

6 Extending a biologically inspired model of choice: multi-alternatives, nonlinearity, and value-based multidimensional choice

Rafal Bogacz, Marius Usher, Jiaxiang Zhang, and James L. McClelland

Summary

The Leaky Competing Accumulator (LCA) is a biologically inspired model of choice. It describes the processes of leaky accumulation and competition observed in neuronal populations during choice tasks and it accounts for reaction time distributions observed in psychophysical experiments. This chapter discusses recent analyses and extensions of the LCA model. First, it reviews the dynamics and it examines the conditions that make the model achieve optimal performance. Second, it shows that nonlinearities of the type present in biological neurons improve performance when the number of choice-alternatives increases. Third, the model is extended to value-based choice, where it is shown that nonlinearities in the value function, explain risk-aversion in risky-choice and preference reversals in choice between alternatives characterised across multiple dimensions.

6.1 Introduction

Making choices on the basis of visual perceptions is an ubiquitous and central element of human and animal life, which has been studied extensively in experimental psychology. Within the last half century, mathematical models of choice reaction times have been proposed which assume that, during the choice process, noisy evidence supporting the alternatives is accumulated (Laming, 1968; Ratcliff, 1978; Stone, 1960; Vickers, 1970). Within the last decade, data from neurobiological experiments have shed further light on the neural bases of such choice. For example, it has been reported that while a monkey decides which of two stimuli is presented, certain neuronal populations gradually increase their firing rate, thereby accumulating evidence supporting the alternatives (Gold and Shadlen, 2002; Schall, 2001; Shadlen and Newsome, 2001). Recently, a series of neurocomputational models have offered an explanation of the neural mechanism

underlying both, psychological measures like reaction times and neurophysiological data of choice. One such model, is the Leaky Competing Accumulator (LCA; Usher and McClelland, 2001), which is sufficiently simple to allow a detailed mathematical analysis. Furthermore, as we will discuss, this model can, for certain values of its parameters, approximate the same computations carried out by a series of mathematical models of choice (Busemeyer and Townsend, 1993; Ratcliff, 1978; Shadlen and Newsome, 2001; Vickers, 1970; Wang, 2002).

Since its original publication, the LCA model (Usher and McClelland, 2001) has been analysed mathematically and extended in a number of directions (Bogacz *et al.*, 2006; Brown *et al.*, 2005; Brown and Holmes, 2001; McMillen and Holmes, 2006). In particular, which values of parameters achieve an optimal performance have been investigated. This matter is important, because if ‘we expect natural selection to produce rational behaviour’, as discussed by Houston *et al.* (2006), then the values of parameters revealed by these analyses should be found in the neural networks mediating choice processes. In this chapter, we will use the word ‘optimal’ to describe the theoretically best possible performance. In some cases, decision networks cannot achieve the optimal performance, e.g., due to some biological constraints, however, it is still of interest to investigate which parameters give best possible performance within the constraints considered – we use the word ‘optimised’ to refer to such performance.

It has been shown that for choices between two alternatives, the LCA model achieves optimal performance for particular values of parameters when its processing is linear (Bogacz *et al.*, 2006) or remains in a linear range (Brown *et al.*, 2005) (the precise meaning of these conditions will be reviewed later). However, it is known that information processing in biological neurons is nonlinear and two questions remain open: (1) is linear processing also optimal for choice between multiple alternatives, and (2) what are the parameters of the nonlinear LCA model that optimise its performance?

This chapter has two aims. First, it reviews the biological mechanisms assumed in the LCA model, and reviews an analysis of the dynamics and performance of the linear and nonlinear LCA models (Section 6.2). Second, it presents new developed extensions connected with the introduction of nonlinearities. In Section 6.3 we show that nonlinearities (of the type present in biological neurons) may improve performance in choice between multiple alternatives. In Section 6.4 we discuss how to optimise the performance of the nonlinear LCA model for two alternatives. Finally, in Section 6.5 we show how nonlinearities in the LCA model also explain counterintuitive results from choice experiments involving multiple goals or stimulus dimensions.

6.2 Review of the LCA model

In this section we briefly review the experimental data on neurophysiology of choice and models proposed to describe them, focusing on the LCA model. We examine the linear version and nonlinear versions of this model and we analyse its dynamics and performance.

6.2.1 Neurophysiology of choice

The neurophysiology of choice processes has been the subject of a number of recent reviews (Schall, 2001; Sugrue *et al.*, 2005). We start by describing a typical task used to study perceptual choice, which makes use of three important processes: representation of noisy evidence, integration of evidence, and meeting a decision criterion.

In a typical experiment used to study neural bases of perceptual choice, animals are presented with a cloud of moving dots on a computer screen (Britten *et al.*, 1993). In each trial, a proportion of the dots are moving coherently in one direction, while the remaining dots are moving randomly. The animal has to indicate the direction of prevalent dot movement by making a saccade in the corresponding direction. There are two versions of this task. The first one is the *free-response* paradigm, in which participants are allowed to respond at any moment of time. The second paradigm is the *interrogation* (or response–signal) paradigm, in which participants are required to continuously observe the stimulus until a particular signal (whose delay is controlled) is provided and which prompts an immediate response.

During the choice process, sensory areas (e.g., motion area MT) provide noisy evidence supporting the alternatives, which is represented in the firing rates of motion-sensitive neurons tuned to specific directions (Britten *et al.*, 1993; Schall, 2001). Let us denote the mean activity of the population providing evidence supporting alternative i by I_i . The perceptual choice problem may be formulated simply as finding which I_i is the highest. However, this question is not trivial, as the activity levels of these input neurons are noisy (Britten *et al.*, 1993), and hence answering this question requires sampling the inputs for a certain period.

It has been observed that in this task neurons in certain cortical regions including the lateral intraparietal area (LIP) and the frontal eye field (FEF) gradually increase their firing rates (Schall, 2001; Shadlen and Newsome, 2001). Furthermore, because the easier the task, the faster is the rate of this increase (Shadlen and Newsome, 2001), it has been suggested that these neurons integrate the evidence from sensory neurons over time (Schall, 2001; Shadlen and Newsome, 2001). This integration averages out the noise present in sensory neurons allowing the accuracy of the choice to increase with time. Moreover, because (in the free-response paradigm) the firing rate, just before the saccade, does not differ between difficulty levels of the task (Roitman and Shadlen, 2002), it is believed that the choice is made when the activity of the neuronal population representing one of the alternatives reaches a decision threshold.

6.2.2 Biologically inspired models of perceptual choice

A number of computational models have been proposed to describe the choice process described above, and their architectures are shown in Figure 6.1 for the case of two alternatives (Mazurek *et al.*, 2003; Usher and McClelland, 2001; Wang, 2002). All of these models include two units (bottom circles in Figure 6.1) corresponding to neuronal populations providing noisy evidence, and two accumulator units (denoted by y_1 and y_2

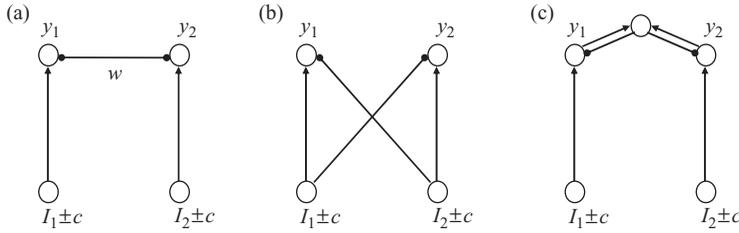


Figure 6.1 Architectures of the models of choice. Arrows denote excitatory connections, lines with filled circles denote inhibitory connections. (a) LCA model (Usher and McClelland, 2001) (b) Mazurek *et al.* (2003) model, (c) Wang (2002) model.

in Figure 6.1) integrating the evidence. The models differ in the way inhibition affects the integration process: in the LCA model (Figure 6.1a) the accumulators inhibit each other, in the Mazurek *et al.* (2003) model (Figure 6.1b) the accumulators receive inhibition from the other inputs, and in the Wang (2002) model (Figure 6.1c) the accumulators inhibit each other via a population of inhibitory inter-neurons. It has been shown that for certain values of their parameters, these models become computationally equivalent, as they all implement the same optimal algorithm for decision between two alternatives (Bogacz *et al.*, 2006). In this chapter, we thus focus on the LCA model, and we review its optimality (analogous analysis for the other two models is described in Bogacz *et al.*, 2006).

6.2.3 Linear LCA model

Figure 6.1a shows the architecture of the LCA model for the two alternative choice tasks (Usher and McClelland, 2001). The accumulator units are modelled as leaky integrators with activity levels denoted by y_1 and y_2 . Each accumulator unit integrates evidence from an input unit with mean activity I_i and independent white noise fluctuations dW_i of amplitude c_i (dW_i denote independent Wiener processes). These units also inhibit each other by way of a connection of weight w . Hence, during the choice process, information is accumulated according to (Usher and McClelland, 2001):

$$\begin{cases} dy_1 = (-ky_1 - wy_2 + I_1) dt + c_1 dW_1 \\ dy_2 = (-ky_2 - wy_1 + I_2) dt + c_2 dW_2 \end{cases}, \quad y_1(0) = y_2(0) = 0. \quad (6.1)$$

In the equations above, the term k denotes the decay rate of the accumulators' activity (i.e., the leak) and $-wy_i$ denotes the mutual inhibition. For simplicity, it is assumed that integration starts from $y_1(0) = y_2(0) = 0$ (cf. Bogacz *et al.*, 2006).

The LCA model can be used to describe the two paradigms described in Subsection 6.2.1. In the free-response paradigm, the model is assumed to make a response as soon as either accumulator exceeds a preassigned threshold, Z . The interrogation paradigm is modelled by assuming that at the interrogation time the choice is made in favour of the alternative with higher y_i at the moment when the choice is requested.

Because the goal of the choice process is to select the alternative with highest mean input I_i , in the following analyses and simulations we always set $I_1 > I_2$. Hence a

simulated choice is considered to be correct if the first alternative is chosen; this will happen in the majority of simulated trials. However, in some trials, due to noise, another alternative may be chosen; such trials correspond to incorrect responses. By simulating the model multiple times expected error rate (ER) may be estimated. In addition, in the free-response paradigm, the average decision time (DT) from choice onset to reaching the threshold can be computed.

