Chapter 13

The neurodynamics of choice, value-based decisions, and preference reversal

Marius Usher

School of Psychology, Birkbeck, University of London, London, UK

Anat Elhalal

School of Psychology, Birkbeck, University of London, London, UK

James L. McClelland

Department of Psychology, Stanford University, Palo Alto, CA, USA

A theory of choice is paramount in all the domains of cognition requiring behavioural output, from perceptual choice in simple psychophysical tasks to motivational *value*-based choice, often labelled as *preferential* choice and which is exhibited in daily decision-making. Until recently, these two classes of choice have been the subject of intensive but separate investigations, within different disciplines. Perceptual choice has been investigated mainly within the experimental psychology and neuroscience disciplines, using rigorous psychophysical methods that examine behavioural accuracy, response latencies (choice-RT), and neurophysiological data (Laming, 1968; Ratcliff & Smith, 2003; Usher & McClelland, 2001; Vickers, 1979). Preferential choice, such as when one has to choose an automobile among a set of alternatives that differ in terms of several attributes or dimensions (e.g., quality and economy) has been investigated mainly within the economics and the social science disciplines, using mainly reports of choice preference. Unlike in perceptual choice, where the dominant models are *process* models that approximate *optimality* (Bogacz *et al.*, 2007; Gold & Shadlen, 2002) based on the Sequential Probability Ratio Test (SPRT; Barnard, 1946; Wald, 1947), the literature on preferential choice has emphasised a series of major deviations from *normativity* (Huber *et al.*, 1982; Kahneman & Tversky, 1979, 2000; Knetch, 1989; Simonson, 1989; Tversky, 1972; Tversky & Kahneman, 1991). This has led to the proposal that decision-makers use a set of disparate heuristics, each addressing some other aspect of these deviations

(LeBoef & Shafir, 2005).[1] Because of this, it is difficult to compare the two types of theories and to discuss their implications for issues related to optimality and to principles of rational choice.

Two recent types of research are providing, however, the opportunity to bridge this gap. First, neurophysiological studies of value-based decisions have been carried out on behaving animals (Glimcher, 2004; Sugrue *et al.*, 2004, 2005). Second, a series of neurocomputational models of preferential choice have been developed, which address the *process* by which preferential choice is issued (Steward, this volume; Roe *et al.*, 2001; Usher & Mclelland, 2004). It is the synthesis of these two lines of work that defines the new field of *neuroeconomics,* whose central aim is to understand the principles that underlie value-based decisions and the neural mechanisms through which these principles are expressed in behaviour. A major insight of neuroeconomics is that these principles can be modelled as implementing an optimising solution to some survival/reproductive challenge in the evolutionary environment (Wikipedia).

The aim of this chapter is to review some of the neurocomputational work by contrasting the various processing assumptions and the way they account for one of the most intriguing patterns in the choice data: *contextual preference-reversal.* We start with introducing a very simple process model of choice, called the leaky competing accumulator (LCA) model (Usher & McClelland, 2001), which has been developed to account for both perceptual and preferential choice data. We will then review some of the data on preference reversal, and on framing effects, which are the main targets of the neurocomputational theories discussed later. Then we will discuss two such theories, the decision-field theory (DFT) developed by Busemeyer and colleagues (Busemeyer & Diederich, 2002, Busemeyer & Johnson, 2003; Diederich, 1997; Roe *et al.*, 2001) and an extension of the LCA. We examine some similarities and contrasts between the models that lead to a set of experimental predictions. Finally, we present experimental data aimed at testing these predictions and consider their implications for the issue of rationality in choice.

Perceptual-choice, optimality, and the LCA model

Consider a simple task that illustrates a typical example of perceptual choice. In this task, the observers (humans or 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 observer's task is to indicate the direction of prevalent motion. The time of the response is either 'up to the observer' (and is measured as a function of the stimulus property and the response accuracy) or is controlled by a response signal (and the accuracy is measured as a function of the observation time). This task presents the basic demand the observer has to accomplish for performing the task: decide which of

[1] One notable exception is work by Tversky and colleagues (e.g., Tversky, 1972; Tversky & Simonson, 1993), who developed a number of mathematically rigorous process models.

a number of options match best a noisy signal. As the noise varies in time, a good way to perform the task is to integrate the evidence from sensory neurons over time (this is consistent with the neural data; Schall, 2001; Shadlen & Newsome, 2001). This integration averages out the noise, allowing the accuracy of the choice to increase with time. The quality of the decision (measured in terms of accuracy per time-taken to make the response, or response-rate) depends on the response-rule and the way in which the evidence is integrated.

A variety of models have been proposed to account for rich data patterns that involve not only accuracy and mean-RTs but also RT-distributions (Smith & Ratcliff 2004). Here we will focus on one such model, the LCA (Usher & McClelland, 2001), which is framed at a neural level and yet is simple enough so that it can be mathematically analysed and its parameters can be manipulated so as to optimise performance. As the main aim of this chapter is to understand value-based decisions, we will only present a brief summary here, but see Bogacz *et al.* (2007) for a detailed and rigorous account.

The model assumes that for each choice option there is a response unit that accumulates evidence in favour of this option. Figure 13.1 illustrates the model, for the case of binary choice. The accumulation of the evidence (or activation) is subject to temporal decay (or leak) and the various response units compete via lateral (and mutual) inhibition. The decay and the lateral inhibition (k and w) are the important parameters that affect the network's choice pattern and performance. In addition each unit's activation is truncated at zero, reflecting the biological constraint that activations correspond to neuronal firing rates, which cannot be negative. This is also reflected in the use of the threshold-linear output function, $f(x)$. Mathematically this can be formulated and simulated as:

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

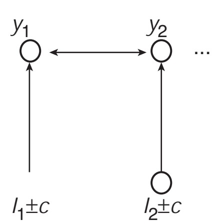

Fig. 13.1. LCA model for binary choice. The variables y_1 and y_2 accumulate noisy evidence (I_1, I_2). The units compete due to lateral inhibition (w). To generalize for n-choice, one adds units, which compete with each other by lateral (all to all) inhibition (Usher & McClelland, 2001).

where the final term (*dW*) corresponds to the differential of a (stochastic) Wiener process (Gaussian noise of amplitude *c*). Finally, the response is issued when the first of the accumulators reaches a common response criterion. This criterion is assumed to be under the control of the observer, reflecting speed-accuracy tradeoffs.

The model's choice pattern depends essentially on the relative values of the decay and the inhibition parameters. When decay exceeds inhibition the model exhibits recency (it favours information that arrives late in time), but when inhibition exceeds decay, it exhibits primacy. Neither of these results produces the highest accuracy possible. However, when decay and inhibition are balanced (decay = inhibition) the network performs optimally. Indeed, in this case the model's behaviour mimics the optimal Sequential Probability Ratio Test (SPRT). This means that among all possible procedures for solving this choice problem, it minimizes the average decision time (DT) for a given error rate (ER).

