

How a biological decision network can implement a statistically optimal test

Rafal Bogacz¹, Eric Brown², Jeff Moehlis³, Philip Holmes^{2,4}, Jonathan D. Cohen^{5,6}

¹Department of Computer Science, University of Bristol, Bristol, BS8 1UB, UK

²Program in Applied and Computational Mathematics, Princeton University, Princeton, NJ 08544, USA

³Department of Mechanical Engineering, University of California, Santa Barbara, CA 93106, USA

⁴Department of Mechanical and Aerospace Engineering, Princeton University, Princeton, NJ 08544, USA

⁵Center for the Study of Brain, Mind and Behavior, Princeton University, Princeton, NJ 08544, USA

⁶Department of Psychology, Princeton University, Princeton, NJ 08544, USA

Abstract

Neurophysiological evidence due to Schall, Newsome and others indicates that decision processes in certain cortical areas (e.g. FEF and LIP) involve the integration of noisy evidence. Within this paradigm, we ask which neuronal architectures and parameter values would allow an animal to make the fastest and most accurate decisions. Since evolutionary pressure promotes such optimality (e.g. in prey capture and predator avoidance), it is plausible that biological decision networks realise or approximate optimal performance. We consider a simple decision model proposed by Usher & McClelland consisting of two populations of neurons integrating evidence in support of two alternatives, and we analyze the dynamics of this model. We show that in order to implement the optimal decision algorithm (sequential probability ratio test) the linearised network must satisfy the following two constraints: (i) it must accumulate the difference between evidence in support of each alternative, as would be implemented by mutual inhibition between the populations; and (ii) the strength of mutual inhibition must be equal to the leak of activity from each population.

1 Introduction

Decision making is a very frequent element of life of humans and animals, and accuracy and speed of the decisions is critical to animal survival. During millions of years of evolution, evolutionary pressure promoted animals whose brains made more efficient decisions. Hence it is plausible that decision circuits in the brain possess architectures and parameters allowing them to implement optimal or nearly optimal algorithms. Therefore, in order to identify the architecture and parameters of decision networks in the brain, it may be informative to ask what is the optimal algorithm for decision making, and what biologically plausible network may implement this optimal algorithm.

This optimality approach is not guaranteed to reveal the true decision network in the brain. But it can produce interesting and counterintuitive experimental predictions, which may be used to test the model suggested by the approach. Furthermore, the algorithm optimally solving a decision problem may uncover (or may inspire) practical computational applications. This report shows how the mathematical analysis of decision processes may help in understanding them and make predictions concerning the architecture of neural networks involved in decision making. In particular, it identifies parameters of the decision making model proposed by Usher & McClelland (2001) under which the computations of the neural decision network are equivalent to an optimal statistical test for decision making (sequential probability ratio test).

In three following sections we briefly review neurophysiology of decision, optimal statistical test, and the model of decision network by Usher & McClelland (2001). Then in Section 5 we identify conditions under which this model achieves optimal performance. Finally, in Section 6 we list other directions in which the theory has been extended.

2 Neurophysiology of decision

The neurobiology underlying decision making has been extensively studied in a task in which monkeys are presented with a visual field of small dots most of which are moving randomly, but a certain fraction of which are moving left on some trials and right on others (Britten et al., 1993). Typically, the animal is trained to respond by making a saccade in the direction of the coherently moving dots. Figure 1a shows schematically the typical patterns of activity observed in area MT of monkeys performing this task (this area is involved in motion processing). When a stimulus with coherent leftward motion is presented, the firing rate of an MT neuron selective for leftward motion is higher than the firing rate of a neuron selective for rightward motion (Britten et al., 1993) (in Figure 1a the grey curve is more often above the black). However, the firing rates for both types of neurons are noisy, hence decisions based on the activity of MT neurons at a given moment in time would

be inaccurate. This reflects the uncertainty inherent in the stimulus and its neural representation.

Figure 1b shows schematically the pattern of activity of neurons in area LIP, which is involved in controlling eye movements. The LIP neurons are believed to integrate the input from MT neurons over the duration of a trial. The decision based on the integrated evidence, namely on activity of LIP neurons after about 0.5s, is much more accurate. This example illustrates that the decision process may be realized in the neural substrate by the integration of noisy information.




