PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B

Phil. Trans. R. Soc. B (2007) **362**, 1655–1670
doi:10.1098/rstb.2007.2059
Published online 11 April 2007

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

Rafal Bogacz[1,*], Marius Usher[2], Jiaxiang Zhang[1] and James L. McClelland[3]

[1]*Department of Computer Science, University of Bristol, Bristol BS8 1UB, UK*
[2]*Department of Psychology, Birkbeck College, University of London, London WC1E 7HX, UK*
[3]*Center for the Neural Bases of Cognition, Carnegie Mellon University, Pittsburgh, PA 15213, USA*

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 paper discusses recent analyses and extensions of the LCA model. First, it reviews the dynamics and 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 characterized across multiple dimensions.

Keywords: decision making; perceptual choice; nonlinear; optimality; utility; preference reversal

1. INTRODUCTION

Making choices is a 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 (Stone 1960; Laming 1968; Vickers 1970; Ratcliff 1978). 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 (Schall 2001; Shadlen & Newsome 2001; Gold & Shadlen 2002). 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 & McClelland 2001), which is sufficiently simple to allow a detailed mathematical analysis. As we will discuss in the following sections, this model can, for certain values of its parameters, approximate the same computations carried out by a series of mathematical models of choice (Vickers 1970; Ratcliff 1978; Busemeyer & Townsend 1993; Shadlen & Newsome 2001; Wang 2002).

* Author for correspondence (r.bogacz@bristol.ac.uk).

Electronic supplementary material is available at http://dx.doi.org/10.1098/rstb.2007.2059 or via http://www.journals.royalsoc.ac.uk.

One contribution of 15 to a Theme Issue 'Modelling natural action selection'.

Since its original publication, the LCA model (Usher & McClelland 2001) has been analysed mathematically and extended in a number of directions (Brown & Holmes 2001; Brown *et al.* 2005; McMillen & Holmes 2006; Bogacz *et al.* 2006). In particular, it has been investigated for which values of parameters it achieves an optimal performance. In this paper, we will use the term 'optimal' to describe the theoretically best possible performance. This is important, because if 'we expect natural selection to produce rational behaviour', as discussed by Houston *et al.* (2007), then the optimal values of parameters should be found in the neural networks mediating choice processes. In some cases, decision networks cannot achieve optimal performance due to biological constraints; however, it is still of interest to investigate which parameters give the best performance within the constraints considered—we use the term 'optimized' 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: (i) is linear processing also optimal for choice between multiple alternatives? and (ii) what are the parameters of the nonlinear LCA model that optimize its performance?

This paper 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 (§2). Second, it presents new developed extensions related to

This journal is © 2007 The Royal Society

the introduction of nonlinearities. In §3, we show that nonlinearities (of the type present in biological neurons) may improve performance in choice between multiple alternatives. In §4, we discuss how to optimize the performance of a nonlinear LCA model for two alternatives. Finally, in §5, we show how nonlinearities in the LCA model also explain counter-intuitive results from choice experiments involving multiple goals or stimulus dimensions.

2. REVIEW OF THE LEAKY COMPETING ACCUMULATOR 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 and nonlinear versions of this model and analyse its dynamics and performance.

(a) *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). On 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 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 the participants are allowed to respond at any moment of time. The second one is the *interrogation* (or response signal) paradigm, in which the participants are required to continuously observe the stimulus until a particular signal (whose delay is controlled) is provided that prompts an immediate response.

During the choice process, sensory areas (e.g. medial temporal (MT) area involved in motion processing) 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 (LIP) area and the frontal eye field gradually increase their firing rates (Schall 2001; Shadlen & Newsome 2001). Furthermore, because the easier the task, the faster the rate of this increase (Shadlen & Newsome 2001), it has been suggested that these neurons integrate the evidence from sensory neurons over time (Schall 2001; Shadlen & Newsome 2001). This integration averages out the noise present in sensory neurons allowing the accuracy of the choice to increase with time. Moreover, since (in the free-response paradigm) the firing rate, just before the saccade, does not differ between difficulty levels of the task (Roitman & Shadlen 2002), it is believed that the choice is made when activity of the neuronal population representing one of the alternatives reaches a decision threshold.

(b) *Biologically inspired models of perceptual choice*

A number of computational models have been proposed to describe the choice process outlined previously, and their architectures are shown in figure 1 for the case of two alternatives (Usher & McClelland 2001; Wang 2002; Mazurek *et al.* 2003). All of these models include two units (bottom circles in figure 1) corresponding to neuronal populations providing noisy evidence, and two accumulator units (denoted by y_1 and y_2 in figure 1) integrating the evidence. The models differ in the way inhibition affects the integration process: in the LCA model (figure 1*a*), the accumulators inhibit each other; in the Mazurek *et al.* (2003) model (figure 1*b*), the accumulators receive inhibition from the other inputs; and in the Wang (2002) model (figure 1*c*), the accumulators inhibit each other via a population of inhibitory interneurons. 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 paper, we thus focus on the LCA model and review its optimality (analogous analysis of the other two models is described in Bogacz *et al.* (2006)).

(c) *Linear LCA model*

Figure 1*a* shows the architecture of the LCA model for the two alternative choice tasks (Usher & 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, $\mathrm{d}W_i$, of amplitude c_i ($\mathrm{d}W_i$ denotes independent Wiener processes). These units also inhibit each other by means of a connection of weight w. Hence, during the choice process, information is accumulated according to the following equations (Usher & McClelland 2001)

$$\begin{cases} \mathrm{d}y_1 = (-ky_1 - wy_2 + I_1)\mathrm{d}t + c_1\mathrm{d}W_1 \\ \mathrm{d}y_2 = (-ky_2 - wy_1 + I_2)\mathrm{d}t + c_2\mathrm{d}W_2 \end{cases}, \; y_1(0) = y_2(0) = 0. \tag{2.1}$$

In the above equations, the term $-ky_i$ 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 outlined in §2*a*. 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

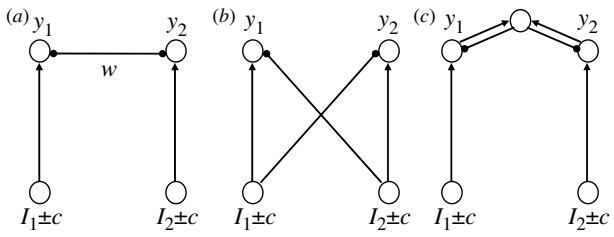

Figure 1. Architectures of the models of choice. Arrows denote excitatory connections, lines with filled circles denote inhibitory connections. (*a*) LCA model (Usher & McClelland 2001). (*b*) Mazurek *et al.* (2003) model. (*c*) Wang (2002) model.

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.

Since 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 on the majority of simulated trials. However, on some trials, due to noise, another alternative may be chosen; such trials correspond to incorrect responses. By simulating the model multiple times, the 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 & McClelland 2001):

$$\mathrm{d}y_i = \left(-ky_i - w\sum_{\substack{j=1\\ j\neq i}}^{N} y_j + I_i\right)\mathrm{d}t + c_i\mathrm{d}W_i, \qquad (2.2)$$

$$y_i(0) = 0.$$

When the decay and inhibition parameters are equal to zero, the terms in equations (2.1) and (2.2) 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 the accumulators integrate noisy evidence independent of one another.