The LCA model can be naturally extended to N alternatives. In this case, the dynamics of each accumulator i is described by the following equation (Usher and McClelland, 2001):

$$dy_i = \left(-ky_i - w \sum_{\substack{j=1 \\ j \neq i}}^N y_j + I_i \right) dt + c_i dW_i, \quad y_i(0) = 0. \quad (6.2)$$

When the decay and inhibition parameters are equal to zero, the terms in Equations (6.1) and (6.2) describing leak and competition disappear, and the linear LCA model reduces to another model known in psychological literature as the *race* model (Vickers, 1970; 1979), in which accumulators integrate noisy evidence independent of one another.

6.2.4 Dynamics of the model

The review of dynamics of the linear LCA model in this subsection is based on Bogacz *et al.* (2006). In the case of two alternatives, the state of the 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 plane whose horizontal and vertical axes correspond to y_1 and y_2 ; the evolution of activities of the accumulator units during the choice process may be visualised as a path in this plane. Representative paths for three different parameter ranges in this plane are shown in Figure 6.2. In each case the choice process starts from $y_1 = 0$ and $y_2 = 0$, i.e., from the bottom left corner of each panel. Initially the activities of both accumulators increase due to stimulus onset, which is represented by a path going in an upper-right direction. But as the accumulators become more active, mutual inhibition causes the activity of the ‘weaker’ accumulator to decrease and the path moves toward the threshold for the more strongly activated accumulator (i.e., the correct choice).

To better understand the dynamics of the model, Figure 6.2 shows its *vector fields*. 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. Note that in all three panels of Figure 6.2 there is a line, indicated by a thick grey line, to which all states are attracted: the arrows point towards this line from both sides. The location along this line represents an important variable: the difference in activity between the two accumulators. As most of the choice-determining dynamics occur along this line, it is helpful to make use of new coordinates rotated clockwise by 45° with respect to the y_1 and y_2 coordinates. These new coordinates are shown in Figure 6.2b: x_1 is parallel to the attracting line and describes the difference

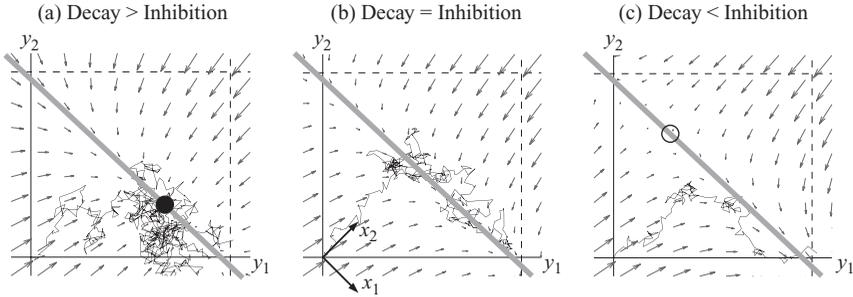


Figure 6.2 Examples of the evolution of the LCA model, showing paths in the state space of the model. The horizontal axes denote the activation of the first accumulator; the vertical axes denote the activation of the second accumulator. The paths show the choice process from stimulus onset (where $y_1 = y_2 = 0$) to reaching a threshold (thresholds are shown by dashed lines). The model was simulated for the following parameters: $I_1 = 4.41$, $I_2 = 3$, $c = 0.33$, $Z = 0.4$. The sum of inhibition (w) and decay (k) is kept constant in all panels, by setting $k + w = 20$, but the parameters themselves have different values in different panels: (a) $w = 7$, $k = 13$; (b) $w = 10$, $k = 10$; (c) $w = 13$, $k = 7$. The simulations were performed using the Euler method with timestep $\Delta t = 0.01$. 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 the normal distribution with mean 0 and variance $c^2 \Delta t$. The arrows show the average direction of movement of LCA model in the state space. The thick grey lines symbolise the attracting lines. The filled circle in panel (a) indicates the attractor. The open circle in panel (c) indicates the unstable fixed point.

between activities of the two accumulators; while x_2 describes the sum of their activities. The transformation from y to x coordinates is given by (cf. Seung, 2003):

$$\begin{cases} x_1 = \frac{y_1 - y_2}{\sqrt{2}}, \\ x_2 = \frac{y_1 + y_2}{\sqrt{2}}. \end{cases} \quad (6.3)$$

In these new coordinates Equations (6.1) become (Bogacz *et al.*, 2006):

$$dx_1 = \left((w - k)x_1 + \frac{I_1 - I_2}{\sqrt{2}} \right) dt + \frac{c_1}{\sqrt{2}} dW_1 - \frac{c_2}{\sqrt{2}} dW_2, \quad (6.4)$$

$$dx_2 = \left((-k - w)x_2 + \frac{I_1 + I_2}{\sqrt{2}} \right) dt + \frac{c_1}{\sqrt{2}} dW_1 + \frac{c_2}{\sqrt{2}} dW_2. \quad (6.5)$$

Equations (6.4) and (6.5) are *uncoupled*; that is, the rate of change of each x_i depends only on x_i itself (this was not the case for y_1 and y_2 in Equations (6.1)). Hence, the evolution of x_1 and x_2 may be analysed separately.

We first consider the dynamics in the x_2 direction, corresponding to the summed activity of the two accumulators, which has the faster dynamics. As noted above, on all panels of Figure 6.2 there is a line to whose proximity the state is attracted, implying that x_2 initially increases and then fluctuates around the value corresponding to the position of the attracting line. The magnitude of these fluctuations depends on the inhibition and decay parameters; the larger the sum of inhibition and decay, the smaller the fluctuation (i.e., the closer the system stays to the attracting line).

Figure 6.2 also shows that the dynamics of the system in the direction of coordinate x_1 . These dynamics are slower than the x_2 dynamics and correspond to a motion along the line. Their characteristics depend on the relative values of inhibitory weight w and decay k . When decay is larger than inhibition, there are also attractor dynamics for the x_1 dynamics, as shown in Figure 6.2a. The system is attracted towards this point and fluctuates in its vicinity. In Figure 6.2a the threshold is reached when noise pushes the system away from the attractor. When inhibition is larger than decay, x_1 -dynamics are characterised by repulsion from the fixed point, as shown in Figure 6.2c.

When inhibition equals decay, the term $(w - k)x_1$ in Equation (6.4) disappears, and Equation (6.4) describing the evolution along the attracting line can be written as:

$$dx_1 = \left(\frac{I_1}{\sqrt{2}}dt + \frac{c_1}{\sqrt{2}}dW_1 \right) - \left(\frac{I_2}{\sqrt{2}}dt + \frac{c_2}{\sqrt{2}}dW_2 \right). \quad (6.6)$$

In the remainder of this chapter we refer to the linear LCA model with inhibition equal to decay as *balanced*. The vector field for this case is shown in Figure 6.2b. In this case, according to Equation (6.6) the value of x_1 changes according to the difference in evidence in support of two alternatives, hence the value of x_1 is equal to the *accumulated* difference in evidence in support of two alternatives.

The three cases illustrated in Figure 6.2 make different predictions about the impact of temporal information on choice in the interrogation paradigm. If inhibition is larger than decay (Figure 6.2c), and the repulsion is high, the state is likely to remain on the same side of the fixed point. This causes a *primacy effect* (Busemeyer and Townsend, 1993; Usher and McClelland, 2001): the inputs at the beginning of the trial determine to which side of the fixed point the state of the network moves, and then due to repulsion, late inputs before the interrogation time have little effect on choice made. Analogously, decay larger than inhibition produces a *recency effect*: the inputs later in the trial have more influence on the choice than inputs at the beginning, whose impact has decayed (Busemeyer and Townsend, 1993; Usher and McClelland, 2001). If the decay is equal to inhibition, inputs during the whole trial (from the stimulus onset to the interrogation signal) influence the choice equally, resulting in a balanced choice (with maximal detection accuracy; see below). Usher and McClelland (2001) tested whether the effects described above are present in human decision makers by manipulating the time flow of input favouring two alternatives, and reported significant individual differences: some participants showed primacy, others showed recency and some were balanced and optimal in their choice.

6.2.5 Performance of linear LCA model

In this subsection we review parameters of the model (w, k) that result in an optimal performance of the linear LCA model in the free-response paradigm for given parameters of the inputs (I_i, c_i). We start with the two alternatives in the free-response paradigm (Bogacz *et al.*, 2006), then we discuss multiple alternatives (see also McMillen and Holmes, 2006), and the interrogation paradigm.

When inhibition and decay are both fairly strong (as in Figure 6.2b), the state evolves very closely to the attracting line (see above) reaching the decision threshold very close to the intersection of the decision threshold and attracting line (see Figure 6.2b). Thus

in this case, the LCA model exceeds one of the decision thresholds approximately when the variable x_1 exceeds a positive value (corresponding to y_1 exceeding Z) or decreases below a certain negative value (corresponding to y_2 exceeding Z).

The above analysis shows that when the LCA model is balanced and both inhibition and decay are high, a choice is made approximately when x_1 representing the accumulated difference between the evidence supporting the two alternatives exceeds a positive or negative threshold. This is the characteristic of a mathematical choice model known as the diffusion model (Laming, 1968; Ratcliff, 1978; Stone, 1960), which implements the optimal statistical test for choice in the free-response paradigm: the Sequential Probability Ratio Test (SPRT) (Barnard, 1946; Wald, 1947). The SPRT is optimal in the following sense: among all possible procedures for solving this choice problem given a certain ER, it minimises the average DT.

In summary, when the linear LCA model of choice between two alternatives is balanced and both inhibition and decay are high, the model approximates the optimal SPRT and makes the fastest decisions for fixed ERs (Bogacz *et al.*, 2006).

In the case of multiple alternatives the performance of the linear LCA model is also optimised when inhibition is equal to decay and both have high values (McMillen and Holmes, 2006). However, in contrast to the case of two alternatives, the LCA model with the above parameters does not achieve as good performance as the statistically (asymptotically) optimal tests: the Multiple SPRT (MSPRT) (Dragalin *et al.*, 1999). The MSPRT tests require much more complex neuronal implementation than the LCA model (McMillen and Holmes, 2006). For example, one of the MSPRT tests may be implemented by the ‘max vs. next’ procedure (McMillen and Holmes, 2006), in which the following quantities are calculated for each alternative at each moment of time: $L_i = y_i - \max_{j \neq i} y_j$, where y_i is the evidence supporting alternative i accumulated according to the race model. The choice is made whenever any of the L_i exceeds a threshold.

