters a bound at $\pm A$. The particular bound that is crossed determines the choice, and the time taken to reach that bound determines the decision time. The important idea is that a single mechanism explains both which choice is made and how long it takes to make it.

These predictions can be described by relatively simple analytical equations, which give rise to the fits in figure 10.1. The psychometric function describes the probability of choosing the positive direction as a function of the motion strength, C:

$$P_+ = \frac{1}{1 + e^{-2kCA}} \tag{10.1}$$

where k and A are fitted parameters. The direction of motion is indicated by the sign of C. The probability of choosing the positive motion direction is P_+. We assume that the subjects are unbiased. Therefore, when $C = 0$, $P_+ = 1 - P_+ = \frac{1}{2}$.

The chronometric function describes the reaction time as a sum of decision and nondecision times. The decision time function shares the same parameters

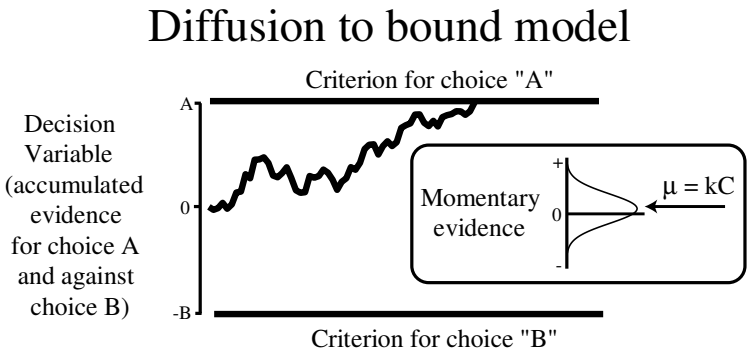

Figure 10.2 Diffusion-to-bound model of the decision process. Momentary evidence in favor of the "A" choice and against the "B" choice is accumulated as a function of time. The process terminates with an "A" or "B" choice when the accumulated evidence reaches the upper or lower bound, respectively, at $+A$ or $-B$. The momentary evidence is distributed as a unit-variance Gaussian whose mean, μ, is proportional to motion strength. The decision variable on a single trial follows a random "diffusion" path, like the one shown. The average of many of these paths would appear as a ramp with slope μ and variance proportional to time. Both decision time and the proportion of "A" and "B" choices are governed by A, B, and μ.

as in the psychometric function:

$$E\left[t\right] = \frac{A}{kC}\tanh\left(kCA\right) \tag{10.2}$$

When C=0, this equation is interpreted as a limit

$$\lim_{C\to 0}\frac{A}{kC}\tanh\left(kCA\right) = A^2 \tag{10.3}$$

We will derive these equations in a later section, but for now, it suffices to say that they capture the data reasonably well. Indeed, this model explains the choices and reaction times of monkey and human subjects on a variety of simple, two-choice discrimination tasks under different speed-accuracy pressures [14, 24]. Before working through the derivations, let's acquaint ourselves with the fitted parameters in equations (10.1) and (10.2).

The simplest form of the diffusion model, as employed in this example, has three parameters. First, there is the bound height, A, which mainly controls the balance between speed and accuracy. We place these bounds equidistant from the starting point because at the beginning of the trial, before any evidence has arrived, the two alternatives are equally likely. We will restrict our analysis to this simple condition. The value of $\pm A$ represents the total amount of evidence

that is required before a decision is made. Because random variations in the momentary evidence tend to average out with larger numbers of samples, a larger absolute value of A results in a greater level of accuracy. But the cost is time: a higher bound takes longer to reach, on average, resulting in a slower reaction time.

The second parameter, k, converts the stimulus strength into the drift rate of the diffusion process. The average drift rate is effectively the average value for the momentary evidence that accumulates per unit time. We will elaborate the concept of momentary evidence and its accumulation in later sections. For now, kC can be thought of as the mean of the momentary evidence normalized by its variance. For the random-dot motion experiments, it provides a very simple conversion from stimulus intensity and direction to the mean of the random number that accumulates sequentially. We are free to think about the momentary evidence or its accumulation. The momentary evidence is described by a distribution of random numbers: at each time step there is a draw. The accumulation has an expected drift rate equal to this mean per time step.

A third parameter is required to fit the reaction time data. The average nondecision time, $\bar{t}_{nd}$, accounts for the sensory and motor latencies outside the decision process per se. On any one trial, t_{nd} is a random number, but for present purposes, we are only concerned with its mean value. The mean reaction time is the sum of the mean decision time and the mean nondecision time.

The three parameters, k, A, $\bar{t}_{nd}$, are chosen to fit the choice and reaction time data in figure 10.1. Clearly they do a reasonable job of capturing the shapes of both functions and their relative position on the motion strength axis. This is by no means guaranteed. It suggests that a single mechanism might underlie the choices and decision times on this task. We emphasize that we do not expect all perceptual decisions to obey this framework. But many probably do.

10.3 Derivation of Choice and Reaction Time Functions

We will now explain how the accumulation of noisy momentary evidence to a positive or negative bound leads to the equations above: a logistic choice function and a $\tanh(x)/x$ decision time function in terms of stimulus strength. The exercise explains the origins of these functions, exposes the key assumptions, and gives us some insight into how these expressions generalize in straightforward ways to other decisions. The mathematics was developed by Wald and summarized in several texts [9, 15]. Here, we attempt to provide a straightforward explanation of the essential steps, because these ideas can be challenging to students and they are not well known to most physiologists. We emphasize that this is a synthesis of well-established results; certainly many important extensions of the theory are omitted. For additional reading, we recommend Link's book [19] and several insightful and didactic papers by Philip Smith [32, 33]. Other key references are cited below.

10.3.1 Overview

Boiled down to its essence, the problem before us is to compute the probability that an accumulation of random numbers will reach an upper bound at $+A$ before it reaches a lower bound at $-A$. This has an intuitive relationship with accuracy in the following sense. If the evidence should favor the decision instantiated by reaching the upper bound, what we will call a positive response, then the expectation of the momentary evidence is a positive number. Obviously, the accumulation of random numbers that tend to be positive, on average, will end at $+A$ more often that at $-A$. We desire an expression that returns the probability of such an accurate "positive" choice based on a description of the random numbers that constitute the momentary evidence. Our primary goal is to develop this formulation. Our second goal is to develop an expression for the number of samples of momentary evidence that it takes, on average, to reach one or the other bound. This is the decision time. Along the way, we will provide some background on the essential probability theory.

10.3.2 Statement of the Problem

Consider a series of random numbers, $X_1, X_2, \ldots, X_n$, each drawn from the same distribution. We are interested in the stochastic process that unfolds sequentially from the sum of these random numbers

$$Y_n = \sum_{i=1}^{n} X_i \tag{10.4}$$

We assume that each value of X is drawn from the same distribution and each draw is independent of the values that preceded it. In other words, the X_i are independent and identically distributed (*i.i.d.*). Like the X_i, the Y_i are also a sequence of random numbers. In fact there is a one-to-one correspondence between the two series. However, unlike the X_i, the Y_i are neither independent nor identically distributed. They are the accumulation of the X_i. Each subsequent value in the sequence of Y depends on the value that had been attained at the previous step. So there is clearly some correlation between Y_n and Y_{n-1}, and we certainly would not say that the distribution of Y on the n^{th} step is the same as it was on the step before. On the other hand, the distribution of Y_n is easy to describe if Y_{n-1} is known. It is just $Y_{n-1} + X_n$. Since it is unnecessary to know how Y_{n-1} attained its particular value, we say that the sequence, Y_i, is a Markov process.

$Y_1, Y_2, \ldots, Y_n$ represents a random path from the origin, the accumulation of random numbers, $X_1, X_2, \ldots, X_n$. The path, Y, can be written as a function of t or as a function of the number of time steps, $n = t/\Delta t$. If the time steps are discrete the stochastic process is termed a random walk, and if t is continuous, Y is a termed a diffusion process. We tend to gloss over these distinctions (but see Appendix 10.1). We can do this because we are always considering time steps that are small enough (with respect to the bounds) so that it takes a

large number to reach the bounds. This means that even if time steps come at irregular intervals, the number of steps multiplied by the average time step is a good approximation to the time to reach the bound, and the increments are small enough that we can treat the last step as if it reached the bound at ±A exactly, without overshoot.

According to our formulation, the process stops when Y_n reaches $+A$ or $-A$. Our goal is to derive expressions for the probability of reaching these bounds and the average stopping times. We want these expressions to be functions of the stimulus strength (sign and magnitude of the evidence). Because the distribution of the values to be accumulated, X_i, depends on the stimulus intensity, our goal is to derive these expressions in terms of the distribution of X_i.