(d) *Dynamics of the model*

The review of the dynamics of the linear LCA model in this subsection is based on the work by 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 visualized as a path in this plane. Representative paths for three different parameter ranges in this plane are shown in figure 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 towards the threshold for the more strongly activated accumulator (i.e. the correct choice).

To better understand the dynamics of the model, figure 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 the three panels of figure 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 2*b*: x_1 is parallel to the attracting line and describes the difference 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)

$$\begin{cases} x_1 = \dfrac{y_1 - y_2}{\sqrt{2}}, \\ x_2 = \dfrac{y_1 + y_2}{\sqrt{2}}. \end{cases} \qquad (2.3)$$

In these new coordinates, equations (2.1) become (Bogacz *et al.* 2006)

$$\mathrm{d}x_1 = \left((w-k)x_1 + \frac{I_1 - I_2}{\sqrt{2}}\right)\mathrm{d}t + \frac{c_1}{\sqrt{2}}\mathrm{d}W_1 - \frac{c_2}{\sqrt{2}}\mathrm{d}W_2, \qquad (2.4)$$

$$\mathrm{d}x_2 = \left((-k-w)x_2 + \frac{I_1 + I_2}{\sqrt{2}}\right)\mathrm{d}t + \frac{c_1}{\sqrt{2}}\mathrm{d}W_1 + \frac{c_2}{\sqrt{2}}\mathrm{d}W_2. \qquad (2.5)$$

Equations (2.4) and (2.5) are *uncoupled*, i.e. 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 equation (2.1)). 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 previously, in figure 2*a–c* 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 2 also shows the dynamics of the system in the direction of coordinate x_1. This is slower than the x_2 dynamics and it corresponds to a motion along the line. Its characteristics depend on the relative values of inhibitory weight, w, and decay, k. When decay is larger than inhibition, attractor dynamics also come into play, as shown in figure 2*a*. The system is attracted towards this point and fluctuates in its vicinity. In figure 2*a*, the

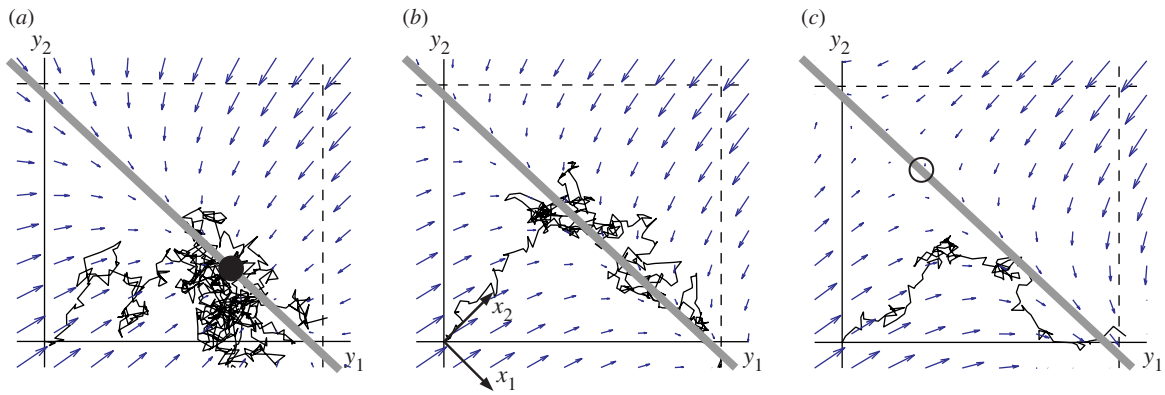

Figure 2. Examples of the evolution of the LCA model, showing paths in the state space of the model. The horizontal axes denote the activation of the first accumulator and the vertical axes denote the activation of the second accumulator. The paths show the choice process from stimulus onset (where $y_1=y_2=0$) 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*) decay > inhibition; $w=7$, $k=13$; (*b*) decay = inhibition; $w=10$, $k=10$; (*c*) decay < inhibition; $w=13$, $k=7$. The simulations were performed using the 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 the normal distribution with mean 0 and variance $c^2\Delta t$. The arrows show the average direction of movement of the LCA model in the state space. The thick grey lines symbolize the attracting lines; filled circle in (*a*) indicates the attractor; open circle in (*c*) indicates the unstable fixed point.

threshold is reached when noise pushes the system away from the attractor. When inhibition is larger than decay, the x_1 dynamic is characterized by repulsion from the fixed point, as shown in figure 2*c*.

When inhibition equals decay, the term $(w-k)\,x_1$ in equation (2.4) disappears and, describing the evolution along the attracting line, it can be written as

$$\mathrm{d}x_1 = \left(\frac{I_1}{\sqrt{2}}\mathrm{d}t + \frac{c_1}{\sqrt{2}}\mathrm{d}W_1\right) - \left(\frac{I_2}{\sqrt{2}}\mathrm{d}t + \frac{c_2}{\sqrt{2}}\mathrm{d}W_2\right). \tag{2.6}$$

In the rest of the paper, we refer to the linear LCA model with inhibition equal to decay as *balanced*. The vector field for this case is shown in figure 2*b*. In this case, according to equation (2.6), 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 the two alternatives.

The three cases shown in figure 2 make different predictions about the impact of temporal information on choice in the interrogation paradigm. If inhibition is larger than decay (figure 2*c*), 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 & Townsend 1993; Usher & 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 & Townsend 1993; Usher & 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 & McClelland (2001) tested whether the effects described previously 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.

(e) *Performance of the linear LCA model*

In this subsection, we review the 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 & Holmes 2006), and the interrogation paradigm.

When both inhibition and decay are fairly strong (figure 2*b*), the state evolves very closely to the attracting line (as mentioned previously) reaching the decision threshold very close to the intersection of the decision threshold and attracting line (figure 2*b*). 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 a negative threshold. This is the characteristic of a mathematical choice model known as the diffusion model (Stone 1960; Laming 1968; Ratcliff 1978), 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

Table 1. Summary of conditions the linear LCA model must satisfy to implement the optimal choice algorithms.

paradigm	no. of the alternatives	
	$N=2$	$N>2$
free response	inhibition = decay and both high	*optimality not attainable*
interrogation (response signal)	inhibition = decay	inhibition = decay

possible procedures for solving this choice problem giving certain ER, it minimizes 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 optimized when inhibition is equal to decay and both have high values (McMillen & 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 & Holmes 2006). For example, one of the MSPRT tests may be implemented by the 'max versus next' procedure (McMillen & 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 & Holmes 2006). Furthermore, as the number of alternatives (N) increases, the best achievable DT for fixed ER of the linear balanced LCA model approaches that of the race model (McMillen & 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; Neyman & Pearson 1933; Bogacz *et al.* 2006) and for multiple alternatives (McMillen & 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 1 summarizes the 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 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 §3.

