

Information Integration in Perceptual and Value-Based Decisions

Konstantinos Tsetsos

Thesis submitted for the degree of Doctor of Philosophy
University College London

December 2011

Abstract

Research on the psychology and neuroscience of simple, evidence-based choices has led to an impressive progress in capturing the underlying mental processes as optimal mechanisms that make the fastest decision for a specified accuracy. The idea that decision-making is an optimal process stands in contrast with findings in more complex, motivation-based decisions, focussed on multiple goals with trade-offs. Here, a number of paradoxical and puzzling choice behaviours have been revealed, posing a serious challenge to the development of a unified theory of choice. These choice anomalies have been traditionally attributed to oddities at the representation of values and little is known about the role of the process under which information is integrated towards a decision. In a series of experiments, by controlling the temporal distribution of the decision-relevant information (i.e., sensory evidence or value), I demonstrate that the characteristics of this process cause many puzzling choice paradoxes, such as temporal, risk and framing biases, as well as preference reversal.

In Chapter 3, I show that information integration is characterized by temporal biases (Experimental Studies 1-2, Computational Studies 1-3). In Chapter 4, I examine the way the integration process is affected by the immediate decision context (Experimental Studies 3-4, Computational Study 4), demonstrating that prior to integration, the momentary ranking of a sample modifies its magnitude. This principle is further scrutinized in Chapter 5, where a rank-dependent accumulation model is developed (Computational Study 5). The rank-dependent model is shown to underlie preference reversal in multi-attribute choice problems and to predict that choice is sensitive, not only to the mean strength of the information, but also to its variance, favouring riskier options (Computational Study 6). This prediction is further confirmed in Chapter 6, in a number of experiments (Experimental Studies 5-7) while the direction of risk preferences is found to be modulated by the cognitive perspective induced by the task framing (Experimental Study 8). I conclude that choice arises from a deliberative process which gathers samples of decision-relevant information, weighs them according to their salience and subsequently accumulates them. The salience of a sample is determined by i) its temporal order and ii) its local ranking in the decision context, while the direction of the weighting is controlled by the task framing. The implications of this simple, microprocess model are discussed with respect to choice optimality while directions for future research, towards the development of a unified theory of choice, are suggested.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text. This work has not been submitted for any other degree or professional qualification except as specified.

Signature:

March 15, 2012

(Konstantinos Tsetsos)

Acknowledgements

I would like to express my sincere gratitude to my supervisors, Nick Chater and Marius Usher for their invaluable guidance, inspiration and encouragement throughout the course of my PhD. I would like to thank Chris Olivola, Juan Gao and Tobias Donner for their constructive collaborations. I am especially grateful to James McClelland for his kind and generous intellectual contribution to part of this work.

I would also like to thank Nick Chater and Dave Lagnado for providing unique and stimulating conditions within their labs. And all the members of the Chater+Lagnado Laboratory, for being great colleagues: Christos Bechlivanidis, Caren Frosch, Tobias Gerstenberg, Adam Harris, Anne Hsu, Petter Johansson, Jens Koed Mansen, Irma Kurniawan, Dave Lagnado, Milena Nikolic, Chris Olivola, Ramsey Raafat, Stian Reimers, Costi Rezlescu, Adam Sanborn, Katya Tentori, Ivo Vlaev and Ro'i Zultan. Many thanks to my participants for enduring my tasks.

I am grateful for the part-time employment and financial support kindly offered to me during these years by: AIG Foundation (Institute for International Education), Decision Technology LTD, ELSE Research Centre (Nick Chater and Chris Olivola), James McClelland and Marius Usher, Ramsey Raafat, and Katya Tentori.

Finally, this work would not have been completed without the courageous and substantial financial aid from my parents. This thesis is dedicated wholeheartedly to them.

Collaborations

Chapter 1: Parts of the text are reprinted from one article, currently submitted in a special topic in *Frontiers in Cognitive Science* and *Frontiers in Neuroscience* with Marius Usher, Juan Gao and James L. McClelland. Other parts are reprinted from two published articles in *Psychological Review* (October 2010) with Marius Usher and Nick Chater. These parts were also presented as a poster in the Mathematical Psychology conference, held in Amsterdam in August 2009.

Chapter 3: Parts of the text are reprinted from one article, currently submitted in a special topic in *Frontiers in Cognitive Science* and *Frontiers in Neuroscience* with Marius Usher, Juan Gao and James L. McClelland. Other parts are reprinted from a currently submitted article with Nick Chater and Marius Usher. These parts were presented as a poster in SPUDM conference, held in London, in August 2011.

Chapter 4: Parts of the text are reprinted from one article, published in *Frontiers in Neuroscience* with Marius Usher and James L. McClelland, in May 2011. Other parts are reprinted from a currently submitted article with Nick Chater and Marius Usher and were also presented as a poster in SPUDM conference, held in London, in August 2011.

Chapter 6: Parts of the text are reprinted from a currently submitted article with Nick Chater and Marius Usher and were also presented as a poster in SPUDM conference, held in London, in August 2011.

Publications arising directly from this submission

Tsetsos, K., Gao, J., Usher, M., & McClelland, J. (submitted). Using time-varying evidence to probe decision dynamics.

Tsetsos, K., Chater, N., Usher, M. (submitted). Value psychophysics: How salience-driven value integration explains decision biases and preference reversal.

Tsetsos, K., Usher, M., & McClelland, J. (2011). Testing multi-alternative decision models with non-stationary evidence. *Frontiers in neuroscience*, 5. doi:10.3389/fnins.2011.00063.

Tsetsos, K., Usher, M., & Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological review*, 117(4), 1275-1291. doi:10.1037/a0020580.

Usher, M., **Tsetsos, K.**, & Chater, N. (2010). Postscript: Contrasting predictions for preference reversal. *Psychological review*, 117(4), 1291-1293. doi:10.1037/0033-295X.117.4.1291.

Table of Contents

List of Figures	11
List of Tables	14
1 Introduction	15
1.1 Order Effects in Perceptual Choice	20
1.1.1 Sequential Probability Ratio Test	21
1.1.2 Psychological Models of Choice: Diffusion and Race	21
1.1.3 Variants of the classical diffusion model	23
1.1.4 A Neural Model of Choice: The Leaky Competing Accumulators	25
1.1.5 An Experimental Study on Order Effects	27
1.1.6 Summary	30
1.2 Context Effects in Value-based Choice	32
1.2.1 Preference Reversal in Multi-attribute Choice	33
1.2.2 Mechanisms for reversal effects	35
1.2.3 Two Neurocomputational Approaches	38
1.2.4 Contrasting DFT and LCA	42
1.2.5 Alternative Models	46
1.2.6 Summary	47
1.3 Summary and Overview of the Thesis	50
2 General Methods and Techniques	53
2.1 Overview	53
2.2 Experimental Methods	53
2.2.1 Stimuli and Tasks	53
2.2.2 Non-stationary Information	57
2.3 Computational Techniques	60

2.3.1	Monte Carlo Model Simulations	60
2.3.2	Optimization Procedure	60
2.3.3	Model Evaluation	61
3	Time-dependent Weighting of Information	62
3.1	Overview	62
3.2	Order Effects and Trial Duration (Computational Study 1)	63
3.2.1	Method	64
3.2.2	Results and Discussion	64
3.3	Order Effects in Evidence Integration (Experimental Study 1)	68
3.3.1	Experiment 1a	69
3.3.2	Experiment 1b	72
3.3.3	Discussion	74
3.4	Order Effects in Value Integration (Experimental Study 2)	75
3.4.1	Method	76
3.4.2	Results	80
3.4.3	Computational Models (Computational Study 2) and Discussion	82
3.5	Optimality in Decisions under Uncertainty (Computational Study 3) .	87
3.5.1	Methods	88
3.5.2	Results and Discussion	88
3.6	Summary and General Discussion	92
4	Context-dependent Weighting of Information	94
4.1	Overview	94
4.2	Context Effects in Evidence Integration (Experimental Study 3) . . .	95
4.2.1	Method	96
4.2.2	Results	100
4.2.3	Computational Models of Multi-alternative Perceptual Choice (Computational Study 4)	104
4.2.4	Discussion	117
4.3	Context Effects in Value Integration (Experimental Study 4)	120
4.3.1	Method	121
4.3.2	Results and Discussion	123
4.4	Summary and General Discussion	126
5	Rank-dependent Leaky Integration	129

5.1	Overview	129
5.2	Rank-dependent Leaky Integration Model (Computational Study 5)	130
5.2.1	Model Implementation	130
5.2.2	Rank-dependency in Perceptual Decisions	131
5.2.3	Rank-dependency in Value Integration	134
5.2.4	Discussion	136
5.3	Predictions for Binary Choice: Relativity of Value and Risk-Attitudes (Computational Study 6)	138
5.3.1	Relativity of Value	138
5.3.2	Risk Attitudes	140
5.3.3	Discussion	140
5.4	Summary and General Discussion	142
6	Value Integration and Risk-attitudes	145
6.1	Overview	145
6.2	Risk-seeking in Gains (Experimental Study 5)	146
6.2.1	Method	147
6.2.2	Results	148
6.2.3	Discussion	150
6.3	Risk-seeking in Losses (Experimental Study 6)	150
6.3.1	Method	150
6.3.2	Results	151
6.3.3	Discussion	152
6.4	Risk-aversion in Mixed Sequences (Experimental Study 7)	152
6.4.1	Experiment 7a	153
6.4.2	Experiment 7b	155
6.4.3	Discussion	158
6.5	Risk-preferences and Task Framing (Experimental Study 8)	159
6.5.1	Experiment 8a	159
6.5.2	Experiment 8b	161
6.5.3	Discussion	164
6.6	Risk-preferences and Self-paced Sampling (Experimental Study 9)	165
6.6.1	Method	166
6.6.2	Results	166
6.6.3	Discussion	168

6.7	Risk-preferences and Rare Events (Experimental Study 10)	168
6.7.1	Method	169
6.7.2	Results	170
6.7.3	Discussion	172
6.8	Summary and General Discussion	172
7	Summary and Conclusions	175
7.1	Summary of Findings	176
7.2	Implications	177
7.3	Future Directions	182
7.4	Conclusion	185
	References	186

List of Figures

1.1	Reverse correlations showing primacy in Kiani et al., 2008	24
1.2	Effect of motion pulse on detection of motion for one observer	28
1.3	The effect of the pulse timing of the shift of the psychometric function for one observer.	28
1.4	Illustration of a choice space for options that vary on two dimensions .	34
1.5	Illustration of decision field theory and leaky competing accumulators models in neural networks	40
1.6	Illustration of the attraction and similarity effects	43
1.7	Dominance reversal in DFT	44
2.1	Stimuli used in the perceptual choice tasks.	55
2.2	The fast-value integration paradigm.	56
2.3	Unfolding in time a 2-dimensional choice problem	58
3.1	Reverse correlations for the LCA model.	65
3.2	Fraction of trials in which the accumulator associated with early evi- dence wins.	66
3.3	Single trial activation of the accumulators for long and short durations	67
3.4	Results in Experiment 1a	71
3.5	Results in Experiment 1b	73
3.6	The timeline of an experimental trial in Experimental Study 2.	77
3.7	The distributions in the unbalanced condition.	79
3.8	The balanced condition experimental condition.	79
3.9	Performance in the unbalanced trials of Experimental Study 2.	80
3.10	Recency bias in the balanced trials of Experimental Study 1.	81
3.11	Logistic-regression weights in the unbalanced trials in Study 1.	82
3.12	Data fits of perfect and leaky integration (without parameter w).	84

3.13	Data fits of perfect and leaky integration with overweighting the last item.	85
3.14	Signal embedded in noise.	89
3.15	ROC detection curves for signal embedded in continuous noise	90
3.16	Signal strength as a function of the signal duration for each leak level	90
3.17	Signal trial activations for short and long time constants	91
4.1	Density distribution that determines the time of switching from one phase to the other	97
4.2	The non-stationary evidence to three alternatives in a single trial of a similarity condition	99
4.3	The time course of an experimental trial.	100
4.4	Accuracy in the filler conditions of Experimental Study 3.	101
4.5	Preference for target and competitor in the attraction condition	102
4.6	Preference for the dissimilar option in the similarity condition	102
4.7	Preference for the stationary option in the compromise and anti-compromise condition	103
4.8	Neural implementation of perceptual choice models for the interrogation paradigm	105
4.9	Single trial activations for race, diffusion and LCA for initial evidence supporting A/B and supporting C	109
4.10	Predictions for the bounded diffusion and race models and the LCA model with zero and high levels of processing noise	112
4.11	Individual choice for the dissimilar alternative, P(C), in the similarity condition of the experiment and in the models.	115
4.12	The time course of an experimental trial in Experimental Study 4. . .	122
4.13	The mean values of the distributions for the 4 experimental conditions	123
4.14	Results for the four conditions in Experimental Study 4	125
5.1	Data fits for Experimental Study 4.	135
5.2	Performance in the unbalanced trials of Experimental Study 2.	139
5.3	Risk-seeking prediction in value integration.	141
6.1	Risk-seeking pattern in Experimental Study 5.	149
6.2	Risk-seeking pattern in Experimental Study 6.	151
6.3	Risk-aversion pattern in Experimental Study 7a.	154

6.4	Results in Experimental Study 7b.	157
6.5	Task-framing effect in Experimental Study 8a.	160
6.6	Task-framing effect in Experimental Study 8b.	163
6.7	Results in Experimental Study 9.	167
6.8	Results in Experimental Study 10.	171

List of Tables

1.1	A summary of DFT and LCA accounts for the contextual preference reversal effects.	42
3.1	Statistical analysis for the four participants of Experiment 1b (main order effect and its interaction with trial duration).	74
3.2	Model parameters and BIC values for perfect and leaky integration). .	85
3.3	Model parameters and BIC values for perfect and leaky integration). .	86
4.1	Experimental conditions. The mean values of the dim option D are omitted here since they were always the same: $\mu_1 = 0.1$ and $\mu_2 = 0.1$.	98
4.2	The mean values of the sequences in the 4 experimental conditions. .	123
5.1	Conditions in <i>Experimental Study 3</i>	132
5.2	Rank ordering of the options of table 5.1.	133
5.3	The ranking of the sequences in each distribution in the the 4 experimental conditions of <i>Experimental Study 4</i>	134
5.4	Optimized parameters of the rank-dependent model for <i>Experimental Study 4</i>	136

Chapter 1

Introduction

Background

Decision-making is ubiquitous across all domains of cognition that require behavioural output, from simple, *perceptual* choices on the basis of sensory evidence, to more complex, *value-based* ones, exhibited in daily life activities. These two classes of choice have been the subject of separate investigations, within different paradigms and disciplines. Perceptual decisions have been traditionally studied within experimental psychology and neuroscience, using choice accuracy and response latencies in simple psychophysical tasks in which the observer has to classify the perceived *sensory input* into one out of two or more competing hypotheses (Laming, 1968; Vickers, 1979; Britten, Shadlen, Newsome, & Movshon, 1993; Smith & Ratcliff, 2004). On the other hand, *value-based* or *preferential* choices, such as when deciding which laptop to buy out of many alternatives differing in several attributes, have been mainly studied within behavioural economics and social sciences, using primarily reports of choice preference in laboratory tasks that emulate real decision problems (Simon, 1982; Kahneman & Tversky, 2000).

Research on these two distinct areas, differing in methods and techniques, has produced divergent proposals regarding the *competence* of the decision-makers. In perceptual choice, the tasks are simple and well-defined, with the observer's responses being either correct or incorrect. In this type of tasks, where the choice accuracy and response times can be objectively measured, *optimality* is defined in statistical terms based on the Sequential Probability Ratio Test benchmark (SPRT; Barnard, 1946; Wald, 1947),

which produces the fastest responses for a given level of accuracy. A central tenet, arising from almost forty years of research on behavioural psychophysics (Audley, 1960; Stone, 1960; Vickers, 1970; Link & Heath, 1975; Ratcliff, 1978; Ratcliff & Rouder, 1998; Ratcliff & Smith, 2004) and complemented by physiological data (Gold & Shadlen, 2001, 2002, 2007; Bogacz & Gurney, 2007; Donner, Siegel, Fries, & Engel, 2009), holds that the process underlying perceptual decisions is an approximation of the SPRT test and hence approaches statistical optimality (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Bogacz, 2007, 2009). On the other hand, in the value-based paradigm, choice cannot be settled on sensory information only and each alternative is evaluated by examining its consequences in relation to internal motivations. The normative benchmark is provided by the expected utility theory, the cornerstone of economic theory and rational choice explanation (Von Neuman & Morgenstern, 1947; Debreu, 1960). The abundance of decision biases, highlighted in the seminal work of Tversky and Kahneman (e.g., Kahneman & Tversky, 1979; Tversky & Kahneman, 1981; Huber, Payne, & Puto, 1982; Tversky & Kahneman, 1986; Knetsch, 1989; Simonson, 1989; Kahneman & Tversky, 2000; Gilovich, Griffin, & Kahneman, 2002), has given rise to the general view that humans are fundamentally irrational and the normative axioms elusive.

This tension, that behaviour is nearly optimal in the perceptual domain and suboptimal in the value-based one, is also reflected at the different theoretical frameworks proposed in the two areas. In perceptual choice, models have evolved around the optimal SPRT test attempting to implement it algorithmically, based on the hypothesis that choice is the result of the accumulation of sequentially sampled sensory evidence, towards a response criterion (Laming, 1968; ?, ?; Ratcliff, 1978). Therefore the prototype model is *dynamic* in nature, explicitly describing the deliberation *process* and determining the response time of the decision. On the other side, value-based models have initially attempted to achieve descriptive adequacy by gradually modifying the normative theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992; Tversky & Simonson, 1993). These models similar to their predecessor, the expected utility theory, are *algebraic* and *static* in nature, providing choice output but saying nothing about the underlying deliberation process. An alternative route has been taken by completely abandoning the mathematical formalisms and assumptions of the rational theory. In this line of research, emphasis was placed on capturing the behavioural regularities independently of the expected utility theory, as algorithms that are not optimal

but can produce fast and reasonable choices under certain situations. This research has led to the generic proposal that decision-makers use a set of disparate heuristics, each addressing some other aspect of choice behaviour (Todd & Gigerenzer, 2000; Gigerenzer, 2006; Harvey, 2007).

Determining the relationship between perceptual and value-based choice is central for our broader understanding of decision-making. One viable way to explore this relationship is through the direct comparison of the two classes of theories. This has been particularly difficult since, as described above, the two types of choice have been analysed at different levels using completely different theoretical frameworks. However, the recent development of mathematically formalized dynamic models, which address the time course of *preference* formation (Busemeyer & Townsend, 1993; Roe, Busemeyer, & Townsend, 2001; Usher & McClelland, 2004; Stewart, Chater, & Brown, 2006; Stewart & Simpson, 2008; Krajbich & Rangel, 2011), promises to initiate a theoretical link between perceptual and value-based choice. These models build upon the *sequential sampling* tradition of perceptual choice, with the assumption that value (and not evidence) is accumulated across time, towards a response criterion. The advantages of the dynamic models of preference over previous approaches are several. First, as opposed to descriptive, algebraic models, they bear explanatory power, accounting not only for the overt choice behaviour but also for its exact time course and for phenomena such as vacillations and changes of mind. Second, contrary to the heuristic framework that assumes different, ad-hoc algorithms for different choice problems (Chater, 2001), dynamic models are predictive and parsimonious, consisting of fixed sets of independently motivated principles¹. Finally, if biologically constrained, these models, can be naturally mapped on physiological data offering that way the possibility of a detailed understanding of the neural underpinnings of decision-making.

The advent of process models of preference has revived the interest for recording information acquisition in the course of value-based decision-making (Russo & Rosen, 1975; Rayner, 1978; Armel, Beaumel, & Rangel, 2008; Reutskaja, Nagel, Camerer, & Rangel, 2011; Glöckner & Herbold, 2011), to track the regularities in the sampling of information and to determine the exact timing and amount of the input that the decision-maker receives. Knowing *what* aspects of choice alternatives are consid-

¹In dynamic models, the choice input is transformed to behaviour through the synergy of predetermined mechanisms. Under the heuristic doctrine, since different algorithms are employed for different problems, it appears that the choice input transforms the mechanisms to produce output that matches the experimental data.

ered during deliberation, can be particularly useful but only if the *way* under which value is processed and integrated is known. Contrary to perceptual choice models whose computational elements have been scrutinized for years, little is known about the micro-process of value-integration. And although process models of evidence and preference share a common conceptual framework, it should not be assumed *a priori* that both theories share the exact same mechanisms. Understanding the way that value is integrated, in comparison to the mechanisms of evidence integration, is essential for the further refinement of process models of preference, which in turn can delineate the relationship between perceptual and value-based choice.