Figure 1. Cartoon of typical peri-stimulus time histograms of neuronal activity during ‘the moving dots task’. The figure does not show the actual data, but it is a sketch based on data described by Britten et al. (1993), Shadlen & Newsome (2001), and Schall (2001). Horizontal axes show time from stimulus onset. Vertical axes indicate firing rate. Representative firing rates are shown for stimulus with coherent leftward motion. a) Firing rate of neurons in the area MT: gray line represents a typical neuron selective to leftward motion, and black line for rightward motion (the curves were generated by adding noise to exponentially decaying functions). b) Firing rate of neurons in the area LIP: gray line represents a typical neuron selective for leftward saccades, and black line for rightward saccades (the curves were generated by integrating the difference between curves from panel a and some noise).

3 The decision problem

Let us formalize the decision problem on the basis of the above example. We assume that there are two populations of neurons whose activities provide evidence in support of the two alternative decisions (e.g., corresponding to the two groups of MT neurons in the above example). We denote the activities of the first populations over a given trial by $y_1^1, y_2^1, \dots, y_n^1$, and of the second population by $y_1^2, y_2^2, \dots, y_n^2$. Let us assume that the samples y_i^1 come from a normal

distribution with mean I_1 and standard deviation c , and samples y_i^2 come from a normal distribution with mean I_2 and standard deviation c . The goal of the decision process is to identify which mean activity, I_1 or I_2 , is higher. Note that this is equivalent to asking whether $I_1 - I_2$ is positive or negative, i.e. whether the differences between input samples have positive or negative mean. Let us denote the differences between activities of input populations by $x_i = y_i^1 - y_i^2$.

Within this framework, the question of optimal decision making is the following: For given signal and noise levels I_1, I_2 and c in the input populations, what is the optimal strategy for integration of evidence (e.g., by LIP neurons) that would allow the most accurate and fastest decisions on average, over the course of many trials? More precisely, there are two sub-questions: (i) Which strategy yields the lowest error rate (ER), when a given (fixed) time for decision is allotted, and (ii) which strategy yields the fastest reaction times (RT) for a given ER?

The two questions above refer to optimality in two different conditions under which decision tasks can be run. The first question relates to a decision process in which participants are presented with the stimulus for a fixed duration, at the end of which they are expected to answer, usually on presentation of response prompt, thus constraining their RTs. The second question refers to a decision process in which participants are asked to respond freely, when they are ready, usually being instructed to be as accurate and as fast as possible. We refer to this situation as the *free-response protocol*. We focus only on this protocol.

The answer to the second question, regarding optimality in the free-response paradigm, is provided by the sequential probability ratio test (SPRT) of Barnard (1946) and Wald (1947). In contrast to classical decision procedures in which a previously fixed number of samples is collected before the decision is rendered, SPRT may be applied to continuously accumulating data. A decision is made as soon as a threshold, which depends upon the required accuracy, is reached. Specifically, let H_1 and H_0 denote the two alternative hypotheses, as above, and assume that samples x_i in support of these are drawn at random from two probability distributions with densities $p_1(x), p_0(x)$. In particular, in the case of the decision problem defined at the beginning of this Section, $H_1: p_1(x)$ is a normal distribution with positive mean μ and standard deviation σ , $H_0: p_0(x)$ is a normal distribution with negative mean $-\mu$ and standard deviation σ . After each sample the ratio of probabilities $p_1(x_i)/p_0(x_i)$ is calculated and the product of these likelihood ratios is accumulated. Observations continue as long as the likelihood ratio lies within the boundaries $Z_0 < Z_1$:

$$Z_0 < \frac{p_1(x_1)p_1(x_2)\dots p_1(x_n)}{p_0(x_1)p_0(x_2)\dots p_0(x_n)} < Z_1 \quad (1)$$

Thus, after each measurement one recomputes the current likelihood ratio, thereby assessing the net weight of evidence in favor of H_1 over H_0 . When the ratio first exceeds Z_1 or falls below Z_0 , sampling ends and either H_1 or H_0 is respectively accepted; otherwise sampling continues.

SPRT is the optimal test for decision-making on the basis of accumulating noisy data in the following sense: Among all fixed or variable sample decision methods that guarantee fixed error probabilities, SPRT requires on average the smallest number of samples to render a decision (Wald & Wolfowitz, 1948). In other words, for given ER, SPRT delivers the fastest RT.