One advantage of this neural model over other (more abstract) models such as the random-walk, is that it is easy to generalise to choice among more than two alternatives. In this case, all the units representing the alternative choices race towards a common response criterion, while they all simultaneously inhibit each other. This mutual inhibition allows the network choice to depend on relative evidence, without the need for a complex readout mechanism that would depend on computing the difference between the activations of the most active unit and the next most active. Moreover, as discussed in detail in Bogacz *et al.* (2007), the threshold-nonlinearity (the truncation of negative activations at zero) is critical in maintaining this advantage when the input favours only a small subset of a larger ensemble of options. The reason for this is that with the threshold-nonlinearity, the activation of all the non-relevant units is maintained at zero, rather than becoming negative and sending uninformative (positive) input into the relevant competing units (see Bogacz *et al.*, 2007). This feature is important in our application of the LCA model to value-based choice.

Value-based decisions, preference-reversal, and reference effects

Although we assume that both perceptual and value/motivational decisions involve a common selection mechanism, the basis on which this selection operates differs. The aim of this section is to review the underlying principles of value-based decisions, focusing on a series of puzzling patterns of preference reversal in multi-attribute choice (see Stewart & Simpson, this volume, for a discussion of risky choice), which raise challenges for an optimal or normative theory of choice.

Unlike in perceptual choice, the preferential choice cannot be settled on the basis of perceptual information alone. Rather, each alternative 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 flats, which vary on their properties as described on a number of dimensions (number of rooms, price, distance from work, etc).

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 preferential choice, 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.

First, the preference between alternatives depends and may reverse depending on a reference, which corresponds either to the present state of the decision-maker, or to an *expected* state, which is subject to manipulation. Consider, for example the following situation (Fig. 13.2a). 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) participants prefer the option that is more similar to the reference (Tversky & Kahneman, 1991).

Second, it has been shown that the preference between two options can be reversed 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 (Tversky, 1972; Simonson, 1989; Huber *et al*., 1982). To illustrate these effects consider a set of options, *A*, *B*, *C*, and *S*, which are characterized 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 (Fig. 13.2b). 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

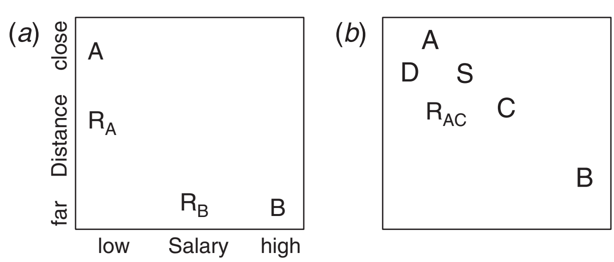

Fig. 13.2. Configurations of alternatives in the attribute space. In each panel the two axes denote two attributes (sample attributes' labels are given in panel a). Capital letters denote the positions of the alternatives in the attribute space, while letters Ri denote the reference points. (a) Reference effect in multi-attribute decision-making (after Tversky & Kahneman, 1991). (b) Contextual preference reversal: similarity, attraction and the compromise effects. Alternatives A, B, C, S lie on the indifference line.

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.

Third, a large number of studies indicate that human (and animal) observers exhibit *loss-aversion*. A typical illustration of loss-aversion was shown in the following study (Knetch, 1989). Three groups of participants are offered a choice between two objects of roughly equal value (a mug and a chocolate bar), labelled here as A and B. One group is first offered the A-object and then the option to exchange it for the B-object. The second group is offered the B-object followed by the option to exchange it for A. The control group is simply offered a choice between the two objects. The results reported by Knetch (1989) are striking. Whereas the control participants chose the two objects in roughly equal fractions (56% vs. 44%), 90% of the participants in either of the groups that were first offered one of the objects prefer to keep it rather than exchange it for the other one (see also Samuelson & Zeckhauser, 1988). This effect is directly explained by Tversky and Kahneman by appealing to an asymmetric value-function, which is steeper in the domains of losses than in that of gains. Because losses are weighted more than gains, participants who evaluate their choices with the already-owned object serving as the reference point decline the exchange. [For the control participants the values may be computed either relative to the neutral reference (Tversky & Kahneman, 1991), or each option can be used as a reference for the other options (Tversky & Simonson, 1993); in both cases, there is no reference bias, consistent with the nearly equal choice fractions in this case.]

Computational models of preferential choice and of reversal effects

Recent work on neurocomputational models on preferential choice (Busemeyer & Diederich, 2002; Busemeyer & Johnson, 2003; Roe *et al.*, 2001; Stewart & Simpson, this volume; Usher & McClelland, 2004) may seem at odds with the emphasis on heuristics of choice that is the most common approach of the field.[2] We believe, however, that the two approaches are complementary (and we will discuss this further in the Discussion section), as best suggested by the pioneering research of Amos Tversky, who in addition to his work on heuristics, was one of the first to develop formal mathematical models for preferential choice. Two of his models are the *elimination by aspects* (*EBA*; Tversky, 1972), which accounts for the similarity effect, and the

[2] For example, our request for funding was rejected by the ESRC because one reviewer has forcefully and eloquently argued that neurocomputational models are inherently opaque and the better way to understand decision-making is via the use of heuristics.

context-dependent-advantage model, which accounts for the attraction and for the compromise effects (Tversky & Simonson, 1993).

Interestingly, however, the properties responsible for accounting for these effects have not been combined within a single model. Moreover, as observed by Roe *et al.* (2001), the context-dependent-advantage model cannot explain the preference reversals in similarity effect situations. The first unified account of all three reversal effects was proposed by Roe *et al.* (2001), using the DFT approach. 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. Our approach is, in fact, a direct combination of principles used by Tversky in his previous models, within a neurocomputational framework. Another computational model, closely related to the LCA, has been proposed by Stewart (this volume).

The DFT and LCA models

Both theories are implemented as connectionist models, in which the decision-maker integrates (with decay) a momentary *preference* towards a response criterion, for each choice-alternative. These theories built upon earlier work by Tversky (1972), in assuming that the decision-makers undergo a stochastic process of switching attention between dimensions or attributes. We start with a brief formulation of the DFT model and its account of preference reversal.