The aim of this thesis is to probe the computations that are performed when samples of information, either corresponding to *sensory input* or to *values*, are integrated towards a choice. Setting the micro-foundations of information integration in decision-making will not only determine the relationship between perceptual and value-based choice but will ultimately lead us to more complete choice models which will encompass accounts for both information acquisition [stemming from the recent advances in recording eye-tracking data (Armel et al., 2008; Reutskaja et al., 2011; Glöckner & Herbold, 2011)] and information processing.

Information Weighting in Decision-making

Two specific aspects of information processing will be addressed using a combined experimental and computational method: a) how does the temporal order of information affect its weighting and b) how does the immediate context influence information integration? Understanding the differential weighting of decision-relevant information as a function of its *temporal order* and of the *immediate context* is theoretically critical since, as summarized below, both types of weighting undermine the ideas of the *optimal* and *rational* decision-maker.

For example the optimality assumption in the perceptual literature, resulting from the partial descriptive success of the SPRT, imposes that choice is the result of *Bayesian inference*, whereby all pieces of evidence confer equally to the making of the decision. In other words a decision that is made based on a particular evidence-set should not change if the exact same evidence appears in a different order. Contrary to this claim, experiments in perceptual choice in both primates (Huk & Shadlen, 2005; Kiani,

Hanks, & Shadlen, 2008) and humans (Pietsch & Vickers, 1997; Usher & McClelland, 2001; Tsetsos, Usher, & McClelland, 2011) have revealed order effects, showing that decisions are weighted by either primacy or recency. Thus, order effects undermine the proposal of the ideal observer and exploring their basis might improve our understanding of the very nature of optimality.

On the opposite pole, within the value-based literature, the falsification of the rational man assumption has been mainly driven by contextual effects. The goodness value assigned to an alternative is dependent on the context in which it is presented. Although the underlying psychological assumption of the rational choice theory, that people should behave consistently, has been significantly undermined by contextual biases, it has not been totally dispensed. The program of ecological rationality has highlighted that decisions that may appear poor in ecologically invalid or unstable contexts (e.g. laboratory tasks), may be highly efficient and rational in the natural environment (Gigerenzer, 1991; Gigerenzer & Hoffrage, 1995; Oaksford & Chater, 1998) and thus, rationality should not be assessed independently of the environmental context and its stability (Oaksford & Chater, 1995; Shanks, Tunney, & McCarthy, 2002). It still remains unclear whether people are inherently inconsistent or whether they behave so under specific contextual peculiarities.

To summarize, order and context effects play a pivotal role in the toggling between optimality/ suboptimality and rationality/ irrationality. The experimental scrutiny of these effects that is pursued in this thesis, is expected to shed light on the cognitive mechanisms that generate them. Admittedly, human decision making is multifaceted, violating sometimes the norms and complying with them others, with the “nearly optimal” and “definitely irrational” popular statements, resembling the “half full or half empty glass” rhetorical argument and reflecting mostly different viewpoints when it comes to the interpretation of the empirical truth. This inconclusiveness, characterizing the psychology of decision-making, has influenced the current project at the theoretical front. In particular, the decision models developed in this thesis are not shaped to *a priori* comply with or violate the normative theory, but are constructed bottom-up from the synthesis of simple mechanisms of information integration that capture the empirical reality.

In the remaining part of the current, introductory chapter, previous research that I did on order effects in perceptual choice (Tsetsos, Gao, Usher, & McClelland, 2011) and on context effects in value-based choice (Tsetsos, Usher, & Chater, 2010; Usher, Tset-

sos, & Chater, 2010), will be presented along with the relevant literature. This will set the scene for the following chapters where order and context effects are experimentally examined, jointly on perceptual and preferential choice, helping to infer the core computational elements of *evidence* and *value* integration.

1.1 Order Effects in Perceptual Choice

Almost sixty years of research on perceptual choice, starting from the early signal detection theory (Stone, 1960; Laming, 1968; Vickers, 1970; Link & Heath, 1975) and ending with recent neuroeconomics studies (Rorie, Gao, McClelland, & Newsome, 2010; Summerfield & Koechlin, 2010; Krajbich & Rangel, 2011), has converged to the idea that the most likely cause of a perceived experience is inferred using multiple samples of noisy evidence (Ratcliff, 1978; ?, ?; Ratcliff & Rouder, 1998; Usher & McClelland, 2001; Bogacz et al., 2006; Smith & Ratcliff, 2004; Wong, Huk, Shadlen, & Wang, 2007; Ratcliff & McKoon, 2008). These samples are accrued up to a response criterion, with the time needed to breach the criterion determining the response latency. This simple principle has been descriptively successful, explaining a fundamental behavioural pattern, the speed-accuracy tradeoff: with more time available, one can take more samples of evidence and thus be more accurate. Further support to the *sequential sampling* hypothesis has been provided by neurophysiological studies of motion discrimination with behaving animals. There, neurons in visual-motor integration areas (e.g. the lateral intraparietal cortex), showed ramped activity which correlated with the amount of the integrated evidence (Gold & Shadlen, 2001, 2002; Roitman & Shadlen, 2002).

One advantage of the integrate-to-threshold principle is that, under specific conditions, it can be statistically optimal, generating the fastest decisions for a given error rate (Wald & Wolfowitz, 1948). The optimality of perceptual decisions can be understood with evolutionary terms, assuming that animals capable of making fast and accurate decisions are promoted. Many theories have been developed within the sequential sampling framework, differing in aspects like the stopping rule or the boundary of integration, and although optimality is not always achieved it helps constrain the existing models and the interpretation of the empirical data. I will start by statistically formulating an abstract binary perceptual decision problem and defining the strategy of the ideal observer. This will be useful for the next subsection where two of the most promi-

ment mathematical models of perceptual choice, the classical drift-diffusion (Laming, 1968) and race models (Vickers, 1970) will be introduced.

1.1.1 Sequential Probability Ratio Test

Lets assume that the observer receives sensory information at discrete time-steps. The sensory evidence at time t , supporting alternative i is denoted as $x_i(t)$. According to Gold and Shadlen (2001) the choice problem can be formalized statistically by assuming that evidence $x_i(t)$ comes from a normal distribution with mean μ_i and standard deviation σ . The observer needs to determine which μ_i is the highest. In other words hypothesis H_i states that the evidence supporting alternative i is stronger. For a binary choice task the two hypotheses become:

$$H_1 : \mu_1 = \mu^+, \mu_2 = \mu^- ; H_2 : \mu_1 = \mu^-, \mu_2 = \mu^+ \quad (1.1)$$

with $\mu^+ > \mu^-$. Wald (1947) provided the optimal procedure to distinguish between the two competing hypothesis, the Sequential Probability Ratio Test. According to the SPRT, at each time t the ratio of the likelihood of the evidence given the hypothesis is computed:

$$R = \frac{P(x(1..t)|H_1)}{P(x(1..t)|H_2)}, \quad (1.2)$$

with $x(1..t)$ corresponding to the sensory evidence presented until time t (i.e. $x_1(1), x_2(1), \dots, x_1(t), x_2(t)$). When R reaches an upper threshold Z_1 then a choice is made in favour of H_1 whereas when it goes below a lower threshold Z_2 then H_2 is the most likely cause of the perceived experience. If neither of the thresholds is crossed then an extra sample of evidence is considered, until the process distinguishes between the two hypotheses.

1.1.2 Psychological Models of Choice: Diffusion and Race

The first two psychological models to capitalize on the sequential sampling aspect of the SPRT were the race (Vickers, 1970) and the diffusion models (Stone, 1960; Laming, 1968; Ratcliff, 1978). The two models, albeit similar in assuming integrate-to-criterion mechanisms, differ in their definition of the criterion that terminates the decision. Lets denote the integrated evidence for alternative i at time t by:

$$y_i(t) = \sum_{\tau=1}^t x_i(\tau), \quad (1.3)$$

and assume a choice problem between two competing hypotheses. In the *race* model a choice is initiated once the integrated evidence, y_i , of any of the two alternatives i exceeds a threshold. In the *diffusion* model the process is terminated once the *difference* between the integrated evidence of the two alternatives exceeds a threshold. An equivalent way to implement the diffusion's stopping criterion is by assuming a single decision variable that accumulates the difference between the sensory evidence supporting the two alternatives, i.e. $d = y_1 - y_2$. In that case, the choice is made once the accumulated difference d crosses a positive (decision in favour of H_1) or a negative (decision for H_2) boundary.

As Bogacz (2009) has shown, the accumulated difference, d , between the two perceptual hypotheses is exactly proportional to the logarithm of R in equation 1.2. Therefore the diffusion model is an implementation of the optimal SPRT, providing the fastest decisions for a given accuracy level. For example if the thresholds of both the race and the diffusion model are set so as to produce correct responses 90% of the time, the diffusion model, on average, will be faster. The optimality of the diffusion model lies on its stopping rule which considers *relative* rather than *absolute* levels of evidence. Intuitively, the accumulation of relative evidence makes the diffusion model adaptable to the difficulty of the task at hand. If the losing alternative is much worse than the winning one, then the diffusion model will generate a fast decision. On the contrary if the evidence for the two alternatives is very ambivalent then the decision will take much longer because discrimination between the two hypotheses will be much harder. This is not the case in the race model which predicts that the stronger the evidence for the losing alternative is the faster the decision will be. This property of the race model, known as statistical facilitation (Raab, 1962; Townsend & Nozawa, 1995), can be conceptualized by thinking the decision process as two athletes running an independent race (i.e. not being able to help or hinder each other). Now two different races are assumed: in the first race, one fast athlete (F) runs against a medium one (M); in the second race, the same fast athlete (F) runs against a very slow one (S). On average, the termination time of the first race will be faster compared to the second race. This happens since the fast runner (F) is just as fast in both races but runner (M) is faster than runner (S). So, runner (F) loses more of his slower runs to runner (M) than to runner (S), resulting in a speed-up of the overall finishing times.

1.1.3 Variants of the classical diffusion model

Overall the diffusion model has outperformed race in fitting the choice patterns and the distributions of the response times in a variety of choice tasks and paradigms (Ratcliff & Rouder, 2000; Ratcliff, Gomez, & McKoon, 2004; Ratcliff & Smith, 2004). Moreover, at the neural level, the growth of information in the decision relevant neurons has been better approximated by the dynamics of the diffusion model (Ratcliff, Cherian, & Segraves, 2003) while microstimulation studies have provided support for the accumulation of relative, rather than absolute (i.e. race), evidence (Ditterich, Mazurek, & Shadlen, 2003; Hanks, Ditterich, & Shadlen, 2006). Despite its success in accounting for behavioural and neural data, the classical drift-diffusion model has been challenged by a striking behavioural pattern; for weak evidence, choice accuracy does not improve with time after some critical interval of about 400 ms but rather saturates at an asymptotic level (Swensson, 1972). This finding contradicts the principle of perfect integration, held within the classical diffusion, according to which the quality of the decision should keep improving with the accumulation of more samples of evidence.

Two assumptions have been proposed, within the drift-diffusion model, in order to account for the saturation in decision accuracy. First, the drift or the mean value, μ_i , of the evidence is subject not only to variability *within* a trial, due to external noise in the sampled evidence or internal noise in the neural responses, but also to variability *between* trials of the same experimental condition (Ratcliff & Rouder, 1998). In other words, the mean drift value is not fixed but sampled in each trial from a normal distribution with mean μ_i and standard deviation σ_d . For example in the motion discrimination task where the observer needs to decide the dominant direction in the movement of a cloud of dots, in a small fraction of trials with rightward dominant motion, the sampled dominant direction will be leftward. Therefore decision accuracy will eventually saturate since in trials such as the aforementioned, observing the signal for longer does not improve the decision quality.

The second assumption that has been employed to account for the accuracy saturation holds that the integration of evidence stops when an *absorbing* boundary is reached even under conditions in which the response time is under the experimenter's control (*bounded diffusion*, Ratcliff, 2006). In practice this means that once the boundary is reached the decision is finalized and later evidence is still perceived but virtually neglected. Note that while the drift-variance diffusion maintains the principle of the

SPRT that all pieces of evidence are equally weighted, the bounded diffusion assigns higher weights to the early evidence, since late evidence will be ignored in those trials where the boundary is reached before the end of the stimulus presentation. Huk and Shadlen (2005) provided empirical support for the bounded integration mechanism by applying small perturbations to the evidence and showing a larger impact of the perturbations that were applied at the beginning of the trial. More recently, the bounded diffusion model and the primacy bias were supported in a neurophysiological study (Kiani et al., 2008). Using the interrogation procedure (i.e. response under experimental control), Kiani et al. (2008) measured behavioral and neural responses from monkeys that engaged in the discrimination of the motion direction of moving-dot stimuli that varied both in the motion coherency (i.e. quality of the evidence) and in the time for inspection. First, the behavioral data confirmed that for coherency levels below 13%, accuracy asymptotes at a level that does not exceed 90%, and that the saturation comes in play for intervals longer than 420 ms. Second, a reverse correlation analysis² revealed that the evidence was weighted by primacy, as predicted by the bounded diffusion model (Figure 1.1).

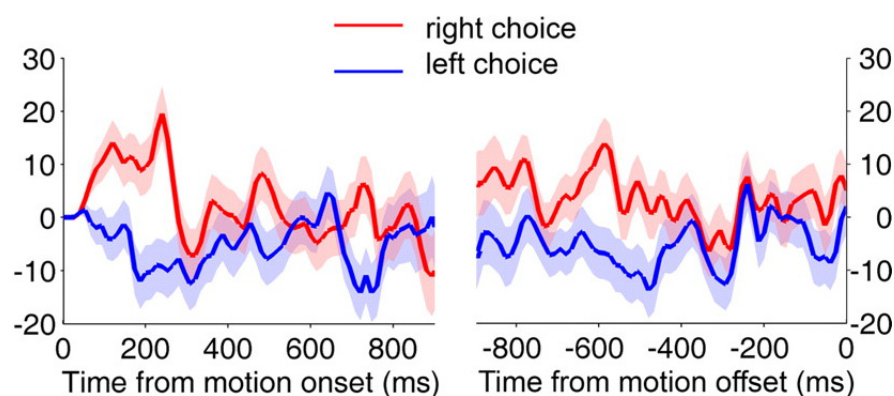


Figure 1.1: *Reverse correlations (reproduced from Kiani et al., 2008) showing primacy. Left, signals aligned with motion onset. Right, signals aligned with motion offset. The difference between the evidence that favours the response (red) and the one that opposes it (blue) is larger at the beginning of the trial.*

These results indicating primacy bias, stand in contrast with an experimental study by Usher and McClelland (2001) (Experiment 3), that was performed to examine the

²In this analysis the zero-coherency trials, where the motion of dots is totally random, are considered. The average evidence that favoured the chosen alternative is compared to the average evidence for the non-chosen alternative. This technique is similar to a logistic regression, showing the relative weight of evidence at different moments.

temporal weighting of evidence. In that experiment participants saw a fast sequence involving two alternative letters (H/S) and had to decide what letter was presented more often. While most of the trials (i.e. regular trials) involved sequences with a majority of either S or H, a subset (i.e. experimental trials) of them was constructed with equal fractions of H and S. In these trials, the time-course of the events was manipulated so that the one alternative was better in the first half and worse in the second. Unlike to the clear primacy results in Kiani et al. (2008), individual differences were observed. Among six observers, some showed primacy bias, others recency bias and others were perfectly balanced, having also the highest accuracy in the regular trials. As discussed above, the diffusion model and its variants is able to account only for uniform or primacy weighting, being unable to capture any recency bias in the integration of evidence. In the next subsection I discuss a more versatile neural model, the leaky competing accumulators (LCA, Usher & McClelland, 2001) which, under different parameters, extrapolates between balanced and temporally weighted (i.e. primacy and recency) evidence integration.

1.1.4 A Neural Model of Choice: The Leaky Competing Accumulators

Recently, a series of neurocomputational models have offered an explanation of the neural mechanism underlying both psychological measures and neurophysiological data of perceptual choice. One such model is the Leaky Competing Accumulator (LCA, Usher & McClelland, 2001), which is sufficiently simple to allow a detailed mathematical analysis. Like in the race model, in LCA a choice between two alternative hypotheses is encoded in the activation states of two accumulators that race against each other. However, unlike to race where the accumulators are independent, in LCA the two decision units compete to each other via lateral inhibition. Another important principle of the LCA model is that it assumes that the decision units are subject to decay or leakage, having a finite effective time constant. The leaky nature of the integration, allows the model to explain the imperfect accuracy levels even at long observation intervals, since integrating evidence for much longer than the integration time constant does not increase the total amount of the accumulated evidence.

For binary choice, the LCA is a stochastic 2 dimensional system, described by two variables, y_1 and y_2 , that correspond to the accumulated evidence in favor of the two

alternatives. Each accumulator unit, y_i , integrates evidence from an input unit with mean activity I_i and independent white noise fluctuations dW_i of amplitude c_i (dW_i denote independent Wiener processes). These units also inhibit each other by way of a connection of weight β . Hence, during the choice process, information is accumulated according to:

$$\begin{aligned} dy_1 &= (I_1 - \kappa y_1 - \beta f(y_2))dt + c_1 dW_1, \\ dy_2 &= (I_2 - \kappa y_2 - \beta f(y_1))dt + c_2 dW_2, \end{aligned} \quad (1.4)$$

with κ denoting the leakage and f a non-linear activation function which truncates negative values of the accumulators back to zero³ before they inhibit each other. The LCA model in the free response paradigm, where the observer is free to respond at any time, assumes that a response is made as soon as either accumulator exceeds a predefined threshold Z . In the interrogation paradigm, where the response time is under experimental control, a choice is made in favour of the accumulator with the highest activation at the moment when the choice is requested.

For positive y -values the LCA behaviour can be approximated by the difference in the activation of the two accumulators, $y = y_1 - y_2$, as a 1-dimensional diffusion. In that case the model becomes equivalent to an Ornstein-Uhlenbeck (OU) diffusion with a leak or expansion coefficient (Busemeyer & Townsend, 1993; Usher & McClelland, 2001):

$$dy = (I - (\kappa - \beta)y)dt + cdW, \quad (1.5)$$

The LCA, depending on the values of leak and inhibition, weighs evidence either differentially, based on its temporal order, or uniformly. First, when the decay exceeds the inhibition parameter ($\kappa > \beta$), a leaky diffusion takes place. In this case the activation difference decays, resulting in both bounded accuracy and recency bias. If, on the other hand, the inhibition parameter exceeds the leak ($\kappa < \beta$), the diffusion coefficient becomes positive, resulting in expanding trajectories and unbounded $y_1 - y_2$

³The non-linearity assumption in the LCA model is biologically motivated: the activation states of the accumulators correspond to firing rates of neuronal populations and therefore cannot be negative. The non-linearity of the LCA is computationally efficient for multi-alternative choice problems (Bogacz, Usher, Zhang, & McClelland, 2007), since it discards, early in the process, poor and uninformative options by inhibiting their accumulators to zero. In a recent study we demonstrated that the non-linearity is also the key mechanism in accounting for context effects in alternatives with non-stationary evidence (Tsetsos, Usher, & McClelland, 2011).

differences; in this case small differences in the activations of the two accumulators that occur early on are expanded, resulting in a primacy bias. Finally, for the special case in which the leak and the inhibition are in balance, the model behaves optimally, mimicking the diffusion model (Bogacz et al., 2006).

1.1.5 An Experimental Study on Order Effects

The *inhibition-dominant LCA* and the *bounded diffusion* both weigh evidence by primacy, consistent with the experimental findings in Kiani et al. (2008). If primacy was a universal property of evidence integration these two models would be equally good candidates for explaining the underlying mechanisms of choice. However, as the study by Usher and McClelland (2001, experiment 3) indicates, evidence weighting is quite diverse involving, apart from primacy and unbiased weighting, also recency. Before delving into model comparison, which would promote the more versatile LCA that is able to account for interpersonal differences (i.e. capturing both primacy and recency), it is important to understand whether the profile of evidence weighting depends on contingencies such as task demands. To achieve so, we conducted two experiments with human observers (Tsetsos, Gao, et al., 2011), using the moving dots paradigm which provides optimal control of the evidence manipulation and a relative long integration-interval.

The first study followed a design similar to Kiani et al. (2008). Observers were asked to discriminate the direction of moving dots displays. The coherency and the duration of the displays were varied and the observers were trained to respond within a window of 300 ms from a response signal that appeared upon stimulus termination. As in Kiani et al. (2008), the duration of the trials was exponentially distributed, ranging from 150 to 1750 ms. The critical manipulation was applied in a subset of trials (80% of the trials with durations 300 ms or longer), in the form of a pulse which resulted in a change in the motion coherency of $\pm 3.2\%$ for 200 ms. All the observers learned to respond within the 300 ms response window and their accuracy increased with motion-coherency. Figure 1.2 shows how the pulse shifts the psychometric function, for one of the observers (similar results were obtained for the rest of the participants). Of special interest is the effect of the timing of the pulse on this shift for all observers (Figure 1.3), with the effect being larger at the start of the trial, indicating primacy and replicating the results of Kiani et al. (2008) with human observers.