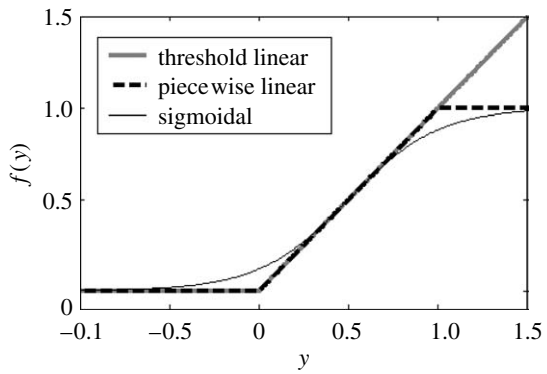

Figure 3. Nonlinear input–output functions used in the LCA model. Threshold linear: $f(y)=y$, for $y \geq 0$ and $f(y)=0$, for $y<0$ (Usher & McClelland 2001). Piecewise linear: $f(y)=0$, for $y<0$, $f(y)=1$, for $y>1$ and $f(y)=y$ otherwise (Brown *et al.* 2005). Sigmoidal: $f(y)=1/(1+e^{-4(y-0.5)})$ (Brown & Holmes 2001; Brown *et al.* 2005).

(f) *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 have been proposed, starting with the original version (Usher & 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:

$$\mathrm{d}y_i = \left(-ky_i - w \sum_{\substack{j=1 \\ j \neq i}}^{N} f(y_j) + I_i \right) \mathrm{d}t + c_i \mathrm{d}W_i, \qquad (2.7)$$

$$y_i(0) = 0.$$

Figure 3 shows three functions $f(y)$ proposed in the literature: threshold linear (Usher & McClelland 2001), piecewise linear (Brown *et al.* 2005) and sigmoidal (Brown & Holmes 2001; Brown *et al.* 2005). The threshold linear function corresponds to the constraint that actual neural activity is bounded (by zero) at its low end. The piecewise linear and sigmoidal functions bound the activity levels of accumulators at both ends (the maximum level of activity being equal to one). In the free-response paradigm, the threshold of the model with piecewise linear activation function (Brown *et al.* 2005) must be lower than one (as otherwise a choice would never be made). Hence, in the free-response paradigm, the nonlinear model with piecewise linear activation function is equivalent to the model with the threshold linear function (Usher & 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 (2.2) (rather than equation (2.7)) and add

reflecting boundaries on y_j at zero, preventing any of y_j from being negative (Usher & McClelland 2001); we refer to such a model as *bounded*. At every step of the simulation, the activity level of an accumulator y_j is being reset to zero if a negative value is obtained. The bounded model behaves very similar to the nonlinear models with threshold linear, piecewise linear and even sigmoidal activation functions, and provides a good approximation for them (see appendix A of Usher & McClelland (2001) for a detailed comparison of the bounded and nonlinear LCA models).

(g) *Performance of the 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 §2*e*) 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 4*a* (but not in figure 4*b*). The distance of the attracting line from the origin of the plane is equal to (Bogacz *et al.* 2006)

$$x_2^* = \frac{I_1 + I_2}{\sqrt{2}(k + w)}. \tag{2.8}$$

According to equation (2.8), the larger the sum of mean inputs I_1+I_2, the further the attracting line is from the origin. Figure 5 compares the performance of bounded LCA models with linear LCA models without boundaries, which we refer to as *unbounded*. Figure 4*a* 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 5*a*. 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 the simulations, the attracting line crosses the threshold before the axes, as shown in figure 4*a*.

Figure 4*b* shows the position of the attracting line for the parameters used in simulations of the unbounded LCA model for $N=2$ alternatives in figure 5*b*. 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 4*b* and hence the state is likely to hit the boundaries before reaching the threshold.

McMillen & Holmes (2006) tested the performance of the bounded LCA model for multiple alternatives, for the following parameters: $I_1=2$, $I_2=\cdots=I_N=0$, $c_1=\cdots=c_N=1$ (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.

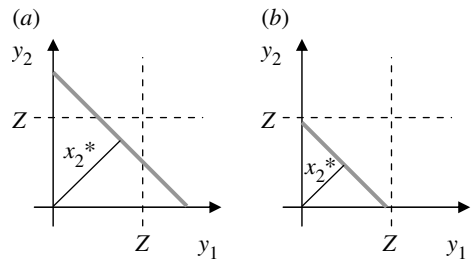

Figure 4. State plane analysis of the LCA model. Thick grey lines symbolize attracting lines in the y_1y_2 plane. (*a*,*b*) The position of the attracting line is shown for parameters used in simulations in figure 5*a*,*b*, respectively. Thus, the distance x_2^* of the attracting line from the origin is equal to 0.26 and 0.12, respectively (from equation (2.8)). 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 5*a*,*b*, respectively, i.e. 0.25 and 0.17.

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 shown in figure 6 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 previously within the free-response paradigm: we will discuss three 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 two more complex cases that reflect progressively more realistic situations.

(a) *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 $I_3=\cdots=I_N=c_3=\cdots=c_N=0$. First, let us examine the dynamics of the bounded LCA model (with $y_1, y_2 \geq 0$). In this case, the other accumulators, $y_3, \ldots, y_N$, do not receive any input but only inhibition from y_1, y_2 and hence they remain equal to zero (i.e. $y_i=0$ for all $i>2$; figure 5*c*). Therefore, the choice process simplifies to a model of two alternatives, as described in equation (2.1). Hence, when the boundaries are present, the performance of the model does not depend on the total number of alternatives, N. This is shown in figure 5*a*,*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.

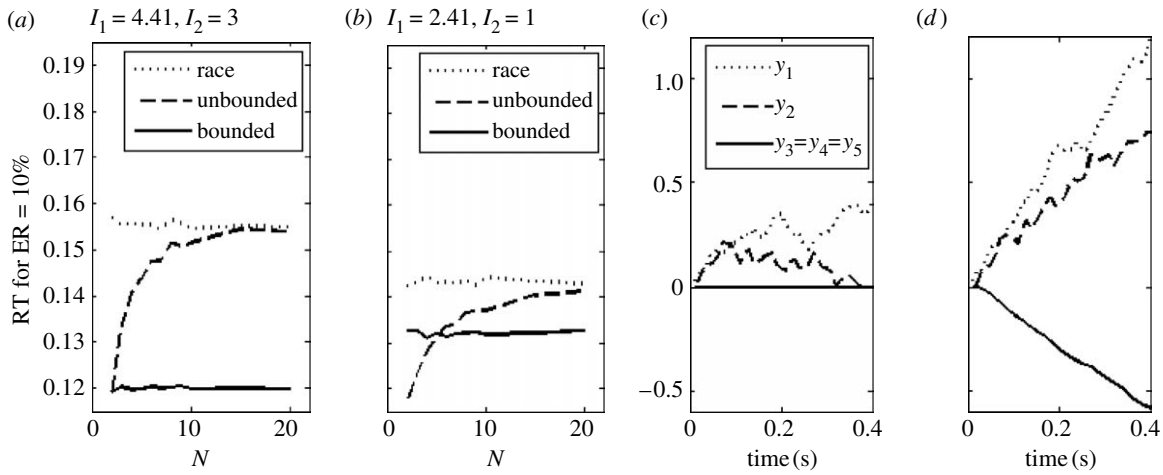

Figure 5. Performance and dynamics of choice models with only two accumulators receiving inputs. All models were simulated using the Euler method with $\Delta t = 0.01$ s. (*a*,*b*) Decision time for a threshold resulting in an ER of 10% of different choice models as a function of the number of alternatives 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, $I_3 = \cdots = I_N = 0$, $c_3 = \cdots = c_N = 0$. The panels differ in the total mean input to the first two accumulators: in (*a*) $I_2 = 3$, while in (*b*), $I_2 = 1$. For each set of parameters, a threshold was found numerically that resulted in ER of $10\% \pm 0.2\%$ (s.e.); this search for the threshold was repeated 20 times. For each of these 20 thresholds, the DT 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 LCA model (*c*) and unbounded LCA model (*d*), showing y_i as functions of time. The models were simulated for the same parameters as in (*a*), and for $N = 5$ alternatives; (*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.