The model can be viewed as a *linear* neural network with four layers (Fig. 13.3). The first layer corresponds to the input attribute values, which feed via weights into units at level 2 that correspond to the two choice alternatives. An attentional mechanism stochastically selects between the attribute units (D1 and D2), so that only one attribute (determined randomly) provides input to level 2 at each time step. Level 3 computes valences for each by subtracting the average level 2 activation of the two other alternatives from its own level 2 activation. As the attention switches between the attributes, the valences vacillate from positive to negative values. Level 4 is the choice layer, which performs a leaky integration of the varying preference-input from level 3. Competition between the options occurs at level 4, mediated by bi-directional inhibitory connections with strengths that are assumed to be distance dependent: the strengths of the inhibitory connections among units in the fourth layer decrease as the distance between them, in the attribute space shown in Fig. 13.2b, increases.

The LCA model is illustrated in Fig. 13.4. It shares many principles with its DFT counterpart but also differs on some. In the DFT the strength of the lateral inhibition is similarity-dependent (stronger for similar alternatives), while in the LCA it is similarity-independent. Furthermore, the DFT is a linear model, where excitation by negated inhibition is allowed, while the LCA incorporates important non-linearities. First, the lateral inhibition is subject to a threshold at 0, so that negative activations do not produce excitation of other alternatives. Second, the LCA incorporates a convex and asymmetric utility-value function (Kahneman & Tversky, 2000).

Like the DFT model, LCA assumes a sequential and stochastic scan of the dimensions. Also like the DFT, the LCA assumes that the inputs to the choice units are (3^{rd} layer in

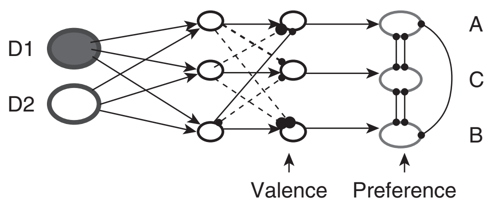

Fig. 13.3. The DFT model for reversal in multi-attribute choice (after Roe *et al.*, 2001). Solid arrows correspond to excitation and the open ones to inhibition. Option C has a stronger inhibition (double arrows) to options A and B due to the fact that the latter are more similar with C than with each other (see Fig. 13.2). At each moment the attention is focussed on one of the two attributes or dimensions.

Fig. 13.4) obtained via a pre-processing stage. In the case of the LCA, however, this involves the computation of relative differences between each option and each of the other options. Specifically, when faced with three alternative choices, participants evaluate the options in relation to each other in terms of gains or losses. Accordingly, the inputs, I, to the leaking accumulators are governed by:

$$I_1 = V(d_{12}) + V(d_{13}) + I_0; \quad I_2 = V(d_{21}) + V(d_{23}) + I_0; \quad I_3 = V(d_{31}) + V(d_{32}) + I_0.$$

where d_{ij} is the differential (advantage or disadvantage) of option i relative to option j, computed on the dimension currently attended; V is the nonlinear advantage function, and I_0 is a positive constant that can be seen as promoting the available alternatives into the choice set. The nonlinear advantage function is chosen (Tversky & Kahneman, 1991) to provide diminishing returns for high gains or losses, and aversion for losses relative to the corresponding gains.[3]

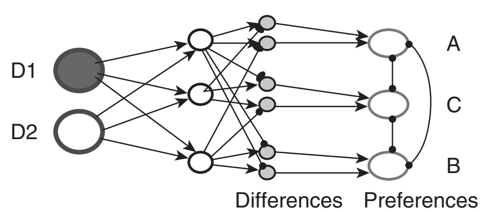

Figure 13.4. LCA model for a choice between three options characterised by two dimensions, D1 and D2. The solid arrows correspond to excitation and the open ones to inhibition. At every time step, an attentional system stochastically selects the activated dimension (D1 in this illustration). The input-item units in the second layer represent each alternative according to its weights on both of the dimensions and project into difference-input units in the 3rd layer. This layer converts the differences via an asymmetric nonlinear value function before transmitting them to choice units in the fourth layer.

[3] In Usher & McClelland (2004) the value function was chosen as: $v(x) = z(x)$ for $x > 0$ and $v(x) = -(z(|x|) + [z(|x|)]^2)$ for $x < 0$, where $z(x) = \log(1+x)$.

These models account for the similarity effect in the same way, and in a way that is similar to Tversky's and Kahneman's Elimination by Aspects (EBA) account: the preferences for the similar alternatives tend to rise as a result of scanning the dimension that supports them, and to fall together when attention is switched to the nonsupporting dimension. As shown in Fig. 13.5, the preferences for the similar alternatives are correlated in time and are anticorrelated with the preference of the dissimilar alternative; correlated alternatives tend to be active (or inactive) together. When they are active together, they have the opportunity to be chosen but they split these opportunities between them. On the other hand, the remaining, anti-correlated alternative is active out of phase with the other two alternatives and does not split its choice opportunities.

The two models offer different accounts of the attraction and compromise effects. While DFT relies on distant-dependent inhibition that decreases with the distance between the alternatives to produce these effects, LCA relies on *loss-aversion* (Tversky & Kahneman, 1991). According to the DFT, the attraction effect is explained as a contrast type effect: the similar options A and D have a strong inhibitory coupling because of the distance-dependent inhibition but this leads to a paradoxical boosting of the support for A because the dominated option D takes on a negative activation, which boosts the activation of A because the negative activation of D times the negative connection weight produces as positive input to A.

In LCA, both the attraction and the compromise effect are explained as a result of an asymmetric value function between gains and losses, according to which losses are weighted more than gains (and this weight asymmetry increases with the magnitude of the loss), consistent with loss-aversion. The LCA accounts for the compromise effect because the compromise has only a small disadvantage relative to either extreme, while the extremes have one small disadvantage (relative to the compromise) and one large disadvantage (relative to each other extreme). Similarly, the attraction effect is obtained because the dissimilar option, B, has two large disadvantages (relative to A and D), while A has only one large disadvantage (relative to B).

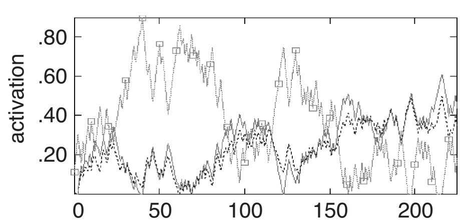

Fig. 13.5. Correlated activation of similar choice options (A, S lines with no symbols) in the LCA model. The dissimilar option, B (line with symbols) has times when it dominates the preference (reproduced from Usher & McClelland, 2004).

Discussion of computational models and further predictions

Both of the models described in the previous section can account for the three contextual reversal effects, promising thus to provide a unified explanation of multiattribute preferential choice. The models also share many properties (the switch of attention between dimensions, the leaky integration of the valences or advantages, and the lateral inhibition between the choice units). There are, however, a few basic differences in the process by which the reversal effects arise. We focus the discussion on these core differences before turning to a set of predictions.