10.3.3 Background: Moment Generating Functions

A concept that will be useful in this endeavor is that of the moment-generating function (MGF) of a probability distribution. The MGF (when it exists) is simply an alternative description of a probability distribution, and it can be thought of as a transformation that is convenient for certain calculations (much like a Laplace or Fourier transform). The MGF of a random variable X is the expectation of $e^{\theta X}$ over all possible values of X:

$$M_X(\theta) \equiv E\left[e^{\theta X}\right] = \int_{-\infty}^{\infty} f(x)e^{\theta x}dx, \tag{10.5}$$

where $f(x)$ is the probability density function for X, and θ can be any arbitrary value. If you are familiar with Fourier analysis or the Laplace transform, think of θ as the frequency variable. If this is a foreign concept, just think of θ as a number; and the MGF is a function of this number.

Figure 10.3B shows the MGF for the normal distribution with mean = 1.5 and standard deviation = 1.5 as shown in figure 10.3a. Interestingly, the slope of the function at θ = 0 is the mean (or first moment) of X. We can see this by taking the derivative of equation (10.5), with respect to θ:

$$M'_X(\theta) = \frac{d}{d\theta}E\left[e^{\theta X}\right] = \frac{d}{d\theta}\int_{-\infty}^{\infty} f(x)e^{\theta x}dx = \int_{-\infty}^{\infty} xf(x)e^{\theta x}dx \tag{10.6}$$

At the point $\theta = 0$,

$$M'_X(0) = \int_{-\infty}^{\infty} x\,f(x)dx = E[x], \tag{10.7}$$

(Some students may need to be reminded that an expectation of a function of x is nothing more than a weighted sum of all possible values of that function; the weights are defined by the probability of observing each of the possible random values of x. The simplest case is the weighted sum of the x values themselves, that is, the expectation of x, which is termed the mean or first moment.)

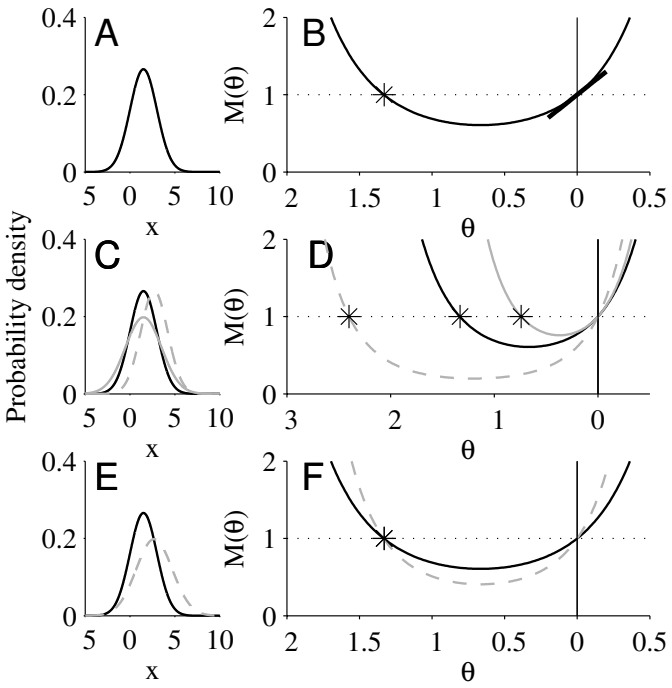

Figure 10.3 Effect of changes in mean and variance of a distribution on its moment-generating functions (MGFs). A. Probability density function (PDF) for a Gaussian distribution with a mean of 1.5 and a variance of 2.25 and standard deviation of 1.5. B. MGF associated with the PDF in A. Tangent line at $\theta = 0$ indicates the first derivative at 0. This is $E[X]$ (the mean), which equals 1.5 (note scaling of ordinate and abscissa). Asterisk indicates the θ_1 zero crossing. C. Same PDF as in A (black trace), alongside a normal PDF with a larger mean and the same variance (gray dashed trace), and a normal PDF with the same mean and a larger variance (gray solid trace). D. MGFs associated with each PDF in C. Line styles are the same as in C. Asterisk indicates the θ_1 zero crossing. E. Same PDF as in A (black trace), alongside a normal PDF with a larger mean and larger variance (gray dashed trace). F. MGFs associated with each PDF in E. Line styles are the same as in E. Asterisk indicates the θ_1 zero crossing.

It is also worth mentioning that the second derivative of the function shown in figure 10.3B evaluated at zero is the expectation of the squared values of the random variable, also known as the second moment. Just differentiate the expression in equation (10.7) again. Now an x^2 appears in the integral. This provides a little insight into the shape of the function shown in figure 10.3B. Since the expectation of a squared random number must be positive, we know that the convexity of the MGF at 0 is positive. It also explains the term "moment-generating": when evaluated at $\theta = 0$, the $1^{st}, 2^{nd}, \ldots, n^{th}$ derivatives of $M(\theta)$ return the expectations of $x, x^2, \ldots, x^n$.

Figure 10.3D and F show examples of MGFs associated with the normal distributions shown in figure 3C and E, each having a different mean and variance. You can see that higher means steepen the slope of the function at $\theta = 0$, whereas higher variance exaggerates the convexity. The figure also highlights a feature of the MGF that will be important to us in a moment. Notice that the $M(0) = 1$ (because e^{0X} is 1 for all X). The MGF then returns to 1 at another point. These are marked in the figure with an asterisk. We refer to this special root of the MGF as θ_1; so

$$M_X(\theta_1) = M_X(0) = 1 \tag{10.8}$$

θ_1 is going to play an important role in the argument we are about to share. You'll want to return to this figure later. For now, notice that θ_1 moves further from 0 when the mean is a larger positive number (figure 10.3C,D, dashed gray trace), it moves toward 0 when the variance is larger (figure 10.3C,D, solid gray trace), and it is unchanged when the ratio of the mean to variance remains the same (figure 10.3E,F).

MGFs are useful for analyzing the sums of random numbers. For example, suppose we add two random numbers to produce a new one:

$$S = X_1 + X_2 \tag{10.9}$$

The distribution of S can be written as a convolution of the distributions for X_1 and X_2:

$$f_S(s) = \int_{-\infty}^{\infty} f_{X_1}(r) f_{X_2}(s - r)\, dr, \tag{10.10}$$

where f_{X_1} and f_{X_2} are the probability density functions for the X_1 and X_2. An intuition for this is that the probability of making a particular sum, $S = s$, is the probability that $X_1 = r$ and $X_2 = s - r$ for all possible values of r. Effectively, the distribution of the new variable, S, is achieved by shifting and blurring the distribution for one of the added variables by the other.

Now suppose X_1 and X_2 have MGFs $M_{X_1}(\theta)$ and $M_{X_2}(\theta)$. The MGF for S is simply the product of these:

$$M_S(\theta) = M_{X_1}(\theta) M_{X_2}(\theta) \tag{10.11}$$

Thus convolution of the probability functions is replaced by multiplication of these transformed functions, which is often a more convenient operation. This general idea is a concept that should be familiar to readers acquainted with Fourier and Laplace transforms. It turns out to play a key role in the derivation of the psychometric function.

Thus, the MGF associated with Y_n, as defined in equation (10.4), is the MGF associated with the momentary evidence, X, multiplied by itself n times.

$$M_{Y_n}(\theta) = M_X^n(\theta) \tag{10.12}$$

This is the MGF associated with the height of the trajectory after n steps, ignoring the bounds.

10.3.4 A Moment Generating Function for the Terminated Accumulation

The decision process ends when the accumulation of evidence reaches one of the bounds, that is, when $\tilde{Y} = \pm A$. The tilde above the Y is there to indicate that we are considering the accumulation at its termination. There is no subscript, because the number of steps is a random number: the distribution of decision times.

Using the concept introduced above, we can also express this distribution in terms of its MGF,

$$M_{\tilde{Y}}(\theta) = E[e^{\theta \tilde{Y}}] = P_+ e^{\theta A} + (1 - P_+)e^{-\theta A}, \tag{10.13}$$

where P_+ is the probability of stopping at the positive bound. Notice the brute force expansion of the expectation as the weighted sum $e^{\theta A}$ and $e^{\theta(-A)}$. Our plan is to use this equation to obtain an expression for P_+. As it is, the equation is not practical because it contains an MGF and the variable θ. We would like instead to have terms that we can relate to stimulus intensity (e.g., motion coherence). To achieve this, we seek another way to represent the MGF for the terminated accumulation.