Figure 5*c*,*d* compares the evolution of bounded and unbounded LCA models for $N = 5$ alternatives. Figure 5*c* shows the evolution of the bounded LCA model in which accumulators y_1, y_2 evolve in the way typical of the LCA model for two alternatives (compare with figure 2*b*): the competition between accumulators y_1, y_2 is resolved and as y_1 increases, y_2 decreases towards zero. Figure 5*d* shows that during the evolution of the unbounded model, the accumulators $y_3, \ldots, y_N$ become more and more negative. Hence, the inhibition received by y_1, y_2 from $y_3, \ldots, y_N$ is actually positive, and increases the value of y_1, y_2. Therefore, in figure 5*d* (in contrast to figure 5*c*), the activation of the 'losing' accumulator, y_2, also increases.

To better illustrate the difference between the bounded and unbounded choice behaviours, consider the dynamics of the unbounded model (equation (2.1)) for $N = 3$ alternatives. In such a case, the state is attracted to a plane (figure 7; McMillen & 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 7. 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 the y_3 axis at the current value of y_3. For example, the double grey lines in figure 7 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 (2.1), accumulator y_3 becomes more and more negative (as it receives more and more inhibition from y_1

Figure 6. 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.

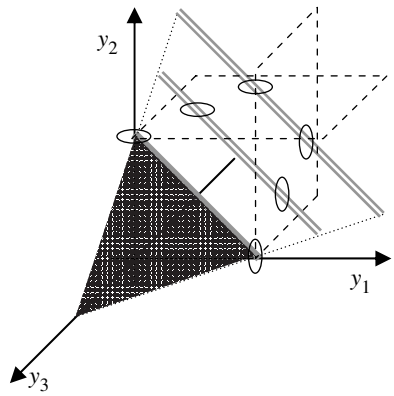

Figure 7. State-space analysis of the LCA model for three alternatives. The grey triangle indicates the attracting plane and dotted lines indicate the intersection of the attracting plane with the $y_1 y_3$ plane and the $y_2 y_3$ plane. The thick grey line symbolizes the attracting line in the $y_1 y_2$ plane; the double grey lines show sample positions of the attracting line in the $y_1 y_2$ 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 $y_1 y_2$ plane with the decision thresholds.

and y_2), as shown in figure 5*d*. Hence, the attracting line in the y_1y_2 plane moves further and further away from the origin of the y_1y_2 plane. For example, the thick grey line in figure 7 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). Owing 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 5*a*,*b*. Also note in figure 7 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.

(b) *Case 2: all accumulators receive equal noise*

In the previous case, accumulators $y_3, \ldots, y_N$ did not receive any input. This assumption is somewhat unrealistic, as the input neurons have a *spontaneous* variable firing rate even if the stimulus does not provide any evidence (Britten *et al.* 1993). In the electronic supplementary material, we extend the previous analysis to the case where all the accumulators are subject to the same level of noise, i.e. $c_1=\cdots=c_N$; as in the previous case we assume that only two accumulators receive positive input, i.e. $I_3=\cdots=I_N=0$. We show that if I_1 and I_2 are sufficiently high, then the conclusion from the previous case generalizes to this case, and the bounded LCA model outperforms the unbounded model. However, if I_1 and I_2 are lower, the performance of the bounded LCA model decreases below that of the unbounded model, but the performance of the bounded LCA model can be improved by increasing inhibition relative to decay.

(c) *Case 3: biologically realistic input parameters for choice with continuous variables*

We assumed previously that only two integrators receive input while the others received none: $I_3=\cdots=I_N=0$. However, in many situations, it might be expected that there is a 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 organization 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 within which stimuli are positioned at various points within 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 distribution function (Snowden *et al.* 1992)

$$I_i = r_{\min} + (r_{\max} - r_{\min})\exp\left(-\frac{d_i^2}{2\sigma}\right), \tag{3.1}$$

where $r_{\min}$ and $r_{\max}$ denote the minimum and the maximum firing rates of the neuron, respectively, and σ describes the width of the tuning curve. In our simulation, we use the parameter values given by Snowden *et al.* (1992), as shown in figure 8*a*.

We made two other changes for the current simulations. First, we assumed in the previous cases that all the accumulators receive the same level of noise, and furthermore that the noise magnitude is independent of the input level. However, a number of studies have shown that the variance in the neuronal firing rate of visual neurons (including neurons in area MT providing inputs in the motion discrimination task) is approximately proportional to the mean firing rate (for review, see Shadlen & Newsome (1998)). On the basis of these studies, Shadlen & Newsome (1998) proposed that a typical relationship between the mean and the variance of the inputs is $c_i^2 \approx 1.5 I_i$, so here we test the performance of the bounded and unbounded LCA models for the levels of noise chosen in this way.

Second, in all simulations described so far, the input to accumulator i at each time-step of numerical integration $I_i\,\mathrm{d}t + c_i\,\mathrm{d}W_i$ was generated from the normal distribution (figure 2) and hence could be negative. The neurons providing input to the accumulators cannot have negative firing rates, hence in the simulations described here, we additionally made the input equal to zero if it was negative on a given integration step.

Figure 8*b* shows the DTs under the assumptions described. 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 approximately 25%. Figure 8*b* also shows that the performance of the nonlinear LCA model of equation (2.7) with sigmoid activation function closely approximates the performance of the bounded LCA (recall from §2*f* that the sigmoid activation function, like the