Non-linearity and value functions

Whereas the DFT is a linear model, there are two types of nonlinearities in the LCA model. The first one is the biological constraint that activation cannot turn negative (unlike valences in DFT) and thus there is no 'excitation by negated-inhibition'. In Usher and McClelland (2004) we have argued that this is an important biological constraint and Busemeyer and colleagues have responded by suggesting possible biological implementations of their scheme (Busemeyer *et al.*, 2005). Here we focus on a set of functional considerations.

In Section 'Perceptual-choice, optimality, and the LCA model' we reviewed the application of the LCA to perceptual choice, showing that its flexible generalisation to multiple choice (allowing it to maintain high performance at large-n) relies to a large extent on the zero-threshold nonlinearity, which makes the units that receive little support drop out of the choice process. In a linear system with many options, as is likely to be the case in daily decision-making, the many options that receive little support will become negatively activated and will send noninformative input into the relevant choice units, reducing the choice quality. This problem is likely to be exacerbated if unavailable options are assumed to compete during the choice process, as the DFT needs to assume to account for reference effects (see below).

The second type of nonlinearity assumed in the LCA but not in DFT involves the nature of the value or utility function. Whereas the DFT has a linear value function, and choice patterns such as loss-aversion are thought to be emergent (or derivative) from the model's behaviour, in the version of the LCA we used, the value-function is explicit. To illustrate this distinction, let us examine how the two approaches can account for one of the cornerstones of preferential choice: the framing effects (see Fig. 13.2a). In the LCA we are following Tversky in assuming that when the choice offers an explicit reference (a present job that is being terminated) the available options are evaluated relative to that reference. Because of the logarithmic nonlinearity of the value function, the observers will prefer A from reference Ra, but will prefer B from reference Rb.[4] To explain the reference effect, protagonists of the DFT

[4] This is the case even without assuming that the value function for losses is steeper than that for gains (see Bogacz *et al.*, 2007, section 5 for details). Assuming a steeper value function for losses will further amplify the effect.

have proposed that the unavailable reference option takes part in the choice process, but is not chosen because it has a low value on a third dimension: the *availability*; this makes the reference Ra, somehow similar to the dominated option D (in Fig. 13.2b) and the reference reversal effect is then explained as a contrast effect, similar to the explanation of the attraction effect. Note, however, that by assuming that unavailable options compete for choice, one has to bring in every choice act a potential large set of unattractive options. As explained above, this is likely (in a linear model) to result in a reduction of the choice quality.

In addition to this consideration, we believe that there is independent evidence in favour of nonlinearities in the value function. Consider the task facing the decision-maker in choice options such as those illustrated in Fig. 13.2. To do this one has to represent the magnitudes that correspond to the various alternatives. Since magnitude evaluation is thought to involve a logarithmic representation and is subject to Weber's law,[5] it is plausible that it also affects the value functions. This idea is not new. In fact it dates back to Daniel Bernoulli (1738/1954), who proposed a logarithmic type of nonlinearity in the value function in response to the so-called St. Petersburg paradox, almost two centuries ago.[6] Bernoulli's assumption—that internal utility is logarithmically related to objective value—offers a solution to this paradox and has been included in the dominant theory of risky choice, the prospect theory (Tversky & Kahneman, 1979). Moreover, a logarithmic function, such as $\log(1+x)$ 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).[7]

In our recent paper (Bogacz *et al.*, 2007) we have explored the consequences of using such a logarithmic value function without the assumption of the asymmetry between gains and losses (the value for negative x, is then defined as $-\log(1 - x)$). We show there that this logarithmic assumption alone can suffice for accounting for some of the reversal effects, such as the reference effect and some aspects of the attraction and the compromise effect. We elaborate here on the latter.

---

[5] 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.

[6] This paradox was first noticed by the casino operators of St. Petersburg (see for example Glimcher, 2004, pp. 188–192 for detailed descriptions of the paradox and of Bernoulli's solution). Here is a brief description. 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 ($E = \Sigma_i{=}1,\ldots, 1/2^i\, 2^i = \Sigma_i{=}1,\ldots, 1 =$), 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*.

[7] While neuronal firing rates 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).

The nature of the compromise effect

There are a few issues to examine in understanding why people prefer compromise options. The first explanation, originally offered by Simonson (1989) is that this is the result of a conscious justification strategy: people choose the compromise because they have an easy way to justify it to themselves and to others. Note, however, that existence of such a heuristic does not imply that there are no *additional* factors that contribute to the effect. In the following we explore a number of potential contributing factors (on top of the justification heuristics) that emerge from neurocomputational models. The first (non-justificatory) factor, we examine, involves to the way in which nonlinear utilities are combined across two or more dimensions. Assuming a logarithmic value function, $v(x) = \log(1 + x)$, one can see that when summing across two dimensions, one obtains: $U(x_1,x_2) = u(x_1) + u(x_2) = \log[1+ (x_1 + x_2) + x_1x_2]$. Figure 13.6 illustrates a contour plot of this 2D utility function.

One can observe that equal preference curves are curved in the x_1–x_2 continuum: the compromise (0.5,0.5) has a higher utility than the (1,0) or (0,1) options. While this cannot (on its own) account for the compromise effect, which requires a change in the shares of the *same* two options when a third option in introduced, it still leads to an interesting prediction. Decision-makers are likely to prefer a middle-range option (0.5, 0.5) to an extreme range one $(1\text{-}x, x)$ for $0 < x < 0.5$, in binary choice, and the preference difference should increase the smaller x is (extreme options are less attractive than middle-range options even in the absence of context). Data supporting this prediction is presented in the following section.

Contextual contributions can further add to this tendency to choose options in the middle of the range. Two possibilities arise from the DFT and the LCA models. According to the DFT, the preference for the compromise is due to the dynamically correlated activations (the choice preferences). According to the DFT the preferences of the two extreme options are correlated with each other but not with the compromise (because of the stronger inhibition between more similar options) and thus, they split their wins, making the compromise option stand out and receive a larger share of choices (Roe *et al.*, 2001).