The SPRT is equivalent to a random walk with thresholds corresponding to alternative decisions. To see this, take the logarithm of both sides of Equation 1:

$$\log Z_0 < \log \frac{p_1(x_1)}{p_0(x_1)} + \dots + \log \frac{p_1(x_n)}{p_0(x_n)} < \log Z_1 \quad (2)$$

If we denote the logarithm of the likelihood ratio defined in Equation 1 by I^n , then Equation 1 implies that we iteratively accumulate I^n after each observation:

$$I^n = I^{n-1} + \log \frac{p_1(x_n)}{p_0(x_n)} \quad (3)$$

Let us evaluate the probability ratio for the hypotheses defined earlier in the section. Equation 3 becomes:

$$\begin{aligned} I^n &= I^{n-1} + \log p_1(x_n) - \log p_0(x_n) = \\ &= I^{n-1} + \log \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x_n-\mu)^2}{2\sigma^2}} - \log \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{(x_n+\mu)^2}{2\sigma^2}} = \\ &= I^{n-1} + \frac{-x_n^2 + 2\mu x_n - \mu^2 + x_n^2 + 2\mu x_n + \mu^2}{2\sigma^2} = \\ &= I^{n-1} + \frac{2\mu}{\sigma^2} x_n \end{aligned}$$

Thus from the above equation, the SPRT is equivalent to a random walk starting at $I^0 = 0$, and continuing until the logarithm of the likelihood ratio I^n reaches one of the thresholds: $\log Z_0$ or $\log Z_1$. During this random walk the samples (i.e. the differences between the two inputs) are accumulated (a more rigorous and complete analysis of the relationship between SPRT and random walks is given by Gold & Shadlen, 2001).

4 Model of decision network

Figure 2a shows the architecture of an abstract neural network (or connectionist) model for the two alternative decision tasks (Usher & McClelland, 2001). The model includes four units representing the mean activities of neuronal populations: two input units representing populations providing evidence in support of the two alternative decisions (e.g., corresponding to groups of movement sensitive MT neurons from the example in Section 2); and two decision units representing populations integrating the evidence




Figure 2. Usher & McClelland (2001) model. a) Architecture of the model: Arrows denote excitatory connections, line with filled circles denotes inhibitory connections. b) An example of the evolution of the model, showing y_1 and y_2 as functions of time. c) The phase- or state space of the mutual inhibition model. Horizontal axis denotes the activation of the first decision unit; vertical axis denotes the activation of the second decision unit. The path shows the decision process from stimulus onset (where $y_1 = y_2 = 0$) to reaching a decision threshold (decision thresholds are shown by dashed lines). The model was simulated for the following parameters: $I_1 = 2$, $I_2 = 1.5$, $c = 0.2$, $w = k = 1.5$, $Z = 1$. The simulations were performed using Euler method with time-step $\Delta t = 0.01$. To simulate the Wiener processes, at every step of integration, each of the variables y_1 and y_2 was increased by a random number from normal distribution with mean 0 and variance $c^2\Delta t$.

(e.g., corresponding to the LIP neurons involved in controlling eye movement).

The decision units are modeled as leaky integrators with activity levels denoted by y_1 and y_2 . Each decision unit accumulates evidence from an input unit with mean activity I_j and independent white noise fluctuations η_i of Root Mean Square (RMS) strength c (η_i denote independent Wiener processes). These units also inhibit each other by way of a connection of weight w . Hence, during the decision process, information is accumulated according to:

$$\begin{cases} \dot{y}_1 = -ky_1 - wy_2 + I_1 + c\eta_1 \\ \dot{y}_2 = -ky_2 - wy_1 + I_2 + c\eta_2 \end{cases}, \quad y_1(0) = y_2(0) = 0. \quad (4)$$

In the equations above, the term k denotes the decay rate of the units' activity (i.e., the leak) and $-wy_i$ denotes the mutual inhibition. Note that terms $-ky_i$ cause the activity to decay to zero in the absence of inputs to the unit (because if y_i were positive in the absence of inhibition, input, and noise, \dot{y}_i would be negative, and y_i would decrease). The scale of the units' activity is chosen so that zero represents the baseline activity of both units in the absence of all inputs, hence integration starts from $y_1(0) = y_2(0) = 0$. As soon as either unit exceeds a preassigned threshold Z , the model is assumed to make a response.

The state of this model at a given moment in time is described by the values of y_1 and y_2 , and may therefore be represented as a point on a *phase plane* whose horizontal and vertical axes correspond to y_1 and y_2 ; the evolution of activities of the decision units during the decision process may be visualized as a path in this plane. An example is shown in Figure 2c, corresponding to the individual time courses of y_1 and y_2 shown in Figure 2b.