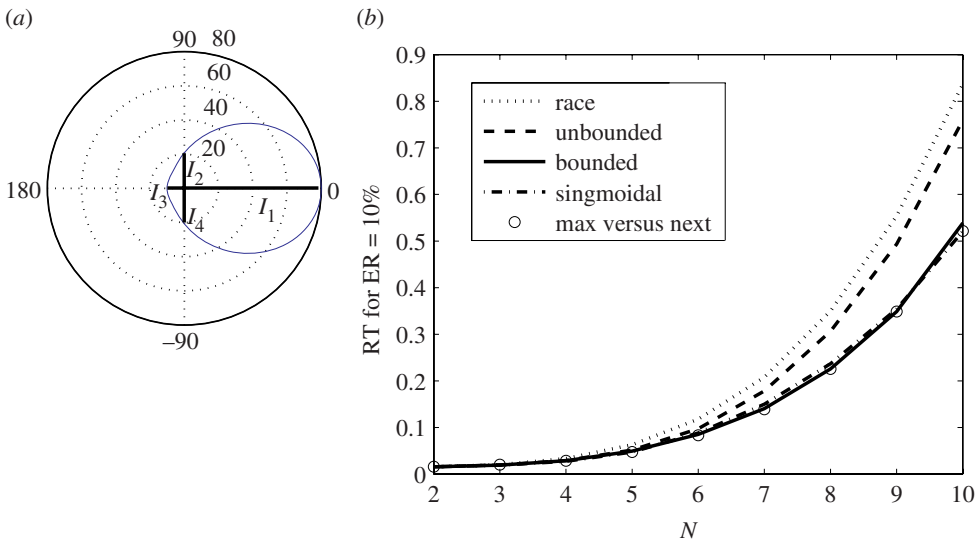

Figure 8. Simulation of the motion discrimination task. (*a*) Tuning curve describing the simulated firing rate of MT neurons (network inputs) as a function of angular difference, d_i, between the direction of coherent motion in the stimulus and neurons' preferred direction. We used the following parameters of the tuning curve: $\sigma = 46.5°$ (the average value over tuning curves fitted to 30 MT neurons by Snowden *et al.* (1992)), $r_{\min} = 10$ Hz, $r_{\max} = 80$ Hz (values from the neuron in right panel of fig. 7 in Snowden *et al.* (1992)). 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 $d_i = 360°(i-1)/N$, then if $d_i > 180°$, we made $d_i = d_i - 360°$, and then we computed I_i from equation (3.1). (*b*) Decision time with a threshold resulting in ER of 10% of different models as a function of the number of alternatives, N (shown on x-axis). Five models are shown: the race model, the unbounded LCA model, the bounded LCA model, the nonlinear LCA model with the sigmoid activation function and max versus next (see key). Methods of simulations are as in figure 5. The parameters of the LCA model are equal to $w = k = 10$. The mean inputs were chosen as described for (*a*), and the standard deviations were chosen as $c_i = \sqrt{1.5 I_i}$ (Shadlen & Newsome 1998). The following scaled sigmoid function was used in the simulation of the nonlinear LCA model: $f(y) = 10/(1 + \exp(-4(y/10 - 0.5)))$.

boundaries, prevents the activity levels of accumulators from decreasing below zero).

Furthermore, figure 8*b* also shows that the performance of the bounded LCA model, and of the nonlinear LCA model with sigmoid activation function, is very close to that of the max-versus-next procedure (that may implement an asymptotically optimal test; see §2*e*). In summary, this simulation shows that introduction of the biologically realistic assumption that 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 closely approximate the optimal performance.

4. OPTIMIZATION OF PERFORMANCE OF BOUNDED LEAKY COMPETING ACCUMULATOR 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 & McClelland 2001). However, a more complex assumption regarding the process that terminates decisions in the interrogation paradigm is also possible. As suggested by Ratcliff (1988), a response criterion may still be in place (as in the free-response paradigm) and participants use this response criterion (again as in the free-response paradigm) so that when the activation reaches this criterion, they make a preliminary decision (and stop integrating input). Accordingly, there are two types of trials: (i) those that reach criterion (as mentioned previously) and (ii) 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 with the introduction of absorbing upper boundaries 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 (Ratcliff 1988; Mazurek *et al.* 2003). 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 the interrogation paradigm, the time courses of average responses from the population of LIP neurons stop increasing after a certain period following the stimulus onset, and are maintained until the interrogation time (Roitman & Shadlen 2002).

In §2*e*, we showed that the unbounded LCA model achieves optimal performance when the decay is equal to inhibition. The following question then arises: does the balance of decay and inhibition still optimize 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 9 shows the ER of bounded LCA model for $N = 2$ alternatives. To make the position of the attracting line stable (cf. equation (2.8)), we fixed parameters $w + k$ but varied $w - k$. The results illustrate that by decreasing inhibition relative to decay the bounded model can achieve lower ER in the interrogation

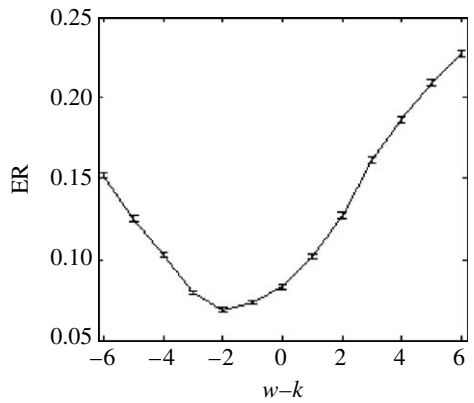

Figure 9. The ER of the bounded LCA model in the interrogation paradigm. The models were simulated with parameters: $I_1=5.414$, $I_2=4$, $c_1=c_2=0.8$, the height of the boundary $B=1.4$, interrogation time $T=2.5$. The sum of decay and inhibition was fixed at $w+k=6$, while their difference changed from -6 to 6. Data are averaged from 10 000 trials.

paradigm. This happens because in this case, there is an attracting point to which the state of the model is attracted, as shown in figure 2*a*, 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 minimized. In summary, in the interrogation paradigm, in contrast to the unbounded model, a balance of decay and inhibition does not optimize ER for the bounded model. Instead, optimal performance within the tested range of parameters is achieved when inhibition is smaller than decay.

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 judgments. 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 many theories have been accumulated (Kahneman & Tversky 2000), more recently, neurophysiological studies of value-based decisions have also been carried on behaving animals (Platt & 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.

(a) *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 among three laptops, which vary in their properties as described on 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 a set of contextual factors.

(b) *Violations of expected value and preference reversals*

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

Second, the preference between alternatives depends on a reference, which corresponds to either the present state of the decision maker, or even an *expected* state, which is subject to manipulation. Consider, for example, the following situation (figure 10*a*). When offered a choice between two job alternatives A and B, described on two dimensions (e.g. distance from home and salary), to replace a 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 & 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 in the *similarity*, the *attraction* and the *compromise* effects.

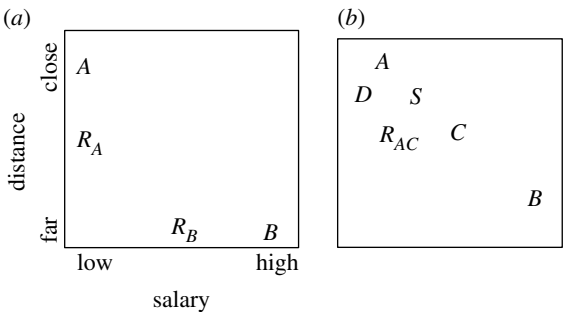

Figure 10. Configurations of alternatives in the attribute space. In each panel, the two axes denote two attributes of the alternatives (sample attributes' labels are given in (*a*)). The capital letters denote the positions of the alternative choices in the attribute space, while letters R_i denote the reference points. (*a*) Reference effect in multi-attribute decision making (after Tversky & Kahneman 1991). (*b*) Contextual preference reversal: the similarity, attraction and compromise effects. Alternatives *A*, *B*, *C* and *S* lie on the indifference line.