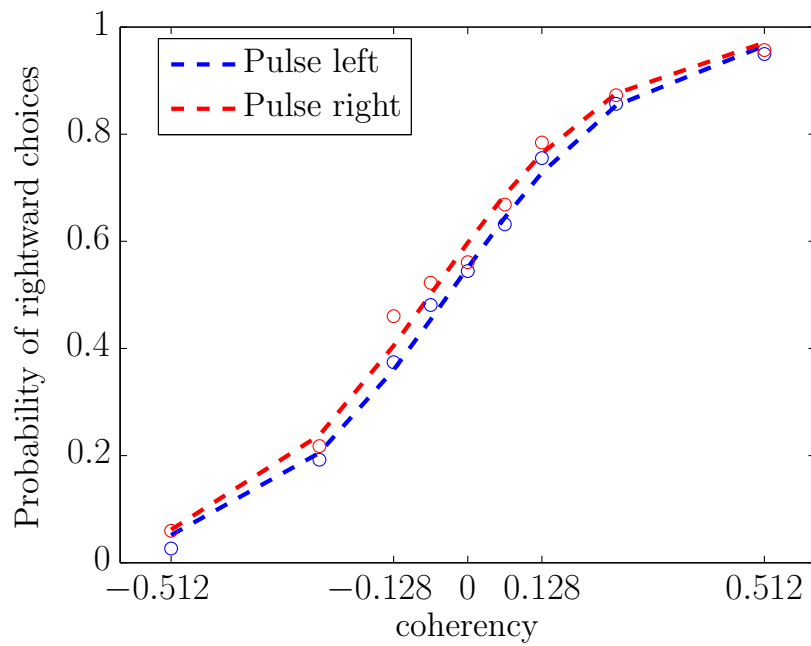


Figure 1.2: *Effect of motion pulse on detection of motion for one observer. Red line shows probability of choosing “right” when a rightward perturbation is introduced and blue curve the same for leftward perturbations. The pulse is equivalent to a change of $\pm 3.2\%$ in coherency.*

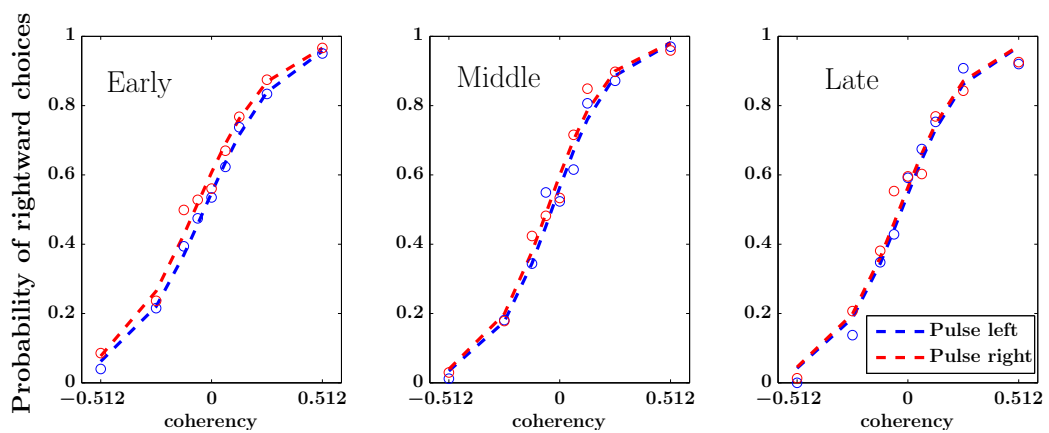


Figure 1.3: *The effect of the pulse timing on the shift of the psychometric function for one observer from early in the trial (left) to late (right).*

The second experiment was carried out in order to obtain a more robust measure for the recency-primacy bias and understand it with regards to the task characteristics. To do so, for each coherency and duration, 3 conditions were created: i) the constant condition corresponding to a fixed coherency during the whole trial (as in Experiment 1), ii) the early condition, corresponding to a fixed coherency during the first half of the stimulus, which is set to zero (random motion) during the second half, iii) the late condition in which the first half has random motion and the second half a fixed coherency. Comparison between the accuracy in early and late conditions provides a measure of the order bias. The experiment was conducted on two groups with the between subjects factor being the task characteristics. In particular for the first group the trials duration followed an exponential distribution from 150 to 1750 ms while the response deadline was 300 ms, identical to Experiment 1 (and to Kiani et al., 2008). For the second group the trial durations were uniformly distributed from 150 to 1750 ms and the response deadline was relaxed to 1000 ms. The primacy score, calculated as the mean accuracy in the early condition minus the mean accuracy in the late condition, was significantly larger for the first group (11% against 2%; $t[8] = 2.98$; $p < 0.02$). Furthermore, while all the observers in the short-deadline condition showed the primacy effect, there was considerable variation in the bias for the observers in the second group.

This reduction in the primacy bias can be understood in relation to the two procedural differences between the two groups. The first difference is the response deadline which was stringent in the first group (i.e. 300 ms) and relaxed in the second (i.e. 1000 ms). A stringent deadline encourages participants to prepare a response before the stimulus termination and response signal, so that they do not miss the deadline. The second procedural difference involves the exponential distribution of stimulus duration in the first group which changed into a uniform distribution in the second group. The reason why Kiani et al. (2008) used an exponential distribution for the stimulus duration was to ensure that observers have no information about the time when the response signal will appear⁴. On the other hand it is possible to argue that knowing that the stimulus is about to end during the course of the trial is less relevant to temporal biases than knowing that the trial duration is likely to be short in advance of the trial. The latter is the case when the trial duration is selected from an exponential distribution, resulting in an optimal policy (assuming a capacity limited resource pool) that allocates more

⁴An exponential distribution has a flat hazard function which means that the observer cannot tell in the course of the trial whether the trial is ending or not. On the other hand the uniform distribution has a peaked hazard function and when the time gets close to the upper bound of the distribution the observer knows that the response cue is approaching.

attention to the evidence at the start of the trial.

No clear recency was found in any of the subjects in the two above experimental studies and also in Kiani et al. (2008). One factor that may explain the lack of recency, is the degree of practice which was quite extensive. As Brown and Heathcote (2005) suggested, practice increases the efficiency of evidence accumulation by reducing the effective leak. Future research is thus needed to better understand the various factors that affect the temporal weighting of evidence. Nevertheless, the conclusion drawn by Kiani et al. (2008) that bounded integration is a universal decision principle that applies not only to self-paced decisions but also to experimentally controlled ones, needs to be reconsidered and understood with respect to the task contingencies. At the theoretical level both the bounded diffusion and the LCA model are able to account for the inefficiency in evidence integration (e.g. imperfect performance even at long durations). However within the bounded diffusion model this inefficiency is exaggerated by assigning to late evidence, not lower (i.e. like in LCA) but zero weight.

1.1.6 Summary

When people make decisions on the basis of sensory information they integrate noisy samples of evidence up to a criterion. This mechanism, motivated by the statistically optimal SPRT which conceptualizes choice as a Bayesian inference problem, is central in a series of mathematical models. The most successful of them has been for years the *diffusion* model which holds that the *relative*, rather than the absolute, evidence is accumulated across time, virtually implementing the optimal bayesian procedure. One direct implication of the perfect integration that the diffusion model assumes, is that all pieces of evidence are equally weighted. Despite its success in accounting for a rich set of psychological and neural data, the classical diffusion model has been challenged by the fact that human choice behaviour is often suboptimal. First, for weak evidence, observers' performance is not perfect even when the observation time is infinite. Second, a series of experimental studies revealed temporal biases in the weighting of evidence.

These two challenges have been addressed by two separate revisions of the diffusion model. The saturation of performance for long intervals was accounted by a *diffusion* with *drift-variance* which, however, maintains the balanced and unbiased weighting of evidence across time. The second challenge of temporal biases was accounted by

a *bounded diffusion* which assumes that evidence integration ceases and the choice is finalized once the decision boundary is reached, even if new evidence keeps flowing in. The bounded diffusion model gives a clear-cut primacy prediction, accounting also for the saturation in accuracy. Nevertheless human choice behaviour is more diverse being subject to both primacy and recency biases. The diffusion framework though is unable to capture this diversity in evidence weighting, producing either unbiased or primacy-biased choices.

This is not the case with the more flexible and neurally inspired LCA model. The LCA model assumes that evidence integration is leaky with competition among the decision units triggered by lateral inhibition. The leakiness of the accumulation process naturally accounts for the accuracy saturation, since integration for a period longer than the effective time constant does not improve the decision quality. Additionally the balance between the leak and inhibition parameters is able to generate three different modes of temporal weighting. When inhibition is larger than leak, then the model exhibits primacy. On the other hand when the model is leak dominant then recent information is overweighted. Finally when leak and inhibition are in balance then the model mimics the drift-variance diffusion.

In a recent neurophysiological study with primates (Kiani et al., 2008) strong primacy biases were obtained and the bounded diffusion model was proposed as the most successful mechanistic account of perceptual choice. In a follow up study we did (Tsetsos, Gao, et al., 2011), we first replicated the primacy bias with human observers and subsequently revealed that it is mainly triggered by task contingencies. In particular, we showed that a stringent response deadline and an anticipation regarding the short length of the trial, prior to the trial onset, urges observers to assign higher weight to early evidence. We concluded that choice dynamics are adaptive to the type of the task and its characteristics and that more experimental work is required to understand the basis of temporal biases. In the forthcoming chapters I will experimentally explore temporal biases, emphasizing the role of recency (which challenges most existing models) and questioning whether the same profile of temporal weighting is maintained across domains, i.e. in both perceptual and value-based decisions. Next I turn into reviewing the other important factor of differential information weighting that this thesis addresses: context-dependent weighting.

1.2 Context Effects in Value-based Choice

As we saw in the previous section, recent advances in the neuroscience and psychology of perceptual choice has led to impressive progress in understanding and characterizing the central underlying process of evidence-integration, indicating a mechanism that is regarded to be nearly optimal. This stands in contrast with the more challenging field of complex, value-based decisions, such as deciding what car to buy or which flat to rent, in which we are faced with a number of puzzling patterns like preference reversals and decision-biases, as highlighted in the seminal work of Tversky and Kahneman (Kahneman & Tversky, 1979; Tversky & Kahneman, 1981; Huber et al., 1982; Tversky & Kahneman, 1986; Knetsch, 1989; Simonson, 1989; Kahneman & Tversky, 2000; Gilovich et al., 2002). The seemingly endless anomalies characterizing human choice behaviour, have directed the field of behavioural economics and psychology into mostly experimental routes, in an attempt to reveal as many as possible regularities and paradoxes that contradict the economic theory (Von Neuman & Morgenstern, 1947; Debreu, 1960). As a consequence of using the normative account as a measure of comparison to human choice behaviour, theorizing preferential choice was greatly influenced by the formalisms of economic theory and the most successful models were originally descriptive modifications of expected utility theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992)⁵. More recently though the field has developed further and several different approaches of increased predictive and explanatory power, that address decision-making at the process level, were developed ⁶.

Despite the recent promising advances from the early descriptive models to process-based and explanatory ones, all theoretical frameworks still face a central puzzle; people's preferences between options can be reversed by the presence of decoy options (that are not chosen) or by the presence of other irrelevant options added to the choice set. Three types of contextual reversal effect reported in the decision-making literature, the attraction, compromise, and similarity effects, have been explained by a number of independent proposals. Yet a major theoretical challenge is capturing all 3 effects simultaneously. In the next sections, I review the range of mechanisms that have been proposed to account for decoy effects and analyse in more detail two computational

⁵In a recent review article by Vlaev, Chater, Stewart, and Brown (2011), these theories are labelled as "value-first".

⁶According to Vlaev et al. (2011) in many of these models comparison is a central mechanism. These models can be further classified into those that assume a representation of value versus others that rely merely on ordinal comparisons, without assuming any internal value scales.

models at the process level, decision field theory (Roe et al., 2001) and leaky competing accumulators (Usher & McClelland, 2004), that aim to combine several such mechanisms into an integrated account. I argue that the LCA framework, which follows on Tversky's relational evaluation with loss aversion (Tversky & Simonson, 1993), provides a more robust account, potentially implying that common mechanisms are involved in both high-level decision making and perceptual choice, for which LCA was originally developed. The latter hypothesis is furthered pursued in this thesis where the information integration mechanism, in both perceptual and value-based decisions, is probed in relation to context effects.

1.2.1 Preference Reversal in Multi-attribute Choice

Confronted with an unusually short dessert menu, Ms. X vacillates between two options, A and B. Finally, she plumps for A, at which point the waiter responds that, in fact, there is also the daily special, Option C. "Thank goodness you told me that," says Ms. X, relieved, "In that case, I'd prefer B." There is something paradoxical about Ms. X's change of heart. How can the availability of a third option, C, possibly affect whether A or B is preferred? The relative pleasure of eating Dessert A or B surely should depend on the properties of A and B alone and not on the properties of any other dessert C, whether that C is an available option or not. To hammer home how paradoxical any influence of C might be, let us push the story a little further. The waiter returns with Dessert B and says, "Actually, the chef has just told me that C is sold out." "In that case, I'd like to switch back to A, please," decides Ms. X.

The puzzling behaviour of Ms X in this situation is a case of contextual preference reversal. It is fascinating that such reversals have been reported to characterize human decision making between alternatives that vary on several dimensions, as illustrated in Figure 1.4, where one has to choose one out of several cars that vary on two attributes (i.e. economy and quality). Three such reversal effects have been reported in the literature. The most puzzling of them are the *attraction* effect (Huber et al., 1982) and the *compromise* effects (Simonson, 1989), which have the form of Ms X's preference reversal. For the attraction effect, the irrelevant option D is a decoy (an inferior or dominated option), similar but of less value than A. For the compromise situation, the option C is of approximately equal value as A and B, but it is such that it places B in the 'middle' within the 2D attribute space, making it a compromise. A third and

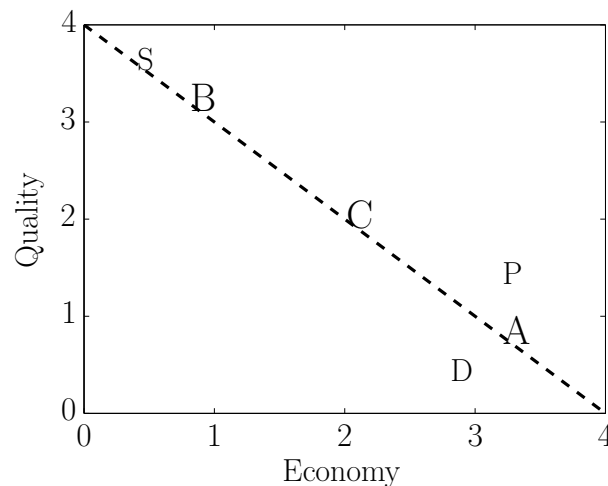


Figure 1.4: *Illustration of a choice space for options that vary on two dimensions. The pattern of preferences between A and B can be affected by the presence of other, irrelevant options (C, D, P, S) in the choice set*

perhaps less puzzling choice reversal is the *similarity* effect (Tversky, 1972). Here the introduction of a new option S, very similar to B (and of equal value), shifts the relative choice between A and B in favour of the dissimilar option, A. More recently, a new type of decoy effects, the *phantom decoy* effects, has been observed (Pratkanis & Farquhar, 1992; Dhar & Glazer, 1996; Choplin & Hummel, 2005; Pettibone & Wedell, 2000, 2007) in which the introduction of an unavailable but dominant option (P in Figure 1.4) can bias the decision towards the similar dominated option (A). These phantom decoy effects raise an additional challenge to the theory of choice (Pettibone & Wedell, 2007).

Such paradoxical preference reversals are, not surprisingly, ruled out by many theories of choice. In particular, they are ruled out by any theory of choice which separately assigns some goodness value to each option and proposes that people always or more likely (if the choice mechanism is stochastic) prefer options with higher goodness values. I shall call such accounts *option-based* theories (also known as simple scalable choice models) where the crucial assumption is that a value is assigned *independently* to the available options and choice is determined by the comparison of values.

Option-based accounts of choice, require that whether A is chosen rather than B depends on the relative value of A and B. And, by assumption, these values are determined by independent consideration of each option. No further option, C, can affect the relative values of A and B. Value-based account of choice include expected utility

theory, the cornerstone of economic theory and rational choice explanation (Von Neuman & Morgenstern, 1947; Debreu, 1960). Moreover, they apply to any variants of such theories which allow noise, either in the assignment of goodness values, or the decision between goodness values [e.g., stochastic expected utility (Blavatsky, 2007)]. This class is broad, and includes many psychological theories of choice, including, for example, prospect theory (Kahneman & Tversky, 1979; but see Tversky & Simonson, 1993) for a prospect theory variant that allows contextual preference reversal).

How can such apparent anomalies be explained? As we shall see, a wide variety of theoretical proposals have been put forward, although no single mechanism accounts for all three decoy effects. What is required is an integration of several mechanisms into a single computational model. Here, we analyse two such models, both based on principles of neural computation: the Decision Field Theory (DFT; Roe et al., 2001) and the Leaky Competing Accumulators (LCA; Usher & McClelland, 2004). The aim of this section is to compare in a systematic way DFT and LCA in their account to reversal effects and to derive novel predictions from these models (see also Pettibone & Wedell, 2007 for a comparison of models focused on phantom decoys).

In the next subsection, I will review the variety of mechanisms that have been proposed to explain preference reversal and will clarify which mechanisms explain which effects. Then, two neurocomputational approaches, DFT and LCA, will be described in relation to the core theoretical mechanisms; and consider the similarities and differences between them while contrasting their predictions. Two apparent problems of the DFT approach will be raised and evaluated against empirical data; first undesired predictions due to local inhibition and linearity and second lack of robustness of the correlational mechanism which accounts for the compromise effect.

1.2.2 Mechanisms for reversal effects

Before plunging into details concerning DFT and LCA, it is worth considering, in general terms, how a third option might influence the choice between two existing options. There are three broad classes of mechanism, based on: i) attentional switching to different choice aspects, ii) relational, rather than independent, evaluation of properties and loss-aversion, iii) value-shifts or contrast effects, mediated by lateral inhibition. I consider these briefly in turn.

1.2.2.1 Attention to choice aspects and temporal correlations

As shown by Tversky in his elimination by aspects (EBA; Tversky, 1972), the similarity effect follows immediately, and fairly uncontroversially, from a stochastic criteria shifting mechanism. Assume that, while struggling to choose between tiramisu and fruit salad at some moments, Ms X is swayed by taste (favoring the tiramisu); at some moments she is swayed by health (favoring the fruit salad). That is, her criterion for choice (or in the language of the EBA, her attention to the choice-aspects) is continually shifting. Suppose that there is a 0.6 probability that she will choose fruit salad. But before she can choose, then waiter points out that there is a third option, “fruit surprise”, which turns out to be almost exactly the same as, and no better or worse than, fruit salad. Ms X resumes her oscillations between taste and health. Now if, as before, there is a 0.6 chance that health will win out, and she will choose fruit, note that she now has a further choice: between fruit salad and fruit surprise. If she makes this choice randomly, then the probability of choosing fruit salad is now 0.3—i.e., less than the 0.4 probability of choosing tiramisu. But before “fruit surprise” was added, the probability of choosing fruit salad was greater than the probability of choosing tiramisu.

The preference reversal described above can also be seen as an instantiation of a more general principle of fluctuating and temporally correlated preference. What happens to Ms X above, is that her preferences fluctuate and that the preference for fruit-salad and fruit-surprise are positive correlated (they raise and fall together). In this case the correlation is caused by the switching of attention to different choice-attributes, but as we will see below, such correlations can be also caused by other mechanisms. The general idea, however, is that when temporal correlations between momentary preference exist, the correlated options split their wins, and hence loose share, relative to the uncorrelated options.

1.2.2.2 Relational evaluation of options and loss-aversion

The impact of relational, rather than independent, evaluation of options or properties is best illustrated by considering the attraction effect. This corresponds to the addition to the menu of a “second-tiramisu”, which is just like tiramisu, but marginally inferior in every way (or, more strictly, marginally inferior in at least one way, and no better in any other way). Now consider the relative goodness of each option. If we are not

sure how to weigh up the different dimensions of desserts, we may feel that fruit salad is roughly as good as tiramisu; and fruit salad is roughly as good as second tiramisu; but however we weigh the dimensions, it is clear that tiramisu is better than second tiramisu. The specific account of why tiramisu is now relatively favoured can take various forms. For example, according to reason-based decision making (Simonson, 1989; Pennington & Hastie, 1993), people choose by searching for a justification for their choice. The choice of tiramisu may be justified by its clear superiority to second tiramisu (i.e., it is clearly relatively better, even if we are not sure how much we like either option, in absolute terms); but fruit salad has no clear justification, being difficult to compare with either alternative option.