10.3.5 Wald's Martingale

This can be done in the following way. First, we create a new stochastic process, $Z_1, Z_2, \ldots, Z_n$ that parallels the sequence $Y_1, Y_2, \ldots, Y_n$:

$$Z_n = M_x^{-n}(\theta) e^{\theta Y_n} \tag{10.14}$$

For any particular choice of θ, the MGF part becomes a number raised to the negative nth power. In that case (and if $\theta \neq 0$) the Z_i form a sequence of random numbers in one-to-one correspondence with the Y_i. In fact, we could say that before the process begins, the accumulation starts at $Y_0 = 0$ and $Z_0 = 1$.

This newly created sequence has the following important property. If at the n^{th} step, the process happens to have attained the value Z_n, then the expectation of the random value that it will attain on the next step is also Z_n

$$E\left[Z_{n+1} \,|Z_n\right] = Z_n : \tag{10.15}$$

Think about what this means. Z_n is a random number—the n^{th} value in a sequence of random numbers. Imagine observing this sequence as it plays out, one step at a time. Suppose we have just seen the n^{th} step and we are now waiting for the next random number in the sequence. At this point, the value of Z_n is known. Now, we know that the next step will produce a random number. But the expectation of this random number (i.e., the average if only we could repeat this next step many times) is the number we have in hand. That is the meaning of equation (10.15). It says that on average, we do not expect the next step in the sequence to change from the value that it happened to attain on the pervious step. Of course, this is only true on average. Any given sequence of

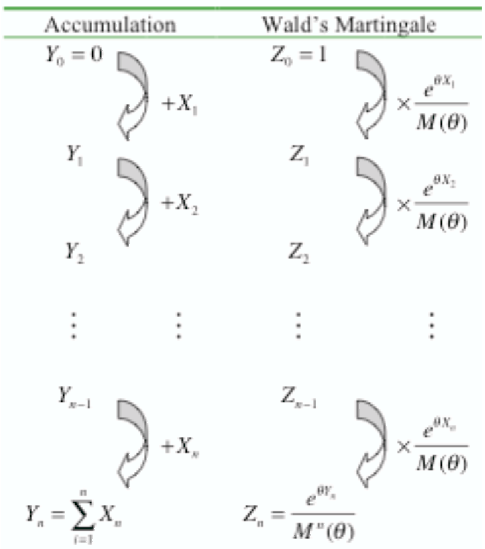

Figure 10.4 Two stochastic processes. The two lists of random numbers are created by updating the previous value in the list using a random draw, X_n. Each, $X_1, X_2, X_3, ...$, is drawn from the same distribution, and each draw is independent from the others. The formulas show how to use the current draw to update the previous value of Y or Z using the previous value attained and the new draw of X.

Z_i will wander about because the next number is unlikely to actually be the last number that was attained—that's just the expectation.

We can appreciate this in another way by considering an alternative recipe for generating the Z_i, illustrated in figure 10.4. To make any Z_{n+1}, begin with Z_n, draw a random value, X_{n+1}; multiply Z_n by $e^{\theta X_{n+1}}$ and divide by $M_X(\theta)$. Since $M_X(\theta)$ is $E\left[e^{\theta X_{n+1}}\right]$, the expectation of Z_{n+1} is Z_n.

A random variable sequence that has such properties is called a *martingale*. In fact, the stochastic process $Z_0, Z_1, \ldots, Z_n$ is a martingale with respect to the Y_n because $E\left[Z_{n+1} | Y_0, Y_1, Y_2, \ldots, Y_n\right] = Z_n$. This property can also be worked out directly from the definition of Y and Z, as follows:

$$\begin{aligned}
E\left[Z_{n+1} | Y_0, Y_1, Y_2, \ldots, Y_n\right] &= E\left[M_X^{-(n+1)}(\theta) e^{\theta Y_{n+1}} | Y_0, Y_1, Y_2, \ldots, Y_n\right] \\
&= E\left[M_X^{-(n+1)}(\theta) e^{\theta(Y_n + X_{n+1})}\right] \text{ (by the rule for generating } Y_{n+1}) \\
&= E[M_X^{-1}(\theta) M_X^{-n}(\theta) e^{\theta Y_n} e^{\theta X_{n+1}}] \\
&= E[M_X^{-1} Z_n(\theta) e^{\theta X_{n+1}}] \text{(using the definition of } Z_n) \\
&= M_X^{-1}(\theta) Z_n E\left[e^{\theta X_{n+1}}\right] \text{ (because } Z_n \text{ and } M_X(\theta) \text{ are known)} \\
&= Z_n
\end{aligned} \tag{10.16}$$

Any particular instantiation of Z is a random process that is one to one with a particular instantiation of a random walk, Y. Yet, the expectation of any Z_n is always the same. In fact it is always 1.

$$E[Z_n] = E\left[M_X^{-n}(\theta)e^{\theta Y_n}\right] = M_X^{-n}(\theta)E\left[e^{\theta Y_n}\right] = M_X^{-n}(\theta)M_{Y_n}(\theta) = 1 \quad (10.17)$$

Notice that the MGF is a function of θ — at any particular value of θ, it is just a number — so it can be removed from the expectation in the first line of equation (10.17). The last equality follows from equation (10.12).

The stochastic process, Z, is known as Wald's martingale and the identity in equation (10.17) is known as Wald's identity, usually written

$$E\left[M_X^{-n}(\theta)e^{\theta Y_n}\right] = 1 \quad (10.18)$$

The usefulness of creating Z is not yet apparent, but in a moment, we will exploit an important property it possesses.

10.3.6 Another Moment Generating Function for the Terminated Accumulation

Equation (10.14) gives us an expression that relates each Z_n to its associated Y_n and therefore to the distribution of each X_i, but remember that our goal is to relate the stopped variable $\tilde{Y}$ to that distribution. Fortunately, an important property of martingales, the optional stopping theorem, will enable us to do so. Suppose the sequence Z is forced to stop when the accumulation Y reaches one of the bounds. Let $\tilde{Z}$ represent the "stopped" value of Z.

$$\tilde{Z} = M_X^{-\tilde{n}}(\theta)e^{\theta\tilde{Y}} \quad (10.19)$$

We placed a tilde over the n in this expression to indicate that we do not know the number of steps it takes to terminate. In fact n is a random number that we will relate to decision time. Because n is a random number, we cannot simply take the stopped accumulation and convert it to a stopped value for Wald's martingale. We do not know how many times we need to divide by $M_X(\theta)$. Fortunately, we can work with the expectation of $\tilde{Z}$.

The optional stopping theorem for martingales states that $E\left[\tilde{Z}\right] = E[Z_n]$ even though the number of steps, n, is a random number. We will not prove this theorem, but it ought to come as no surprise that this would hold, because the expectation of Z is always the same, in this case 1. Thus,

$$E\left[M_X^{-\tilde{n}}(\theta)e^{\theta\tilde{Y}}\right] = 1 \quad (10.20)$$

The left side of this expression almost looks like the definition of an MGF for $\tilde{Y}$, which is what we desire. But alas, there is the $M_X^{-\tilde{n}}(\theta)$ bit inside the expectation. This expression is especially problematic because it introduces a random variable, $\tilde{n}$. If only we could make this bit disappear.

Here is how. Equation (10.20) holds for any value of θ that we choose. So let's choose θ such that $M_X(\theta) = 1$. Then it no longer matters that $\tilde{n}$ is a random number, because 1 raised to any power is always 1. Recall from the discussion of MGFs above that there are two values for θ that satisfy this identity: 0 and the point marked θ_1(the asterisk in figure 10.3B). For all MGFs, $M(0) = 1$, but this is not useful (Z is not a stochastic process — it is always 1). However, this other value away from the origin is profoundly interesting. It is guaranteed to exist if (i) Xhas a MGF, (*ii*)X has finite variance, (*iii*) $E[X] \neq 0$, and (*iv*) the range of X includes positive and negative values. The first two conditions are formalities. The third omits a special case when the slope of the MGF is zero at 0 (effectively, $\theta_1 = 0$). The last condition implies that the random increments can take on positive and negative values — in other words, the conditions leading to a random walk or diffusion process. Suffice it to say that for the conditions that interest us, θ_1 exists. At this special value, $\theta = \theta_1$, equation (10.20) simplifies to

$$E\left[e^{\theta_1 \tilde{Y}}\right] = 1 \tag{10.21}$$

Equation (10.21) is not an MGF for the stopped random value $\tilde{Y}$, but it is a point on that MGF.