5 Model parameters resulting in optimal performance

As illustrated in Section 4, the behaviour of the model may be visualized by plotting states on the phase plane. Figure 2c shows a representative path in state space: initially the activities of both decision units increase due to stimulus onset, but as the units become more active, mutual inhibition causes the activity of the 'weaker' unit to decrease and the path moves toward the threshold for the more strongly activated unit (i.e., the correct choice).

To better understand the dynamics of the model, Figure 3 shows its *vector fields* for three different parameter ranges. Each arrow shows the average direction in which the state moves from the point indicated by the arrow's tail, and its length corresponds to the speed of movement (i.e., rate of change) in the absence of noise. In Figure 3, as for most other simulations described in this article, we set $I_1 > I_2$; that is, we assume that the first alternative is the correct one (the opposite case is obtained simply by reflecting about the diagonal $y_1 = y_2$).

Note that in all three panels of Figure 3 there is a line (an eigenvector), sloping down and to the right, to which system states are attracted: The arrows point towards this line from both sides. The orientation of this line represents an important dimension: the difference in activity between the two decision units. Note that the evolution *along* the line differs for different values of decay and inhibition, as does the strength of attraction toward the line, and its location in the phase plane. Most of the interesting dynamics determining decisions occur along this line. Therefore, it is easier to understand these in terms of new coordinates rotated clockwise by 45° with respect to the y_1 and y_2 coordinates, so that one of the new axes is parallel to the attracting line. These new coordinates are shown in Figure 3b, denoted by x_1 (parallel to the attracting line) and x_2 . The transformation from y to x coordinates is given by (cf. Seung, 2003):

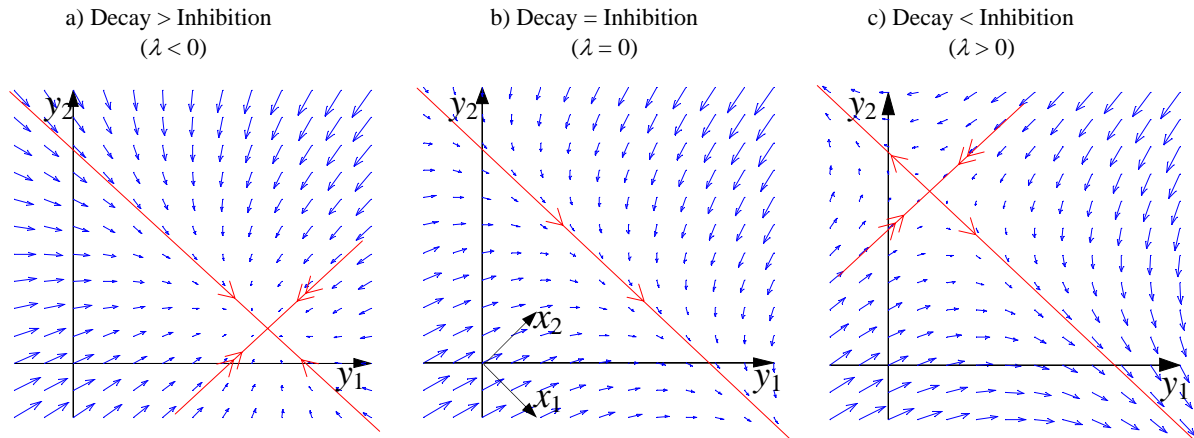


Figure 3. Vector fields for the model. In all plots $I_1 = 2, I_2 = 1$. Inhibition (w) and decay (k) have different values in different panels: a) $w = 0.5, k = 1.5$; b) $w = 1, k = 1$; c) $w = 1.5, k = 0.5$.

$$\begin{cases} x_1 = \frac{y_1 - y_2}{\sqrt{2}}, \\ x_2 = \frac{y_1 + y_2}{\sqrt{2}}. \end{cases} \quad (5)$$

Equations 5 derive from the geometry shown in Figure 3b: x_1 describes the difference between activities of the two decision units, while x_2 describes the sum of their activities. The square root of two in the denominators of Equations 5 is a normalization factor, included to ensure that y and x coordinates have the same scale.