Alternatively, both the attraction and the compromise effects could be explained, without appealing to a justification process ⁷, by assuming that values are computed via pairwise comparisons. For example, we might assume that each option is compared with each other option and that the differences, advantages or disadvantages (on each dimension, separately) are transformed into utilities via a value function (Tversky & Simonson, 1993), characterized by loss-aversion (a steeper slope in the domain of losses than in that of gains, so that losses loom larger than gains)⁸.

1.2.2.3 Inhibition as contrast-enhancement between similar options

An alternative way to explain the attraction effect is a type of local contrast enhancement, as observed in visual perception (e.g. a circle appears larger when surrounded by smaller circles; Massaro & Anderson, 1971). One mechanism that can mediate such a process is lateral inhibition between similar items, so that only alternatives that are similar inhibit each other. To cause an enhancement of the dominating option, one needs to assume that the local inhibition operates on a relational attribute evaluation function (inferior options, A' have negative values, while superior options A have positive values; thus A' causes an enhancement in the value of A since passing negative activation via an inhibitory link results in excitation).

The mechanisms described above are not the only ones that can account for reversal effects. Other mechanisms, such as dimensional weight change, ranking, grouping,

⁷Note that decoy effects have been found in other non-human species (Hurly & Oseen, 1999; S. Shafir, Waite, & Smith, 2002) suggesting that justification is not crucial for such effects to occur.

⁸Loss aversion explains the endowment effect (Knetsch, 1989) reflecting the fact that people tend to stick with the current choice because they overweight losses incurred from switching, relative to gains.

etc, have been proposed in various models (Guo & Holyoak, 2002; Stewart et al., 2006; Pettibone & Wedell, 2007). I focus here on these three mechanisms, because they are used in the models that are contrasted. In particular, the first two are used in the LCA, which implements key elements of two of Tversky's models, EBA (Tversky, 1972) and the context-dependent advantage model (Tversky & Simonson, 1993), while the first and the last are used in DFT.

1.2.3 Two Neurocomputational Approaches

Although the mechanisms described above can explain the various decoy effects, no single mechanism appears to explain the full range of effects. A computational account integrating several mechanisms appears to be required to provide an adequate explanation of the effects and make parametric predictions for choice as a function of how the options are situated in the attribute-space. Recently, a number of dynamical theories of value-based decision making have been proposed accounting not only for the choice outcome, but also for the dynamics of the decision process as it unfolds over time. In contrast to heuristics and computational theories with static parametrisation, dynamical models can make predictions on temporal aspects of decision-making such as vacillations and decision times, and they are also in the position make contact with recent neurophysiological studies of perceptual choice. Here we focus on two such theories, the decision-field theory for multi-attribute choice (DFT; Roe et al., 2001) and the leaky-competing accumulators (LCA; Usher & McClelland, 2004), which account simultaneously, for all the three contextual reversal effects.

Both DFT and LCA conceptualize choice as an Ornstein-Uhlenbeck (OU) diffusion process (see also equation 1.5 in the previous section for similar dynamics), or in other words, a leaky integration of preference subject to choice competition and driven by attentional shifts. This allows both models to account for the similarity effect, following Tversky (1972) as a result of a stochastic attention shift. Despite many processing similarities between the models there are also few important differences. While DFT is a linear model, which has the appeal of mathematical tractability, the LCA assumes two types of non-linearity. The first concerns value of the activations which (corresponding to firing-rates) are not allowed to go negative. The second non-linearity is carried over from prospect theory, in the form of an asymmetric value function with loss-aversion (losses weighted higher than gains), which is taken by LCA as a primitive. Unlike the

LCA, which maintains most of the aspects of Tversky's theories, DFT does not assume loss-aversion as a primitive but rather derives it as an emergent property. To do so, it assumes that the inhibition between the choice alternatives is a decreasing function of their similarity in the attribute space.

Below the exact instantiation of the two models is reviewed. Both models are instantiated in four-layered connectionist networks as illustrated in Figure 1.5. The first layer corresponds to the choice-attributes (two attributes are illustrated here). In both models, it is assumed that the attention of the decision maker switches stochastically across dimensions (D_1, D_2), according to a Bernoulli process hence at any time step, only one of the attributes is active. The 2D-characterization of each alternative on the D_1/D_2 space (Figure 1.4), is given by the connectivity between the first and the second layer (i.e., a 2×3 matrix m_{ij}). Each node in the second layer corresponds to the integrated attribute values of each choice-alternative according to the following equation:

$$U_i(t) = \sum_{j=1,2} w_j(t) m_{ij} + \varepsilon_i(t), \quad (1.6)$$

where ε is the probability of attending irrelevant dimensions, and w_j is 0 or 1 depending on which dimension is attended.

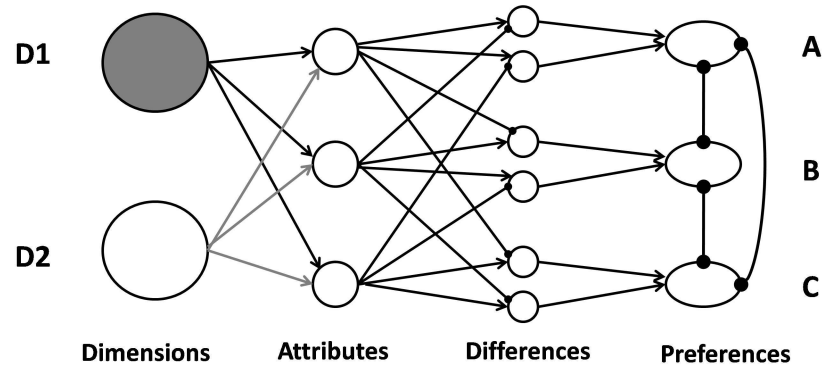
The two models differ slightly on the intermediate computations performed in the third layer and on the way in which the preferences are integrated in the fourth layer. In DFT the third layer computes contrasts between each option and the other alternatives (also mentioned as valences), as the difference between the value of the option and the mean value of the other options, with respect to the active dimension:

$$v_i(t) = U_i(t) - \frac{\sum_{k \neq j} U_k(t)}{n-1} \quad (1.7)$$

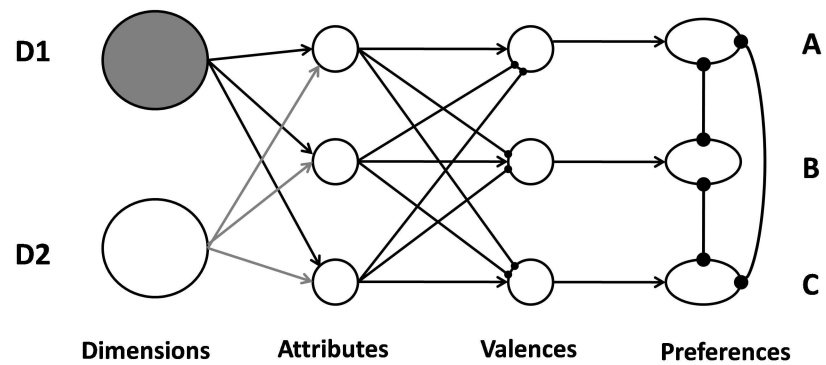
In LCA, the third layer computes advantages and disadvantages between all pairs of options, which are transformed by a non-linear, asymmetric (loss-averse) value function:

$$I_i(t) = \sum_{j \neq i} V(d_{ij}) + I_0, \quad (1.8)$$

with d_{ij} being the advantage or disadvantage of option i to option j on the active dimension, V is a non-linear value function with loss-aversion and I_0 is a positive



(a)



(b)

Figure 1.5: Illustration of decision field theory and leaky competing accumulators models in neural networks; circle arrow heads correspond to inhibition. a: Connectionist network for decision field theory. b: Connectionist network for leaky competing accumulators.

constant that promotes the alternatives to the choice process, namely, prevents the I_i of the inferior options from being negative. Finally, in both models, the fourth layer integrates the contrasted differences (valences or sum of advantages/disadvantages in DFT and LCA respectively), as preferences, across time. For DFT:

$$P_i(t+1) = v_i(t) + \sum_j s_{ij} \cdot P_j(t) + \xi(t), \quad (1.9)$$

and for LCA:

$$P_i(t+1) = I_i(t) + \sum_j s_{ij} \cdot P_j(t) + \xi(t), \quad (1.10)$$

with ξ standing for additive, Gaussian noise.

The integration of preference for each option is imperfect (leaky) and subject to competition with the preferences of the other options (see equations 1.9 and 1.10). The leaky integration of preferences and the competitive interactions between the options are implemented in a connectivity matrix, whose diagonal term corresponds to a self connectivity coefficient (or the leak parameter), and whose off-diagonal elements s_{ij} , correspond to inhibitory connections. While in LCA all the off-diagonal elements are constant (global inhibition), in DFT their magnitude depends on the distance between the alternatives i and j , in the 2D attribute space). Finally, as mentioned above, DFT is linear and thus preference states can take both positive and negative values, as opposed to LCA where negative activations at the fourth layer, are truncated to zero.

While the two models explain identically the similarity effect, their explanation for the attraction and compromise effects are very different. In DFT it is the contrast enhancement mediated by local inhibition that accounts for the attraction effect; the value of the dominating option, A, is enhanced by the similar decoy. In particular, the similarity between nearby alternatives (A and D in Figure 1.4) results in their being coupled by strong local inhibition. As option D is inferior to both A and B, it has negative valence. Therefore, option D boosts the preference of option A by passing its negative activation value through a negative connection (I will call this activation by *negated inhibition*). The function that specifies the local inhibition relates the psychological distance (i.e., similarity) of the options and the degree they compete by lateral inhibition.

The compromise effect is also accounted by DFT due to the distance-dependent inhibition; however, the key mechanism is correlation, not contrast enhancement. In this case, the extremes (A and B) and the compromise (C) interact via strong inhibitory

links whereas the extremes, A and B, are too distant from each other to compete. As the extremes do not inhibit each other, while they inhibit the compromise option, their momentary preference becomes decorrelated from the compromise, but correlated with each other. Thus the correlated extremes split their wins making the compromise option stand out and take a larger share of choices (see Roe et al., 2001 for details).

Unlike in DFT, the LCA account of the attraction and the compromise effects is similar to the context-dependent advantage model (Tversky & Simonson, 1993), and does not require a distance-dependent inhibitory mechanism. Instead, it follows the principles suggested by Tversky and Simonson (1993), according to which the value for each option is evaluated in relation to all other options in the choice set (so far, this is not fundamentally different from DFT) via a nonlinear loss-aversion value-function. In particular, for the attraction effect when option D is introduced, option B is penalized more by having two large disadvantages (relative to A and D, when dimension of Economy is attended) relative to A (which has one large disadvantage only). The same principle, helps the LCA account for the compromise effect; the extreme options (A and B) have one large and one small disadvantage each, whereas the compromise option has two small disadvantages. Due to the asymmetry of the value function, large disadvantages are penalized more, favouring that way the compromise option. A summary of the accounts that each model gives for each effect is given in Table 1.1.

Table 1.1: A summary of DFT and LCA accounts for the contextual preference reversal effects.

Effect	Model	
	DFT	LCA
Similarity	Attentional switching across dimensions	Attentional switching across dimensions
Attraction	Excitation by negated local inhibition	Loss-averse value function
Compromise	Correlations due to local inhibition	Loss-averse value function

1.2.4 Contrasting DFT and LCA

In this section, I explore parametrically how the choices depends on the locations of the choice alternatives in the attribute space. Specifically, two options (i.e., A and B) remain constant, while the third option (i.e., C) moves across the two dimensional

space with an increment of .05 at each step. I consider only results at or below the diagonal between A and B, as above the diagonal option C is always chosen. Details about the parameters used for each model are given in Tsetsos et al. (2010). In Figure 1.6 I present the magnitude of the attraction and similarity effects with respect to option A, as the difference between the probability of choosing A and the probability of choosing B, for different locations of the option C in the 2-d lattice. I use a gray scale, where brighter points correspond to a stronger enhancement of the preference of A by the introduction of C. For both models we can see the similarity effect illustrated as a thin white line close to option B (1,3) and adjacent to the diagonal (i.e., the introduction of an option C similar- neither dominating nor dominated by- option B results in boosting the preference for the dissimilar option A).

However, predictions for the attraction effect diverge. For the LCA model (Figure 1.6(b)) the attraction effect is present in the triangular white area close to option A. The magnitude of the effect gradually decreases as the distance between the decoy (option C) and the target (option A) increases. On the other hand DFT gives a more dichotomous prediction regarding the magnitude and the location of the attraction effect (Figure 1.6(b)). More importantly, the DFT predicts that the magnitude of the effect is relatively flat within the area where it takes place. This discontinuity directly stems from the relatively abrupt distance inhibition functions (i.e., step function see Tsetsos et al., 2010 for details).

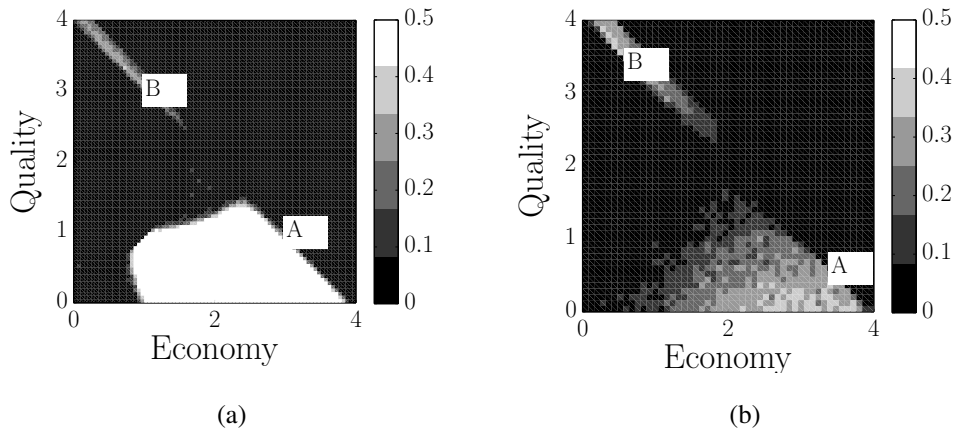


Figure 1.6: *Illustration of the attraction and similarity effects as the boost that A gets relative to B by the introduction of C (i.e., $P[A|A, B, C] - P[B|A, B, C]$) in various places of the two-dimensional lattice. a: Predictions for DFT, sigmoidal inhibition function. b: Predictions for LCA.*

The attraction effect in DFT is a type of contrast effect, in which the decoy enhances the dominating option with which it is contrasted. While this works well in the attraction situation, this mechanism has the danger of causing dominance-reversals for options that are in a strict domination order, as illustrated in Figure 1.7(a) (C dominates B and B dominates A). Such reversals may take place (depending on the magnitude of the inhibition), when the distance between the options is such that A and B inhibit each other, while C is more distant and outside the inhibition range of the two dominated options.

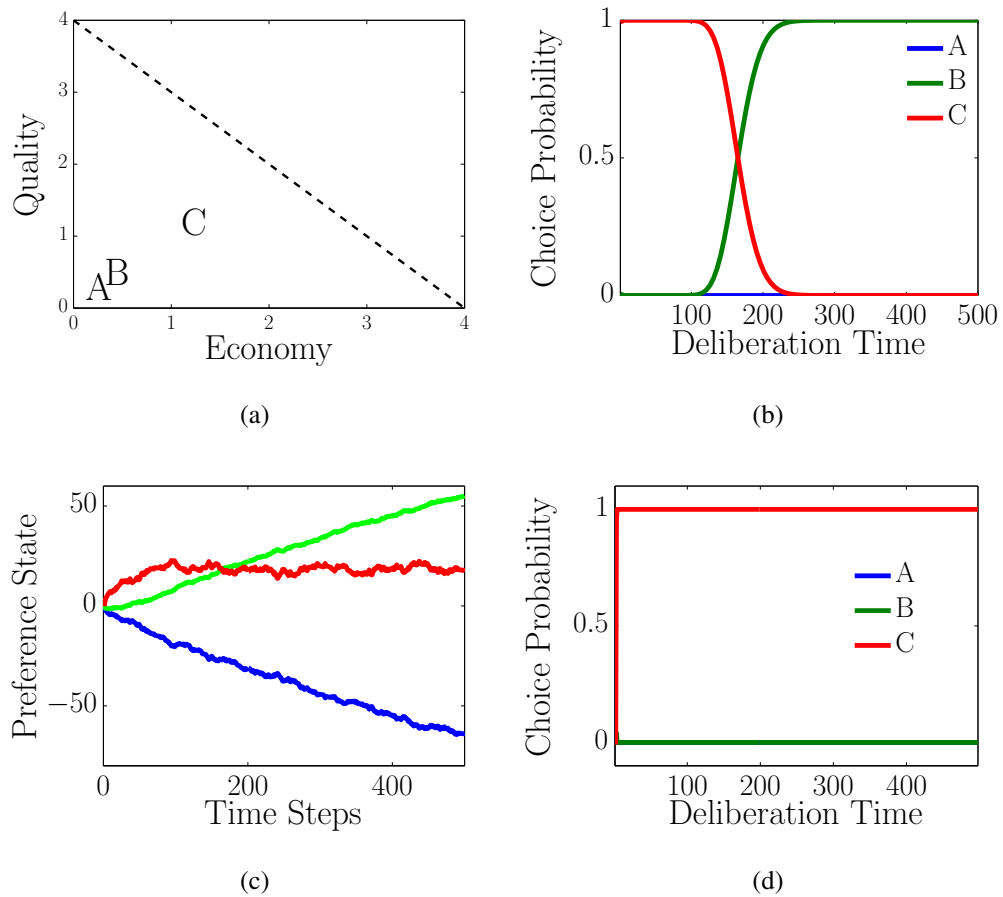


Figure 1.7: *a: A choice scenario where Option C has the highest additive utility. b: Probability of choice for the three options in decision field theory (DFT); after approximately 200 time steps, the inferior Option B outplays C due to the sharp boundaries of the inhibition. c: Single-trial trajectory for DFT. d: Predictions for the leaky competing accumulators model.*

In the simulation result in Figure 1.7(b), a localized distance function that allows DFT to reproduce all three reversal effects was used. As we can see, although the activations

are bounded (all eigenvalues of the s-matrix are smaller than 1) after approximately 200 time steps, the dominated option B emerges as the choice winner (Figure 1.7(c) shows a single trial trajectory for DFT). Intuitively this prediction results from the fact that the superior option, C, does not benefit from the boosting by negated inhibition from any option since it does not interact with either A or B. On the other hand the inferior decoy, A, which has a negative valence, confers excitation on B. On the contrary as illustrated in Figure 1.7(d), the LCA gives the correct prediction, since due to the non-linearity in the preference states, uninformative options are deactivated (stuck at zero) at early stages of the decision process. In Appendix C in Tsetsos et al. (2010) we show how one needs to constrain the DFT parameters in order to avoid dominance reversals and at the same time account for all three effects.

Another diverging point in the two models is the way they account for the compromise effect. According to the original DFT model (Roe et al., 2001), the compromise effect occurs because the preferences of the extremes are correlated in time. This account for the nature of the effect is very different from the one offered by the LCA, which follows Tversky and Simonson (1993) proposal that it is a result of the pairwise comparisons between the options, and the large penalties applied on the disadvantages. If the compromise effect arises from the temporal correlation of the extremes, it should be possible, in principle, to detect a signature of this correlation. One way to investigate this, was recently explored experimentally (Usher, Elhalal, & McClelland, 2008). In this study, participants were presented with a three-choice compromise choice-set, and in some cases, following the participants' choice of an extreme option, this option was announced to be unavailable and a speeded 2nd choice was requested for one of the remaining two options. If the two extremes are indeed correlated, one may predict that, at the moment when one of them reaches a response-criterion, the other extreme is also high in its activation and therefore, is more likely to be selected very fast. The experimental results, also reported in Tsetsos et al. (2010), showed that after the choice of an unavailable extreme, the participants had an overwhelming tendency to choose the compromise, rather than the other extreme. Furthermore the selection times were longer when participants chose the other extreme than when they chose the compromise.