10.3.7 The Psychometric Function

We have now found a value of θ for which the expression $E[e^{\theta \tilde{Y}}]$ is known, and we can plug this back into equation (10.13), yielding:

$$M_{\tilde{Y}}(\theta_1) = E[e^{\theta_1 \tilde{Y}}] = P_+ e^{\theta_1 A} + (1 - P_+)e^{-\theta_1 A} = 1 \tag{10.22}$$

Solving for P_+

$$P_+ = \frac{1 - e^{-\theta_1 A}}{e^{\theta_1 A} - e^{-\theta_1 A}} = \frac{1 - e^{-\theta_1 A}}{e^{-\theta_1 A}\left(e^{\theta_1 A} + 1\right)\left(e^{\theta_1 A} - 1\right)} = \frac{1}{1 + e^{\theta_1 A}} \tag{10.23}$$

The probability of terminating a decision in the upper vs. lower bound is a logistic function of the argument $\theta_1 A$.

We have very nearly achieved an expression for the psychometric function. All that remains is to relate θ_1 to the stimulus intensity variable (e.g., motion coherence). We know that θ_1 is a root of the MGF associated with the distribution of X, the momentary evidence that accumulates toward the bounds. Suppose X obeys a normal distribution with mean μ and standard deviation σ. The MGF associated with the normal distribution is

$$M_X(\theta) = e^{\theta\mu + \frac{1}{2}\theta^2\sigma^2}, \tag{10.24}$$

which is 1 when

$$\theta_1 = -\frac{2\mu}{\sigma^2} \tag{10.25}$$

Suppose that stimulus intensity leads to a proportional change in the mean of the momentary evidence but no change in the variance. Then, we can write

$$\theta_1 = -2kC/\sigma^2 \qquad (k > 0), \tag{10.26}$$

which gives us a logistic function resembling equation (10.1). The only difference is the σ^2 term here. This term, which we set to 1, provides a scaling for A and μ. We will justify these assumptions (at least partially) in a later section and consider the meaning of the fitted parameters, k and A, in terms that connect to noisy neurons in the visual cortex.

The form of the equation will depend ultimately on whatever particular assumptions we make about the relationship between motion coherence and the distribution of momentary evidence. The important point of the exercise is not to justify these assumptions but to recognize what it is about the distribution of momentary evidence that will lead to a prediction for the psychometric function (PMF). We prefer to leave this section with the PMF in its most general form prescribed by equation (10.23). This is a remarkable insight with many fascinating implications [19], some of which we will describe in detail later.

10.3.8 Decision Time

Now that we have computed the probability of terminating a decision in the upper vs. lower bound, we will derive an expression for the average time to reach the bound. This is the length of a time step multiplied by the number of steps taken to reach the bound, what we called $\tilde{n}$. We will derive an expression for the mean of this quantity, $E[\tilde{n}]$. The first step is to take the derivative of both sides of equation (10.20) with respect to θ:

$$E\left[e^{\theta\tilde{Y}}\tilde{Y}M_X^{-\tilde{n}}(\theta) - e^{\theta\tilde{Y}}\tilde{n}M_X^{-1-\tilde{n}}(\theta)M_X'(\theta)\right] = 0 \tag{10.27}$$

Evaluating this expression at $\theta = 0$ greatly simplifies matters. Recall that an MGF is always 1 at $\theta = 0$ and the first derivative is the first moment (i.e., the mean). Therefore

$$E\left[\tilde{Y} - \tilde{n}\mu\right] = 0, \tag{10.28}$$

which can be rearranged

$$E[\tilde{n}] = \frac{E[\tilde{Y}]}{\mu} \qquad (\text{for } \mu \neq 0) \tag{10.29}$$

This is a deceptively simple expression. It tells us that the mean number of steps to reach one of the bounds is the average accumulation at the time of the decision divided by the average increment in the evidence. Now $\tilde{Y}$ is either +A or –A, but it is more often +A when $\mu > 0$, and it is more often –A when $\mu < 0$. The expectation is just a weighted sum of these two values. Therefore

$$E[\tilde{Y}] = P_+A + (1 - P_+)(-A) = (2P_+ - 1)\,A \tag{10.30}$$

and

$$E[\tilde{n}] = \frac{(2P_+ - 1)\,A}{\mu} \tag{10.31}$$

Substituting the expression for P_+ from equation (10.23), yields

$$E[\tilde{n}] = \frac{A}{\mu}\left(\frac{2}{1+e^{\theta_1 A}} - 1\right) = \frac{A}{\mu}\left(\frac{1-e^{\theta_1 A}}{1+e^{\theta_1 A}}\right) \tag{10.32}$$

Multiplying numerator and denominator by $e^{-\frac{\theta_1 A}{2}}$ yields a simple expression based on the hyperbolic tangent function:

$$E\left[\tilde{n}\right] = \frac{A}{\mu}\left(\frac{e^{-\frac{\theta_1 A}{2}} - e^{\frac{\theta_1 A}{2}}}{e^{-\frac{\theta_1 A}{2}} + e^{\frac{\theta_1 A}{2}}}\right) = \frac{A}{\mu}\tanh\left(-\frac{\theta_1 A}{2}\right) \tag{10.33}$$

This general expression tells us that the number of steps is related to the bound height, the mean of the momentary evidence, and θ_1. The value for θ_1 depends on the distribution that gives rise to the values of momentary evidence that are accumulated. We offer some helpful intuition on this below, but it is a number that tends to scale directly with the mean and inversely with the variance of that distribution. For the normal distribution, this tendency holds exactly. If we assume that the X_i are drawn from a normal distribution with mean μ and variance σ^2, then

$$E\left[\tilde{n}\right] = \frac{A}{\mu}\tanh\left(\frac{\mu A}{\sigma^2}\right) \tag{10.34}$$

In the special case where the momentary evidence favors the two decision outcomes equally, $\mu = 0$ and

$$E\left[\tilde{n}\right] = \lim_{\mu \to 0}\frac{A}{\mu}\tanh\left(\frac{\mu A}{\sigma^2}\right) = \frac{A^2}{\sigma^2} \tag{10.35}$$

Again, suppose a change in the motion strength, C, leads to a proportional change in the mean of the momentary evidence, but no change in the variance. Then substituting $-2kC/\sigma^2$ for θ_1 into equation (10.33) yields

$$E\left[\tilde{n}\right] = \frac{A}{\mu}\tanh\left(\frac{kCA}{\sigma^2}\right) \tag{10.36}$$

Equation (10.36) resembles equation (10.2), which we fit to the reaction time data in figure 10.1. The only difference is the σ^2 term here. This term, which we set to 1, provides a scaling for A and μ. The average decision time at $C = 0$ is therefore $E\left[\tilde{n}\right] = A^2$. This implies that the variance associated with the momentary evidence at each time step is 1. This implies that the variance associated with the momentary evidence at each time step is 1, and the coefficient, k, should be interpreted as follows: $\mu = kC$ is the mean of a normal distribution

whose variance is 1 in a single time step. The momentary evidence is drawn from this distribution in each time step. This gives us a way to reconcile the fitted parameters, k and A, with measurements from neurons. We will resume this thread below.

Before leaving this section, we should point out one other aspect of the equations for decision time: they apply to both choices. In the equations above, $E[\tilde{n}]$ is the expected number of steps to reach *either* bound. It is the expected number of steps to finish, regardless of which bound is reached. However, when the upper and lower bounds are equidistant from the starting point, it is often the case that the expected number of steps to one or the other bound is exactly the same. This holds when the momentary samples (X_i) are drawn from the normal distribution. This is obvious when the mean of the momentary evidence is 0 — the case of a purely random walk. For any path of steps that ends in the upper bound, there is an equally probable path to the lower bound. It may be less obvious when there is a tendency to drift in the positive or negative direction. In that case, the claim is that the average number of steps to reach the "correct" bound (e.g., the positive bound when the evidence is positive, on average) is the same as the average number of steps to reach the "incorrect" bound. In fact, the entire distribution of stopping times should be identical and merely scaled by P_+ and its complement.