The process that is responsible for the effect in the LCA model is not *correlational*, but rather due (as Tversky suggested) to the nature of the nonlinearity and framing in value evaluations. Consider first, the simple logarithmic nonlinearity described above (without the asymmetry for losses) and note that the value of options defined in a

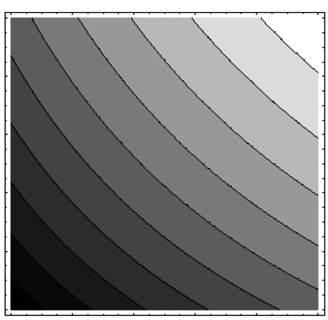

Fig. 13.6. 2D logarithmic value-function, $U(x_1,x_2) = u(x_1)+u(x_2)$ $= \log[1+ (x_1+x_2) + x_1x_2]$

2D parametric space depends on a reference (see below). If this reference changes with the choice set, a reversal effect arises. Consider, for example, a trinary choice between options: $A = (1,0)$, $B=(0.5,0.5)$, and $C = (0,1)$ and another binary choice between options $A = (1,0)$, $B = (0.5,0.5)$. If we assume that decision makers use the minimum value on both dimensions as reference (i.e., (0,0) in the trinary choice and (0.5) in the binary choice, we obtain that decision-makers will be indifferent between *A* and *B* in binary choice, but they will prefer *B* (the compromise) in trinary choice (see, Bogacz *et al.*, 2007). Second, the existence of an asymmetric value function, which is steeper in the domain of losses, provides another explanation, without the need to assume that the reference changes from the binary to the trinary set. In this case, all we need to assume is (following Tversky & Simonson, 1993) that decision-makers use each option as a reference to each other option (i.e., they evaluate differences rather than the options themselves). As shown in the previous section, this leads to a compromise effect because the extremes (but not the compromise) have large disadvantages that are penalised by the loss-aversive value-function.

In the following, we present some preliminary data that are aimed at testing these predictions. The first experiment examines binary and trinary preferences among options defined over two attributes with tradeoffs, by comparing a middle option (both attribute values in middle of the range) with an extreme one (one of the attributes at the high-end and the other one at the low-end). The second experiment, examines the correlational hypothesis. This is done by announcing, immediately after the decision-maker has chosen an extreme option, that this option is now unavailable and a speeded choice between the remaining two options has to be made.

Preliminary experimental investigations of the compromise effect

Experiment 1—A parametric study of the value of 2D trade-off options

The experiment was conducted on choice between options defined over two dimensions with values that create a tradeoff (Fig. 13.7). The aim was to compare the likelihood of choosing an alternative with attribute values in the middle of the range, relative to alternatives with extreme values, as a function of the distance between the extremes in the attribute space (small vs. large separation). Two groups of participants were tested. The first group were tested on trinary choice (Experiment 1A) and the second group on binary choice (Experiment 1B).

Method

Participants. Seventy-eight subjects participated in the trinary study (Experiment 1A) and 28 in the binary choice (Experiment 1B).

Design

The experiment was within-subjects with two levels: small separation (B,C,D) or (B,C) versus large separation (A,C,E), or (A,C) as illustrated in Fig. 13.7. The dependent variable was the probability of choosing the compromise and the extreme options.

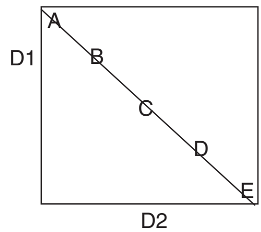

Fig. 13.7. Choice options in Experiment 1. The letters represent the alternatives on the two dimensions D1 and D2. For example, A is lowest on D2 and highest on D1. C represents the compromise and is in the middle range of both dimensions.

Materials

Experiment 1A. Each decision-problem consisted of three alternatives that differed on two dimensions. The values for extreme choices were symmetric relative to the compromise. For instance, car A is high on riding quality, but low on reliability, whereas car E is the opposite. Ten different decision problems were designed and for each problem a small separation (B,C,D) and a large separation (A,C,E) were created (materials are available online at www.bbk.ac.uk/psyc/staff/academic/musher/documents/DMproblems.pdf). These two conditions were counterbalanced so that one group of subjects were presented with the a first half of the problems in the (B,C,D) condition and the other half in the (A,C,E) condition, while the other half received the reversed pairing. The actual order of presentation of the problems was randomised.

Experiment 1B. The same materials were used, except that the 10 trinary choice problems were used to create 20 problems, which included the middle option and only one of the extremes. Each participant was presented with 10 choice problems (only one from each domain, such as laptops (below). Five of the problems involved a small separation and five of them involved a large separation. The allocation of problems to conditions was counterbalanced (four sets of problems were presented to four groups of participants).

Procedure

Experiment 1A. Participants were presented with a booklet of 10 pages, each containing one problem. They were instructed to imagine having to make a choice among three available options, which are identical on all the other properties except of the two described, and to make a selection that reflects their preference. One example is: '*Imagine you want to buy a laptop. You have a selection of three laptops, which have the same characteristics except for weight and screen size. If you had to choose one laptop out of the three, which laptop would you select?*'

	Laptop I	**Laptop II**	**Laptop III**
Weight (kg)	2.7	1.9	3.5
Screen size (")	14	13	15

Experiment 1B. The procedure was identical to experiment 1A, except for the fact that the participants were emailed the choice problems in a Word-attachment and they sent it back with their marked responses.

Results

The fraction of choices for the compromise and the extreme options are shown in Fig. 13.8 (trinary choice: left panel and binary choice right panel), for the small and the large separation conditions. [In trinary choice (left panel) the extreme conditions were averaged, thus the normalisation is: 2P(extreme) + P(compromise) =1].

The choice probabilities in trinary choice (Experiment 1A) were analysed using a 2 (extreme vs. compromise) by 2 (small vs. large separation) within subjects ANOVA. This yielded a significant main effects of compromise, $F(1,77) = 69.84$, $p < 0.001$, and of degree of separation, $F(1,77) = 42.066$, $p < 0.001$, and a highly significant interaction of compromise vs. degree of separation, $F(1,77)=42.066$, $p < 0.001$. This indicates that the compromise options receive a higher share than the extreme options and that this effect is larger at high separation (by 21%). Interestingly, a similar pattern is found in the binary choice (Experiment 1B), where participants chose the option (C) (in the middle of the value range) more than the extreme option (B, C, A or E), and this effect increases at large separations (A, E).

This suggests that the main factor that contributes to the participants' preference of the compromise option is the fact that this option is within the middle of the preference range, which has a higher 2D-value (Fig. 13.6). It is possible that a further contextual effect contributes to the compromise effect in trinary choice.[8] As the present experiments were between-participants (and there were some differences in procedure), further experiments are required to evaluate accurately the magnitude of these contributions and their dependence on the distance between the extremes. Although this experiment does not distinguish between the LCA/DFT accounts of the compromise

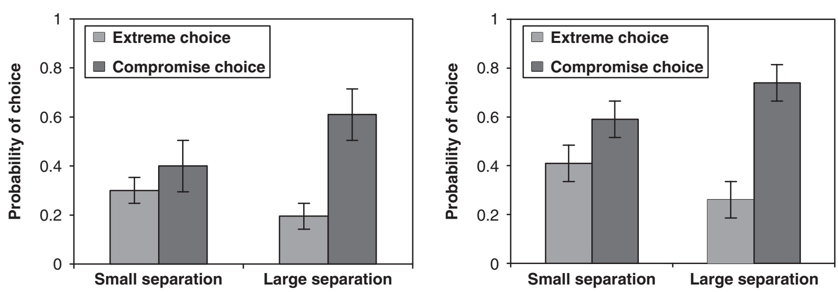

Fig. 13.8. Proportion of choices in trinary (Experiment 1A; left panel) and in binary choice (Experiment 1B; right panel). Error bars are SEM.