Although the linear and balanced LCA with high inhibition and decay achieves shorter DT for fixed ER than the linear LCA model with other values of parameters (e.g., inhibition different from decay, or both equal to zero), it is slower than MSPRT (McMillen and Holmes, 2006). Furthermore, as the number of alternatives for N increases, the best achievable DT for a fixed ER of the linear balanced LCA model approaches that of the race model (McMillen and Holmes, 2006).

In the interrogation paradigm, the LCA model achieves optimal performance when it is balanced both for two alternatives (it then implements the Neyman–Pearson test (Bogacz *et al.*, 2006; Neyman and Pearson, 1933)) and for multiple alternatives (McMillen and Holmes, 2006). However, by contrast to the free-response paradigm, in the interrogation paradigm, the high value of decay and inhibition is not necessary for optimal performance and the balanced LCA model (even with high inhibition and decay) achieves the same performance as the race model.

Table 6.1 summarises conditions necessary for the linear LCA model to implement the optimal algorithm for a given type of choice problem. Note that the linear LCA model can implement the algorithms achieving best possible performance for all cases except of choice between multiple alternatives in the free-response paradigm. Hence this is the only case in which there exists room for improvement of the LCA model – this case is addressed in Section 6.3.

Table 6.1 Summary of conditions the linear LCA model must satisfy to implement the optimal choice algorithms.

Paradigm	# of alternatives	
	$N = 2$	$N > 2$
Free response	Inhibition = Decay and both high	<i>Optimality not attainable</i>
Interrogation (response-signal)	Inhibition = Decay	Inhibition = Decay

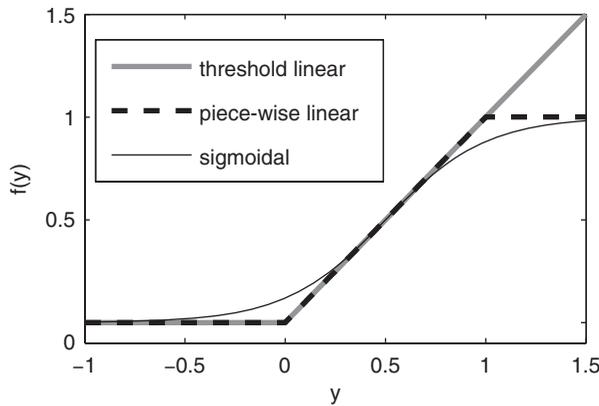


Figure 6.3 Nonlinear input–output functions used in the LCA model. Threshold linear: $f(y) = y$ for $y \geq 0$ and $f(y) = 0$ for $y < 0$ (Usher and McClelland, 2001). Piece-wise linear: $f(y) = 0$ for $y < 0$, $f(y) = 1$ for $y > 1$, and $f(y) = y$ otherwise (Brown *et al.*, 2005). Sigmoidal: $f(y) = 1 / (1 + e^{-4(y-0.5)})$ (Brown *et al.*, 2005; Brown and Holmes, 2001).

6.2.6 Nonlinear LCA model

In the linear version of the LCA model described so far, during the course of the choice process, the activity levels of accumulators can achieve arbitrarily large or small (including negative) values. However, the firing rate of biological neurons cannot be negative and cannot exceed a certain level (due to the refractory period of biological neurons). A number of ways of capturing these limits in the LCA model has been proposed, starting with the original version (Usher and McClelland, 2001), where the values of y_1 and y_2 are transformed through a nonlinear activation function $f(y)$ before they influence (inhibit) each other:

$$dy_i = \left(-ky_i - w \sum_{\substack{j=1 \\ j \neq i}}^N f(y_j) + I_i \right) dt + c_i dW_i, \quad y_i(0) = 0. \quad (6.7)$$

Figure 6.3 shows three functions $f(y)$ proposed in the literature: threshold linear (Usher and McClelland, 2001), piece-wise linear (Brown *et al.*, 2005), and sigmoidal (Brown *et al.*, 2005; Brown and Holmes, 2001). The threshold linear function corresponds to the

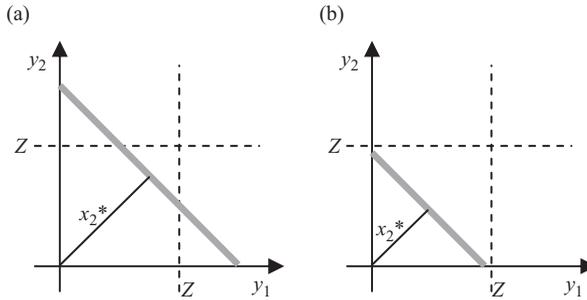


Figure 6.4 State plane analysis of the LCA model. Thick grey lines symbolise attracting lines in the y_1y_2 plane. The position of the attracting line is shown for parameters used in simulations in Figure 6.5a and b, respectively. Thus the distance x_2^* of the attracting line from the origin is equal to 0.26 and 0.12, respectively (from Equation (6.8)). The dashed lines indicate the thresholds. The values of the threshold are shown that produce $ER = 10\%$ in simulations of the unbounded (linear) LCA model for $N = 2$ alternatives in Figure 6.5a and b respectively, i.e., 0.25 and 0.17.

constraint that actual neural activity is bounded (by zero) at its low end. The piece-wise linear and sigmoidal functions bound the activity levels of accumulators at both ends (the maximum level of activity being equal to 1). In the free-response paradigm, the threshold of the model with piece-wise linear activation function (Brown *et al.*, 2005) must be lower than 1 (as otherwise a choice would never be made). Hence, in the free-response paradigm the nonlinear model with piece-wise linear activation function is equivalent to the model with the threshold linear function (Usher and McClelland, 2001) (the upper boundary cannot be reached); these models only differ in the interrogation paradigm.

One way to simplify the analysis is to use linear Equation (6.2) (rather than (6.7)) and add reflecting boundaries on y_j at 0, preventing any of y_j from being negative (Usher and McClelland, 2001), and we refer to such model as *bounded*. In every step of the simulation of the bounded LCA model, the activity level of an accumulator y_j is being reset to 0 if a negative value is obtained. The bounded model behaves very similarly to the nonlinear models with threshold linear, piece-wise linear and even sigmoidal activation functions and provides a good approximation for them (see Appendix A of Usher and McClelland, 2001, for detailed comparison between the bounded and nonlinear LCA models).

6.2.7 Performance of bounded LCA model

For two alternatives, the bounded model implements the optimal choice algorithm, as long as decay is equal to inhibition and both are large (see Subsection 6.2.5) and the model remains in the linear range (i.e., the levels of accumulators never decrease to zero; cf. Brown *et al.*, 2005). Since during the choice process the state of the model moves rapidly towards the attracting line, the levels of y_j are likely to remain positive, if the attracting line crosses the decision thresholds before the axes as shown in Figure 6.4a

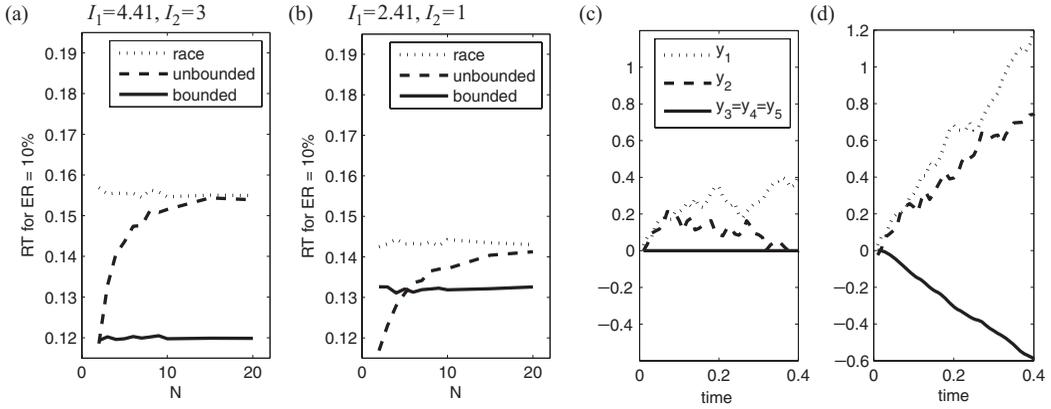


Figure 6.5 Performance and dynamics of choice models with only two accumulators receiving inputs. All models were simulated using the Euler method with $\Delta t = 0.01$ s. (a), (b) Decision time for a threshold resulting in an error rate (ER) of 10% of different choice models as a function of the number of alternatives N (shown on x-axis). Three models are shown: the race model, the unbounded (i.e., linear) LCA model, and the bounded LCA model (see key). The parameters of the LCA model are equal to $w = k = 10$. The parameters of the first two inputs were chosen such that $c_1 = c_2 = 0.33$, $I_1 - I_2 = 1.41$ (values estimated from data of a sample participant of Experiment 1 in the study of Bogacz *et al.*, 2006), while the other inputs were equal to 0, $I_3 = \dots = I_N = 0$, $c_3 = \dots = c_N = 0$. The panels differ in the total mean input to the first two accumulators: in panel (a) $I_2 = 3$, while in panel (b) $I_2 = 1$. For each set of parameters, a threshold was found numerically that resulted in $ER = 10 \pm 0.2\%$ (s.e.); this search for the threshold was repeated 20 times. For each of these 20 thresholds, the decision time was then found by simulation and their average used to construct the data points. Standard error of the mean was lower than 2 ms for all data points hence the error bars are not shown. (c), (d) Examples of the evolution of the bounded (c) and the unbounded (d) LCA model, showing y_i as functions of time. The models were simulated for the same parameters as in panel (a), and for $N = 5$ alternatives. Panels (c) and (d) were simulated for the same initial seed of the random number generator hence in both cases the networks received exactly the same inputs.

(but not in Figure 6.4b). The distance of the attracting line from the origin of the plane is equal to (Bogacz *et al.*, 2006):

$$x_2^* = \frac{I_1 + I_2}{\sqrt{2}(k + w)}. \quad (6.8)$$

According to Equation (6.8), the larger the sum of mean inputs $I_1 + I_2$, the further the attracting line is from the origin. Figure 6.5 compares the performance of the bounded LCA model and the linear LCA model without boundaries, which we refer to as *unbounded*. Figure 6.4a shows the position of the attracting line relative to thresholds for the parameters used in the simulations of the unbounded LCA model for $N = 2$ alternatives in Figure 6.5. For $N = 2$, adding the reflecting boundaries at $y_i = 0$ does not affect the performance of the model (the left end of the solid line coincides with the left end of the dashed line). This can be expected since for the parameters used in simulations, the attracting line crosses the threshold before the axes, as shown in Figure 6.4a.