Consider a "path" that ends in the upper bound after n steps. The likelihood of observing this path is a product of likelihoods of obtaining the n values, $X_1, X_2, \ldots, X_n$ from $N\{\mu, \sigma\}$. Any such path has a mirror image that ends in the lower bound. The likelihood of observing this path is the product of the likelihoods of drawing $-X_1, -X_2, \ldots, -X_n$ from the same normal distribution. The ratio of these likelihoods is

$$\frac{P(X_1, X_2, \ldots, X_n \,|\mu, \sigma)}{P(-X_1, -X_2, \ldots, -X_n \,|\mu, \sigma)} = \frac{\prod_i \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(X_i-\mu)^2}{2\sigma^2}}}{\prod_i \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(-X_i-\mu)^2}{2\sigma^2}}} = e^{\frac{2\mu\sum_i X_i}{\sigma^2}} = e^{\frac{2\mu A}{\sigma^2}}, \tag{10.37}$$

where A is the sum of the increments, X_i. Since all paths sum to this same value, every path that ends in the upper bound (+A) has a corresponding path that ends at the lower bound. The probability of observing any path to the upper bound is the same as the one to the lower path multiplied by a constant, which happens to be the odds of hitting the upper bound. We can see this by computing the odds from equation (10.23);

$$\frac{P_+}{1-P_+} = \frac{\left(\frac{1}{1+e^{\theta_1 A}}\right)}{\left(1-\frac{1}{1+e^{\theta_1 A}}\right)} = e^{-\theta_1 A} = e^{\frac{2\mu A}{\sigma^2}} \tag{10.38}$$

From the perspective of the present exercise, the equality of mean decision time regardless of which bound terminates the decision is convenient on the one hand and damaging on the other. It is convenient because our equation for

fitting the reaction time is appropriate even when we fit just the correct choices, as in figure 10.1. It is damaging because it predicts that the mean reaction time for error trials ought to be the same as for correct choices. In our experiments, they are close, but the reaction time on error trials is typically slightly longer than on correct trials [27, 24]. A variety of solutions have been proposed to remedy this deficiency. The most obvious, but the least explored, is to use a non-Gaussian distribution for the momentary evidence [20]. It turns out that if the slopes of the MGF at θ_1 and at 0 are not equal and opposite, then $E[\tilde{n}]$ are not the same when the process terminates at the upper or lower bound. Link and Heath illustrate this for the Laplacian distribution. Other solutions include variability in E[X] and starting point [17, 26] and a nonstationary bound. For example, if an "urgency" signal lowers the bounds as a function of time (or an equivalent addition is made to the accumulation) then longer decisions will be less accurate [10]. Variants on the nonstationary bound idea include adding a cost of computing time.

If the bounds are not equidistant from the starting point of the accumulation, then the decision times for the two choices are wildly different. Expressions for these conditionalized stopping times can be found in [24] and [31] for Gaussian increments, and in [20] for non-Gaussian increments based on θ_1.

10.3.9 Summary

We derived general forms of expressions for P_+, the probability that an accumulation of random numbers will reach a positive bound before it reaches a negative bound, and $E[\tilde{n}]$, the number of random numbers that are added, on average, to reach either bound. The most general forms of these expressions are functions of the absolute value of the criterion and a special root of the MGF associated with the random increments. To achieve our fitting functions, we assumed that the random increments were drawn from a normal distribution. The special root is proportional to the mean divided by the variance of the increments: $\theta_1 = -2\mu/\sigma^2$. To construct our psychometric and chronometric functions, we made the additional simplifying assumption that this term is proportional to our stimulus intensity variable, motion strength. Recall that motion strength is a signed quantity; the sign indicates the direction. Thus, positive motion strengths give rise to positive increments, on average, and negative values of θ_1.

The real dividend of this argument, we hope, is in the general expression that revolves around θ_1. The random numbers that accumulate in the brain may not be normally distributed (see Appendix 10.1) and the mapping between strength of evidence and the parameterization of momentary evidence may not be as simple as a proportional relationship to the mean. The connection between quality of evidence (e.g., stimulus intensity) and behavior (i.e., choice and decision time) are mediated via θ_1. If we know how intensity affects θ_1, then we can predict the shape of the psychometric and chronometric functions. It is reasonable to consider this term proportional to the mean/variance ratio for increments drawn from a variety of distributions. And it is not too hard to

appreciate how neurons could give rise to such a quantity.

10.4 Implementation of Diffusion-to-Bound Framework in the Brain

Our enthusiasm for the random walk to bound framework emanates from physiology experiments. For the random-dot motion task, there are neural correlates for both momentary evidence and its accumulation. Although there is much to learn about many of the most important details of the mechanism, we understand enough to believe that this mechanism, or something very close to it, actually underlies decisions on the random-dot motion task.

10.4.1 The Momentary Evidence for Direction is Represented by Neurons in area MT/V5

MT/V5 is an area of extrastriate visual cortex first identified by Zeki [12] and Allman and Kaas [2] and appears to be specialized for processing visual motion [1, 5, 23, 4]. The random-dot motion stimulus used in our task was tailored by Newsome, Britten, and Movshon [22] to match the receptive field preferences of neurons in the middle temporal area (MT). We know from a variety of stimulation and lesion experiments that area MT plays an essential role in allowing a monkey to perform the discrimination task. Indeed microstimulation of a cluster of rightward-preferring MT neurons causes the monkey to decide that motion is rightward more often [28] and to do so more quickly than when there is no stimulation [11]. Moreover, stimulating these same right-preferring neurons affects leftward decisions negatively. When the monkey decides that the motion is leftward, he does so more slowly when right-preferring neurons have been activated by microstimulation.

MT neurons respond to random-dot motion stimuli by elevating their firing rates. As shown in figure 10.5, the firing rate rises the most when highly coherent motion is in the neuron's preferred direction, and it is lowest when highly coherent motion is in the opposite (null) direction. Notice that the response is elevated relative to baseline when low coherence motion is shown in *either* direction, but the firing rate is slightly greater when the motion is in the neuron's preferred direction. After a brief transient associated with the onset of the random dots, the response reaches a fairly steady (but noisy) firing rate while the random-dot motion is displayed. It is the firing rate during this steady period that constitutes the momentary evidence. Notice that there is no sign of accumulation in these responses. They look nothing like the trajectories of a random walk. Rather, they provide the momentary evidence that is accumulated elsewhere in the brain.

10.4.2 A Crude Estimate of θ_1 from the Neural Recordings

The experiments suggest that the momentary evidence for rightward, say, is based on a difference in firing rates between right-preferring and left-preferring

MT neurons [11]. There is good reason to believe that a decision is not based on a pair of neurons but on the activity from ensembles of neurons. For the random-dot displays used in our experiments, a typical MT neuron will change its average firing rate by ∼40 spikes per second (sp/s) as motion strength increases from 0% to 100% in its preferred direction, and it will decrease its response by ∼10 sp/s over this same range of motion strengths when the direction is opposite to its preferred direction [7]. If this difference in spike rates between ensembles of right- and left-preferring MT neurons constitutes the momentary evidence for rightward,

$$\mu_{MT} \equiv E\left[X_R - X_L\right] = 50C, \tag{10.39}$$

where once again C is motion coherence on a scale from -1 to 1.

What happens to the variance of this difference variable? For a single neuron, the variance of the number of spikes emitted by an MT neuron in a fixed epoch is typically a constant times the mean of this number, termed the Fano factor, ϕ. For a single MT neuron, $\phi_1 \approx 1.5$. Suppose the spike rate is 20 sp/s when C = 0. Then the expected number of spikes in 1 second is $E\left[s\right] = 20$ and the variance of the spike count in a 1 second epoch is $\phi_1 E\left[s\right] = 30$ sp^2. The average of N spike counts is the sum of N random numbers, $\frac{s_1}{N}, \frac{s_2}{N}, , \frac{s_N}{N}$. So the expected mean count is the same as the expectation of the count from a single neuron:

$$E\left[\bar{s}\right] = NE\left[\frac{s}{N}\right] = E[s] \tag{10.40}$$

However, the variance of the mean count is

$$Var\left[\bar{s}\right] = NVar\left[\frac{s}{N}\right] = N\frac{Var\left[s\right]}{N^2} = \frac{\phi_1 E\left[s\right]}{N} \tag{10.41}$$

This expression for the variance of a mean holds when the counts that are averaged are statistically independent. For sums of independent variables, the variance of the sum is the sum of the variances associated with each random variable.

Unfortunately, MT neurons are not independent in their responses. In that case, the variance of the sum is the sum of all the terms in the N by N covariance matrix depicted here,

$$\begin{pmatrix} \sigma_1^2 & \cdots & r_{1n}\sigma_1\sigma_n \\ \vdots & \ddots & \vdots \\ r_{n1}\sigma_n\sigma_1 & \cdots & \sigma_n^2 \end{pmatrix} \tag{10.42}$$

where r_{ij} is the correlation coefficient between the count from the i^{th} and j^{th} neurons. In addition to N variance terms along the main diagonal, there are the other $N^2 - N$ covariance terms (when the correlation is nonzero). Assuming that all the variances are the same and all the coefficients are the same, we can

specify these terms as follows. When we take the average spike count, each of the diagonal terms is

$$Var\left[\frac{s}{N}\right] = \frac{Var\left[s\right]}{N^2} = \frac{\phi_1 E\left[s\right]}{N^2} \tag{10.43}$$

and each off-diagonal term is

$$rVar\left[\frac{s}{N}\right] = \frac{rVar\left[s\right]}{N^2} = \frac{r\phi_1 E\left[s\right]}{N^2} \tag{10.44}$$

The variance of the mean is the sum of the N terms in equation (10.43) plus the $N^2 - N$ terms in equation (10.44)

$$Var\left[\bar{s}\right] = N\left(\frac{\phi_1 E\left[s\right]}{N^2}\right) + \left(N^2 - N\right)\left(\frac{r\phi_1 E\left[s\right]}{N^2}\right) \tag{10.45}$$

From this expression, it is easy to see that in the limit of large N,

$$\lim_{N\to\infty} Var\left[\bar{s}\right] = r\phi_1 E\left[s\right] \tag{10.46}$$

This limit is a reasonable approximation to the variance of the mean count from as few as N = 100 neurons in the ensemble.