[8] This can be computed as $P_3(C)/[P_3(C) + P_3(A)] - P_2(C)/[P_2(C) + P_2(A)]$, where P3 and P2 are the trinary and binary choice probabilities.

effect, it suggests that the nonlinearity of the value function is an important contributor of the decision-makers' preference of middle-range options.

Is the compromise a dynamic correlation effect?

If the compromise effect is caused, as predicted by the DFT, by the fact that the preferences of the extremes are correlated in time, it should be possible, in principle, to detect a signature of this correlation. One way to investigate this is by presenting participants with the three-choice compromise option, and in some of the cases, following the participant's choice, announce that the option chosen is unavailable but a speeded choice is possible (under deadline) for one of the other two options. The idea is that, if the participant chose one of the extremes in her 1st choice, and if the preferences of the two extremes are correlated, then at the moment of the response, there should be a high likelihood that the preference of the other extreme is also high (relative to the compromise, which is not correlated). Thus one may predict that the participant will choose the 2nd extreme in her 2nd choice, following her 1st choice of the other extreme option. One caveat to this prediction is that, following the announcement of the unavailability of her preferred choice, the participant will restart the choice from scratch, in which case the advantage of the correlated alternative becomes immaterial. Such restart, however, is expected to lead to longer choice latencies, leading to a 2nd prediction: the choice latencies of the 2nd response should be faster when the other extreme is chosen than when the compromise is chosen (as in the latter case a restart is more likely). The LCA model makes the opposite prediction. Here the extreme options are anti-correlated (due to the scan of aspects), which together with the fact that the compromise receives more activation (due to the nonlinear value function) leads to a strong preference the compromise when a second choice is offered after the first choice (of an extreme option) is announced as unavailable.

Method

Participants. One hundred and forty-three participants volunteered to take part in this experiment (60% females; age-range 20–61 with a mean of 37).

Materials

The material consisted of 30 choice problems with three alternatives that varied on two dimensions, of a similar type to those shown in Table 13.1. Out of these problems, 25 were of the compromise form. Of those, 14 had a 2nd choice required after announcing the unavailability of the 1st chosen option. (The other 11 did not involve unavailable options so as not to create a strategy that prepares the two preferred options in advance.) A table that includes those 25 problems is available on: www.bbk.ac.uk/psyc/staff/academic/musher/documents/DMproblems.pdf

Procedure

The experiment was run on the web, using the 'express' psychology experiments package (Yule & Cooper 2003), which recorded the responses and the latencies.

The participants were instructed to imagine the situations described, as real choice situations they encounter in daily life and to indicate their preference. They were also instructed that sometimes (as in real life) an option they chose may be unavailable and that in such a case they will be able to make a 2nd choice (it was emphasized that this 2nd choice should be fast).

Results

The fraction of choices in favour of the compromise and extreme options (out of 14) in the 1st choice made is shown in Fig. 13.9 (left panel).

One can see (left panel) that, consistent with the data from Experiment 1, the participants chose the compromise option more than they chose the extreme options ($t(142) = 7.185$, $p < 0.001$). The fraction of choices made after a first choice, which was an extreme option announced to be unavailable is shown in the right panel. One can see that in such a situation the participants chose more than 90% the compromise option ($t(142) = 35.354$, $p < 0.001$).

Finally, we examine the response latencies of these 2nd choices and compare the latency of choosing the compromise over the extreme. As shown in Fig. 13.9 (right), the probability of choosing the extreme is less than 10% (only 58 participants out of the 143 made such choices, so we could only compare reaction time data for them). Although most of the responses were faster than 7 sec, indicating the participants followed the instructions of making a speeded response (this time includes the reading of the un-availability), there are also some slow RTs, that are as long as 30 sec. Eliminating such outliers (10 out of about 2000 in the whole data set), one finds that the average choice is significantly faster in the compromise 2nd choice (4.3 sec), than in the extreme 2nd choice (6.3 sec), $t(57) = 4.092$, $p < 0.001$. The full RT density-distributions (collapsed over all the participants) for the compromise (blue) and the extreme 2nd choices is shown in Fig. 13.10 (left). One can observe that extreme choices are slower in their mode and have a longer tail. This is consistent with RT-data from individual participants (right panel), which show slower responses for extreme (green) than to compromise (red) 2nd choices.

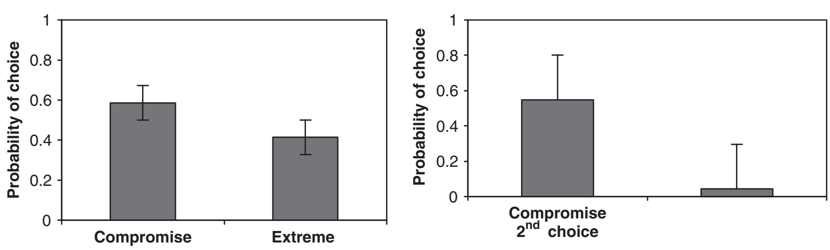

Fig. 13.9. The probability to choose the compromise and the extreme options in the first choice opportunity (left), and to choose compromise/extreme in the second choice.

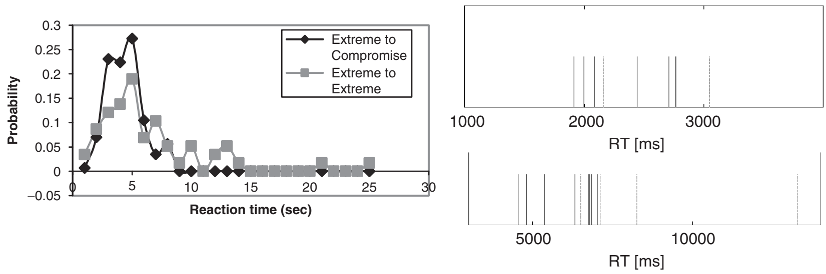

Fig. 13.10. Response latencies for 2nd choices in Experiment 2. Left: RT density-distribution of all responses of 58 participants; Right: RTs of two participants.

This result does not support the correlational hypothesis of the compromise effect, however, one cannot totally rule out such an account, because one can explain the preference for the compromise in the 2nd choice, within the DFT (J. Busemeyer, personal communication), by assuming that once an option becomes unavailable it acts as a dominated decoy and enhances the likelihood to chose the compromise (which is more similar with it), as in the attraction effect. Further studies, which contrast the magnitude of the compromise effect in situations where the extremes are available or unavailable, could help to test this proposal.