To illustrate these effects, consider a set of options, *A*, *B*, *C* and *S*, which are characterized by two attributes (or dimensions) and located on a decision maker's indifference curve: the person is of equal preference on a choice between any two of these options (figure 10*b*). The similarity effect is the finding that the preference between *A* and *B* can be modified in 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* and *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 & Shafir 2005). One notable exception is the work by Tversky and colleagues, who developed a mathematical, context-dependent advantage model that accounts for reference effects and preference reversal in multi-dimensional choice (Tversky & Simonson 1993). 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 & Townsend 1993). More recently, Usher & 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 LCA models not only share many principles but also differ on some. While DFT is a linear model (where excitation by negated inhibition, of the type described in §2, is allowed) and 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 §3), while the second one involves a convex utility–value function (Kahneman & Tversky 2000).

It is beyond the scope of this paper to compare detailed predictions of the two models (but see Usher & 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, in the following one, we show how a simple LCA-type model, taking these principles on board, can address value-based decisions.

(c) *Nonlinear utility functions and the Weber law*

The need for a nonlinear relation between internal utility and objective value was noticed by 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, Glimcher (2004), pp. 188–192 and Martin (2004) for good descriptions of the paradox and of Bernoulli's solution.) Owing 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 it comes up 'heads'. 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 ($E = \sum_{i=1,\ldots,\infty}(1/2)^i 2^i = \sum_{i=1,\ldots,\infty} 1 = \infty$), 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% win less than £5), but in a few cases one can win a lot. Paying a large amount for playing the game results in a high probability of making a loss and a small probability for 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 & Kahneman 1979). A logarithmic value function $u(x) = \log(x)$, used as the expected utility, gives a value of approximately £4 for the St Petersburg

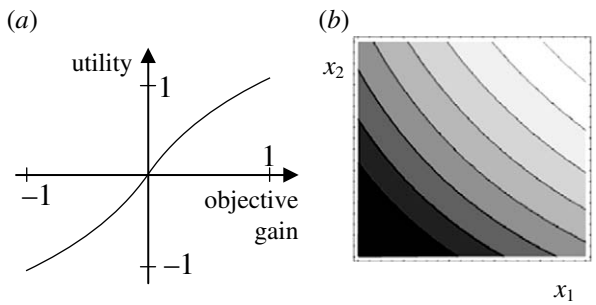

Figure 11. (*a*) Utility function, $u(x)=\log(1+kx)$, for $x>0$ and $-\gamma\log(1-kx)$, for $x<0$. ($k=2$, $\gamma=1$). (*b*) Combined two-dimensional utility function for gains ($x_1>0$, $x_2>0$).

game. A more complex version of the game and resulting paradox described in the electronic supplementary material.

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 (figure 10) where the options are characterized 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 which 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 $x+\mathrm{d}x$, the just noticeable difference, $\mathrm{d}x$, is proportional to x itself.

One simple way to simultaneously satisfy the Weber law and the Bernoulli diminishing return intuition 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 c^2, as shown in the electronic supplementary material.

As proposed by Bernoulli, a logarithmic nonlinearity accounts simultaneously for risk aversion and the Weber law. 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 one in the logarithm corresponds to a baseline of present value before any gains or losses are received). (In prospect theory (Tversky & Simonson 1993; Kahneman & Tversky 2000), one chooses $\gamma>1$, indicating a higher slope for losses than for gains. This is also assumed by Usher & McClelland (2004). Here, we use $\gamma=1$ in order to explore the simplest set of assumptions that can result in these reversals effects; increasing γ strengthens the effects.) As shown in figure 11*a*, function $u(x)$ starts linearly and then is subject to diminishing returns, which is a good approximation to neuronal input–output response function of neurons at low to intermediate firing rates (Usher & Niebur 1996). While neuronal firing rates eventually saturate, it is possible that a logarithmic dependency exists on a wide range of gains and losses, with an adaptive baseline and range (Tobler *et al.* 2005).

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), $U(x_1, x_2)=u(x_1)+u(x_2)=\log[1+k(x_1+x_2)+k^2x_1x_2]$. Note that to maximize this utility function one has to maximize a combination of linear and multiplicative terms. The inclusion of a multiplicative term in the utility optimization is supported by a survival rationale: to survive, animals need to ensure the joined (rather than separate) possession of essential resources (like food and water). Figure 11*b* shows a contour plot of this two-dimensional 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 §5*b*, the reference depends on the subjective expectations and on the information accessible to the decision maker (Kahneman 2003). In §5*d*, the combination of nonlinear utility and reference dependence explains the presence of contextual preference reversals. Finally, when choice alternative are characterized over multiple dimensions, we assume (following Tversky's (1972) elimination by aspects and the various DFT applications; Busemeyer & 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 regards to one of the dimensions and the preference is integrated by the LCAs. In §5*d*, these components of utility evaluations are introduced into an LCA model and applied to the value-based decision patterns described previously.

(d) *Modelling value-based choice in the LCA framework*

To allow for the switching between the alternatives–dimensions, the LCA simulations are done using a discretized version of the LCA model of equation (2.2) (single step of Euler method; note a threshold nonlinearity at zero is imposed: only $y_i>0$ are allowed)

$$y_i(t+\Delta t)=y_i(t)+\Delta t\left(-ky_i-w\sum_{\substack{j=1\\ j\neq 1}}^{N}(y_j)+I_i+I_0+\text{noise}\right), \tag{5.1}$$

where I_i is evaluated according to the utility function described previously 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 s.d. 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 shown in

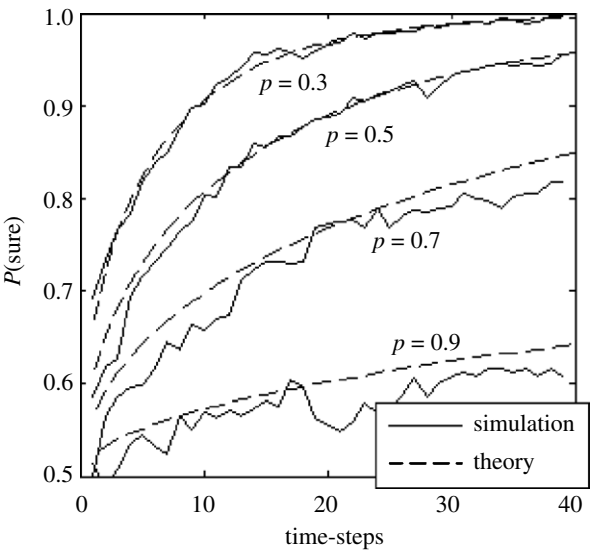

Figure 12. Probability of choosing the sure option as a function of deliberation time for five values of risk (indicated in the plot). Solid lines were obtained from simulations of the LCA model for the following parameters: $W=1$, $I_0=0.6$, s.d. $=0.5$, and the utility function from figure 11. Dashed lines come from the equation derived in the electronic supplementary material.

figure 10), 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; Usher & McClelland 2004) via an interrogation-like procedure. The response units are allowed to accumulate their preference evaluation for T time-steps. Five-hundred 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 demonstration of risk aversion in probabilistic monetary choice and then we turn to preference reversals in multidimensional choice.

(i) *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 corresponds to a probabilistic win of W/p, with probability p (note that the two have an equal expected objective value, W, and that p provides a measure of risk: lower p are 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 (s.d. $=0.5$) 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 12, we show the probability to choose 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 probable the bias of choosing the sure option (this bias starts at a value approximately proportional to $1-p$ and increases due to time integration to an asymptotic value). This is consistent with experimental data, except for low p where, as explained by the prospect theory (Tversky & 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). In the electronic supplementary material, the probability time-curve is approximated analytically, and it is shown that it increases with the square root of deliberation time until it saturates. Risk seeking for losses can be simulated analogously.