Pairs of MT neurons with overlapping receptive fields exhibit correlation coefficients of $r \approx 0.15$ to 0.2. This implies that the average spike count from an ensemble of rightward motion sensors in MT is a random number whose expectation is the mean spike count from a neuron and whose variance is 0.225 to 0.3 times this expectation. We will use the top end of this range for our exercise below. We will also assume that the right-preferring MT neurons respond independently from left-preferring MT neurons. In other words, neurons with common response properties share some of their variability, but neurons with different preferences do not [38, 3, 16].

When C = 0, right- and left-preferring MT neurons both respond at about 20 sp/s. In 1 second, we expect the mean count from N right-preferring neurons to be 20 spikes per neuron with a variance of 6 spikes^2 per neuron^2. The same numbers apply for the left-preferring neurons. Thus, the expected difference is 0. The variance associated with this difference is 12 spikes^2 per neuron^2, assuming that the right- and left-preferring MT neurons are independent. For the strongest rightward motion condition shown in figure 10.5, C = 0.51. In that case, we expect the mean and variance of the average count from right-preferring neurons to be 40 spikes per neuron and 12 spikes^2 per neuron^2. For the left-preferring neurons, the mean and variance are 15 spikes per neuron and 4.5 spikes^2 per neuron^2. So the difference in the ensemble averages has expectation and variance of 25 and 16.5, respectively. The signal-to-noise ratio is $25/\sqrt{16.5} \approx 6.2$. For motion in the opposite direction, the expectation of the difference reverses sign, but the variance remains the same. If C = 1, then the expected difference is $60 - 10 = 50$ spikes per neuron with variance $(0.3)(60 + 10) = 21$. So the signal-to-noise ratio is $50/\sqrt{21} \approx 11$.

Based on these assumptions, we plotted in figure 10.6 the expected signal-to-noise ratio that right- and left-preferring MT neurons furnish in 1 second. We can see from this example that the signal-to-noise ratio is approximately linear over the range of motion strengths used in our experiments. The best line in this figure has a slope of 12.7. Now relate this to the equation for the psychometric and chronometric functions shown in figure 10.1. The fitting equation presumes that in each millisecond, a random number representing momentary evidence is drawn from a normal distribution with mean $\mu = kC$ and variance 1. What is the value of k that would provide the signal-to-noise ratios graphed in figure 10.6 after 1000 ms? After 1000 samples, the numerator must be 1000μ and the denominator must be the square root of the accumulated variance from 1000 samples. Therefore

$$\begin{aligned} \frac{1000kC}{\sqrt{1000}} &= 12.7C \\ k &= \frac{12.7}{\sqrt{1000}} \approx 0.4 \end{aligned} \tag{10.47}$$

This value is remarkably close to the value for k estimated by fitting the behavioral data in figure 10.1 (0.43 $\pm$ 0.01). Of course, we can also express this number in terms of θ_1. Under the assumption that the momentary evidence in each millisecond is a random draw from a normal distribution with unit variance and mean equal to *kC*, $\theta_1 = -2\mu/\sigma^2 \approx -0.8C$.

10.4.3 The Accumulated Evidence for Direction is Represented by Neurons in the Lateral Intraparietal Area (LIP)

Several observations argue that the LIP might be a suitable place to mediate decisions during the motion task. First, neurons in LIP receive a major input from direction selective neurons in area MT. Second, neurons in LIP project to the superior colliculus, an area that likely generates the choice response (an eye movement) [13, 25, 18]. Third, many neurons in LIP discharge in a sustained, spatially selective fashion when a monkey is instructed to make a delayed eye movement. This persistent activity can last for seconds before the eye movement is made, and it does not require the continued presence of a visual target. It is enough to flash a visual cue and to ask the monkey to make an eye movement to its remembered location some time later. Thus LIP neurons should be capable of representing the outcome of the decision about motion — a plan to make an eye movement to one of the targets. Indeed their capacity to emit persistent elevations in firing rate motivates the hypothesis that they can represent the integral of their inputs — the integral of a pulse is a step in activity. Thus, we suspected that these neurons might play a role in the conversion of momentary evidence to a binary decision (for review, see [30]).

During decision-making in the reaction time motion task, neurons in area LIP undergo ramplike changes in spike rate (figure 10.7). These changes are evident from $\sim$225 ms after onset of the random-dot motion until the monkey initiates its eye movement response. The neural activity represents the process leading to the monkey's choice, rising when the monkey chooses the target in

the neuron's response field (T_{in}) and declining when the monkey chooses the target outside the neuron's response field (T_{out}). The graphs shown in figure 10.7 were obtained in the choice-reaction time experiments that produced the behavioral data shown earlier. The responses on the left side of the graph begin when the motion is turned on and stop at the median reaction time for each coherence. The responses on the right side of the graph are aligned to the end of the trial when the monkey initiates an eye movement response either to the choice target in the response field or to the other target.

During the epoch of decision formation, there is a clear effect of motion strength on the rate of change in activity. Strong (easy) stimuli cause a rapid and consistent increase or decrease in the spike rate. Weak (difficult) stimuli cause smaller increases and decreases that can change direction from moment to moment and thus meander like a particle in Brownian motion. For all stimuli, however, when the spike rate reaches a critical value, a decision for T_{in} is reached, and an eye movement ensues ∼70 ms later [27]. LIP neurons represent which choice the monkey will make, the evolution of this decision in time, and the quality of the evidence upon which this decision is based.

With one important caveat, the responses in figure 10.7 can be related to the diffusion-to-bound model. The decision variable we focused on in this chapter is the accumulation of momentary evidence from MT, what we think of as a difference in spike rates from two ensembles of neurons selective for motion toward T_{in} and T_{out}, respectively. Suppose that beginning 225 ms after onset of motion, the firing rate of the LIP neuron represents the accumulation of this spike rate difference. Then its spike rate on any trial will meander like a sample trajectory shown in figure 10.2 until it reaches a bound. The solid curves in figure 10.7 are the expected averaged trajectories that end in the upper bound. They are averages conditionalized on choice. This is why even when the motion strength is 0%, the average response (solid cyan curve) seems to have a positive drift. This is true for error trials too (not shown).

The dashed curves force the caveat mentioned above. These are also conditionalized averages, but not to the response reaching a lower bound. For every LIP neuron with a right-choice target in its receptive fields, there is an LIP neuron with the left-choice target in its receptive fields. According to our ideas, these neurons also accumulate momentary evidence: a difference in spike rates from two ensembles of neurons selective for motion toward T_{in} and T_{out}. Of course these neurons tend to accumulate the evidence with an opposite sign. If T_{in} and T_{out} refer to right and left for the neuron we are recording, then T_{in} and T_{out} refer to left and right for these other LIP neurons. Suppose these neurons terminate the decision process with a left choice if their firing rates reach an upper bound. On these trials, the neurons we recorded (which signal right choices) will tend to decrease their responses because the evidence is against right. However, the end of the trial is highly variable relative to T_{in} neurons because the recorded neurons are not controlling the end of the trial when motion is in their null direction. This explains why the response averages shown by the dashed curves on the right side of figure 10.7 do not come to a stereotyped lower value before the saccade.

Rather than a single diffusion process (or random walk) that ends in two bounds, the physiology forces us to consider a process that is more like a race between two processes. One gathers evidence for right and against left, say. The other gathers evidence for left and against right. If these sources of momentary evidence were exactly the same and if they account for all of the variability in the LIP firing rates, then the two accumulations would be copies of one another, simply mirrored around the starting firing rate. If this inverse correlation were perfect, we could represent the two accumulations on a single graph with stopping bounds at $\pm$A. Of course, the two LIP neurons (or ensembles) are not perfectly anticorrelated. We see evidence for this in the conditionalized response averages in figure 10.7. But it turns out that so long as the race is between processes that are to a large extent anticorrelated, the equations we developed above hold. This is perfectly sensible because although we call this a race, there is no real competition. When one accumulator takes a positive step, the other takes a negative step. The correlation is not perfect, but this tendency renders moot any potential for the two mechanisms to actually compete to reach their respective upper bounds.