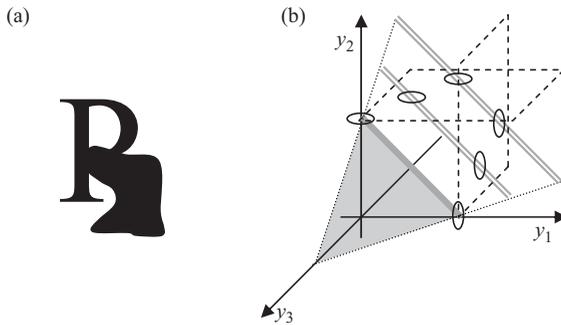


Figure 6.6 (a) Example of a stimulus providing strong evidence in favour of two letters (P and R) and very weak evidence in favour of any other letter. (b) State space analysis of the LCA model for three alternatives. The grey triangle indicates the attracting plane, and dotted lines indicate the intersection of the attracting plane with the y_1y_3 plane and the y_2y_3 plane. Thick grey line symbolises the attracting line in the y_1y_2 plane. The double grey lines show sample positions of the attracting line in the y_1y_2 plane for two negative values of y_3 . The two planes surrounded by dashed lines indicate positions of the decision thresholds for alternatives 1 and 2. The ellipses indicate the intersections of the attracting lines in the y_1y_2 plane with the decision thresholds.

Figure 6.4b shows the position of the attracting line for the parameters used in simulations of the unbounded LCA model for $N = 2$ alternatives in Figure 6.5b. For $N = 2$, adding the reflecting boundaries at $y_i = 0$ degrades the performance of the model (the left end of the solid line lies above the left end of the dashed line). This happens because the attracting line reaches the axes before crossing the threshold, as shown in Figure 6.4b and hence the state is likely to hit the boundaries before reaching the threshold.

McMillen and Holmes (2006) tested the performance of the bounded LCA model for multiple alternatives, for the following parameters: $I_1 = 2$, $I_2 = \dots = I_N = 0$, $c_1 = \dots = c_N = 1$ (all accumulators received noise of equal standard deviation), $w = k = 1$, and N varying from 2 to 16. They found that the DT of bounded LCA for ER = 10% was slower than that of the unbounded LCA model. However, it will be shown here that this is not the case for more biologically realistic types of inputs.

6.3 The advantage of nonlinearity in multiple choice

Most real-life decisions involve the need to select between multiple alternatives, on the basis of partial evidence that supports a small subset of them. One ubiquitous example could correspond to a letter (or word) classification task, based on occluded (or partial) information. This is illustrated in Figure 6.6a for a visual stimulus that provides strong evidence in favour of P/R and very weak evidence in favour of any other letter (a simple analogue for the case of word-classification would consist of a word-stem consistent with few word completions). Note the need to select among multiple alternatives, based on input that supports only a few of them.

We compare the performance of the bounded and unbounded LCA models in the tasks of type described above within the free-response paradigm: we will discuss two cases (with regards to the type of evidence and noise parameters), which may arise in such situations. We start with a simplified case, which is helpful for the purpose of mathematical analysis, followed by a more realistic situation.

6.3.1 Case 1: only two accumulators receive input and noise

Consider a model of N accumulators y_i (corresponding to N alternatives), two of which receive input (supporting evidence; with means I_1, I_2 and standard deviation c), while other accumulators do not, so that $I_3 = \dots = I_N = c_3 = \dots = c_N = 0$. Let us examine first the dynamics of the bounded LCA model (with $y_1, y_2 \geq 0$). In this case, the other accumulators y_3, \dots, y_N do not receive any input but only inhibition from y_1, y_2 and hence they remain equal to 0 (i.e., $y_i = 0$ for all $i > 2$; see Figure 6.5c). Therefore, the choice process simplifies to a model of two alternatives, as described by Equation (6.1). Hence, when the boundaries are present, the performance of the model does not depend on the total number of alternatives N . This is illustrated in Figure 6.5a and b for sample parameters of the model. Note that DTs for fixed ER in each panel (shown by solid lines) do not differ significantly between different values of N .

Figure 6.5c and d compare the evolution of bounded and unbounded LCA models for $N = 5$ alternatives. Figure 6.5c shows the evolution of the bounded LCA model in which accumulators y_1, y_2 evolve in the way typical for the LCA model for two alternatives (compare with Figure 6.2b): the competition between accumulators y_1, y_2 is resolved and as y_1 increases, y_2 decreases towards 0. Figure 6.5d shows that during the evolution of the unbounded model, the accumulators y_3, \dots, y_N become more and more negative. Hence the inhibition received by y_1, y_2 from y_3, \dots, y_N is actually positive, and increases the value of *both* y_1, y_2 . Therefore, in Figure 6.5d (by contrast to Figure 6.5c) the activation of the ‘losing’ accumulator, y_2 , also increases.

To better illustrate the difference between the bounded and unbounded choice behaviour, consider the dynamics of the unbounded model (Equation (6.1)) for $N = 3$ alternatives. In such a case, the state is attracted to a plane (indicated in Figure 6.6b; McMillen and Holmes, 2006). However, since only alternatives 1 and 2 can be chosen, it is still useful to examine the dynamics in the $y_1 y_2$ plane. In the $y_1 y_2$ plane the state of the model is attracted to a line, and the position of this line is determined by the value of y_3 . For example, if $y_3 = 0$, then the attracting line in the $y_1 y_2$ plane is the intersection of the attracting plane and the $y_1 y_2$ plane, i.e., the thick grey line in Figure 6.6b. For other values of y_3 , the attracting line in the $y_1 y_2$ plane is the intersection of the attracting plane and the plane parallel to the $y_1 y_2$ plane intersecting y_3 axis in the current value of y_3 . For example, the double grey lines in Figure 6.6b show the attracting lines in the $y_1 y_2$ plane for two negative values of y_3 .

During the choice process of unbounded LCA of Equation (6.1), accumulator y_3 becomes more and more negative (as it receives more and more inhibition from y_1 and y_2), as illustrated in Figure 6.5d. Hence the attracting line in the $y_1 y_2$ plane moves further and further away from the origin of the $y_1 y_2$ plane. For example, the thick grey line in

Figure 6.6b shows the position of the attracting line in the y_1y_2 plane at the beginning of the choice process and the double grey lines show the positions at two later time points. Therefore, the choice involves two processes: evolution along the attracting line (the optimal process) and evolution of this line's position (which depends on the total input integrated so far). Due to the presence of the second process the performance of the unbounded LCA model for $N = 3$ departs from that for $N = 2$, which is visible in Figure 6.5a and b. Also note in Figure 6.6b that as y_3 becomes more and more negative, the relative positions of the decision thresholds and the attracting line change, and the part of the attracting line between the thresholds becomes shorter and shorter. Hence relative to the attractive line, the thresholds move during the choice process. This situation is in contrast to the case of the bounded LCA model, in which y_3 is constant (as stated above), and hence the position of the attracting line in the y_1y_2 plane (and thus its relation to the thresholds) does not change.

In summary, in the case of choice between multiple alternatives with only two alternatives receiving supporting evidence, the boundaries allow the LCA model to achieve the performance of the LCA model for two alternatives (close to the optimal performance). The performance of the unbounded LCA model is lower – approaching that of the race model as the number of alternatives increases.

6.3.2 Case 2: biologically realistic input parameters for choice with continuous variables

We assumed above that only two integrators receive input while the others received none: $I_3 = \dots = I_N = 0$. However, in many situations, it might be expected that there is a more graded similarity among the different inputs, with the strength of the input falling off as a continuous function of similarity. This would be the case, for example, in tasks where the stimuli were arranged along a continuum, as they might be in a wavelength or length discrimination task. Here we consider the case of stimuli arranged at N equally spaced positions around a ring, an organisation that is relevant to many tasks used in psychophysical and physiological experiments, where the ring may be defined in terms of positions, orientations, or directions of motion. We use the motion case since it is well studied in the perceptual decision-making literature but the analysis applies equally to other such cases as well, and may be instructive for the larger class of cases in which stimuli are positioned at various points in a space.

Considering the motion discrimination case, motion-sensitive neurons in area MT are thought to provide evidence of the direction of stimulus motion. Neurons providing evidence for alternative i respond with a mean firing rate that is a function of the angular distance d_i between the direction of coherent motion in the stimulus and their preferred direction. This function is called a tuning curve, and can be well approximated by a Gaussian (Snowden *et al.*, 1992):

$$I_i = r_{\min} + (r_{\max} - r_{\min}) \exp\left(-\frac{d_i^2}{2\sigma^2}\right) \quad (6.9)$$

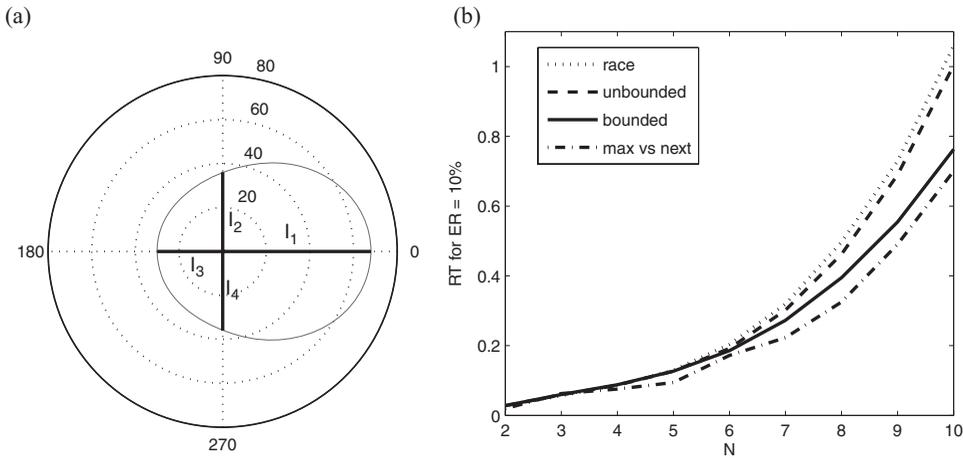


Figure 6.7 Simulation of the motion discrimination task. (a) Tuning curve describing the simulated firing rate of MT neurons (network inputs) as a function of angular difference d_i between the direction of coherent motion in the stimulus and neurons' preferred direction. We used the following parameters of the tuning curve $\sigma = 46.5^\circ$ (the average value over tuning curves fitted to 30 MT neurons by Snowden *et al.*, 1992), $r_{min} = 30$ Hz, $r_{max} = 68$ Hz (values from a neuron analysed in Britten *et al.*, 1993). Thick lines show sample values of I_i in case of $N = 4$, which were computed in the following way. Since we assume that the first alternative is correct and alternatives are equally distributed around 360° , we computed $d_i = 360^\circ(i - 1)/N$, then if $d_i > 180^\circ$, we made $d_i = d_i - 360^\circ$, and then we computed I_i from Equation (6.9). (b) Decision time with a threshold resulting in error rate of 10% of different models as a function of the number of alternatives N (shown on x-axis). Four models are shown: the race model, the unbounded LCA model, the bounded LCA model, and max vs. next (see key). Methods of simulations as in Figure 6.5. The parameters of LCA model are equal to $w = k = 4$; this value was chosen as it optimised performance of the bounded LCA model for the inputs described in panel (a).

where r_{min} and r_{max} denote the minimum and the maximum firing rate of the neuron, and σ describes the width of the tuning curve. In our simulation we use the parameter values that generate the tuning curve function shown in Figure 6.7a.