(ii) *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 10*a*, 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 13*a*) that the R_A reference increases the probability to choose 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 $u(0)+u(0.2)=\mathrm{u}(0.2)>u(0.6)-u(0.4)=u(0.6)+u(-0.4)$) 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 did 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 10*b*. 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 13*b*) that the compromise alternative is preferred among the three. This is a direct result of two-dimensional utility function (figure 11*b*). 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 & McClelland 2004). This maintains an equal preference between C and the

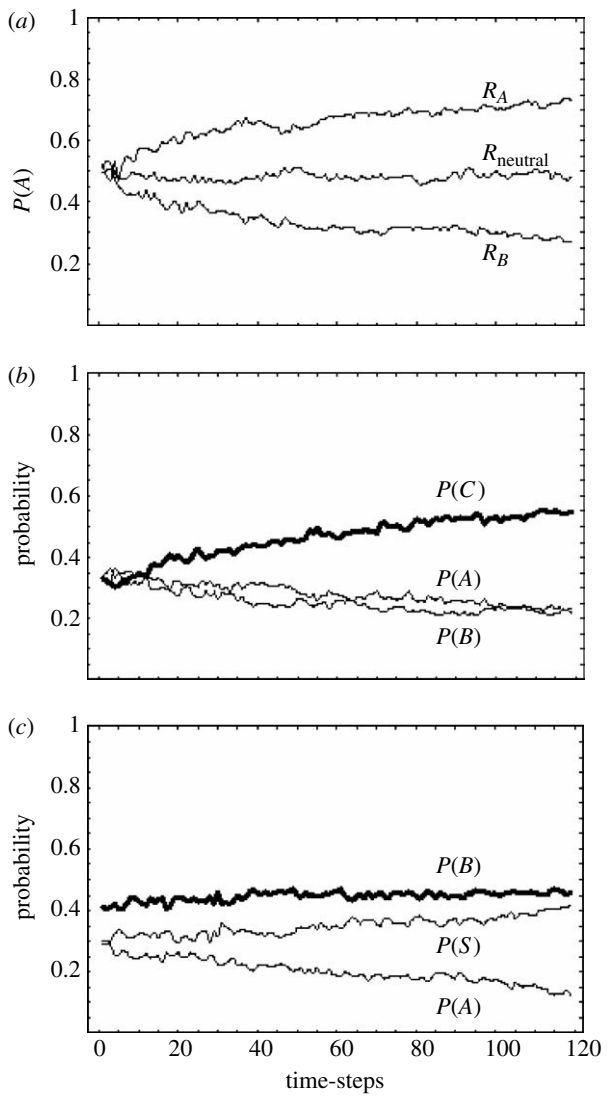

Figure 13. Contextual preference reversal. (*a*) Reference effects in binary choice. (*b*) Compromise effect. (*c*) Similarity effect.

extremes in binary choice. Note also the dynamics of the compromise effect. This takes time to develop; for short times, the preference is larger for the extremes, depending on the dimension evaluated first. Experimental data indicate 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 13*c*, 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 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 reference for every other option; Tversky & Simonson 1993; Usher & McClelland 2004.)

To summarize, 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 & Kahneman 1979), displaying risk aversion (it prefers the sure option to 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. DISCUSSION

In this paper, 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 amount of supporting evidence in the free-response paradigm. Moreover, we have shown that for choices involving multiple alternatives in the free-response paradigm, the nonlinearities of type present in biological decision networks can improve the performance, and in fact may allow the networks to closely approximate the optimal choice algorithm. This raises an intriguing possibility, that these nonlinearities are not just a result of biological constraints, but may rather be a result of evolutionary pressure for speed of decision.

We have also identified conditions (see §§3*b* and 4), in which the performance can be optimized by an elevation/decrease in the level of lateral inhibition relative to the leak (this may be achieved via neuromodulation; Usher & Davelaar 2002). It will be interesting to test whether the behavioural manifestations of unbalance of decay and inhibition (Usher & 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 & Krebs 1974). In the simulations of the bounded LCA model in figure 5*a*,*b*, the DT does not depend on the number of potentially available alternatives. Note, however, this simulation was designed to model the task described in the beginning of §3 (figure 6) 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 the accumulators received an equal level 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 a linear range, it is equivalent to the linear model, and the linear model

satisfies Hick's law when accumulators receive an equal level of noise (McMillen & Holmes 2006). However, it has recently been 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 recently been 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 & 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 the MSPRT algorithm prevents an obvious cortical implementation; hence it is of great interest to investigate the parameters optimizing the LCA model describing the 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 issue of *Philosophical Transactions*. In particular, it is very similar to the model of action selection in the cerebral cortex of Cisek (2007) which also includes accumulation of evidence and competition between neuronal populations corresponding to different alternatives.

This work was supported by EPSRC grant EP/C514416/1. We thank Andrew Lulham for reading the manuscript and very useful comments. MATLAB codes for simulation and finding DT of the LCA model are included in the electronic supplementary material.

REFERENCES

Barnard, G. 1946 Sequential tests in industrial statistics. *J. R. Stat. Soc. Suppl.* **8**, 1–26. (doi:10.2307/2983610)

Bernoulli, D. [1738] 1954 Exposition of a new theory on the measurement of risk. *Ekonometrica* **22**, 23–36. (doi:10.2307/1909829)

Bogacz, R. & Gurney, K. 2007 The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput* **19**, 442–477.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J. D. 2006 The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced choice tasks. *Psychol. Rev.* **113**, 700–765.

Britten, K. H., Shadlen, M. N., Newsome, W. T. & Movshon, J. A. 1993 Responses of neurons in macaque MT to stochastic motion signals. *Vis. Neurosci.* **10**, 1157–1169.

Brown, E. & Holmes, P. 2001 Modeling a simple choice task: stochastic dynamics of mutually inhibitory neural groups. *Stoch. Dyn.* **1**, 159–191. (doi:10.1142/S0219493701000102)

Brown, E., Gao, J., Holmes, P., Bogacz, R., Gilzenrat, M. & Cohen, J. D. 2005 Simple networks that optimize decisions. *Int. J. Bifurc. Chaos* **15**, 803–826. (doi:10.1142/S0218127405012478)

Busemeyer, J. R. & Townsend, J. T. 1993 Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* **100**, 432–459. (doi:10.1037/0033-295X.100.3.432)

Busemeyer, J. R., Townsend, J. T., Diederich, A. & Barkan, R. 2005 Contrast effects or loss aversion? Comment on Usher & McClelland (2004). *Psychol. Rev.* **111**, 757–769.

Cisek, P. 2007 Cortical mechanisms of action selection: the affordance competition hypothesis. *Phil. Trans. R. Soc. B* **362**, 1585–1599. (doi:10.1098/rstb.2007.2054)

Dhar, R., Nowlis, S. M. & Sherman, S. J. 2000 Trying hard or hardly trying: an analysis of context effects in choice. *J. Consum. Psychol.* **9**, 189–200. (doi:10.1207/S15327663JCP0904_1)

Dragalin, V. P., Tertakovsky, A. G. & Veeravalli, V. V. 1999 Multihypothesis sequential probability ratio tests—part I: asymptotic optimality. *IEEE Trans. Inf. Theory* **45**, 2448–2461. (doi:10.1109/18.796383)

Glimcher, P. W. 2004 *Decisions, uncertainty, and the brain: the science of neuroeconomics*. Cambridge, MA: MIT Press.