It is useful to think of a race between accumulators because this architecture is likely to extend to decisions among more than two options. It is hard to see how the model in figure 10.2 would extend to three, four, or more choices (see [17]). But a race among N accumulators is straightforward [35]. We are conducting experiments to test these ideas, and they look promising [8].

10.5 Conclusions

We have described a framework, diffusion to bound, for understanding how simple decisions are made in the brain. We have explained the mathematical tools that allow us to make predictions about the speed and accuracy of such decisions and have shown that these predictions match the behavior we observe. Lastly, we have provided evidence that the neural basis for the machinery we have described is in the parietal cortex.

We hope the mathematical tutorial will be useful to students interested in applying the equations. Indeed, the more general forms of the equations developed in terms of θ_1 may provide deeper insight into processes that are less straightforward than our motion experiments. Insights of fundamental importance to perception and decision-making can be found in [20, 19].

Decision-making in the random-dot motion task certainly takes place in a rarefied context: only two choices are present, they are equally probable *a priori*, there is only one source of relevant evidence, and all correct answers are of the same reward value. Yet, even these simplified conditions provide insights into the basic computations involved in decision-making. By suggesting the simple framework of accumulation of noisy momentary evidence to bound, we have a firm foundation from which we can begin to ask questions about more complicated decisions.

Are there certain kinds of decisions for which simply accumulating evidence

over time is a poor strategy? Certainly. To integrate, the brain must know when the relevant information should be gathered. It would be detrimental to accumulate noise when the stimulus is not present. For example, many detection experiments in vision show little evidence for temporal integration beyond ~80 ms, the limits of Bloch's law [36, 37]. Even if there is no uncertainty about when to integrate, it may make little sense to do so if sufficient information can be acquired in a short epoch. We suspect that many discrimination problems in vision involve analyses of spatial correspondences (e.g., curvature, texture, elongation) that are best achieved on a stable representation of the image in the visual cortex. These analyses are typically performed in a snapshot, so to speak, before the pieces of the representation have time to shift about on the cortical map. While the fidelity of any decision variable could improve with further temporal integration, we suspect that often the benefit in accuracy is not sufficient to overcome natural tendencies to trade accuracy against speed.

On the other hand, organisms are faced with many decisions where the evidence is noisy or arrives in piecemeal fashion over time. For decision-making under these conditions, accumulating evidence in favor of one choice or the other may be an optimal strategy, and may be well described by a bounded accumulation model. Ultimately, we would like to know how the brain combines evidence from multiple sources over time with factors associated with prior probability, predicted costs and benefits of the outcomes, and the cost of time. We would like to understand the neural mechanisms that underlie the incorporation of these factors into the decision process. We are optimistic that the principles of sequential analysis will guide this research program. Just how far we can push the kind of simple perceptual decisions described in this chapter toward achieving this understanding remains to be seen.

Appendix 10.1: Discrete Increments or Summation of Infinitesimals?

From one perspective, there is no particular reason to conceive of the accumulation of increments and decrements in discrete steps. According to this argument, each Δt is infinitesimally small: $\Delta t \rightarrow \delta t$. The random increments are drawn from a distribution with mean μ and variance $\sigma^2 \delta t$. Indeed, the process can be written as a stochastic differential equation for the change in the accumulation, Y:

$$\frac{dY}{dt} = \mu + N\left\{0, \sigma\sqrt{dt}\right\} \tag{10.48}$$

The second term is the noise term, a normal distribution with mean and standard deviation given by the arguments in the curly braces. The peculiar term for the standard deviation ensures that the variance of the accumulation is σ^2 when t = 1. Remember, it is the variance that accumulates, not the standard deviation. Equation (10.48) gives a glimpse of how one might set up a variety of accumulation problems, for example, those that involve some leakiness of the accumulation. For additional reading on this approach, we recommend [33, 34].

There are two important points to be made here. First, by moving to the infinitesimal, we tacitly assume that the random increments are drawn from a normal distribution. This is simply a consequence of the central limit theorem: in any finite Δt, there are so many infinitesimal increments that they must add to a random number that is Gaussian [33]. So, according to this formulation, we can forget about other distributions. In short, forget about generalizations revolving around θ_1; according to this formulation, there is only one formula for θ_1 (the one in equation (10.25)). The second point counters this position.

A common assumption behind all of the formulations we have considered is that the increments that are accumulated to form the random walk are independent of one another. That is why the variance of Y is the sum of the variances of the increments and not the sum of the covariances. This assumption is suspect in the brain, especially over very short time scales. Responses from neurons are weakly correlated over a time scale of $\sim$10 to 50 ms [3]. To obtain independent samples of the spike rate from an ensemble of neurons in the visual cortex, it is essential to wait a minimum of 10 ms between samples (see [21]). Over very short time scales, and certainly for infinitesimal increments, independent samples are unlikely to be available in the brain for accumulation. Therefore, when we map the momentary evidence to spike rates from ensembles of neurons in the brain, it is difficult to embrace the assumption of independence over very short time scales. For this reason, we think it is useful to consider the random walk process as if it occurs in discrete time with increments that are not necessarily drawn from a Gaussian distribution.

Acknowledgments

This study was supported by the Howard Hughes Medical Institute (HHMI) and the National Eye Institute. T.D.H. is also supported by a HHMI predoctoral fellowship.

References

[1] Albright TD (1993) Cortical processing of visual motion. In Miles FA, Wallman J, eds, *Visual Motion and Its Role in the Stabilization of Gaze*, pages 177-201, New York: Elsevier.

[2] Allman JM, Kaas JH (1971) A representation of the visual field in the caudal third of the middle temporal gyrus of the owl monkey (*Aotus trivirgatus*). *Brain Research*, 31:85-105.

[3] Bair W, Zohary E, Newsome WT (2001) Correlated firing in macaque visual area MT: time scales and relationship to behavior. *Journal of Neuroscience*, 21:1676-1697.

[4] Born RT, Bradley DC (2005) Structure and Function of Visual Area MT. *Annual Review of Neuroscience*, 28:157-189.

[5] Britten K (2003) The middle temporal area: motion processing and the link to perception. In Chalupa WJ, Werner JS, eds., *The Visual Neurosciences*, volume 2, pages 1203-1216, Cambridge, MA: MIT Press.

[6] Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *Journal of Neuroscience*, 12:4745-4765.

[7] Britten KH, Shadlen MN, Newsome WT, Movshon JA (1993) Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, 10:1157-1169.

[8] Churchland AK, Kiani R, Tam M, Palmer J, Shadlen MN (2005) Responses of LIP neurons reflect accumulation of evidence in a multiple choice decision task. 2005 Abstract Viewer/Itinerary Planner Washington, DC: Society for Neuroscience Program No. 16.8.

[9] Cox D, Miller H (1965) *The theory of stochastic processes*. London: Chapman and Hall.

[10] Ditterich J (2006) Stochastic models of decisions about motion direction: Behavior and physiology. *Neural Networks*, in press.

[11] Ditterich J, Mazurek M, Shadlen MN (2003) Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience*, 6:891-898.

[12] Dubner R, Zeki SM (1971) Response properties and receptive fields of cells in an anatomically defined region of the superior temporal sulcus. *Brain Research*, 35:528-532.

[13] Gnadt JW, Andersen RA (1988) Memory related motor planning activity in posterior parietal cortex of monkey. *Experimental Brain Research*, 70:216-220.

[14] Gold JI, Shadlen MN (2003) The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *Journal of Neuroscience*, 23:632-651.

[15] Karlin S, Taylor HM (1975) *A First Course in Stochastic Processes, 2nd edition*. Boston: Academic Press.

[16] Kohn A, Smith MA (2005) Stimulus dependence of neuronal correlation in primary visual cortex of the macaque. *Journal of Neuroscience*, 25:3661-3673.

[17] Laming DRJ (1968) *Information Theory of Choice-Rreaction Times*. London: Academic Press.

[18] Lewis JW, Van Essen DC (2000) Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *Journal of Comparative Neurology*, 428:112-137.

[19] Link SW (1992) *The Wave Theory of Difference and Similarity*. Hillsdale, NJ: Lawrence Erlbaum Associates.

[20] Link SW, Heath RA (1975) A sequential theory of psychological discrimination. *Psychometrika*, 40:77-105.

[21] Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A role for neural integrators in perceptual decision making. *Cerebral Cortex*, 13:1257-1269.