In deciding between two alternatives, it is natural that the *difference* between the activities of the units selective for the alternatives should be a useful descriptor of the decision process (see Section 3). However, the new coordinates do more than merely emphasize this point. They allow us to factor the two Equations 4 that describe the decision process into two decoupled processes, separating the evolution of the difference in the activity of the two units (x_1) from the change in their overall (summed) activity (x_2). If we can show that the latter has minimal impact on the decision process, then we can reduce the description of this process from one that is two-dimensional to a simpler one that is one-dimensional. As we will show, for certain parameters this one-dimensional description reduces to the diffusion model (Ratcliff, 1978).

To transform Equations 4 into the new coordinates, we first calculate the derivative (rate of change) of x_1 . Substituting Equations 4 into the first of Equations 5, we obtain:

$$\begin{aligned} \dot{x}_1 &= \frac{\dot{y}_1 - \dot{y}_2}{\sqrt{2}} = \\ &-k \frac{y_1 - y_2}{\sqrt{2}} + w \frac{y_1 - y_2}{\sqrt{2}} + \frac{I_1 - I_2}{\sqrt{2}} + \frac{1}{\sqrt{2}}(c\eta_1 - c\eta_2) \end{aligned} \quad (6)$$

We assumed earlier that the noise processes for the input units are independent. Since the standard deviation of the sum (or difference) of two independent random variables is equal to the square root of the sum of their variances, the noise process in x_1 may be written:

$$\frac{1}{\sqrt{2}}(c\eta_1 - c\eta_2) = \frac{\sqrt{c^2 + c^2}}{\sqrt{2}}\eta_1 = c\eta_1. \quad (7)$$

In Equation 7, η_1 again denotes a noise process with mean equal to 0 and an RMS strength of 1. Substituting Equation 7 and the definition of x_1 from Equation 5 into Equation 6, we obtain Equation 8. Following analogous calculations for x_2 , we have:

$$\dot{x}_1 = (w - k)x_1 + \frac{I_1 - I_2}{\sqrt{2}} + c\eta_1, \quad (8)$$

$$\dot{x}_2 = (-k - w)x_2 + \frac{I_1 + I_2}{\sqrt{2}} + c\eta_2. \quad (9)$$

Equations 8 and 9 are *uncoupled*; that is, the rate of change of each x_i depends only on x_i itself (this was not the case for the decision units in Equations 4). Hence, the evolution of x_1 and x_2 may be analyzed separately, and in fact each is described by an Ornstein-Uhlenbeck (O-U) process (Busemeyer & Townsend, 1993). In particular, Equation 8, for the x_1 process, involves a drift term proportional to the *difference* between the inputs I_1 and I_2 . This process may be stable or unstable, depending upon the relative magnitudes of k and w . Equation 9, for the x_2 process, always gives a stable O-U process (corresponding to attraction to the line in Figure 3), since $-k - w < 0$.

We first consider the dynamics in the x_2 direction, corresponding to the summed activity of the two decision units. As noted above, on all panels of Figure 3 there is a line to which the noise-free state is attracted, implying that x_2 approaches a limiting value as time increases. The rate of this (exponential) approach is $\lambda = k + w$, and it is kept constant in the three cases of Figure 3 by setting $k + w = 2$.

Figure 3 also shows that the dynamics of the system in the direction of coordinate x_1 depends on the relative values of inhibitory weight w and decay k . This dependence is due to the fact that the dynamics of x_1 are described in Equation 8 by a O-U process with coefficient $\lambda = w - k$. When decay is larger than inhibition, then $\lambda < 0$, and there is also an attractor for the x_1 dynamics, as shown in Figure 3a. When inhibition is larger than decay, then $\lambda > 0$, and there is repulsion from the fixed point in the x_1 direction, as shown in Figure 3c. The fixed point is a saddle in this case.

Since $|k+w| > |k-w|$ for *all* parameter values $k > 0$ and $w > 0$, the average state of the system approaches the attracting line faster (and often considerably faster) than it moves along it (e.g., see Figure 2c). Hence, the decision process divides into two phases: an initial phase in which the activity of both units increases quickly, and there is rapid equilibration to a neighborhood around the attracting line; followed by slower motion along the line, governed by an O-U process in which the difference between the activities of the two units grows as one of them prevails and the other subsides.

Most relevant to the current discussion, when inhibition equals decay the term $(w - k)x_1$ in Equation 8 disappears. The vector field for this case is shown in Figure 3b. When inhibition and decay are both fairly strong (as in Figure 3b), the attraction toward the line dominates diffusion along it. Hence, typical paths migrate quickly toward the attracting line and then move slowly along (or near) it.