Gold, J. I. & Shadlen, M. N. 2002 Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron* **36**, 299–308. (doi:10.1016/S0896-6273(02)00971-6)

Hacking, I. 1980 Strange expectations. *Phil. Sci.* **47**, 562–567. (doi:10.1086/288956)

Hertwig, R., Barron, G., Weber, E. U. & Erev, I. 2004 Decisions from experience and the effect of rare events in risky choice. *Psychol. Sci.* **15**, 534–539. (doi:10.1111/j.0956-7976.2004.00715.x)

Houston, A. I., McNamara, J. & Steer, M. 2007 Do we expect natural selection to produce rational behaviour? *Phil. Trans. R. Soc. B* **362**, 1531–1543. (doi:10.1098/rstb.2007.2051)

Kahneman, D. 2003 Maps of bounded rationality: psychology for behavioral economics. *Am. Econ. Rev.* **93**, 1449–1475. (doi:10.1257/000282803322655392)

Kahneman, D. & Tversky, A. (eds) 2000 *Choices, values and frames*. Cambridge, UK: Cambridge University Press.

Kveraga, K., Boucher, L. & Hughes, H. C. 2002 Saccades operate in violation of Hick's law. *Exp. Brain Res.* **146**, 307–314. (doi:10.1007/s00221-002-1168-8)

Laming, D. R. J. 1968 *Information theory of choice reaction time*. New York, NY: Wiley.

LeBoef, R. & Shafir, E. B. 2005 Decision-making. In *Cambridge handbook of thinking and reasoning* (eds K. J. Holyoak & R. G. Morisson). Cambridge, UK: Cambridge University Press.

Martin, R. 2004 The St. Petersburg paradox. In *The Stanford encyclopedia of philosophy* (ed. E. Zalta). Stanford, CA: Stanford University.

Mazurek, M. E., Roitman, J. D., Ditterich, J. & Shadlen, M. N. 2003 A role for neural integrators in perceptual decision making. *Cereb. Cortex* **13**, 1257–1269. (doi:10.1093/cercor/bhg097)

McMillen, T. & Holmes, P. 2006 The dynamics of choice among multiple alternatives. *J. Math. Psychol.* **50**, 30–57. (doi:10.1016/j.jmp.2005.10.003)

Neyman, J. & Pearson, E. S. 1933 On the problem of the most efficient tests of statistical hypotheses. *Phil. Trans. R. Soc. A* **231**, 289–337. (doi:10.1098/rsta.1933.0009)

Platt, M. L. & Glimcher, P. W. 1999 Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238. (doi:10.1038/22268)

Ratcliff, R. 1978 A theory of memory retrieval. *Psychol. Rev.* **83**, 59–108. (doi:10.1037/0033-295X.85.2.59)

Ratcliff, R. 1988 Continuous versus discrete information processing: modeling accumulation of partial information. *Psychol. Rev.* **95**, 238–255. (doi:10.1037/0033-295X.95.2.z238)

Roe, R. M., Busemeyer, J. R. & Townsend, J. T. 2001 Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol. Rev.* **108**, 370–392. (doi:10.1037/0033-295X.108.2.370)

Roitman, J. D. & Shadlen, M. N. 2002 Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* **22**, 9475–9489.

Schall, J. D. 2001 Neural basis of deciding, choosing and acting. *Nat. Rev. Neurosci.* **2**, 33–42. (doi:10.1038/35049054)

Seung, H. S. 2003 Amplification, attenuation, and integration. In *The handbook of brain theory and neural networks* (ed. M. A. Adbib), pp. 94–97, 2nd edn. Cambridge, MA: MIT Press.

Shadlen, M. N. & Newsome, W. T. 1998 The variable discharge of cortical neurons: implications for connectivity, computation, and information coding. *J. Neurosci.* **18**, 3870–3896.

Shadlen, M. N. & Newsome, W. T. 2001 Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936.

Slovic, P. 1995 The construction of preference. *Am. Psychol.* **50**, 364–371. (doi:10.1037/0003-066X.50.5.364)

Snowden, R. J., Treue, S. & Andersen, R. A. 1992 The response of neurons in areas V1 and MT of the alert rhesus monkey to moving random dot patterns. *Exp. Brain Res.* **88**, 389–400. (doi:10.1007/BF02259114)

Stone, M. 1960 Models for choice reaction time. *Psychometrika* **25**, 251–260. (doi:10.1007/BF02289729)

Sugrue, L. P., Corrado, G. S. & Newsome, W. T. 2004 Matching behavior and the representation of value in the parietal cortex. *Science* **304**, 1782–1787. (doi:10.1126/science.1094765)

Sugrue, L. P., Corrado, G. S. & Newsome, W. T. 2005 Choosing the greater of two goods: neural currencies for valuation and decision making. *Nat. Rev. Neurosci.* **6**, 363–375. (doi:10.1038/nrn1666)

Teichner, W. H. & Krebs, M. J. 1974 Laws of visual choice reaction time. *Psychol. Rev.* **81**, 75–98. (doi:10.1037/h0035867)

Tobler, P. N., Fiorillo, C. D. & Schultz, W. 2005 Adaptive coding of reward value by dopamine neurons. *Science* **307**, 1642–1645. (doi:10.1126/science.1105370)

Tversky, A. 1972 Elimination by aspects: a theory of choice. *Psychol. Rev.* **79**, 281–299. (doi:10.1037/h0032955)

Tversky, A. & Kahneman, D. 1979 Prospect theory: an analysis of decision under risk. *Econometrica* **47**, 263–292. (doi:10.2307/1914185)

Tversky, A. & Kahneman, D. 1991 Loss aversion in riskless choice: a reference-dependent model. *Q. J. Econometr.* **106**, 1039–1061. (doi:10.2307/2937956)

Tversky, A. & Simonson, I. 1993 Context-dependent preferences. *Manage. Sci.* **39**, 1179–1189.

Usher, M. & Davelaar, E. J. 2002 Neuromodulation of decision and response selection. *Neural Netw.* **15**, 635–645. (doi:10.1016/S0893-6080(02)00054-0)

Usher, M. & McClelland, J. L. 2001 The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592. (doi:10.1037/0033-295X.108.3.550)

Usher, M. & McClelland, J. L. 2004 Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol. Rev.* **111**, 759–769.

Usher, M. & Niebur, N. 1996 Modeling the temporal dynamics of IT neurons in visual search: a mechanism for top-down selective attention. *J. Cogn. Neurosci.* **8**, 311–327.

Vickers, D. 1970 Evidence for an accumulator model of psychophysical discrimination. *Ergonomics* **13**, 37–58.

Vickers, D. 1979 *Decision processes in perception*. New York, NY: Academic Press.

Wald, A. 1947 *Sequential analysis*. New York, NY: Wiley.

Wang, X. J. 2002 Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968. (doi:10.1016/S0896-6273(02)01092-9)