Furthermore, we simulated the input to the accumulators as spikes (rather than values chosen from a Gaussian distribution). In particular, we assumed that the input to accumulator i comes from the Poisson process with mean I_i , because the Poisson process captures many aspects of firing of cortical neurons (Shadlen and Newsome, 1998). Thus the input to accumulator i within a very short interval dt is chosen stochastically such that it is equal to 1 with probability $I_i dt$ (that corresponds to a spike being produced by sensory population i), and 0 with probability $1 - I_i dt$.

Figure 6.7b shows the DTs under the assumptions described above. The DT grows rapidly as N increases, because as N grows, the difference between the largest input (I_1) and the next two largest inputs (I_2 and I_N) decreases. Importantly, in the simulation, introduction of boundaries to the LCA model reduce DT (for a fixed ER of 10%) very significantly, as N increases. For example, for $N = 10$ the boundaries reduce the DT by

about 25%. Figure 6.7b also shows that the performance of bounded LCA model is close to that of the max vs. next procedure (that implements asymptotically optimal test; see Subsection 6.2.5).

In summary, this simulation shows that the introduction of the biologically realistic assumption that the firing rate of accumulator neurons cannot be negative, may not only improve the performance of choice networks for biologically realistic parameters of inputs, but it also allows the LCA model to approximate the optimal performance.

6.4 Optimisation of performance of bounded LCA model in the interrogation paradigm

It is typically assumed that in the interrogation paradigm the decision threshold is no longer used to render a choice. Instead, the alternative with the highest activity level is chosen when the interrogation signal appears (Usher and McClelland, 2001). However, a more complex assumption regarding the process that terminates decisions in the interrogation paradigm is also possible. As suggested by Roger Ratcliff (1988), a response criterion is still in place (as in the free-response paradigm) and participants use a response criterion (like in free-response paradigm) and when the activation reaches this criterion, they make a preliminary decision (and stop integrating input). Accordingly there are two type of trials: (1) those that reach criterion (as above), and (2) those that do not reach criterion until the interrogation signal is received and where the choice is determined by the unit with highest activation. This is mathematically equivalent to the introduction of an absorbing upper boundary on the accumulator trajectories; once an accumulator hits the upper boundary, it terminates the decision process, so that the state of the model does not change from that moment until the interrogation time (Mazurek *et al.*, 2003; Ratcliff, 1988). Mazurek *et al.* (2003) point out that the dynamics of the model with absorbing upper boundaries is consistent with the observation that in the motion discrimination task under interrogation paradigm, the time courses of average responses from population of LIP neurons cease increasing after a certain period following the stimulus onset, and are maintained until the interrogation time (Roitman and Shadlen, 2002).

In Subsection 6.2.5, we showed that the unbounded LCA model achieves optimal performance when the decay is equal to inhibition. Then the following question arises: does the balance of decay and inhibition still optimise the performance of the bounded LCA model in the interrogation paradigm, when an absorbing upper boundary is assumed (to account for pre-interrogation decisions)? Figure 6.8 illustrates the ER of bounded LCA model for $N = 2$ alternatives. To make the position of the attracting line stable (cf. Equation (6.8)), we fixed parameters $w + k$ but varied $w - k$. The results illustrate that by increasing inhibition relative to decay the bounded model can achieve lower ER in the interrogation paradigm. This happens because in this case, there is an attracting point to which state of the model is attracted, as shown in Figure 6.2a, and this attraction prevents the model from hitting the absorbing boundary prematurely due to noise; thus the biasing effect of early input leading to premature choice is minimised. In summary,

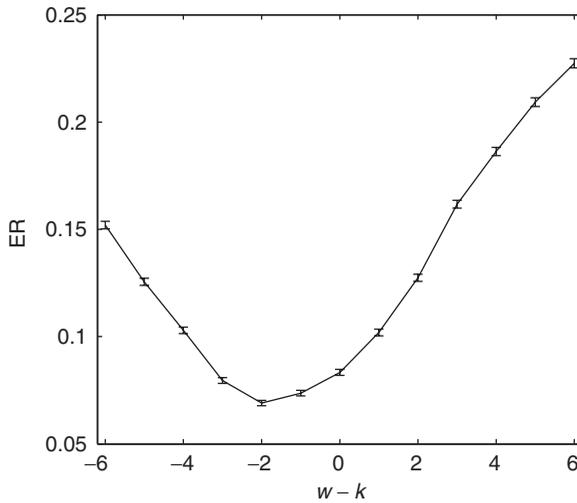


Figure 6.8 The ER of bounded LCA model in the interrogation paradigm. The models were simulated with parameters: $I_1 = 5.414$, $I_2 = 4$, $c_1 = c_2 = 0.8$, $B = 1.4$, $T = 2.5$. The sum of decay and inhibition was fixed $w + k = 6$, while their difference changed from -6 to 6 . Data is averaged from 10 000 trials.

in contrast to the unbounded model, a balance of decay and inhibition did not optimise ER in the interrogation paradigm for the bounded model. Instead, optimal performance within the tested range was achieved when inhibition was smaller than decay.

6.5 Value-based decisions

The LCA model and its extensions discussed so far are targeting an important, but special type of choice; the type deployed in perceptual classification judgements. A different type of choice, of no less importance to humans and animals, is deciding between alternatives on the basis of their match to a set of internal motivations. Typically, this comes under the label of decision making. While human decision making is a mature field, where much data and theories have been accumulated (Kahneman and Tversky, 2000), more recently neurophysiological studies of value-based decisions have also been conducted on behaving animals (Platt and Glimcher, 1999; Sugrue *et al.*, 2004).

Although both the perceptual and the value/motivational decisions involve a common selection mechanism, the basis on which this selection operates differs. The aim of this section is to discuss the underlying principles of value-based decisions and to suggest ways in which a simple LCA type of mechanism can be used to explain the underlying cognitive processes. We start with a brief review of these principles and of some puzzling challenges they raise for an optimal theory of choice, before we explore a computational model that addresses the underlying processes.

6.5.1 Value and motivation-based choice

Unlike in perceptual choice, the decisions we consider here cannot be settled on the basis of perceptual information alone. Rather, each alternative (typically an action, such as purchasing a laptop from a set of alternatives) needs to be evaluated in relation to its potential consequences and its match to internal motivations. Often, this is a complex process, where the preferences for the various alternatives are being constructed as part of the decision process itself (Slovic, 1995). In some situations, where the consequences are obvious or explicitly described, the process can be simplified. Consider, for example, a choice between three laptops, which vary in their properties as described by a number of dimensions (screen size, price, etc.) or a choice between lotteries described in terms of their potential win and corresponding risks. The immediate challenge facing a choice in such situations is the need to convert between the different currencies, associated with the various dimensions. The concept of *value* is central to decision making, as a way to provide such a universal internal currency.

Assuming the existence of a value function, associated with each dimension, a simple normative rule of decision making, the *expected-additive-value*, seems to result. Accordingly, one should add the values that an alternative has on each dimension and compute expectation values when the consequences of the alternatives are probabilistic. Such a rule is then bound to generate a fixed and stable preference order for the various alternatives. Behavioural research in decision making indicates, however, that humans and animals violate expected-value prescriptions and change their preferences between a set of options depending on the way the options are described and on a set of contextual factors.

6.5.2 Violations of expected-value and preference reversals

First, consider the pattern of *risk-aversion* for gains. Humans and animals prefer the less risky of two options that are equated for expected value (Kahneman and Tversky, 2000). For example most people prefer a *sure* gain of £100 to a lottery with a probability 0.5 of winning £200 and nothing otherwise. An opposite pattern, *risk-seeking* is apparent for losses: most people prefer to play lottery with a chance 0.5 of losing £200 (and nothing otherwise) to a *sure* loss of £100.

Second, the preference between alternatives depends on a reference, which corresponds either to the present state of the decision maker, or even to an *expected* state, which is subject to manipulation. Consider, for example the following situation (Figure 6.9a). When offered a choice between two job alternatives A and B , described on two dimensions (e.g., distance from home and salary) to replace an hypothetical job that is being terminated – the *reference* (R_A or R_B , which is manipulated between groups) – participants prefer the option that is more similar to the reference (Tversky and Kahneman, 1991).