These results contradict the correlational account of the compromise effect assumed by the DFT model. Within the DFT framework an alternative explanation for this effect has been provided (Busemeyer & Johnson, 2004). According to this account, *avail-*

ability can be seen as a third choice-attribute, which makes an unavailable option less desirable, but allows it to compete for selection. According to this, the unavailable option bears negative valence due to its low attribute value in the availability dimension, and boosts the compromise, due to their mutually inhibition. The *availability* mechanism can be tested by considering a choice set with three options, as in the attraction case. During the deliberation we announce that the decoy is unavailable. According to the availability assumption, this will make the valence of the decoy option even more negative and thus the boost it should give to the dominant option should be further enhanced. On the other hand, if unavailable options are simply eliminated from the choice-set, we should expect that the attraction effect will diminish towards the baseline for a binary choice. To test this in Tsetsos et al. (2010) we presented to 30 participants, three choice problems, all of which involving the same two alternatives A and B, which create a tradeoff between two choice attributes. The first problem was a binary choice between A and B. The second problem, was a ternary choice, in which a decoy dominated by A and similar to it was added, and the third problem was identical to the second, except that after 15 seconds of deliberation, the decoy was announced to the participants as unavailable. The results, did not fit with this alternative explanation of the decoy effects. The decoy induces a strong attraction effect in favour of the dominating option ($\chi^2(1, N = 30) = 6.67, p < .01$ between ternary and binary). When the decoy is announced as unavailable during the deliberation, its impact disappears and the choice between A and B option reverses very close to baseline ($\chi^2(1, N = 30) = .07, p > .78$ between ternary-unavailable and binary). These results rule out the *unavailability* hypothesis and shows that unavailable options do not continue competing for choice but are rather removed from the selection process.

1.2.5 Alternative Models

So far I focused only on DFT and LCA as they are the only two theories which have accounted for the three contextual preference reversal effects simultaneously. Alternative theories have been proposed for multi-alternative, multi-attribute choice, namely Decision by Sampling (DbS; Stewart et al., 2006) and the ECHO model (Guo & Holyoak, 2002) with the latter accounting for a subset of the reversal effects. Three particular mechanisms stand out of the two models as promising; ranking, grouping and bidirectional connections in the neural network.

In DbS no underlying psychoeconomic scales are assumed. Instead the subjective value of an attribute is its rank in the decision sample which consists of attribute values both present in the decision context and drawn from memory. Thus, the value of a given option is constructed online using basic cognitive tools such as binary comparisons and frequency accumulation. Drawing from simple psychological principles DbS accounts for a large set of decision phenomena such as loss-aversion, temporal discounting and the overestimation of small probabilities. Being explanatory robust in several domains, the DbS and its mechanisms (ranking and ordinal comparisons) appear to be promising for the case of contextual preference reversal effects. Recently, the DbS was integrated with leaky competing accumulators in a dynamical model for decisions under risk (Stewart & Simpson, 2008). This model can also be extended for multi-attribute decisions and its descriptive power in that domain can be the subject of future computational explorations.

The second alternative model, the ECHO model proposed by Guo and Holyoak (2002), has been applied for the similarity and the attraction effects. Its central assumption is that decisions follow a sequential two stage process. At the first stage the two similar options are grouped and processed together. The first-stage, preference states of the similar options are carried over as initial activations at the second stage, where all the three alternatives are compared together. Thus the similar, grouped options receive more processing overall. Note that the mechanism of grouping can be comparable to the step sigmoid inhibitory function in DFT, which involves competition only between the similar options. Another assumption in the ECHO model is that the preference states of the alternatives are passed backwards to the attribute nodes, providing positive feedback. Therefore it is predicted that during deliberation, the attribute values of the option that is dominating the preference, will be enhanced and thus appear to be more important, which has been tested experimentally (Holyoak & Simon, 1999). It is interesting to test what further predictions the LCA and DFT models would yield, by changing their connectionist networks from feedforward to bidirectional.

1.2.6 Summary

Contextual reversal effects have been challenging decision theories for years. Two computational models, DFT and LCA, provided accounts for all the effects under single frameworks. The two models share many properties and use similar connectionist

framework, but they differ in the way they account for the attraction and compromise effects. While the LCA follows the more traditional account offered by Tversky and Simonson (1993), in which the effects arise from the asymmetry of the value function and the fact that options are compared with each other, DFT does not assume asymmetric loss-aversion value functions. Instead it derives the attraction and the compromise effects from the emergent properties of the local inhibition within a linear network. The attraction effect is viewed as a contrast effect, which results from the fact that the decoy boosts the preference of the similar dominating alternative by the mechanism of activation-by-negated-inhibition. The compromise effect is the outcome of an emergent correlation between the extremes, which share their wins in the choice, favouring the compromise option.

Simulations showed that, as a result of the local inhibition boundary, the range of the attribute space in which DFT produces reversal effects has also relatively sharp boundaries, which stand in contrast with the more continuous effects obtained in the LCA model. As a result, the predictions of DFT are less robust to the introduction of new options in the choice set (see Figure 1.7), resulting also in a smaller parameter space which accounts for all the effects simultaneously. The second point that was addressed was the correlational explanation of the compromise effect, which is probably the most original mechanism in the DFT account of multi-attribute decision-making. To examine the correlational prediction, decision-makers were presented with a choice between three alternatives that form a compromise situation, and following the choice of an extreme option, that option was announced as unavailable and a second speeded choice was requested. The overwhelming fraction of speeded choices went to the compromise option, rather than to the other extreme which rules out the correlational explanation. In one version of DFT, such a result can be accounted by assuming that “unavailability” is a third attribute, which does not eliminate an option from the choice process, but only reduces its valence, making it less attractive. Under such a mechanism, the unavailable extreme would activate the compromise via activation by negated inhibition. To test this, we carried out an experiment, which compared the attraction effect in a normal situation, to that in a situation in which the decoy is announced as unavailable after 15 seconds of deliberation. We found, that conversely to the prediction that the unavailability of the decoy reduces its valence enhancing the attraction effect, this effect is reduced towards the baseline of binary choice. This suggests that unavailability should not be viewed as a third choice-dimension (making unavailable options slightly

less desirable), but rather that it maintains their desirability, while eliminating them from the decision-process.

While the contrast between two state of the art models of context effects, presented here and in Tsetsos et al. (2010) and Usher et al. (2010), was quite useful for the refinement and improvement of a specific theory (Hotelling, Busemeyer, & Li, 2010), here we face a paradox; in the field of multi-attribute choice there is an abundance of theoretical models which try to simultaneously account for three basic phenomena (i.e., attraction, compromise and similarity effects). Most of our comparison focused on ruling out some of the DFT mechanisms that were computationally erratic. Although small experiments were conducted, they were tailored on specific model predictions aiming mostly to rule out particular mechanisms. Importantly, the three preference reversal effects have never been replicated under the same experimental paradigm or within subjects. It is therefore conceivable that the attempt to capture the three phenomena under a unique parametrization of a single model of choice is elusive. And yet, the psychological basis and realism of the process-based approach is not experimentally corroborated on value-based experiments but rather heavily inspired by the sequentially sampling models of evidence accumulation in perceptual choice.

Ruling out the DFT account does not prove on its own that the LCA account is correct. And in fact the latter model, although descriptively robust and potentially explanatory, involves a variety of neural (i.e. leaky integration and global competition) and psychological (i.e. asymmetric value functions as in prospect theory) mechanisms, which are too detailed to be tested using conventional behavioural experiments. Facing a reality where the computational models are more advanced than our understanding of the behaviour we try to explain, I take in this thesis a step back and question the very principle of dynamic models of preference, i.e. that preferential choice, sharing the same underlying assumption with perceptual choice where *evidence* is integrated across time, is driven by *value integration*. Thus, the main question I will address is whether preference is shaped via integration of value and whether this process is subject to context effects. Additionally, by examining the presence of context effects in perceptual choice, I will also examine the similarities in the mechanisms that underlie value and evidence integration.

1.3 Summary and Overview of the Thesis

The working hypothesis in this thesis is that decision-making is driven by sampling of information which can either correspond to sensory input or to values. My aim is to examine the presence of order and context effects in both perceptual and preferential choice. Through this examination, and using computational modelling, I will determine the mechanistic interplay between sensory and motivational choice, expecting to ultimately deduce the principles that underlie decision making in both domains.

Chapter 2 will introduce the main experimental methodology that will be used throughout the thesis. In the perceptual literature, experimental paradigms involve the presentation of *dynamic* evidence and hence results can be readily related to process models. On the contrary, in the value-based literature experiments traditionally involve the static presentation of alternatives and as a result the deliberation stage of the decision is covert to the experimenter. In order to overcome this problem, I introduce a new experimental paradigm which interpolates between psychophysics and preferential decisions. There, participants are presented with two or three fast sequences of numerical values and need to indicate which sequence had the highest average. I label this paradigm *fast value integration*. One additional novelty in the experimental techniques I used in this thesis, is the presentation of *non-stationary information*. As I will discuss in Chapter 2, non-stationary information provides a way to externally induce dynamics similar to deliberation phenomena that are covert to the experimenter during complex decisions (e.g., vacillations or changes of focus to different aspects).

In Chapter 3, I will focus on order effects in evidence and value integration. I will start with a surprising prediction that the LCA model gives due to its non-linearity; within the same parameter set, the type of the temporal bias depends on the length of the trial. In particular, and as we will see in more detail, for short trial durations the model predicts primacy while for longer durations recency (*Computational Study 1*). This surprising prediction is supported in a perceptual experiment using the moving dots paradigm where participants show a transition from primacy to recency as the trial duration increases (*Experimental Study 1*). This, however, is not the case when a similar experiment is conducted using the fast value integration paradigm. There, participants still show an increase in recency with trial duration but they do not show any sign of primacy (*Experimental Study 2* and *Computational Study 2*). Finally I perform a formal optimality analysis for decisions under uncertainty, showing that

the optimal model in a dynamically changing environment is *leaky* and not *perfect* integration (*Computational Study 3*). Therefore the fact that human choice is subject to order effects and thus to imperfect integration mechanisms, does not necessarily violate the principle of optimality since the latter must be defined afresh for each different environment. On the contrary, the fact that humans evolved in uncertain environments squares well with the claim that information integration is subject to leakage .

In Chapter 4, I move on to examining the presence of context effects in decision behaviour. My key approach departs from traditional ones and involves the presentation of non-stationary information. By manipulating the dynamics of the information I induce temporal correlations between the alternatives, reducing that way multi-attribute choice problems into a single attribute, and looking for decoy effects; the attraction, similarity and compromise effects. I start with a perceptual experiment where attraction and compromise effects are not obtained, but participants show a very strong similarity effect (*Experimental Study 3*). An analysis using the prominent sequential sampling models of perceptual choice follows, to reveal that the key mechanism to account for the observed effects is the zero non-linearity of the LCA (*Computational Study 4*). A similar experiment is performed using the fast-value integration paradigm and there participants show systematically the patterns predicted by the attraction and similarity effects (*Experimental Study 4*).

The latter finding leads to Chapter 5 where a computational model, that accounts for the fast-value integration context effects, is developed. The model holds that in value integration, the absolute value of an alternative and its relative rank in the current context are combined (*Computational Study 5*). In particular, the absolute value of each alternative is weighted by its momentary rank and this product is subsequently integrated in a leaky accumulator. I label this model *rank-dependent leaky integration* and demonstrate how it accounts for decoy effects, in a much simpler way compared to DFT and LCA. Additionally, I will bring to light a surprising prediction of this model for the fast-value integration paradigm; risk-seeking in the domain of gains (*Computational Study 6*).

In Chapter 6, the predictions of the rank-dependent model regarding the way that risk is dealt within the fast-value integration, is tested experimentally. In a series of experiments (*Experimental Studies 5-10*), I confirm that people when asked to make choices on the basis of fast numerical sequences, are risk-seeking (Experimental Study 5). This pattern persists even when the numerical values correspond to losses (*Experimen-*

tal Study 6). However it switches to risk-aversion, when the sequences are mixed i.e. involving both gains and losses (*Experimental Study 7*). These findings are quite surprising and contradict decisions by description and prospect theory. In order to better understand the basis of these findings, I conducted another series of experiments which revealed that choice behaviour switches from risk-seeking to risk-averse, depending on the framing of the problem at hand (*Experimental Study 8*). When people select a sequence they are risk-seeking but when they have to reject one sequence they are risk-averse. These results integrate well with previous research on high-order decision making and the reason-based framework and help the further refinement of the rank-dependent model. Finally, the relationship between the fast value integration paradigm and decisions-by-experience is further explored (*Experimental Studies 9-10*).

Chapter 2

General Methods and Techniques

2.1 Overview

In this chapter, I give a general description of the experimental and computational methods and techniques used throughout this thesis. I start by describing the stimuli and protocols employed in perceptual (i.e., moving dots and brightness discrimination tasks) and value-based (i.e., fast-value integration task) experiments. Next, I outline a novel aspect of most of the experimental designs presented here, the use of non-stationary information. As I discuss below, using non-stationary information in combination with an interrogation response protocol is ideal for the study of temporal biases. Furthermore, it allows the experimenter to induce complex temporal correlations between the options, in a way that emulates the deliberative process over multi-attribute alternatives with trade-offs. Finally, I outline the basic techniques of simulating and evaluating computational models against the empirical data.

2.2 Experimental Methods

2.2.1 Stimuli and Tasks

Perceptual Choice

For the study of perceptual decisions, I used two different psychophysical tasks. First, when examining order effects (Chapter 3) in binary choice, I adopted a paradigm

extensively used in the neuroscience of decision making, the *moving dots* paradigm (Britten, Shadlen, Newsome, & Movshon, 1992; Shadlen & Newsome, 2001; Roitman & Shadlen, 2002; Gold & Shadlen, 2003; Kiani et al., 2008). In each trial of this task, participants observe a display of noisy, dynamic moving dots. A fraction of dots (depending on the difficulty of the trial) moves coherently towards one direction and the observer has, when prompted by a response cue, to discriminate whether the direction of the coherent motion is left or right (i.e. 2 alternative forced choice). The time course of a typical trial in this task is depicted in Figure 2.1(a) (for a detailed description see also Roitman & Shadlen, 2002).

Although the moving dots paradigm has been applied to neuroscience studies of multi-choice (Niwa & Ditterich, 2008; Churchland, Kiani, & Shadlen, 2008), when examining context effects in multi-alternative decisions, I used a brightness discrimination task (Caspi, Beutter, & Eckstein, 2004; Ludwig, Gilchrist, McSorley, & Baddeley, 2005). There, participants saw four patches whose brightness fluctuated across the course of the trial (the brightness of each spot was normally distributed) and when prompted they had to indicate which of the four spots was the brightest overall (Figure 2.1(b)). The advantage of the brightness discrimination task is that it allows the *independent* manipulation of the evidence for each alternative. On the contrary, in the motion detection task the evidence for the different hypotheses is interdependent. Assume a 4-choice alternative task, with four possible motion directions: up, down, left and right. The up-down and left-right directions are not orthogonal e.g., support for upward motion subtracts evidence for the downward direction (and same for the left-right directions). In previous studies, this lack of independence among the different hypotheses was explicitly taken into account in the developed computational models (e.g., Churchland et al., 2008). However, the purpose of the contextual effects study I conducted (Chapter 4) was to conclude whether there is interaction among alternative options at the decisional (i.e. integration) level. Consequently, due to the pre-decisional, confounding competition (i.e. at the input level) in the motion detection task, the brightness paradigm was preferred.

Value-based Choice

Value integration is an essential process in decision-making between alternatives that are characterized by multiple values (Hertwig, Barron, Weber, & Erev, 2004; Ludvig &

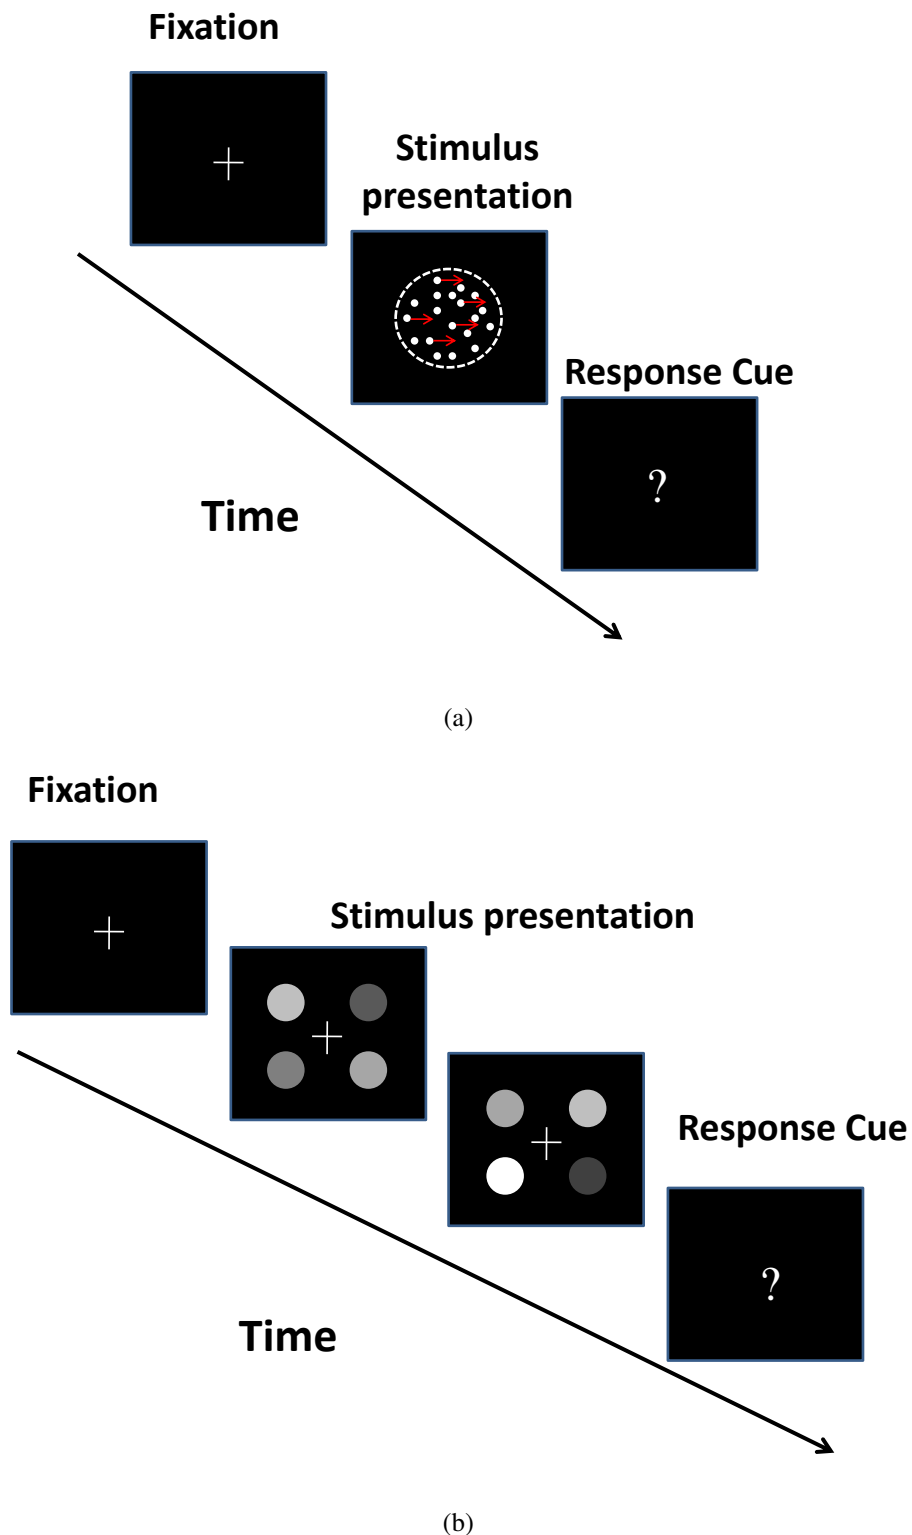


Figure 2.1: Stimuli used in the perceptual choice tasks. *a*: The time course of a trial in the moving dots paradigm (Chapter 3); red arrows indicate the dots with coherent rightward motion, *b*: The time course of a trial in the brightness discrimination task (Chapter 4); the brightness of each spot fluctuates across time and the observer needs to indicate which spot is the brightest overall.

Spetch, 2011) or attributes (Tversky, 1972; Huber et al., 1982; Tversky & Simonson, 1993). For example, to decide what car to buy or which flat to rent, the cognitive system needs to integrate a multitude of goodness values across different dimensions. Our understanding of this mechanism, although of central importance in recent dynamic models of preference (Roe et al., 2001; Usher & McClelland, 2004; Tsetsos et al., 2010), is impeded due to covert processes involved in complex decisions. For instance when trying to decide among alternatives that differ in several dimensions, the decision maker might internally switch focus to different choice aspects or allocate a different amount of processing to each alternative.