General discussion

In this chapter, we have examined how the LCA model, which was originally developed to account for perceptual choice, can be extended to preferential choice between alternatives that vary on several dimensions or attributes (see Stewart & Simpson, this volume, for a model of risky choice that treats probability and value as independent dimensions, and makes decisions on the basis of integrated comparisons). The LCA is a neurocomputational model, from a similar family with the DFT. As such, they share a similar approach, and vary on a number of important but secondary mechanisms. In the previous sections we discussed the differences between the LCA and the DFT accounts to preference reversal and we suggested possible ways to examine them.[9] Here we highlight their common approach by contrast with heuristic approaches (LeBoef & Shafir, 2005; Todd & Gigerenzer, 2000; Gigerenzer 2006) and we briefly address some implications to the debate on *rationality* in human choice.

Neucomputational models vs. heuristics

It is customary to understand choice heuristics as algorithms that are not optimal, but which can produce fast and reasonable choices for given situations (Gigerenzer, 2000, 2006). Understood in this minimal way, neurocomputational models such as DFT and

[9] The LCA is more similar in its approach with the decision-by-sample model (Stewart, this volume), but they vary on their assumptions about the order of attentional switches (between attributes in the LCA) and between alternatives first, in the decision-by-sample).

LCA are heuristics, as they indeed *can* produce reasonable and fast choices. There is, nevertheless, a feeling that something distinguishes these models from typical heuristics. We believe that the main contrast between the two resides in the nature of the algorithm. While the typical heuristics involve rules that can be *verbally* formulated in a *propositional* format and which have a *sequential* nature, the neurocomputational algorithms are mathematical rather than propositional and they involve some degree of parallel (rather than sequential) processing. This stems from the distinction between parallel distributed processes (PDP) and symbolic ones, which has been discussed in detail in other domains of cognition (McClelland *et al.*, 1986; Rumelhart *et al.*, 1986). In the light of this distinction, a number of considerations need to be addressed [see further discussion in the BBS replies to Todd and Gigerenzer (2000), in particular: Chater, 2000; Cooper, 2000; Oaksford, 2000; Shanks & Lagnado, 2000].

The LCA approach relies on some elements of prospect theory, in particular the form of the value function and loss-aversion. In their recent heuristic proposal to preferential choice under risk, the *priority* heuristic, Brandstatter *et al.* (2006) argue against the various variants of the prospect theory on several grounds. First, they maintain that all *weighing* models based on Bernoulli type corrections of the EV principle (such as prospect theory) are overly complex, as they rely on all the information available and require both summations and multiplications. Second, they suggest that heuristics, but not the prospect theory, provides a process model of preference. Third, they argue that such heuristics are ecologically rational and that, therefore, the so-called 'violations of rationality' are a misleading outcome of our over-reliance on unrealistic EV-type rationality norms. While we are sympathetic to the rational-ecological approach, we believe that the use of mathematically defined value functions (a la Bernoulli) within neurocomputational process models can achieve more than the use of disparate verbal heuristics, on all the three grounds.

Consider simplicity considerations first. On the one hand, while weighing models are complex from a symbolic perspective (where one actually multiplies and adds values) they are not so within a PDP one, as there is nothing more straightforward than computing weighted averages that implement EV (see Figs. 13.3 and 13.4)[10] and as logarithmic nonlinearities come for free (they are assumed anyway within basic psychophysical principles of magnitude evaluations, such as Weber's law). On the other hand, there is complexity hidden in the symbolic form of some heuristics, which require fundamentally different computations for different problems within the same domain. As an example, the priority heuristic (a lexicographic type heuristic) requires different computations for choices between options that differ in EV (thus some computation of EV needs to be done anyway) and for gains vs. losses.[11] Similarly, one can

[10] In fact we face the inverse puzzle: what are the processes that limit a straightforward EV computation.

[11] The rationale stated for the priority heuristic is that the first concern is the lowest possible outcome (i.e., the lowest gain). By this principle, the first concern in the domains of losses should be the highest possible loss, rather than the lowest one as the heuristics assumes to account for loss-aversion. In terms of complexity this assumption is not simpler than prospect's theory asymmetry of value functions for gains and losses.

explain the preference reversal effects in 2D attribute choice with verbal heuristics of the type: in a compromise choice, choose the middle option, but in a similarity situation choose the different (extreme) option. As there is a continuity of choice options over the 2D attribute space, one has to decide when to switch from one version of the heuristic to the other; this seems thus to require a meta-model to decide on the heuristic to use. We believe that a neurocomputational approach that employs continuous functions has a better prospect of obtaining a unified account of preference. Moreover, in some conditions, such a model may be *approximately* described by a verbal heuristic. For example, the LCA has some properties in common with the EBA heuristic (the shift of attention from attribute to attribute) and one of its neural precursors (Usher & Zakay, 1993) was shown to extrapolate among a large number of heuristics in multi-attribute choice.

Second, the neural models are in essence process models, which make dynamic predictions (Usher & McClelland, 2004). While such models can obtain fast and reasonable choices like the verbal heuristics, they can also do something that, paradoxically, heuristics don't do but people (unfortunately) do: vacillate and procrastinate in their decision! For example, unlike lexicographic type heuristics, models such as LCA and DFT vacillate in their preferences. One interesting possibility is that neurocomputational models, of the type discussed here, are at the interface between purely conscious (rule based and capacity limited) lexicographic type of decision-making and intuitive/gut-feeling type of choices, which are not subject to the capacity limitations of conscious thought and are able to weigh large number of attributes in parallel (Dijksterhuis & Nordgren, 2006). Finally, we examine a few considerations on optimality and rationality principles.

Optimality and rationality

The domain of multi-attribute preference does not possess an external criterion for the value (or quality) of choice. Nevertheless, some of its ingredients, the leaky integration, the logarithmic value function, and the loss-aversion, can be interpreted in relation to adaptive principles.

In Section 'Perceptual-choice, optimality, and the LCA model', we summarised analysis that indicates optimal performance in choice under stationary conditions, when leak and inhibition are perfectly balanced. Our own experiments in perceptual choice (Usher & McClelland, 2001) and those by Hertwig *et al.* (2004) in feedback-driven value decisions, indicate that some participants show a *recency* bias (leak dominance) that results in suboptimality. One interesting possibility is that this suboptimality is the cost one needs to pay for enabling agents to maintain sensitivity to changes in their environment (Daw *et al.*, 2006)—a type of the exploitation/exploration tradeoff. Consider next the 2D value function (Fig. 13.6), which results in preference for middle-of-the-range options. This value function involves 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). In addition, the higher slope of the value function for losses (relative to gains) can be justified by the 'asymmetry between pleasure and pain' (Tversky & Kahneman, 1991). This, in turn, could be an adaptive outcome of the fact that preventing losses is more critical for survival than getting gains (a single large loss can be critical; attaining large gains is not).