[22] Newsome WT, Britten KH, Movshon JA (1989) Neuronal correlates of a perceptual decision. *Nature*, 341:52-54.

[23] Orban GA, Fize D, Peuskens H, Denys K, Nelissen K, Sunaert S, Todd J, Vanduffel W (2003) Similarities and differences in motion processing between the human and macaque brain: evidence from fMRI. *Neuropsychologia*, 41:1757-1768.

[24] Palmer J, Huk AC, Shadlen MN (2005) The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5:376-404.

[25] Paré M, Wurtz RH (1997) Monkey posterior parietal cortex neurons antidromically activated from superior colliculus. *Journal of Neurophysiology*, 78:3493-3498.

[26] Ratcliff R, Rouder JN (1998) Modeling response times for two-choice decisions. *Psychological Science*, 9:347-356.

[27] Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of Neuroscience*, 22:9475-9489.

[28] Salzman CD, Britten KH, Newsome WT (1990) Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346:174-177.

[29] Sartre JP (1984) Existentialism and Human Emotions. New York: Carol Publishing Group.

[30] Shadlen MN, Gold JI (2004) The neurophysiology of decision-making as a window on cognition. In Gazzaniga MS, eds., *The Cognitive Neurosciences, 3rd edition*, pages 1229-1241, Cambridge, MA: MIT Press.

[31] Smith PL (1990) A note on the distribution of response times for a random walk with Gaussian increments. *Journal of Mathematical Psychology*, 34:445-459.

[32] Smith PL (1995) Psychophysically principled models of visual simple reaction time. *Psychological Review*, 102:567-593.

[33] Smith PL (2000) Stochastic dynamic models of response time and accuracy: A foundational primer. *Journal of Mathematical Psychology*, 44:408-463.

[34] Usher M, McClelland JL (2001) The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108:550-592.

[35] Usher M, Olami Z, McClelland JL (2002) Hick's law in a stochastic race model with speed-accuracy tradeoff. *Journal of Mathematical Psychology*, 46:704-715.

[36] Watson AB (1979) Probability summation over time. *Vision Research*, 19:515-522.

[37] Watson AB (1986) Temporal sensitivity. In Boff KR, Kaufman L, Thomas JP, eds., *Handbook of Perception and Human Performance*, pages 6.1-6.43, New York: Wiley.

[38] Zohary E, Shadlen MN, Newsome WT (1994) Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370:140-143.

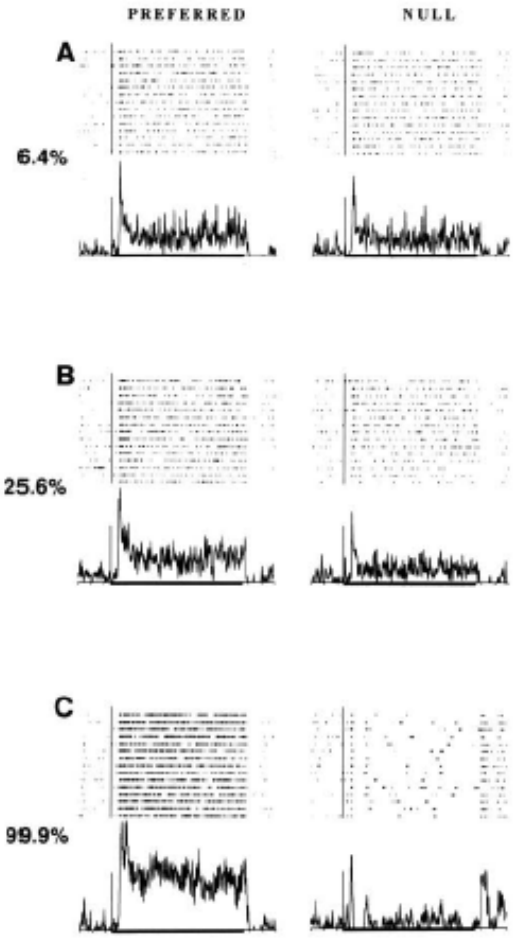

Figure 10.5 Response of MT neurons to stimuli with different motion strengths. A. Responses to motion at 6.4% coherence. Left column shows responses to motion in the preferred direction, and the right column shows responses to motion in the opposite direction. Top: Rasters of neural spike times. Solid line indicates onset of stimulus motion. Bottom: Average firing rate. The vertical bar at stimulus onset represents 100 sp/s, and the horizontal bar under the abscissa shows the 2-second duration of stimulus presentation. B. Same as A for 25.6% coherence. C. Same as A for 99.9% coherence. (Reproduced with permission from K. H. Britten, Shadlen, W. T. Newsome, J. A. Movshon, Response of neurons in macaque MT to stochastic motion signals. Vis. Neurosci. 10:1157-1169. ©1993 by the Cambridge University Press.)

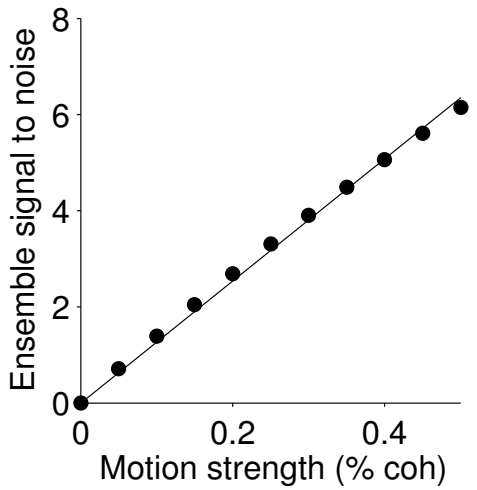

Figure 10.6 Estimated signal-to-noise ratio from an ensemble of MT neurons at different motion strengths. The calculations used to make these estimates are described in the text, and they are based on a stimulus duration of 1 second. The relationship is approximately linear over the range of motion strengths used in our experiments.

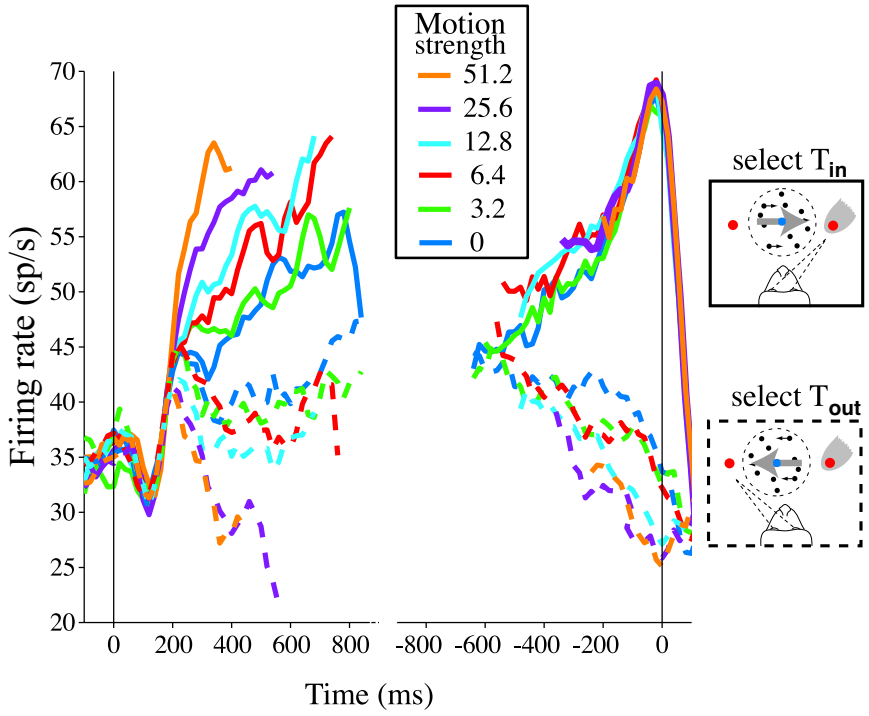

Figure 10.7 Time course of LIP activity during the reaction time version of the direction discrimination task. Average spike rates for 54 LIP neurons. Responses are grouped by choice and motion strength. Solid and dashed lines show responses for trials that were terminated with a T_{in} and T_{out} choice, respectively. The colors indicate the motion strength. The responses are aligned on the left to the onset of motion stimulus and drawn only up to the median RT for each motion strength, excluding any activity within 100 ms of the saccade. The responses on the right are aligned to the initiation of the saccadic eye movement response, and they are plotted backward in time to the median RT, excluding any activity occurring within 200 ms of motion onset. Only correct choices are included except for the 0% coherence case. Adapted with permission from J. D. Roitman and M. N. Shadlen, [27].

Index