Third, it has been shown that the preference order between two options can be modified by the introduction of a third option, even when this option is not being chosen. Three such situations have been widely discussed in the decision-making literature, resulting

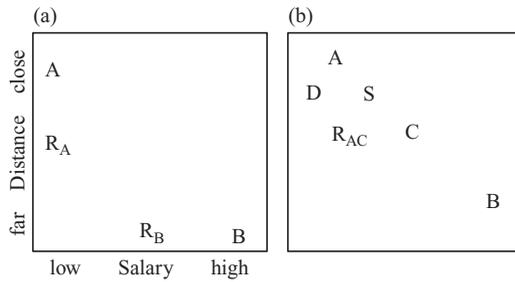


Figure 6.9 Configurations of alternatives in the attribute space. In each panel the two axes denote two attributes of the alternatives (sample attributes' labels are given in panel (a)). The capital letters denote the positions of the alternative choices in the attribute space, while letters R_i denote the reference points. (a) Reference effect in multi-attribute decision making (after Tversky and Kahneman, 1991). (b) Contextual preference reversal: similarity, attraction and the compromise effects. Alternatives A, B, C, S lie on the indifference line.

in the *similarity*, the *attraction* and the *compromise* effects. To illustrate these effects consider a set of options, A, B, C , and S , which are characterised by two attributes (or dimensions) and which are located on a decision-maker indifference curve: the person is of equal preference on a choice between any two of these options (Figure 6.9b). The similarity effect is the finding that the preference between A and B can be modified in the favour of B by the introduction of a new option, S , similar to A in the choice-set. The attraction effect corresponds to the finding that, when a new option similar to A , D , and dominated by it (D is worse than A on both dimensions) is introduced into the choice set, the choice preference is modified in favour of A (the similar option; note that while the similarity effects favours the dissimilar option, the attraction effect favours the similar one). Finally, the compromise effect corresponds to the finding that, when a new option such as B is introduced into the choice set of two options A and C , the choice is now biased in favour of the intermediate one, C , the compromise.

The traditional way in which the decision-making literature addresses such deviations from the normative (additive-expected-value) theory is via the introduction of a set of disparate heuristics, each addressing some other aspect of these deviations (LeBoef and Shafir, 2005). One notable exception is work by Tversky and colleagues (Tversky, 1972; Tversky and Simonson, 1993), who developed a mathematical, context-dependent-advantage model that accounts for reference effects and preference reversal in multidimensional choice. However, as observed by Roe *et al.* (2001), the context-dependent-advantage model cannot explain the preference reversals in similarity effect situations (interestingly, a much earlier model by Tversky (1972), the elimination by aspects, accounts for the similarity effect but not for the attraction, the compromise or other reference effect). In turn, Roe *et al.* (2001), have proposed a neurocomputational account of preference reversal in multidimensional choice, termed the Decision Field Theory (DFT; see also Busemeyer and Townsend, 1993). More recently, Usher and McClelland (2004) have proposed a neurocomputational account of the same findings, using the LCA framework extended to include some assumptions regarding nonlinearities in value functions and reference effects introduced by Tversky and colleagues.

The DFT and the LCA models share many principles but also differ on some. While DFT is a linear model (where excitation by negated inhibition, of the type described in Section 6.2, is allowed) and where the degree of lateral inhibition depends on the similarity between the alternatives, in the LCA account the lateral inhibition is constant (not similarity dependent) but we impose two types of nonlinearity. The first type corresponds to a zero-activation threshold (discussed in Section 6.3), while the second one involves a convex utility-value function (Kahneman and Tversky, 2000).

It is beyond the scope of this chapter to compare detailed predictions of the two models (but see Usher and McClelland, 2004, and reply by Busemeyer *et al.*, 2005). We believe, however, that there is enough independent motivation for nonlinearity and reference dependency of the value functions. In the next subsection we discuss some principles underlying value evaluation and then we show how a simple LCA type model, taking these principles on board, can address value-based decisions.

6.5.3 Nonlinear utility functions and the Weber law

The need for a nonlinear relation between internal utility and objective value was noticed by Daniel Bernoulli (1738 [1954]), almost two centuries ago. Bernoulli proposed a logarithmic type of nonlinearity in the value function in response to the so-called St. Petersburg Paradox. (The paradox was first noticed by the casino operators of St. Petersburg. See for example Martin, 2004, and Glimcher, 2004, pp. 188–92 for good descriptions of the paradox and of Bernoulli’s solution). Due to its simple logic and intuitive appeal, we reiterate it here.

Consider the option of entering a game, where you are allowed to repeatedly toss a fair coin until ‘head’ comes. If the ‘head’ comes in the first toss you receive £2. If the ‘head’ comes in the second toss, you receive £4, if in the third toss, £8, and so on (with each new toss needed to obtain a ‘head’ the value is doubled). The question is what is the price that a person should be willing to pay for playing this game. The puzzle is that although the expected value of the game is infinite ($E = \sum_{i=1, \dots, \infty} \frac{1}{2}^i 2^i = \sum_{i=1, \dots, \infty} 1 = \infty$), as the casino operators in St. Petersburg discovered, most people are not willing to pay more than £4 for playing the game and very few more than £25 (Hacking, 1980). Most people show *risk-aversion*. (In this game, most often one wins small amounts (75% to win less than £5), but in few cases one can win a lot. Paying a large amount to play the game results in a high probability of making a loss and a small probability of a high win. Hence the low value that people are willing to pay reflects risk-aversion.)

Bernoulli’s assumption, that internal utility is nonlinearly (with diminishing returns) related to objective value, offers a solution to this paradox (the utility of a twice larger value is less than twice the utility of the original value) and has been included in the dominant theory of risky choice, the prospect theory (Tversky and Kahneman, 1979). A logarithmic value function $u(x) = \log_{10}(x)$, used as the expected utility, gives a value of about £4 for the St. Petersburg game.

Note that the need to trade between the utility associated with different objective values arises, not only in risky choice between options associated with monetary values but also in cases of multidimensional choice (as illustrated in Figure 6.9) where the

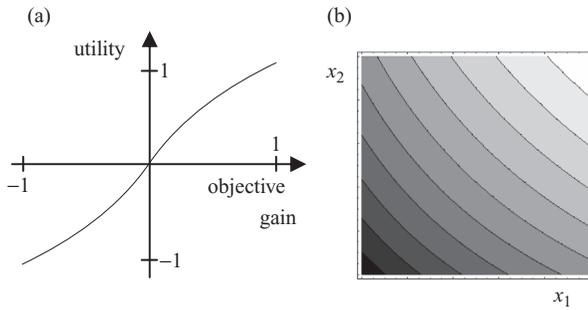


Figure 6.10 (a) Utility function, $u(x) = \log(1 + kx)$ for $x > 0$, and $-\gamma \log(1 - kx)$, for $x < 0$. ($k = 2, \gamma = 1$). (b) Combined 2D-utility function for gains ($x_1 > 0, x_2 > 0$).

options are characterised by their value on two or more dimensions. Moreover, as such values are examples of analogue magnitude representations, one attractive idea is to assume that their evaluation obeys a psychophysical principle that applies to magnitude judgements, in general: the Weber law. The Weber law states that to be able to discriminate between two magnitudes (e.g., weights), x and $x + dx$, the just-noticeable-difference, dx , is proportional to x itself.

One simple way to satisfy the Weber law is to assume that there are neural representations that transform their input (which corresponds to objective value) under a logarithmic type of nonlinearity and that the output is subject to additional independent noise of constant variance. This explanation for the Weber law is consistent with neurophysiological data from a task in which monkeys discriminated between stimuli differing in the number of dots, that suggest that prefrontal neurons represent the logarithms of numbers (Dehaene, 2003).

As proposed by Bernoulli (1738 [1954]) a logarithmic nonlinearity also accounts for risk aversion. Here we assume a logarithmic nonlinearity of the type, $u(x) = \log(1 + kx)$ for $x > 0$, and $u(x) = -\gamma \log(1 - kx)$, for $x < 0$ ($x > 0$ corresponds to gains and $x < 0$ to losses); the constant of 1 in the logarithm corresponds to a baseline of present value before any gains or losses are received). [In prospect theory (Kahneman and Tversky, 2000; Tversky and Simonson, 1993) one chooses, $\gamma > 1$, indicating a higher slope for losses than for gains. This is also assumed in Usher and McClelland (2004). Here we use $\gamma = 1$ in order to explore the simplest set of assumptions that can result in these reversal effects; increasing γ strengthens the effects.] As shown in Figure 6.10a, function $u(x)$ starts linearly and then is subject to diminishing returns, which is a good approximation to the neuronal input–output response function of neurons at low to intermediate firing rates (Usher and Niebur, 1996).

There is a third rationale for a logarithmic utility function, which relates to the need to combine utilities across dimensions. When summing such a utility function across multiple dimensions, one obtains (for two dimensions), $U(x_1, x_2) = u(x_1) + u(x_2) = \log[1 + k(x_1 + x_2) + k^2 x_1 x_2]$. Note that to maximise this utility function one has to maximise a combination of linear and multiplicative terms. The inclusion of a multiplicative term in the utility optimisation is supported by a survival rationale: to survive animals

needs to ensure the joined (rather than separate) possession of essential resources (like food and water). Figure 6.10b illustrates a contour plot of this 2D utility function. One can observe that equal preference curves are now curved in the x_1 - x_2 continuum: the compromise (0.5,0.5) has a much better utility than the (1,0) option.

Another component of the utility evaluation is its reference dependence. Moreover, as discussed in Subsection 6.5.2, the reference depends on the subjective expectations and on the information accessible to the decision maker (Kahneman, 2003). As we show below, the combination of nonlinear utility and reference dependence explains the presence of contextual preference reversals. Finally, when choice alternatives are characterised over multiple dimensions, we assume (following Tversky's elimination by aspects, Tversky, 1972, and the various DFT applications, Busemeyer and Townsend, 1993; Roe *et al.*, 2001) that decision makers switch their attention, stochastically, from dimension to dimension. Thus at every time step the evaluation is performed with regard to one of the dimensions and the preference is integrated by the leaky competing accumulators. In the following subsection, these components of utility evaluations are introduced into an LCA model and applied to the value-based decision patterns described above.

6.5.4 Modelling value-based choice in the LCA framework

To allow for the switching between the alternative dimensions, the LCA simulations are done using a discretised version of the LCA model of Equation (6.2) (single step of Euler method; note a threshold nonlinearity at zero is imposed: only $y_i > 0$ are allowed)

$$y_i(t + \Delta t) = y_i(t) + \Delta t \left(-ky_i - w \sum_{\substack{j=1 \\ j \neq i}}^N y_j + I_i + I_0 + noise \right) \quad (6.10)$$

where I_i were evaluated according to the utility function described above and I_0 is a constant input added to all choice units, which is forcing a choice (in all simulations reported here this value is chosen as 0.6). To account for the stochastic nature of human choice each integrator received the noise that was Gaussian distributed (with standard deviation (SD) of 0.5). During all simulations the following parameters were chosen, $\Delta t = 0.05$, $k = w = 1$ (balanced network). When a reference location is explicitly provided (as in the situation depicted in Figure 6.9) the utility is computed relative to that reference. When no explicit reference is given, a number of possibilities for implicit reference are considered.