As task optimality is environment-dependent, an important insight is that rational strategies of choice require the flexibility to modify the choice parameters in response to the environment and demands (under some situations inhibition dominance or linear value function are advantageous). Such flexibility, however, may be subject to limitations. Unlike Gigerenzer and colleagues (but like Tversky and Kahaneman, 2000; Kahaneman, 2003), we think that contextual reversal effects of the type discussed here demonstrate a limitation of rationality in choice preference. After all, intransitive preferences and contextual reversals can be used to manipulate agents' choice and even transform them into money pumps; even if we evolved within a different environment, it is *rational* to do well in the present one. The notion of rationality, requiring flexible responses to the present environment/tasks, should not be diluted to a fixed repertoire of strategies adaptive to past environments and tasks. The power of choice models is to account for the sources of the adaptive powers of decision-makers and their limitations, and equally important to quantify the degree of these limitations[12] and find ways to eliminate them. A recent suggestion that requires further investigation is that one can overcome limitations in the ability to integrate across dimensions by relying on intuitive[13]/implicit decion-making (Dijksterhuis, *et al*, 2006).

Acknowledgments

Thanks are due to Claudia Sitz and Yvonne Lukaszewicz for running participants in Experiments 1–2, and to David Lagnado for a critical reading.

References

Barnard, G. (1946). Sequential tests in industrial statistics. *Journal of Royal Statistical Society Supplement*, **8**, 1–26.296

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

Bogacz, R., Usher, M., Zhang, J., & McClelland J. L. (2007). Extending a biologically inspired model of choice: Multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society, B* (in press).

Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review*, **113**, 409–432.

Britten, K. H., Shadlen, M. N., Newsome, W. T., & Movshon, J. A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, **10**(6), 1157–1169.

[12] In this regard, lexicographic type models are the worse off. Fortunately, only a minority of human participants exhibit intransitive preferences as predicted by such models (Tversky, 1969).

[13] But see Slovic *et al.* (2004), for an insightful discussion of the limitations and dangers of intuitive/affective preferential choice.

Busemeyer, J. R., & Johnson, J. G. (2003). Computational models of decision making. In D. Koehler & N. Harvey (Eds.), *Handbook of judgment and decision making*. Blackwell Publishing Co (To appear).

Busemeyer, J. R., & Diederich, A. (2002). Survey of decision field theory. *Mathematical Social Sciences*, **43**, 345–370.

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

Chater, N. (2000). How smart can simple heuristics be? *Behavioural and Brain Sciences*, **23**, 745–746.

Cooper, R. (2000). Simple heuristics could make us smart; but which heuristic do we apply when? *Behavioural and Brain Sciences*, **23**, 746–747.

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in Humans. *Nature*, **441**, 876–879.

Diederich, A. (1997). Dynamic stochastic models for decision making under time constraints. *Journal of Mathematical Psychology*, **41**, 260–274.

Dijksterhuis, A., Bos, M. W., Nordgren, L. F., & Baaren, R. B. (2006). On making the right choice: The deliberation without-attention effect. *Science*, **311**, 1005–1007.

Dijksterhuis, A., & Nordgren, L. F. (2006). A theory of unconscious thought. *Perspectives on Psychological Science*, **1**, 95–109.

Gigerenzer, G. (2006). Bounded and rational. In R. J. Stainton (Ed.), *Contemporary debates in cognitive science* (pp. 115–133). Oxford, UK: Blackwell.

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**(2), 299–308.

Hacking, I. (1980). Strange expectations. *Philosophy of Science*, **47**, 562–567.

Hertwig, R., Barron, G., Weber, E. U., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, **15**, 534–539.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, **9**, 90–98.

Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *The American Economic Review*, **93**, 1449–1475.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision making under risk. *Econometrica*, **XLVII,** 263–291**.**

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

Knetch, J. L. (1989). The endowment effect and evidence of nonreversible indifference curves. *American Economic Review*, **79**, 1277–1284.

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

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

McClelland, J. L., Rumelhart, D. L., & the PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 2). Cambridge, MA: MIT Press.

Oaksford, M. (2000). Speed, frugality and the empirical basis of take-the-best. *Behavioural and Brain Sciences*, **23**, 760–761.

Roe, R. M., Busemeyer, J. R., & Townsend, J. T. (2001). Multi-alternative decision field theory: A dynamic connectionist model of decision-making. *Psychological Review*, **108**, 370–392.

Rumelhart, D. L., McClelland, J. L., & the PDP Research Group. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1). Cambridge, MA: MIT Press.

Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, **1**, 7–59.

Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews. Neuroscience*, **2**(1), 33–42.

Shadlen, M. N., & Newsome, W. T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Neurophysiology*, **86**(4), 1916–1936.

Shanks, D. R., & Lagnado, D. (2000). Sub-optimal reasons for rejecting optimality. *Behavioural and Brain Sciences*, **23**, 761–762.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, **16**, 158–174.

Slovic, P. (1995). The construction of preference. *American Psychologist*, **50**, 364–371.

Slovic, P., Finucane, M. L., Peters, E., & MacGregor, D. G. (2004). Risk as analysis and risk as feelings: Some thoughts about affect, reason, risk, and rationality. *Risk Analysis*, **24**(2), 311–322.

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions, *Trends in Neurosciences*, **27**(3), 161–168.

Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science,* **304**(5678), 1782–1787.

Sugrue, L. P., Corrado, G. S., & Newsome, W. T. (2005). Choosing the greater of two goods: neural currencies for valuation and decision making. *Nature Review and Neuroscience,* **6**(5), 363–375.

Tobler, P. N., Fiorillo, C. D., & Schultz, W. (2005). Adaptive coding of reward value by dopamine neurons. *Science*, **307**(5715), 1642–1645.

Todd, P. M., & Gigerenzer, G. (2000). Precis of *Simple heuristics that make us smart*. *Behavioral and Brain Sciences*, **23**, 727–780.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review* **79**, 281–299.

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

Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *The Quarterly Journal of Econometrics*, **106**, 1039–1061.

Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, **39**, 1179–1189.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, **108**(3), 550–592.

Usher, M., & McClelland, J. L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, **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. *Journal of Cognitive Neuroscience*, **8**, 311–327.

Usher, M., & Zakay, D. (1993). A neural network model for attribute based decisions processes. *Cognitive Science*, **17**, 349–396.

Yule, P., & Cooper, R. P. (2003): Express: A web-based technology to support human and computational experimentation. *Behavior Research Methods, Instruments, & Computers,* **35**, 605–613.

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

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