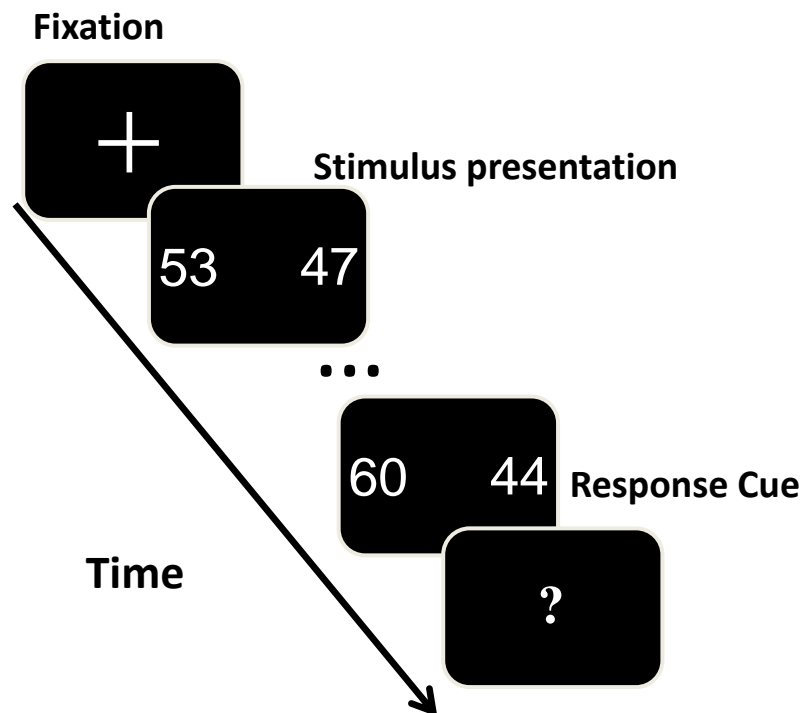


Figure 2.2: *The timeline of stimulus in the binary choice version of the fast value integration task. Participants observe a rapid stream of pairs of numerical values and at the end of the presentation have to decide which sequence, left or right, had the highest average value or which sequence they would like to draw an extra sample from.*

In this thesis, I propose a way to directly probe the micro-mechanisms of value integration, by introducing a decision task of reduced complexity, where the values that the decision-maker receives are defined on a common currency and their time-course is externally controlled by the experimenter. Participants are presented with rapid, varying sequences of pairs or triples of numerical values and are asked to select the one asso-

ciated with the best overall value or the one they would like to draw an extra sample from (Figure 2.2). Importantly, the structure that underlies this paradigm is identical to dynamic perceptual tasks (e.g. brightness discrimination) and the only difference is the presentation of symbolic, numerical values instead of sensory evidence. Based on the remarkable capacity of the cognitive system to make numerosity judgments (Barth et al., 2006) and to integrate affect associated with numerical rewards (Bechara, Damasio, Tranel, & Damasio, 2005), I expected that humans would be able to integrate values across time and select the alternative with the highest payoff, even at a fast presentation rate. This dynamic decision paradigm, which I label *fast value integration*, lies at the intersection of low and higher order decisions and can be used as a proxy to access the underlying process of more complex decisions (e.g. multi-attribute choice; Chapter 4, reason-based choice; Chapter 6).

2.2.2 Non-stationary Information

In most of the psychophysical tasks with dynamic evidence reported in the perceptual literature, the evidence is stationary with its distributional characteristics remaining fixed throughout the experimental trial. This is quite appropriate in the free-response protocol, where both the choice and the response time are of interest and where core aspects of the decision mechanism are examined. Recently though, in studies that followed the interrogation protocol (i.e. the observer responds when prompted and not freely) temporal biases were probed using non-stationary evidence (Usher & McClelland, 2001; Huk & Shadlen, 2005). The merit of using non-stationary evidence with regards to order effects, is that the direction of the evidence can switch midway the trial. For example, in the motion detection task, the dominant motion might be left in the first half of the trial and right during the second half. The observer's responses in trials like this, reveal which half, the first or the second, is more strongly weighted or in other words the type of the temporal bias.

Information with non-stationary characteristics bears ecological validity, since it reflects the dynamics of volatile environments where the underlying structure of the world can unexpectedly change (Summerfield, Behrens, & Koechlin, 2011). Probably mirroring the volatility of the real world, people's views of the environment are often quite flexible. This flexibility manifests itself also in decision-making; decisions are rarely ballistic and humans are known to vacillate, change their minds or endlessly

procrastinate their decisions. What triggers these dynamic effects in the preference state of the decision-maker is of particular interest especially when the environment is stable and the choice is made under certainty. In that case, changes of mind should be attributed to overt (e.g. reflected in the visual fixations) or covert (e.g. internal states) switches of focus to different choice aspects.

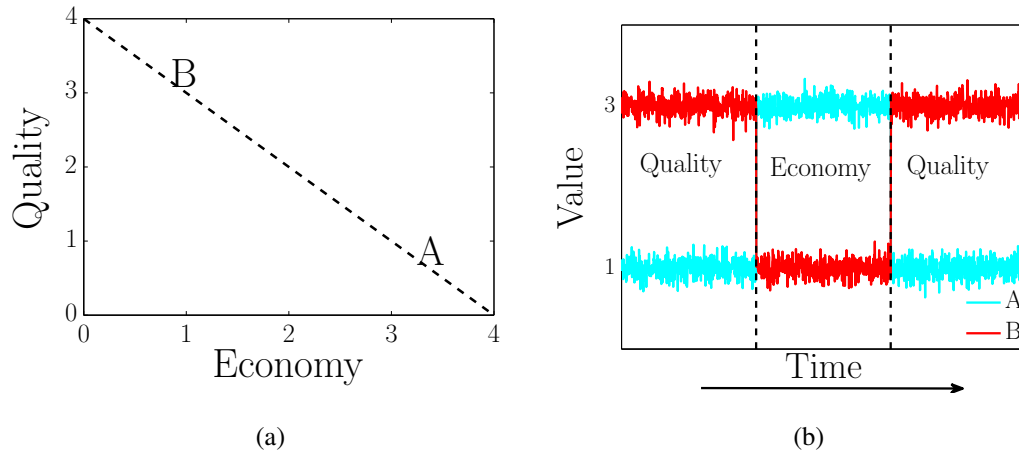


Figure 2.3: *Unfolding a multi-attribute choice problem (a) in time (b). Assuming that each of the two dimensions can be translated into a common currency (i.e. value) and that the decision-maker switches focus among the choice aspects, the time course and amount of the processed information for alternatives A and B can be reconstructed (b).*

For example, when having to decide which car to buy among many alternatives that differ in two dimensions, i.e. economy and quality (Figure 2.3(a)), the decision maker's attention might fluctuate between the two dimensions. In moments, she might consider the quality of the cars, in which case B is favoured over A, while in others the economy might stand out as important, in which case A should look better. This switching of attention squares well with the trade-off inconsistency observed in humans (Stewart, Chater, Stott, & Reimers, 2003); when deciding among qualitatively different options, such as the luxurious and expensive car B and the more basic and economical car A, it is impossible to compare them holistically, on a single dimension of value. In other words the integrated value of complex, multi-attribute alternatives is not readily available to the decision-maker but instead needs to be constructed afresh in any given situation. A plausible mechanism of online value construction for multi-dimensional options, is the fluctuation of attention, the within-attribute evaluation of the alternatives and the subsequent integration of the momentary values across time. This mechanism

is central to dynamic models of preference (Roe et al., 2001; Usher & McClelland, 2004) which, following Tversky's original idea (Tversky, 1972), assume the sequential sampling of different choice attributes until a decision is reached.

Given that the attentional switching to different choice aspects can be external but also internal, it is hard to experimentally observe it and hence test its validity. The route that I will take in this thesis is, instead of measuring the process of attentional switching, to induce temporal fluctuations in dynamic stimulus which mimic this switching. To understand this technique better, I demonstrate in Figure 2.3 how a decision between car A and B can be collapsed and unfolded across time. The decision-maker first considers the quality of the cars and car A receives strong support (i.e. labelled as *value* in the y-axis at the bottom panel of Figure 2.3) while car B a much weaker input. Subsequently, the focus switches from the quality of the cars to their economy and the situation reverses; now A appears quite disadvantageous and B receives stronger input. After considering economy the attention switches again back to quality and eventually the decision-maker finalizes her decision. All these changes of focus can be translated into two input signals, one for A and one for B, which are anti-correlated in time as the two alternatives are in the 2-D choice space (Figure 2.3(a)). If despite this reduction, well-known phenomena in multi-attribute choice, such as decoy effects, persist (Chapter 4) with one-dimensional and temporally manipulated input, then two aims will be achieved. First the attentional switching hypothesis will be corroborated and second the micro-mechanism leading to contextual reversal will be further understood.

To summarize, the merits of presenting non-stationary information will be fully exploited in the rest of the thesis. First when examining temporal biases, the trials will be often divided in two halves with the support to the alternatives changing direction from the one half to the other. Second, in multi-alternative choice problems information will be finely manipulated so as temporal correlations are induced among alternatives. That way the process of attentional switching between choice attributes will be mimicked in the decision input, causing specific internal mental states that are presumably underlying contextual reversal effects. This technique is expected to shed light on the underlying process of multi-attribute choice and also validate the assumption that decisions are driven by the sequential scanning of different choice aspects.

2.3 Computational Techniques

2.3.1 Monte Carlo Model Simulations

All models were simulated using Matlab 7.11.0 (Mathworks Inc., Natick, MA, USA) as stochastic Monte Carlo algorithms that approximated the probability of a certain choice outcome under a specific parameter set, by running multiple trial runs. The core element in all models was a stochastic difference equation (updated over several simulation steps) that described the dynamics of the preference states (the decision variable) for each alternative. The mapping of the simulation steps to the actual experimental time was arbitrary (given for each model in the corresponding sections) in perceptual models where the flow of the perceptual experimental input was very fast. On the contrary, in the fast value integration models the number of the model time-steps coincided with the actual (discrete) experimental frames. In all experiments presented here, an interrogation response protocol was used with a response requested by the experimenter at specific times. Accordingly, a choice was issued in a given model, in favour of the alternative whose decision variable was the highest at the moment of the interrogation. The model choices were collected over several simulation runs and at the end averaged together, to derive the model's choice-probability prediction for each alternative.

2.3.2 Optimization Procedure

For data fitting I used the Matlab toolbox developed by Bogacz and Cohen (2004), which estimates the parameters of a model based on least squares. This toolbox is cast as a Matlab function ("fitparam") which receives as a primary argument the name of the script where the model is implemented. The model script takes as input arguments the model parameters and generates as output data points (predictions) which are evaluated against the empirical data points (see next subsection). The advantage of the method developed by Bogacz and Cohen (2004) is that it extends the multi-dimensional simplex algorithm in order to better handle noisy functions. The fact that the models used throughout this thesis are implemented using Monte-Carlo simulations, renders this technique more robust compared to others that are designed to deal with deterministic functions (e.g. the simplex method in Matlab). The cost function that the optimization routine minimizes is defined as: $cost = \sum_{i=1}^N \left(\frac{e_i - m_i}{n_i} \right)^2$, where m_i

are the statistics of the model, e_i the statistics obtained from the experiment and N the number of the statistics that are fitted. A normalization factor, n_i , is introduced for each statistic i . This is to ensure that all data points contribute equally to the cost function despite differences in the scale across the statistics. As described in detail in Bogacz and Cohen (2004), the selection of the value of the normalization factor varies across the different stages of the optimization process to maximize efficiency. During the initial stages of the optimization process (i.e., searching for starting points and first optimization), n_i takes the value of the average value of the empirical statistics (e_i). At the final stage of the process (i.e., tuning of parameters) n_i becomes the standard deviation of statistic i , obtained after running the model with the same parameters 10 times.

2.3.3 Model Evaluation

In order to evaluate the quantitative fits of a model, I used the Bayesian information criterion (BIC, Schwarz, 1978), which takes into account both the goodness of fit and the complexity of the model. The BIC penalizes the extra free parameters much more strongly than other similar measures (e.g., Akaike information criterion; Akaike, 1974). The BIC is computed as: $-2[\sum_i N e_i \ln(m_i)] + M \ln(N)$, where M is the number of the free parameters of the model.

Chapter 3

Time-dependent Weighting of Information

3.1 Overview

When people make decisions on the basis of dynamic evidence, they often weigh information differentially depending on its temporal order. In a recent neurophysiological study, Kiani et al. (2008) concluded that perceptual decisions are characterized by primacy and that this is a universal property of evidence integration. In follow-up experiments (Tsetsos, Gao, et al., 2011) we showed that the strong primacy pattern, found in Kiani et al. (2008), can be attributed to special procedural aspects of the task, which encouraged observers to have increased attention at the beginning of the trial. Despite our showing the importance of task contingencies in the study of evidence weighting, it yet remains unclear what drives the direction of order effects, with both primacy and recency being reported in different participants of the same task (Usher & McClelland, 2001). In the first section of this chapter (*Computational Study 1*), I will bring to light a novel prediction of the LCA, due to the model's non-linearity: the direction of the order effect interacts with the trial duration; for moderate inhibition dominance, the model predicts primacy weighting for trials of short duration and recency for longer trials. Therefore, according to this prediction, different types of temporal biases might coexist within a single decision strategy (i.e. single parameter set) and be triggered bottom-up, beyond the observer's control (i.e. by the evidence duration). In the second section (*Experimental Study 1*), this prediction is tested in the motion detection task

and its signature is found in two participants. In the third section (*Experimental Study 2*), I examine the influence of the sequence length on the direction and magnitude of the order effect, in the fast value integration paradigm. There, no sign of primacy is found, but consistent with the prediction obtained in Computational Study 1, recency increases with the length of the trial. This result is realized using a leaky diffusion model (*Computational Study 2*). Finally in the last section of this chapter, I discuss the principle of optimality in decisions under uncertainty, showing that in volatile environments the optimal strategy is the imperfect (i.e. leaky) integration of information (*Computational Study 3*).

3.2 Order Effects and Trial Duration (Computational Study 1)

As demonstrated in Tsetsos, Gao, et al. (2011), in perceptual choice tasks, the direction of the temporal bias in evidence integration might be explained by the task characteristics and demands. For example in the study of Kiani et al. (2008), where a strong primacy effect was found, the distribution of the duration of the trials was skewed (i.e. exponential), containing mostly short trials which encouraged the observers to allocate more attention to the beginning of the trial. The dependence of order phenomena on task characteristics is reminiscent of the influential work of Hogarth and Einhorn (1992) who examined how the response mode (i.e. step by step or at the end of the sequence responses), the task complexity and the length of the series of evidence items affect information processing in sequential belief updating tasks. Their proposed anchoring and adjustment model, captured a lot of variance in existing studies and derived novel predictions regarding the interaction of task characteristics with temporal biases.

Contrary to belief updating and higher order reasoning, perceptual choice tasks are much simpler. As a result, from the factors examined in Hogarth and Einhorn (1992) the response mode and task complexity are not relevant since they do not vary across experiments. On the other hand, the length of the presented evidence is a factor that is often varied within the same perceptual task. For instance, the effect of the trial duration on response accuracy has been extensively studied, revealing that accuracy levels off for longer durations (Swensson, 1972; Usher & McClelland, 2001; Kiani et

al., 2008). However, the effect of evidence length on the direction of the order effect has never been addressed. In the current Computational Study, I examine the dynamics of information processing as a function of the duration of the trial, using the non-linear LCA model (Usher & McClelland, 2001). The reason I rely on the LCA model is that, for different parameters, it can operate under three different modes of information weighting: i) uniform weighting, ii) primacy, iii) recency. This versatility of the LCA stands in contrast with other models of perceptual choice such as the diffusion model which can predict either uniform weighting or clear-cut primacy (Ratcliff, 2006), by assuming that when a decision boundary is breached new information that arrives is ignored. In the following, I first demonstrate that, for different parameters, LCA can weigh information by primacy or recency and subsequently I show that, for fixed parameters, information weighting switches from primacy to recency, as the length of the trial increases.

3.2.1 Method

The LCA model was implemented with three free parameters. The first two parameters, β and κ stood for the values of inhibition and leak respectively:

$$\begin{aligned} dy_1 &= (I_1 - \kappa y_1 - \beta f(y_2))dt + \sigma dW_1, \\ dy_2 &= (I_2 - \kappa y_2 - \beta f(y_1))dt + \sigma dW_2, \end{aligned} \quad (3.1)$$

A decision among the two perceptual hypotheses was only made at the end of the stimulus presentation on the basis of the unit with the highest integrated evidence at that moment. The function f is a non-linear function, which truncates negative activations to zero (equivalent to a reflecting boundary; see Usher & McClelland, 2001, Equation 4 and Appendix-A). The last free parameter was the standard deviation of the noise σ . The variables I_1 and I_2 denoted the stimulus strength for each perceptual hypothesis.

3.2.2 Results and Discussion

The LCA model can operate in three modes of evidence weighting: i) $\beta = \kappa$ results in uniform weighting, ii) $\beta > \kappa$ gives primacy and ii) $\beta < \kappa$ generates recency weighting. To illustrate the recency/primacy bias of the LCA, for different parameters, the model was simulated with zero input (i.e. $I_1 = I_2 = 0$) and noise (i.e. $\sigma = 0.1$) for 30000 trials of duration 200 time-steps each. In Figure 3.1, the average input for the winning

accumulator versus the average input for the losing accumulator is presented for two instantiations of the model; the inhibition dominant and leak dominant LCA. For the inhibition dominant LCA ($\beta > \kappa$), what drives the choice is the noise fluctuations in favour of the winning accumulator at the beginning of the trial (Figure 3.1(a)). On the contrary, for the leak dominant LCA ($\beta < \kappa$) when an alternative is chosen (i.e., the winning accumulator) this is due to the noise favouring that alternative towards the end of the trial (Figure 3.1(b)).

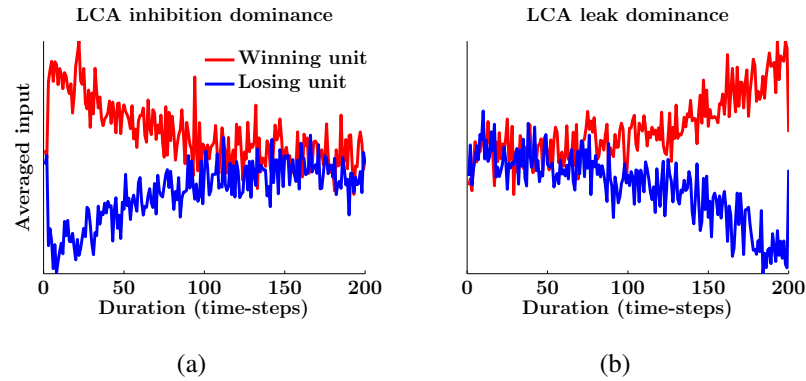


Figure 3.1: *Reverse correlations for the LCA model, showing the average input of the winning and losing accumulators in 30000 trials. a: Inhibition dominance resulting in primacy, b: Leak dominance resulting in recency.*

Under some parameters, the nonlinear inhibition-dominant LCA can predict both primacy and recency, depending on stimulus duration. To illustrate this, I simulate LCA in trials where the evidence favours the one alternative in the first half and the other in the second half, so that the overall evidence is equal for both alternatives. The parameters for the model are $\beta = 0.748$, $\kappa = 0.172$ and $\sigma = 0.1$. The input that the accumulators received were $I_1 = 0.11$ and $I_2 = 0.08$ during the first half (i.e. from $t = 1 \dots \frac{T}{2}$) and $I_1 = 0.08$ and $I_2 = 0.11$ during the second half (i.e. from $t = \frac{T}{2} + 1 \dots T$). In Figure (3.2) the preference for the 1st accumulator that is associated with I_1 is shown for five different trial durations ($N = 30000$ trials for each data point). One can see that, while for short stimulus durations the choice favours the alternative associated with primacy (early evidence), the situation reverses at longer durations.