In all the simulations we present, the decision is monitored (as in Roe *et al.*, 2001, and in Usher and McClelland, 2004) via an interrogation-like procedure. The response units are allowed to accumulate their preference-evaluation for T -time steps. A total of 500 trials of this type are simulated and the probability of choosing an option as a function of time, $P_i(t)$ is computed by counting the fraction of trials in which the corresponding unit has the highest activation (relative to all other units) at time- t . We start with a simple

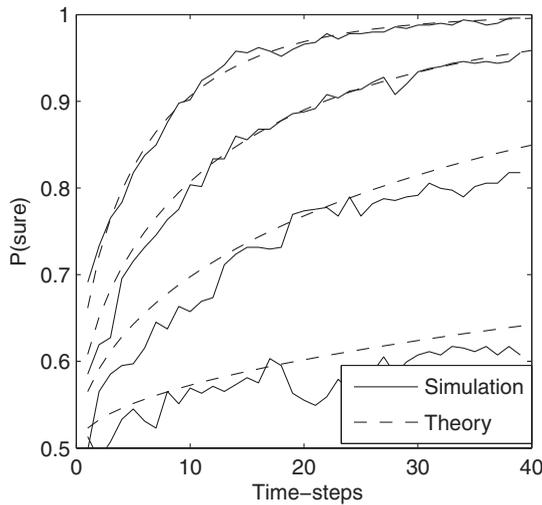


Figure 6.11 Probability of choosing the sure option as a function of deliberation time for five values of risk (indicated in the plot). Solid lines were obtained from simulations of the LCA model for the following parameters $W = 1$, $I_0 = 0.6$, $SD = .5$, and the utility function from Figure 6.10. Dashed lines show corresponding probabilities derived in Bogacz *et al.* (2007).

demonstration of risk-aversion in probabilistic monetary choice and then we turn to preference reversals in multidimensional choice.

6.5.4.1 Risk-aversion in probabilistic choice

We simulate here a choice between two options. The first one corresponds to a ‘sure’ win, W , while the second one to a probabilistic win of W/p , with probability p (note that the two have equal expected objective value, W , and that p provides a measure of risk: lower p is more risky). The model assumes that decision makers undergo a ‘mental simulation’ process, in which the utility of the gain drives the value accumulator, thus the sure unit receives a constant input $I_0 + u(W)$, while the probabilistic unit receives probabilistic input, chosen to be $I_0 + u(W/p)$ with probability p , and I_0 otherwise. In addition, a constant noise input ($SD = 0.5$) is applied to both units at all time steps. Note that due to the shape of utility function u , the average input to the sure unit ($I_0 + u(W)$) is larger than to the probabilistic unit ($I_0 + u(W/p)p$). In Figure 6.11 we show the probability of choosing the sure option as a function of deliberation time for five risk levels, p (small p corresponds to large risk and p close to 1 to low risk). Thus the higher the risk the more likely is the bias of choosing the sure option (this bias starts at value approximately proportional to $1 - p$ and increases due to time integration to asymptotic value). This is consistent with experimental data, except for low p , where as explained by the Prospect Theory (Tversky and Kahneman, 1979), decision makers show an overestimative discrepancy between subjective and objective probability, which we do not address here (but see Hertwig *et al.*, 2004). Risk seeking for losses can be simulated analogously.

6.5.4.2 Multidimensional choice: reference effects and preference reversal

Three simulations are reported. In all of them, at each time step, one dimension is probabilistically chosen (with $p = 0.5$) for evaluation. The preferences are then accumulated across time and the choices for the various options are reported as a function of deliberation time.

First, we examine how the choice between two options, corresponding to A and B in Figure 6.9a is affected by a change of the reference, R_A versus R_B . The options are defined on two dimensions as follows: $A = (0.2, 0.8)$, $B = (0.8, 0.2)$, $R_A = (0.2, 0.6)$ and $R_B = (0.6, 0.2)$. Thus, for example, in simulations with reference R_A , when the first dimension is considered, the inputs I_A and I_B are $I_0 + u(0)$ and $I_0 + u(0.6)$ while when the second dimension is considered the inputs are $I_0 + u(0.2)$ and $I_0 + u(-0.4)$ (this follows from the fact that $A - R_A = (0, 0.2)$ and $B - R_A = (0.6, -0.4)$). We observe (Figure 6.12a) that the R_A reference increases the probability of choosing the similar A -option (top curve) and that the choice preference reverses with the R_B reference (the middle curve corresponds to a neutral $(0, 0)$ reference point). This happens because with reference R_A the average input to A is larger than to B (as $u(0) + u(0.2) = u(0.2) > u(0.6) - u(0.4) = u(0.6) + u(-0.4)$) and vice versa. [If $I_0 = 0$, the net advantage in utility for the nearby option is partially cancelled by an advantage for the distant option due to the zero-activation boundary (negative inputs are reflected by the boundary). The value of I_0 does not affect the other results (compromise or similarity)].

Second, we examine the compromise effect. The options correspond to a choice situation with three alternatives A , B , and C differing on two dimensions as shown in Figure 6.9b. A and B are defined as before and C is defined as $(0.5, 0.5)$. We assume that when all three choices are available the reference is neutral $(0, 0)$. We observe (Figure 6.12b) that the compromise alternative is preferred among the three. This is a direct result of 2D utility function (Figure 6.10b). For binary choice between A and C we assume that the reference point is moved to a point of neutrality between A and C , such as $R_{AC} = (0.2, 0.5)$, which corresponds to a new baseline relative to which the options A and C can be easily evaluated as having only gains and no losses (alternatively, one can assume that each option serves as a reference for the evaluation of the other ones; Usher and McClelland, 2004). This maintains an equal preference between C and the extremes in binary choice. Note also the dynamics of the compromise effect. This takes time to develop; at short times the preference is larger for the extremes, depending on the dimension evaluated first. Experimental data indicates that, indeed, the magnitude of the compromise effect increases with the deliberation time (Dhar *et al.*, 2000).

Third, we examine the similarity effect. In this situation, the option $S = (0.2, 0.7)$ (similar to A) is added to the choice set of A and B . The reference is again neutral $(0, 0)$. We observe that the dissimilar option, B (Figure 6.12c, solid curve), is preferred. This effect is due to the correlation in the activation of the similar alternatives (A and S), which is caused by their co-activation by the same dimensional evaluation. When the supporting dimension is evaluated both of the similar options rise in activation and they split their choices, while the dissimilar option peaks at different times and has a relative advantage. Note also a small compromise effect in this situation. Among the similar options, S (which is a compromise) has a higher choice probability. The attraction effect

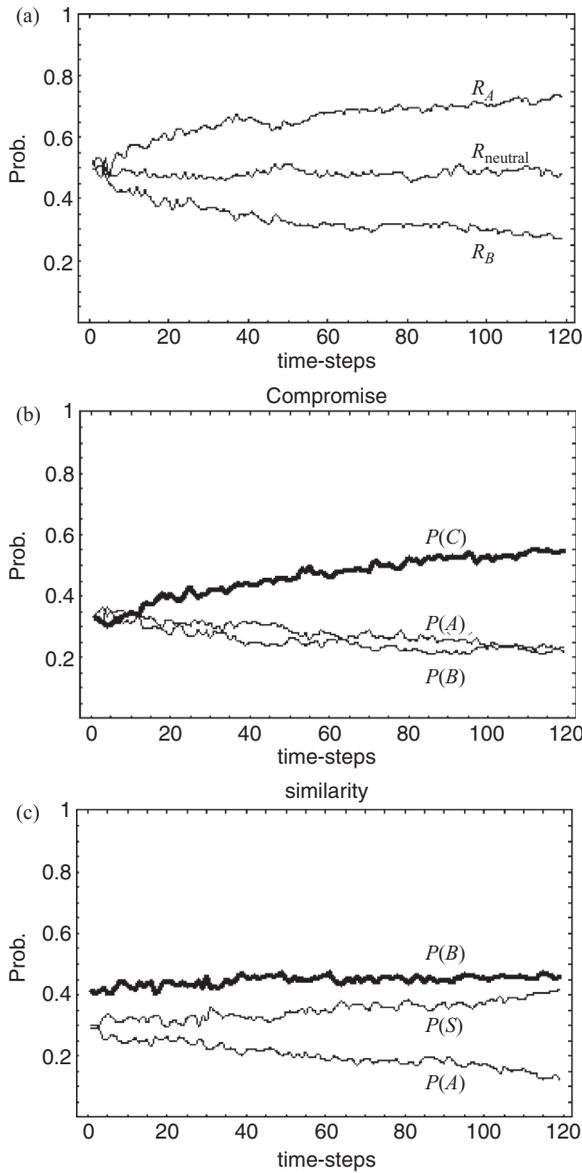


Figure 6.12 Contextual preference reversal. (a) Reference effects in binary choice. (b) Compromise effect. (c) Similarity effect.

is similar to the reference effect. One simple way to explain it is to assume that the reference moves towards the dominated option. (Alternatively, each option may serve as a reference for every other option; Tversky and Simonson, 1993; Usher and McClelland, 2004).

To summarise, we have shown that when the input to LCA choice units is evaluated according to a nonlinear utility function of the type proposed by Bernoulli, which is

applied to differences in value between options and a referent, the model can account for a number of choice patterns that ‘appear’ to violate normativity. For example, the model provides a plausible neural implementation and extension of the Prospect Theory (Tversky and Kahneman, 1979), displaying risk aversion (it prefers the sure option on a risky one of equal expected value) and a series of preference reversals that are due to the effect of context on the choice-reference.

6.6 Discussion

In this chapter we have reviewed the conditions under which various versions of the LCA model (linear and nonlinear) achieve optimal performance for different experimental conditions (free-response and interrogation). We have also shown how the LCA model can be extended to value-based decisions to account for risk aversion and contextual preference reversals.

We have shown that the linear LCA model can implement the optimal choice algorithm for all tasks except the choice between multiple alternatives receiving similar amounts of supporting evidence in the free-response paradigm. Moreover, we have shown that for the choice involving multiple alternatives in the free-response paradigm, the nonlinearities of type present in a biological decision network can improve the performance, and in fact may allow the networks to approximate the optimal choice algorithm. This raises an intriguing possibility, that these nonlinearities are not a result of constraints of biological neurons, but may rather be a result of evolutionary pressure for speed of decisions.