In this case when inhibition is equal to decay, the position of the system in direction x_1 simply accumulates the difference between the evidence in support of first decision and in support of the second decision, and thus undergoes the diffusion process (Stone, 1960; Lamming, 1968; Ratcliff, 1978). Hence when inhibition is equal to decay, the Usher & McClelland model implements the optimal sequential probability ratio test described in Section 3.

Thus one can expect that the Usher & McClelland (2001) model makes the fastest decisions for fixed error rates when it is closest to the diffusion model, namely when the decay equals inhibition. This is indeed the case, as illustrated in Figure 4.




Figure 4. Performance of the model. The following parameters were kept fixed: $I_1 = 1$, $I_2 = 0$, $c = 1$, $w = 1$. Decay (k) is varied (shown on X-axis). Y-axis shows decision time (DT) for the threshold set such that error rate (ER) = 10%. For each set of parameter values, the threshold was increased from zero in steps of 0.01 until the model reached an ER less than or equal to 10%. For each value of the threshold 10,000 trials were simulated. The dots below and above the line indicate the standard error; that is, the standard deviation of DTs across trials divided by the square root of the number of trials (100).

To summarize, when decay is equal to inhibition and both are relatively large, the Usher & McClelland model (2001) approximates the sequential probability ratio test and thus achieves the optimal performance. Thus we predict that in cortical decision network effective decay is also equal to effective inhibition.

6 Further directions

We have extended the theory of neural bases of decision optimization in a number of directions:

- We have shown how more biologically realistic model by Wang (2002) may implement the optimal test (Bogacz et al., 2005).
- We analyzed the values of optimal threshold maximizing the reward rate, which yields a simple relationship between error rates and reaction times, which has been confirmed in a behavioural experiment (Bogacz et al., 2005).
- We analyzed biased decisions in which one of alternatives is more frequent or more rewarded (Bogacz et al., 2005).
- We analysed non-linear version of Usher & McClelland model (2001), and role of gain modulation (Brown et al., 2005)

- We investigated how the optimality generalizes to multiple alternatives (McMillen & Holmes, 2005).

Acknowledgments

This work was supported by the following grants: NIH P50 MH62196, DOE DE-FG02-95ER25238 (P.H.), EPSRC EP/C514416/1, NSF Mathematical Sciences Postdoctoral Research Fellowship (held by J.M.), and the Burroughs-Wellcome Program in Biological Dynamics and Princeton Graduate School (E.B.). We thank anonymous reviewers for very useful comments.

References

- Barnard, G.A. (1946). Sequential tests in industrial statistics. *Journal of Royal Statistical Society Supplement*, 8, 1-26.
- Bogacz, R., Brown, E.T., Moehlis, J., Hu, P., Holmes P., & Cohen, J.D. (2004). The physics of optimal decision making: A formal analysis of performances in two-alternative forced choice tasks. *Psychological Review* (under review).
- Britten, K.H., Shadlen, M.N., Newsome, W.T., & Movshon, J.A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, 10, 1157-1169.
- Brown, E., Gao, J., Holmes, P., Bogacz, R., Gilzenrat, M., & Cohen J.D. (2004). Simple networks that optimize decisions. *International Journal of Bifurcations and Chaos*, in press.
- Busemeyer, J.R., & Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in uncertain environment. *Psychological Review*, 100, 432-459.
- Gold, J.I., & Shadlen, M.N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, 5, 10-16.
- Laming, D.R.J. (1968). *Information theory of choice reaction time*. New York: Wiley.
- McMillen, T., Holmes, P. The dynamics of choice among multiple alternatives. (under review).
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 83, 59-108.
- Schall, J.D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience*, 2, 33-42.
- Seung, H.S. (2003). Amplification, attenuation, and integration. Adbib, M.A. (Ed.) *The Handbook of Brain Theory and Neural Networks*, 2nd edition. Cambridge, MA: MIT Press, pp. 94-97.
- Shadlen, M.N., & Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, 86, 1916-1936.

- Stone, M. (1960). Models for choice reaction time. *Psychometrika*, 25, 251-260.
- Usher, M., & McClelland, J.L. (2001). On the time course of perceptual choice: the leaky competing accumulator model. *Psychological Review*, 108, 550-592.
- Wald, A. (1947). *Sequential Analysis*. New York: Wiley.
- Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19, 326-339.
- Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, 36, 1-20.