In order to explain this pattern I show typical single trial activations of the two accumulators, for one short and one long trial (Figure 3.3). For short duration (Figure 3.3(a)), the accumulator associated with early evidence (blue-line) wins. This happens because at the time of the swap the inhibition is quite strong and although the second

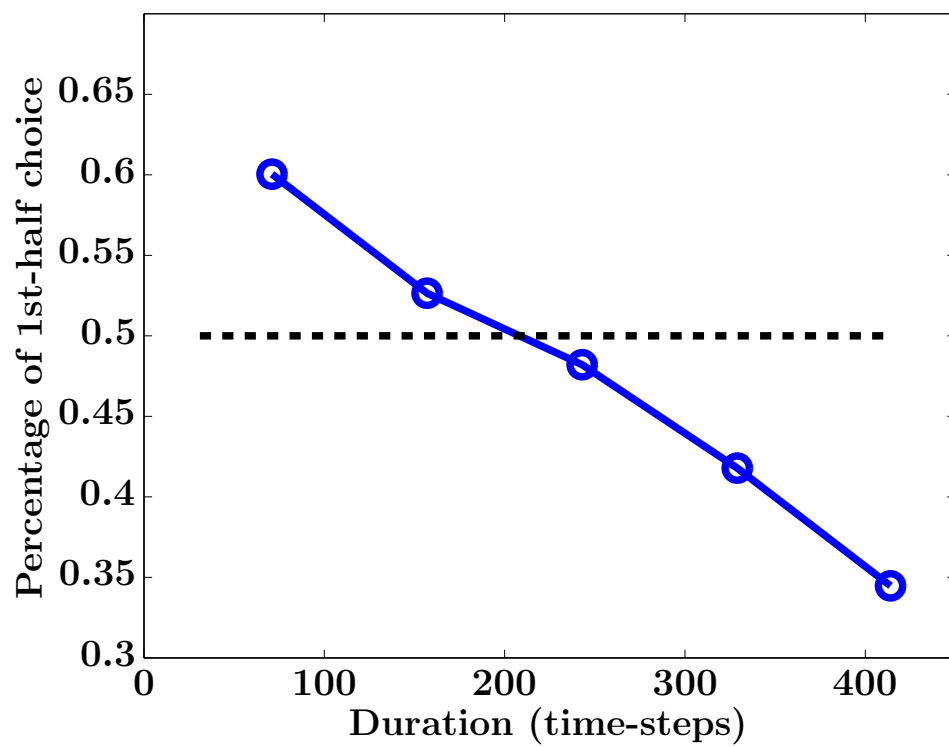
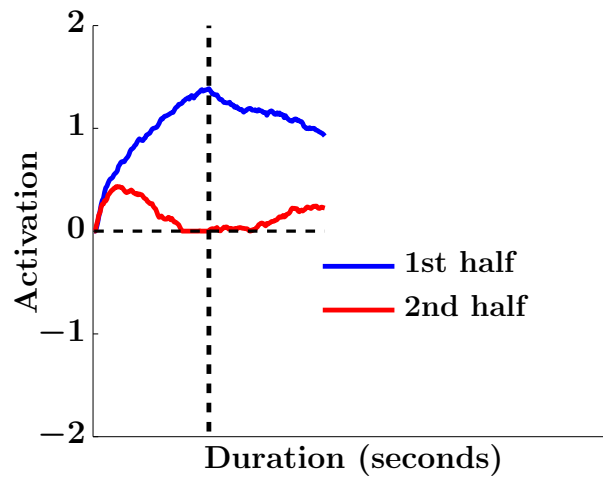
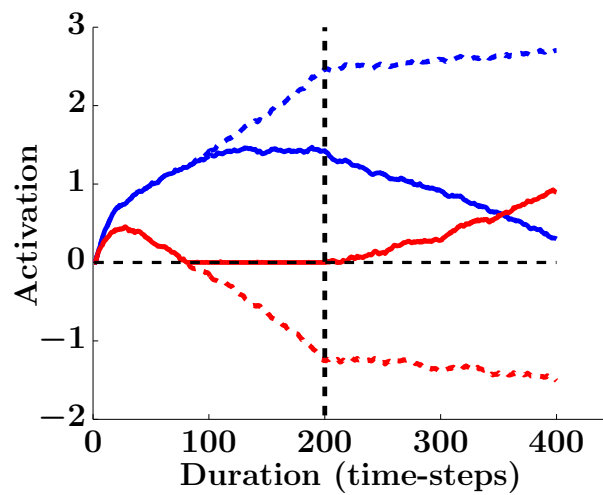


Figure 3.2: *Fraction of trials in which the accumulator associated with early evidence wins.*



(a)



(b)

Figure 3.3: Single trial activation of the accumulators in the simulation in Figure 3.2 for short (a) and for long (b) durations. The accumulator associated with 1st half evidence is shown with blue line and the one associated with 2nd half evidence with red. For the long duration the linear model activations (dashed lines) are also plotted.

accumulator (red) starts to rise, it does not turn over the situation. For longer duration (Figure 3.3(b), solid lines), the second accumulator (red) hits the zero-activation boundary relatively early ($t = 100$ time steps). From that moment and on, the activation of the first accumulator stops increasing (while in a linear system, depicted with dotted lines, it keeps increasing) and the LCA becomes leak dominated, maintaining the same level of activation until $t = T/2$ when the evidence swaps. After the swap, the accumulator associated with the strong evidence in the second half (red) moves up from the zero boundary and the system becomes inhibition dominant again. This unit now has enough time to overtake the first accumulator (note that this would not happen in the linear system; dotted lines).

The key mechanism for the interaction of the temporal bias with the length of the trial, is the non-linearity of the LCA model. In particular, the non-linearity imposes a restriction on the activation of the units and keeps their difference bounded. As a result in any long trial, there will be a point whereby the dynamics will stop evolving and the winning unit will cease benefiting from the extra evidence it receives. Keeping the difference of the two accumulators bounded means that when the evidence swaps direction the, up to the point, losing unit will be close enough (being stuck at zero) and will have enough time (if the trial is long) to reverse the situation and dominate its competing unit. The intuition behind this prediction is that the longer the trial the more likely is the early evidence to be forgotten (i.e. to decay) and therefore the stronger the recency. Interestingly, no other model, including Hogarth and Einhorn (1992) model of belief updating, has ever generated a similar transition from primacy to recency as the length of the sequence increases. In the next section this unique prediction of LCA is probed in a motion detection perceptual task.

3.3 Order Effects in Evidence Integration (Experimental Study 1)

Data from two experiments, presented in Tsetsos, Gao, et al. (2011), are re-examined in an attempt to find support for the intriguing prediction of the LCA, regarding the interaction of the temporal bias with the trial duration. Note that (as also described in the Introduction of this thesis) the purpose of the experiments in Tsetsos, Gao, et al. (2011) was to investigate the dependence of the primacy bias reported in Kiani et al. (2008)

on the task contingencies such as the distribution of the trial lengths and the duration of the response deadline. Therefore the design of the experiments, although the trial duration was varied systematically and the temporal biases were explicitly examined, has some extra characteristics that might be confounding. For example in Experiment 1a (presented below) the distribution of the trial durations was exponential, containing mostly short trials which enhanced primacy patterns. Despite the limitations due to the task design, the data of these experiments will be considered and the signature of the non-linear LCA will be looked for.

3.3.1 Experiment 1a

3.3.1.1 Method

Stimulus

The moving dot stimuli were created following the method described in Kiani et al. (2008). The motion stimulus consisted of circular dots of radius 2 pixels, moving horizontally at a speed of 5 degrees per second. Total dot density was 16.7 dots per degree squared per second. The stimulus was viewed through a circular aperture of radius 5 degrees. Within these parameters, the total number of dots was divided into three sets. One set of dots was displayed per frame. Each set of dots appeared on the monitor once every three frames (frame-triples). In addition, the coherence of the motion stimulus varied between trials. On every displayed frame, each dot had a $(1 - coherence)$ probability of being redrawn at random coordinates within the circular aperture. For instance, at 100% coherence, every dot would be redrawn to move horizontally in the direction specified by the trial, left or right. At 0% coherence, every dot would be redrawn randomly on every frame.

Procedure

Each trial began with a fixation cross at the center of the screen. The moving dots stimulus was displayed 1000 ms later. The coherence levels employed were 6.4%, 12.8%, 25.6%, and 51.2%. Stimulus duration followed an exponential distribution taking values from 100 to 1750 with an increment of 50 ms. The stimulus termination occurred simultaneously with an auditory go signal. In order to receive rewards, participants had to respond by pressing keys on the keyboard of a standard computer, within a 300 ms response window following the go-cue.

Conditions

For each coherency and duration four conditions were created: i) the constant condition corresponding to a fixed coherency during the whole trial, ii) the early-condition, corresponding to a fixed coherency during the first half of the stimulus, which is set to zero (random motion) during the second half, iii) the late condition (the first half with zero-coherence and the second half with fixed coherency) and iv) the switch condition in which the coherency values stays constant in magnitude but the direction of motion switches in the middle of the trial.

Observers

Three participants (two male, one female) with normal or corrected-to-normal vision were tested in several one-hour sessions over several weeks. In each session, participants completed 9 blocks of 100 trials. A self-paced break occurred between blocks to allow rest. Initial sessions where participants familiarized themselves with the task (and hence in which their performance had not fully stabilized) were excluded from the analysis, leaving 14, 17 and 12 sessions for CS, MT and SC, respectively.

3.3.1.2 Results

There are two critical tests regarding the temporal bias. The first one is the accuracy order between the early and the late conditions and the second one is the preference in the swap-condition where the one alternative is supported in the first half of the trial and the other alternative in the second half. The accuracy (averaged across coherency levels) is displayed as a function of condition (constant/early/late) and duration, for the 3 observers in the left panels of Figure 3.4. In all observers, performance improves with stimulus duration (with saturation at longer trials) and the accuracy is higher for the early (green) compared to the late (red) condition. This confirms the Kiani et al. (2008) result, of a primacy bias. The amount of this bias, however, varies among the 3 observers. It is very high in one observer (MT), who totally discarded late evidence, but is smaller in the other two. In one of the observers (SC) one can see a pattern that is indicative of the recency-primacy interaction with the trial duration; the advantage of the early condition is larger for short duration and disappears at longer durations. This is also apparent in the swap condition of the SC observer where primacy in the short trials changes to recency in the longer durations.

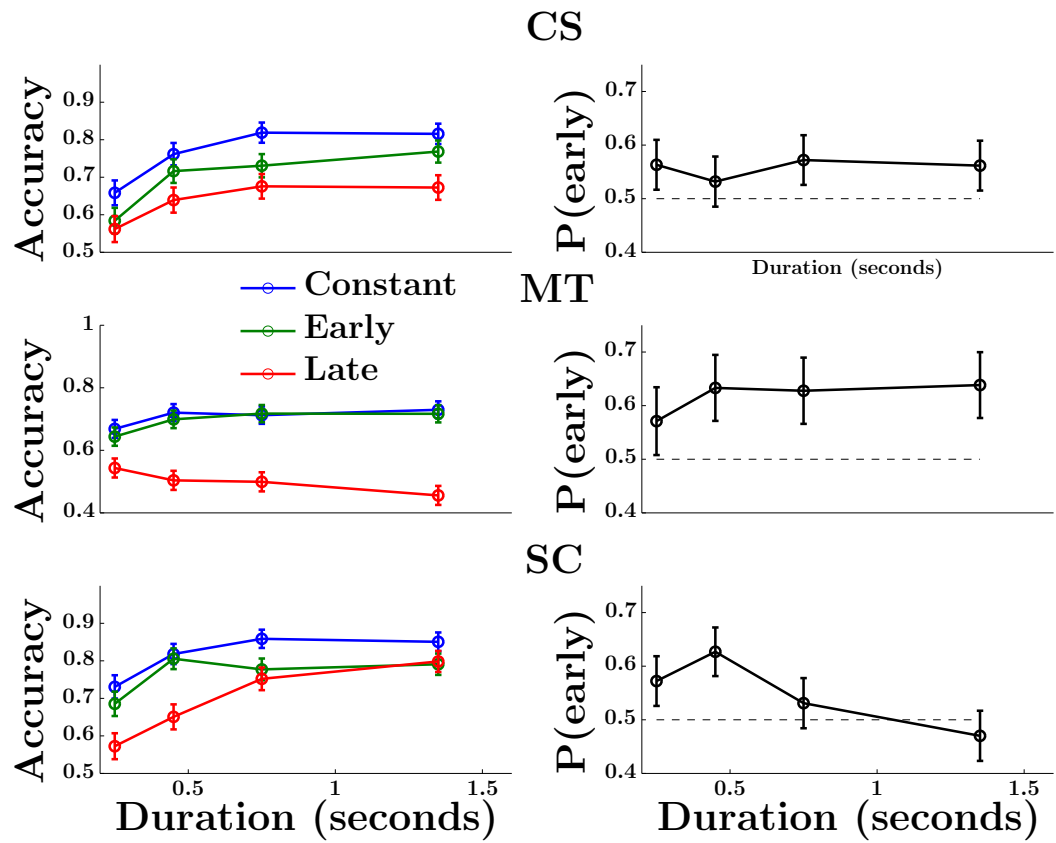


Figure 3.4: *Left: accuracy as a function of stimulus duration and condition (blue-constant; green-early; red-late;). Right: proportion of choice for the direction supported early in the trial, as a function of stimulus duration and condition. Error bars correspond to 95% CI.*

3.3.2 Experiment 1b

3.3.2.1 Method

Stimulus

Same as in Experiment 1a.

Procedure

Each trial began with a fixation cross at the center of the screen. The moving dots stimulus was displayed 1000 ms later. The coherence levels employed were 6.4%, 12.8%, 25.6%, and 51.2%. Stimulus duration followed a uniform distribution taking values from 100 to 1750 with an increment of 50 ms. The stimulus termination occurred simultaneously with an auditory go signal. In order to receive rewards, participants had to respond by pressing keys on the keyboard of a standard computer, within a 1000 ms response window following the go-cue.

Conditions

Same as in Experiment 1a, except the swap conditions which was omitted in this version.

Observers

Four participants (one male, three female) with normal or corrected-to-normal vision were tested in 11 to 25 1-hour sessions over several weeks. The sessions in which participants familiarized themselves with the experiment were not included in the analysis, resulting in 16, 19, 25, 11 sessions respectively, for DG, LK, MM, WW. In each session, participants completed 9 blocks of 100 trials. A self-timed break occurred between blocks.

3.3.2.2 Results

The choice accuracy, for the four observers of Experiment 1b, as a function of stimulus duration and condition (constant/early/late), averaged over coherency is shown in Figure 3.5. As a first observation, the procedural changes in this experiment resulted in a reduced difference in the early-late conditions (i.e. reduced primacy bias). The more balanced distribution of the trial durations (i.e. uniform distribution compared to the exponential in Experiment 1a), afforded a statistical analysis of the weighting profile

for each participant. This was made possible due to the multiple sessions that allowed enough trials for each condition and each trial duration. To carry out this analysis, the data of each observer were divided into mini-sessions or “quasi-subjects” that corresponded to all the *session* \times *coherency* combinations. Each such “quasi-subject” contributed equal number of trials to the relevant dependent variables (i.e. performance for four durations in early and performance for four durations in late), factoring out common variability related to fatigue/practice or performance levels. Thus, one repeated measure (4x2) ANOVA, with 4 temporal durations and 2 conditions (early/ late) was performed.

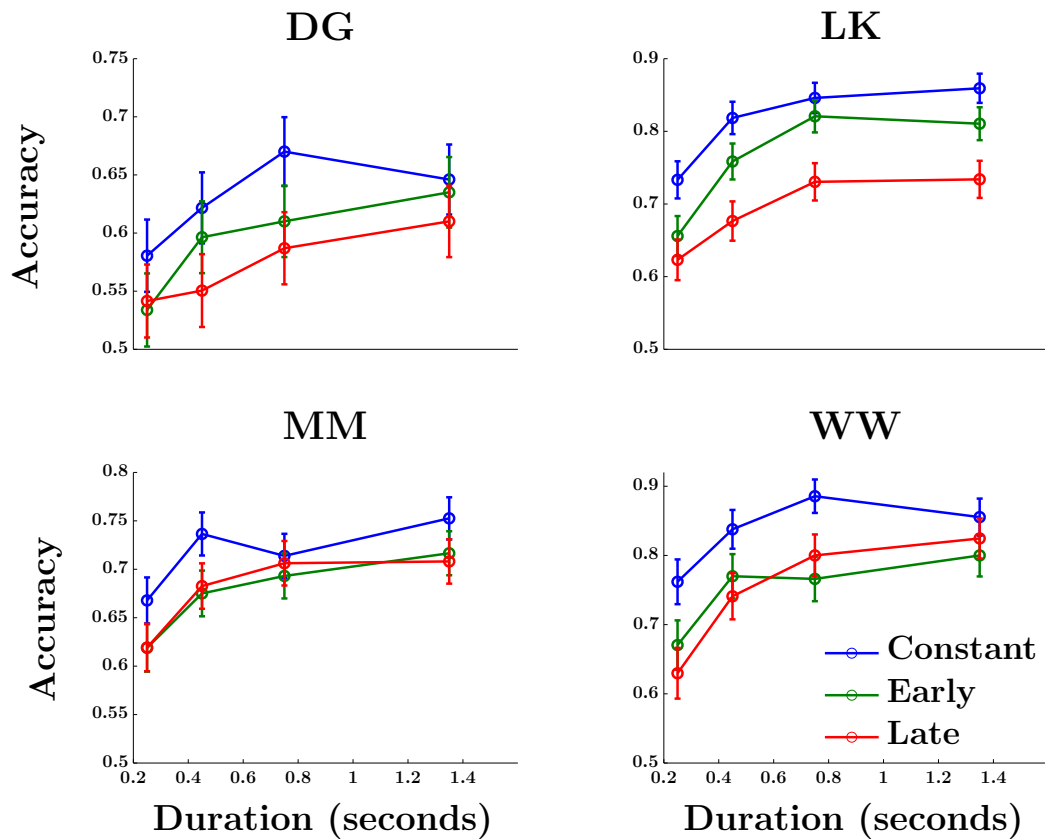


Figure 3.5: Accuracy as a function of stimulus duration and condition in Experiment 1b. Error bars correspond to 95% CI.

The ANOVA results are given in Table 3.1. As opposed to Experiment 1a, where all three observers showed primacy bias, only one out of the four observers (LK) showed a significant primacy effect. Interestingly, one observer (WW) did not show a main effect of primacy, but exhibited a significant interaction between the temporal bias and duration (i.e., primacy at short durations and recency at longer durations), as predicted

by the non-linear inhibition dominant LCA.

Table 3.1: Statistical analysis for the four participants of Experiment 1b (main order effect and its interaction with trial duration).

<i>Subject</i>	<i>Main effect (early/ late)</i>	<i>Interaction: Duration X Bias</i>
DG	$F(1, 63) = 2.715; p = 0.104$	$F(3, 63) = 1.353; p = 0.259$
LK	$F(1, 75) = 40.527; p < 0.001$	$F(3, 75) = 1.297; p = 0.276$
MM	$F(1, 59) = 0.001; p = 0.982$	$F(3, 75) = 1.392; p = 0.247$
WW	$F(1, 43) = 0.062; p = 0.805$	$F(3, 43) = 3.410; p = 0.020$

3.3.3 Discussion

The results of the experiments 1a and 1b, beyond their original scope which was the reduction of the primacy bias due to procedural changes (see Tsetsos, Gao, et al., 2011 for details) revealed individual differences. Of particular interest was to test if any of the observers showed the unique prediction of the non-linear, inhibition-dominant LCA. This was the case for one observer in each of the experiments (CS and WW) who exhibited the predicted signature of the nonlinear inhibition dominant LCA, i.e. the interaction between the temporal bias and the length of the trial.

Although the bias-duration interaction was not universal across all participants, it still poses a major challenge to models of evidence integration. The coexistence of two different weighting profiles within a single decision strategy and the triggering of each profile by bottom-up factors (i.e. trial duration), can be accounted only within the non-linear LCA. Note that the non-linearity of the LCA model is not arbitrary but has biological motivation, capturing the fact that neural firing rates cannot be negative. One interim conclusion that can be drawn, is that this mechanism has the advantage of prioritizing early information, in a flexible, reversible manner.

One caveat of the current Experimental Study is that the two experiments were not designed to explicitly test the interaction between the temporal bias and the trial length. On the contrary, there were factors, especially in Experiment 1a, which encouraged primacy weighting and which might have suppressed the bias-duration interaction. Given that in a number of other experimental paradigms strong recency patterns (Pietsch & Vickers, 1997; Usher & McClelland, 2001; Newell, Wong, Cheung, & Rakow, 2009)

have been obtained, it is important to explicitly test the interaction pattern with different experimental stimuli and protocols. This is done in the next section where the effect of the sequence length on the temporal weighting profile is tested in the value integration paradigm (*Experimental Study 2*).

3.4 Order Effects in Value Integration (Experimental Study 2)

Order effects have been extensively studied in sensory decisions but also in higher-order inference problems where information is presented sequentially (Walker, Thibaut, & Andreoli, 1972; Hogarth & Einhorn, 1992; Furnham, 1986; Lagnado & Channon, 2008; McKenzie, Lee, & Chen, 2002; Trueblood & Busemeyer, 2010; Gerstenberg, Lagnado, Speekenbrink, & Cheung, 2011). Temporal biases are also relevant in cases where the information presentation is not sequential or externally controlled, like when people decide among value-based alternatives which are presented to them simultaneously. There, assuming that during deliberation a covert sequential scanning of the relevant choice-aspects takes place, the order with which different dimensions are processed might determine the choice outcome.

For instance, imagine a choice between two cars differing in quality and economy. The first car is luxurious and expensive (A) while the second is basic and much cheaper (B). The alternatives are equivalent in terms of overall subjective value and in the absence of temporal biases the decision-maker is indifferent. However, if the decision maker exhibits primacy weighting and starts by considering the quality dimension first then, even if the two dimensions are equally important to her, she will end choosing the luxurious car (A). On the contrary if the decision-maker weighs information by recency then the cheap car (B) will be chosen because it is the last one to be favoured. And to push the story a little further, assuming an interaction between information weighting and deliberation length, fast decisions will lead to alternative A while more elongated ones to B, even if the choice-aspects are processed with the exact same order in both cases, starting by considering quality first.