We have also identified a condition (see Section 6.4) in which performance can be optimised by an elevation/decrease in the level of lateral inhibition relative to the leak (this may be achieved via neuromodulation, e.g., Usher and Davelaar, 2002). It will be interesting to test whether the behavioural manifestations of unbalance of decay and inhibition (Usher and McClelland, 2001) can be experimentally observed under these conditions.

One interesting comment relates to Hick’s law, according to which the DT is proportional to the logarithm of the number of alternatives (Teichner and Krebs, 1974). In the simulations of bounded LCA model in Figure 6.5a and b, the DT does not depend on the number of potentially available alternatives. Note, however, this simulation was designed to model the task described at the beginning of Section 6.3 (Figure 6.6a) in which the choice is mainly between two alternatives, which match the ambiguous input (in this simulation only two accumulators receive any input or noise). If all accumulators received equal levels of noise and the bounded LCA model remained in the linear range, it would satisfy Hick’s law, because when the bounded LCA model is in linear range, it is equivalent to the linear model, and the linear model satisfies Hick’s law when accumulators receive equal level of noise (McMillen and Holmes, 2006). However, it has been recently reported that in tasks where one of the alternatives receives much more support than all the others, Hick’s law is indeed violated and the DT does not depend on the number of alternatives (Kveraga *et al.*, 2002). Thus it would be interesting to

investigate the prediction of our theory that a similar independence may occur when two alternatives receive much larger input than the others.

It has been recently proposed that if the balanced LCA model projects to a complex network with architecture resembling that of the basal ganglia, the system as a whole may implement the MSPRT (Bogacz and Gurney, 2007) – the optimal algorithm for this condition. The system involving the basal ganglia may thus optimally make choices between motor actions. However, many other choices (e.g., perceptual or motivational) are likely to be implemented in the cortex. The complexity of MSPRT prevents any obvious cortical implementations, hence it is still of great interest to investigate the parameters optimising the LCA model which can be viewed as an abstraction of cortical processing.

The extension to value based decisions brings the model in closer contact with the topic of action selection. Actions need to be selected according to the value of their consequences, and this requires an estimation of utility and its integration across dimensions. The LCA model is also related to many models of choice on the basis of noisy data presented in this book. In particular, it is very similar to the model of action selection in the cerebral cortex by Cisek (2006), which also includes accumulation of evidence and competition between neuronal populations corresponding to different alternatives.

Acknowledgements

This work was supported by EPSRC grant EP/C514416/1. We thank Andrew Lulham for reading the manuscript and for his very useful comments. Matlab codes for simulation and finding decision time of LCA model can be downloaded from: <http://www.cs.bris.ac.uk/home/rafal/optimal/>.

References

- Barnard, G. (1946). Sequential tests in industrial statistics. *J. Roy. Stat. Soc. Suppl.* **8**: 1–26.
- Bernoulli, D. (1738 [1954]). Exposition of a new theory on the measurement of risk. *Ekonometrica* **22**: 23–36.
- Bogacz, R., E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. *Psychol. Rev.* **113**: 700–65.
- Bogacz, R. and K. Gurney (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput.* **19**: 442–77.
- Bogacz, R., M. Usher, J. Zhang, and J. L. McClelland (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Phil. Trans. Roy. Soc. B.* **362**: 1655–70.
- Britten, K. H., M. N. Shadlen, W. T. Newsome, and J. A. Movshon (1993). Responses of neurons in macaque MT to stochastic motion signals. *Vis. Neurosci.* **10**(6): 1157–69.

- Brown, E., J. Gao, P. Holmes, *et al.* (2005). Simple networks that optimize decisions. *Int. J. Bifurcat. Chaos* **15**: 803–26.
- Brown, E. and P. Holmes (2001). Modeling a simple choice task: stochastic dynamics of mutually inhibitory neural groups. *Stoch. Dynam.* **1**: 159–91.
- Busemeyer, J. R. and J. T. Townsend (1993). Decision field theory: a dynamic-cognitive approach to decision making in uncertain environment. *Psychol. Rev.* **100**: 432–59.
- Busemeyer, J. R., J. T. Townsend, A. Diederich, and R. Barkan (2005). Contrast effects or loss aversion? Comment on Usher and McClelland (2004). *Psychol. Rev.* **111**: 757–69.
- Cisek, P. (2006). Cortical mechanisms of action selection: the affordance competition hypothesis. *Phil. Trans. Roy. Soc. B.* **362**: 1585–600.
- Dehaene, S. (2003). The neural basis of the Weber–Fechner law: a logarithmic mental number line. *Trends Cog. Sci.* **7**: 145–7.
- Dhar, R., S. M. Nowlis, and S. J. Sherman (2000). Trying hard or hardly trying: an analysis of context effects in choice. *J. Con. Psych.* **9**: 189–200.
- Dragalin, V. P., A. G. Tertakovskiy, and V. V. Veeravalli (1999). Multihypothesis sequential probability ratio tests – part I: asymptotic optimality. *IEEE Trans. I.T.* **45**: 2448–61.
- Glimcher, P. W. (2004). *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics*. Cambridge, MA: MIT Press.
- Gold, J. I. and M. N. Shadlen (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* **36**(2): 299–308.
- Hacking, I. (1980). Strange expectations. *Phil. Sci.* **47**: 562–7.
- Hertwig, R., G. Barron, E. U. Weber, and I. Erev (2004). Decisions from experience and the effect of rare events in risky choice. *Psych. Sci.* **15**: 534–9.
- Houston, A. I., J. McNamara, and M. Steer (2006). Do we expect natural selection to produce rational behaviour? *Phil. Trans. Roy. Soc. B* **362**: 1531–44.
- Kahneman, D. (2003). Maps of bounded rationality: psychology for behavioral economics. *Am. Econ. Rev.* **93**: 1449–75.
- Kahneman, D. and A. Tversky, eds. (2000). *Choices, Values and Frames*. Cambridge: Cambridge University Press.
- Kveraga, K., L. Boucher, and H. C. Hughes (2002). Saccades operate in violation of Hick’s law. *Exp. Brain Res.* **146**(3): 307–14.
- Laming, D. R. J. (1968). *Information Theory of Choice Reaction Time*. New York: Wiley.
- LeBoeuf, R. and E. B. Shafir (2005). Decision-making. In *Cambridge Handbook of Thinking and Reasoning*, ed. K. J. Holyoak and R. G. Morisson. Cambridge: Cambridge University Press, pp. 243–66.
- Martin, R. (2004). The St. Petersburg Paradox. In *The Stanford Encyclopedia of Philosophy*, ed. E. Zalta. Stanford, CA: The Metaphysics Research Lab. Available at <http://plato.stanford.edu/archives/fall2008/entries/paradox-stpetersburg/>.
- Mazurek, M. E., J. D. Roitman, J. Ditterich, and M. N. Shadlen (2003). A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13**(11): 1257–69.
- McMillen, T. and P. Holmes, (2006). The dynamics of choice among multiple alternatives. *J. Math. Psych.* **50**: 30–57.
- Neyman, J. and E. S. Pearson (1933). On the problem of the most efficient tests of statistical hypotheses. *Phil. Trans. Roy. Soc. A* **231**: 289–337.
- Platt, M. L. and P. W. Glimcher (1999). Neural correlates of decision variables in parietal cortex. *Nature* **400**(6741): 233–8.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychol. Rev.* **83**: 59–108.

- Ratcliff, R. (1988). Continuous versus discrete information processing: modeling accumulation of partial information. *Psychol. Rev.* **95**: 238–55.
- Roe, R. M., J. R. Busemeyer, and J. T. Townsend (2001). Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol. Rev.* **108**: 370–92.
- Roitman, J. D. and M. N. Shadlen (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**(21): 9475–89.
- Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.* **2**(1): 33–42.
- Seung, H. S. (2003). Amplification, attenuation, and integration. In *The Handbook of Brain Theory and Neural Networks*, 2nd edn, ed. M. A. Adbib. Cambridge, MA: MIT Press, pp. 94–7.
- Shadlen, M. N. and W. T. Newsome (1998). The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**(10): 3870–96.
- Shadlen, M. N. and W. T. Newsome (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**(4): 1916–36.
- Slovic, P. (1995). The construction of preference. *Am. Psychol.* **50**: 364–71.
- Snowden, R. J., S. Treue, and R. A. Andersen (1992). The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp. Brain Res.* **88**(2): 389–400.
- Stone, M. (1960). Models for choice reaction time. *Psychometrika* **25**: 251–60.
- Sugrue, L. P., G. S. Corrado, and W. T. Newsome (2004). Matching behavior and the representation of value in the parietal cortex. *Science* **304**(5678): 1782–7.
- Sugrue, L. P., G. S. Corrado, and W. T. Newsome, (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* **6**(5): 363–75.
- Teichner, W. H. and M. J. Krebs (1974). Laws of visual choice reaction time. *Psychol. Rev.* **81**: 75–98.
- Tversky, A. (1972). Elimination by aspects: a theory of choice. *Psychol. Rev.* **79**: 281–99.
- Tversky, A. and D. Kahneman (1979). Prospect theory: an analysis of decision under risk. *Econometrica* **47**: 263–92.
- Tversky, A. and D. Kahneman (1991). Loss aversion in riskless choice: a reference-dependent model. *Q. J. Economet.* **106**: 1039–61.
- Tversky, A. and I. Simonson (1993). Context-dependent preferences. *Manage. Sci.* **39**: 1179–89.
- Usher, M. and E. J. Davelaar (2002). Neuromodulation of decision and response selection. *Neural Networks* **15**: 635–45.
- Usher, M. and J. L. McClelland (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**(3): 550–92.
- Usher, M. and J. L. McClelland (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol. Rev.* **111**: 759–69.
- Usher, M. and N. Niebur (1996). Modeling the Temporal Dynamics of IT Neurons in Visual Search: A Mechanism for Top-Down Selective Attention. *J. Cog. Neurosci.* **8**: 311–27.
- Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics* **13**: 37–58.
- Vickers, D. (1979). *Decision Processes in Perception*. New York: Academic Press.
- Wald, A. (1947). *Sequential Analysis*. New York: Wiley.
- Wang, X. J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**(5): 955–968.