In this section I will examine the profile of temporal weighting in value integration, using the fast value integration paradigm (see Chapter 2). Participants see rapid, varying sequences of pairs of numerical values and choose the sequence with the highest

overall value. In some trials, the two alternatives will have equal means but the temporal distribution of the numbers will be controlled such that the one sequence will appear better in the beginning and worse towards the end of the presentation. The choice preference in these trials will provide a direct measure of temporal weighting. Additionally, building on the interaction between bias and duration, demonstrated in Computational Study 1, and following on the results of Experimental Study 2, the sequence length will be varied in order to examine the way it influences the order effect.

3.4.1 Method

Participants

Sixteen adults (9 females; aged 19-35; mean age 23.4 years) were recruited from UCL's subject pool and were paid for their participation. The experiment was conducted in two sessions with a maximum of 3 days lag between the two sessions.

Stimuli and Experimental Task

At each trial, participants saw pairs of numbers presented sequentially and had to decide, within 1500 ms, which of the 2 sequences had the highest average value. Each trial started with a presentation of a white fixation cross for 1000 ms, which was positioned at the centre of a black background screen. Afterwards, sequences of pairs of white numbers were presented at a rate of 2 or 4 items per second. The presentation of the last pair of numbers was followed by a green question mark at the centre of the screen for 1500 ms, which prompted the participants to indicate their response (left or right sequence) by pressing the left or the right arrow on the QWERTY keypad of a standard PC. After the response of the participant a black screen stayed on for 250 ms and then the next trial started. For incorrect responses error-feedback (a beep sound) was provided. Failure to respond within 1500 seconds after the response cue's appearance was followed by a "deadline missed" message and a beep sound. Stimulus display and response recording were controlled by Matlab 7.11.0 (Mathworks Inc., Natick, MA, USA) using the COGENT 2000 toolbox (<http://www.vislab.ucl.ac.uk/cogent.php>). The time course of an experimental trial is given in Figure 3.6.

Procedure

Participants were assigned to two different groups. The "slow" group (N=8) per-

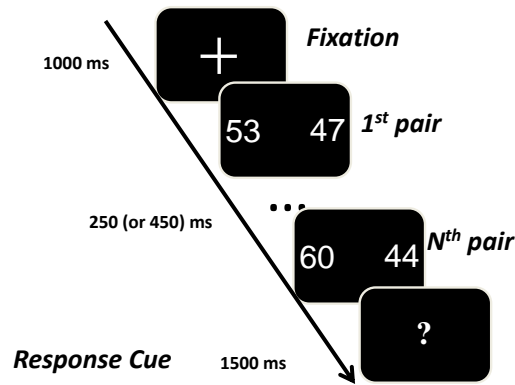


Figure 3.6: *The timeline of an experimental trial in Experimental Study 2. Participants saw pairs of numbers which alternated rapidly, and at the end of the presentation they had to decide which sequence had the highest average.*

formed the task at a presentation rate of 2 pairs/ second while the “fast” group at a rate of 4 pairs/ second. Before the experiment a 10 minutes calibration process was conducted in order to adjust the difficulty of the task for each participant. At each trial of the process, participants saw two sequences of 10 numbers each. One of the sequences was “high” and the other “low”, with their numerical values being generated from Gaussian distributions with mean 50 and 42 respectively. The position of the options was always randomized. The standard deviation of the Gaussians was adapted through a staircase procedure such that the standard deviation of the Gaussians, resulting in 79% accuracy, was estimated at the end (3 up 1 down procedure and step=0.5 units, see also Levitt (1971)). The process started with a standard deviation randomly chosen between 5 and 20 and was terminated after 30 swaps in the direction of the staircase. The obtained standard deviation was used throughout the main experiment for each participant. There was no significant difference in the values of the estimated standard deviations between the participants of the two presentation rate groups ($t(14) = -0.71, p = 0.49$). The mean value of the standard deviation was 11.68 ($SD = 2.78$).

Experimental Conditions

The main experiment consisted of 3 conditions and the repeated measure factor was the sequence length which was varied at 3 levels: 6, 12 and 24 pairs. Each session consisted of 300 trials (600 overall), 100 (overall 200) for each sequence length. The

trials were fully randomized and there were 10 (overall 20) blocks of 30 trials each. At the end of each block the participants were informed about their accuracy score up to that point.

The two first conditions, called hereafter “unbalanced”, involved the selection between two sequences which were generated by one “high” and one “low” distribution (the standard deviation of both distributions was tailored to each participant through the staircase procedure described above, see also Figure 3.7). In both conditions the dependent variable was the decision accuracy (i.e. the fraction of choices for the “high” distribution sequence). In the 1st condition the mean value of the high distribution was between 45 and 55 (randomly determined at each trial) while the mean value of the low distribution was always 8 units smaller. Overall 150 Condition-1 trials (50 for each sequence length) were presented. The 2nd condition was identical to the 1st however the highest number was always placed in the “low” sequence. To match the two conditions for difficulty we created Condition-2 trials by taking Condition-1 trials and modifying the low sequences. We did so by adding a constant to one of the low sequence numbers in order to make it the largest of the trial. To maintain the summed difference between the two sequences equal, we subtracted from another number of the low sequence the same constant that was added. Condition-2 trials were necessary to examine whether participants applied a rule of thumb, according to which the sequence where the largest number appeared is chosen. One hundred and fifty Condition-2 trials (50 for each sequence length) were presented. The sequences of the two first conditions were re-sampled until they were equated in difficulty, for all sequence lengths (to diminish the effect of sampling error).

The third condition, hereafter called “balanced”, consisted of two sequences generated from the same distribution whose mean was randomly determined by a uniform distribution in the 45-55 range. Crucially, in the first half of each trial the values of the 1st option (labelled “high-first”) were sampled from a truncated Gaussian, clipped one standard deviation below the mean (Figure 3.8(a)). At the second half of the trial the values of that option were sampled from the truncated Gaussian clipped one standard deviation above the mean. The values of the second option, labelled “low-first”, followed the exactly opposite time-course to those of the “high-first” option, with lower values appearing in the 1st half and higher in the 2nd (Figure 3.8(b)). Overall both options had, by definition, the same mean value and no feedback was provided in these trials (also these trials were excluded from the calculation of accuracy scores presented

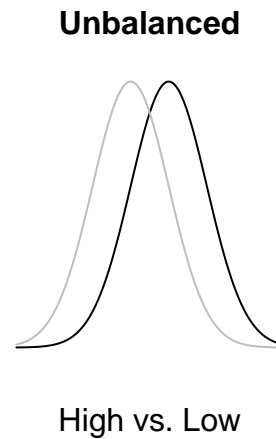


Figure 3.7: In the unbalanced conditions the numbers of the one sequence are always drawn from a distribution with higher (i.e. 8 units) mean.

at the end of the blocks). Overall 300 balanced trials (100 per sequence length) were presented.

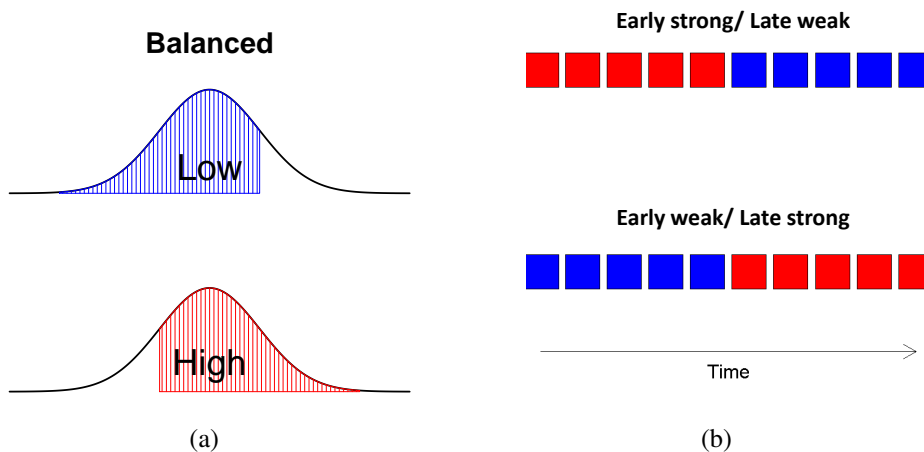


Figure 3.8: The balanced condition experimental condition. a: Clipping of a normal distribution one standard deviation above the mean (blue) and one standard deviation below the mean (red), b: Outline of a balanced trial; for one option (early first) the values of the first half are sampled from the high end of the distribution (red area) and from the low end during the second half. The time course of the other option is exactly opposite.

3.4.2 Results

The participants' ability to perform the task is reflected in their accuracy in the unbalanced conditions. At both presentation rates, participants were able to select the alternative associated with the highest overall value ($t(14) = 18.93, p < 0.001$) and their accuracy increased with the length of the sequence ($F(1, 14) = 91.09, p < 0.001$, Figure 3.9), indicating integration. This conclusion is further supported by the rejection of a simple heuristic decision rule according to which participants choose the sequence where the highest number appeared. In particular, Condition 2 trials were identical to those in Condition 1. However in these trials the maximum number appeared in the low-average sequence. Yet, in Condition 2, participants were able to choose the high-average sequence above the chance level (red line in Figure 3.9, $t(15) = 16.43, p < 0.001$) which rules out a maximum-number choice strategy. The difference in the accuracy between the two conditions is marginally significant ($F(1, 15) = 4.96, p = 0.042$) indicating that the presence of the maximum number in the low-sequence affected but did not reverse the choice outcome.

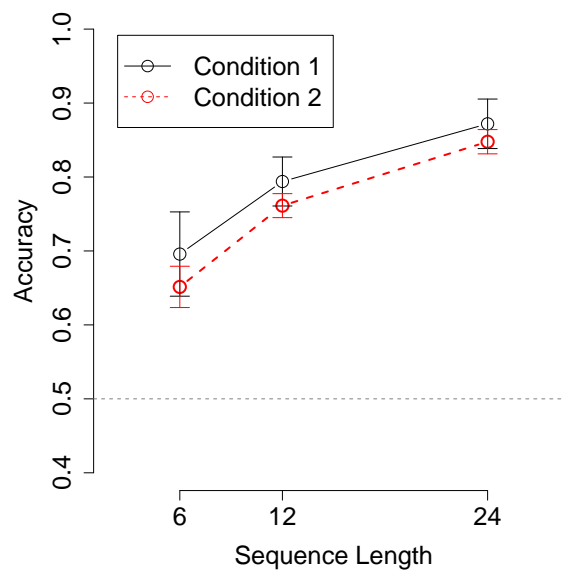


Figure 3.9: Performance in the two conditions of the unbalanced trials. Condition 2 differs from Condition 1 in that the maximum number appears always in the low-average sequence. In both conditions accuracy keeps increasing with the sequence length. Error bars correspond to 95% CI.

To understand the properties of the value integration mechanism, I subsequently questioned the presence of order effects. This was done by examining the choice preference for alternatives with the same mean (balanced sequences), but which differ in the temporal distribution of values, such that one option appears better in the first half and worse in the second. Both presentation rates show a clear recency effect: the values of recent pairs are more strongly weighted ($t(15) = 7.76, p < 0.001$). Moreover, recency increases with sequence length ($F(1, 14) = 15.89, p < 0.005$, Figure 3.10), as the impact of earlier values decays, consistent with leaky (decay-based) integration. The interaction between the temporal bias and the trial length, predicted in Computational Study 1, is clearly found; the longer the trial the higher the recency. However, as opposed to Computational Study 1 and the two participants in Experimental Study 2, no primacy pattern is found and thus there is no qualitative transition of the order effect from primacy to recency.

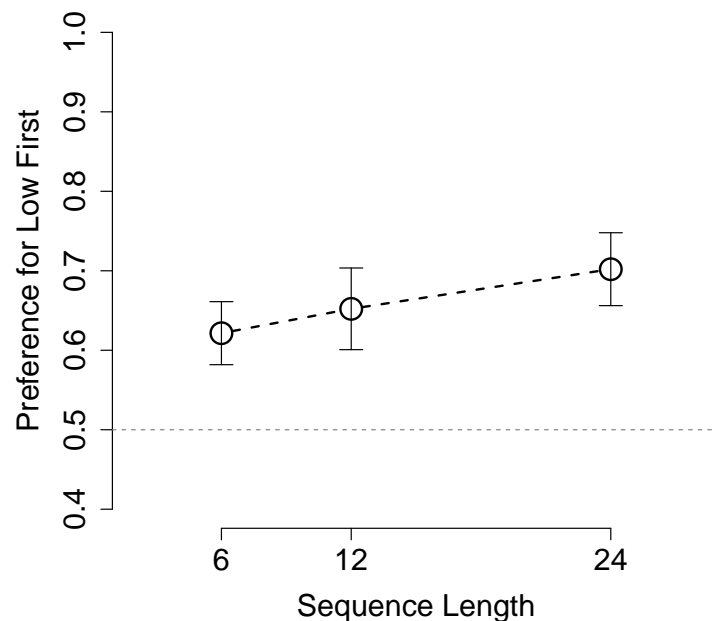


Figure 3.10: Choice preference in the balanced trials showing increased preference for the sequence that begins with low numbers and ends up with high (low-first). The recency bias increases with the sequence length. Error bars correspond to 95% CI.

In order to further confirm the strong recency bias, a logistic regression of the actual input the participants saw in the unbalanced sequences was conducted. As Figure 3.11

shows, the last items in the sequence receive higher weights, for all the three different sequence lengths.

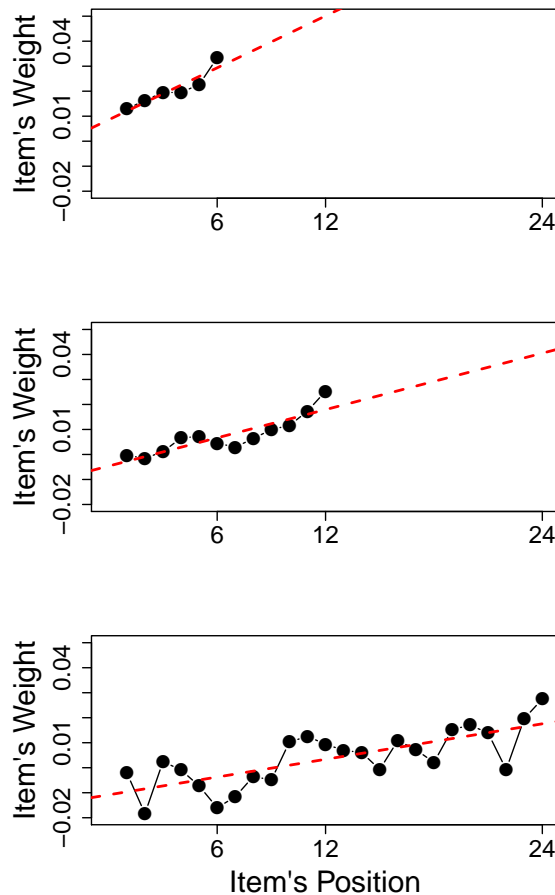


Figure 3.11: Logistic-regression weights in the unbalanced trials in Study 1 for all participants. The weights of the items at different positions in the trial are shown, for trials with different sequence lengths (length= 6, 12, 24 in A, B and C). The dashed red regression lines show the increasing trend in the weights of the later items (i.e. recency effect).

3.4.3 Computational Models (Computational Study 2) and Discussion

When dealing with novel choice alternatives, the cognitive system must integrate information about the features of these alternatives. A central question I addressed in Experimental Study 2, is how such “value integration” occurs in a simplified context where many instances of values on the same currency are presented sequentially. The

main questions of interest were the dependence of the evaluation accuracy on sequence length (the integration bound), and the temporal profile with which decision-makers weigh the values (recency/ primacy). The results demonstrated a significant range of integration since accuracy improved, even in the range of 12-24 items, and a significant recency which increases with sequence length.

In order to characterize the mechanism of value integration I test several models that could account for the data. These are: i) a perfect integrator, ii) a leaky integrator, iii) a model which samples k -pairs (randomly) and forms a decision based on that sample, iv) a model in which each pair has a probability of p to be ignored, and v) a model which considers items only if they are above a threshold. All these models assume some type of integration, since simpler heuristic strategies are rejected by the data (see Figure 3.9).

First I consider models iii-v) on a qualitative basis. The k -samples model is not in a position to account for the increase in performance, unless it assumes that k is close to the maximum sequence length, in which case it becomes equivalent to perfect integration. Furthermore, this model does not make any prediction regarding order effects and recency. Model iv) assumes that the probability of encoding and thus integrating a pair is a free parameter p . This model, can predict the increase of performance with sequence length. The longer the sequence the more items will be integrated and thus more noise will be average out. However, no prediction regarding the temporal order of information is provided. This is also the case with model v) which assumes that the observers adopt a strategy whereby items below a threshold are ignored. Again, this model predicts the performance improvement but says nothing about order effects ¹.

Next, I focus on the two first models the perfect and leaky integrator. These models assume continuous integration in a fashion similar to the diffusion model of perceptual choice. In order to model the value integration task, I assume two sequence of numbers, V_A and V_B , presented sequentially for N frames ($t = 1..N$ with $step = 1$). The distributional characteristics of the sequence values were identical to those of the actual experiment (see *Methods*) The preference state at time t is defined as:

$$P(t) = \lambda \cdot P(t-1) + [V_A - V_B] + N(0, \sigma). \quad (3.2)$$

At the end of the trial ($t = N$), if the preference state is positive a decision is made

¹Variations of these models have been suggested to account for a similar paradigm where decisions are based on samples of values, actively obtained by the participant (Hertwig et al., 2004)

in favour of A and otherwise for B . There are 2 parameters; λ is the decay parameter which is fixed at 1 for the perfect integration and is a free parameter (between 0 and 1) in the leaky integration, determining the time constant of the integration. The other parameter σ defines the additive/internal noise. By definition, the perfect integration model assigns uniform weights to all items and thus cannot produce order effects. On the other hand the leaky integration model can generate recency weighting since, due to the decay, early information dissipates.

In order to assess the descriptive adequacy of the two models, I fitted them on the average data in the balanced and unbalanced trials (Figure 3.12). As expected, both models captured perfectly the performance improvement with sequence length. Regarding the recency pattern, the perfect integration model missed both the main effect and the interaction with the duration (Figure 3.12(b)). On the contrary, the leaky integration accounted well for recency, predicting also that it increases with sequence length. A summary of the fitted parameters and the BIC values of the two models is given in table 3.2.

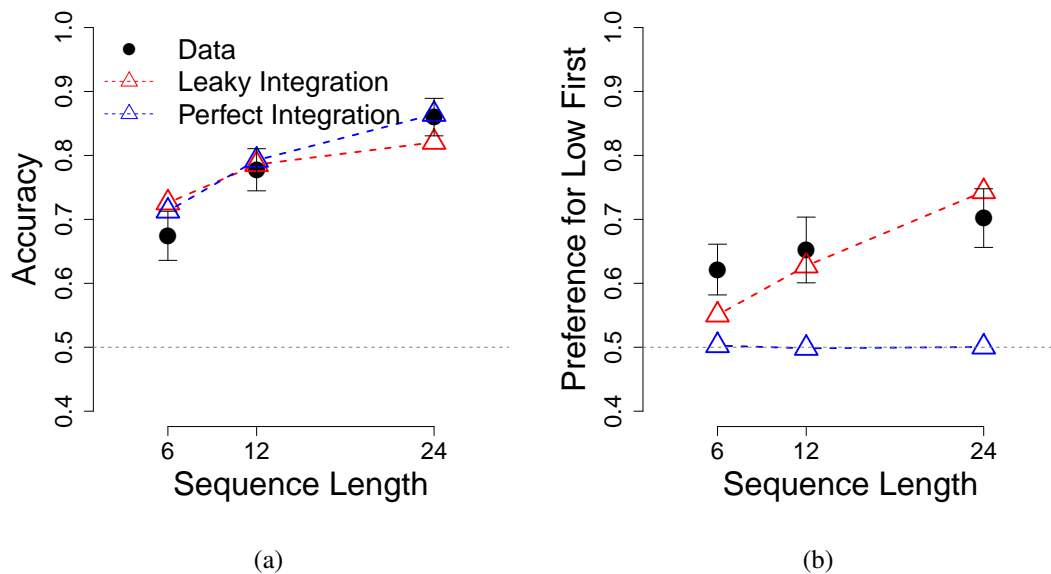


Figure 3.12: Data fits of perfect (blue) and leaky (red) integration with two parameters (Equation 3.2). a: Unbalanced trials, b: Balanced trials.

Despite that the leaky integration captured all the range of patterns in the data, the quantitative fit on the balanced trials was not satisfactory. To improve that and to also allow the perfect integration model to capture the recency effect, I assumed that the last

Table 3.2: Model parameters and BIC values for perfect and leaky integration).

<i>Model</i>	<i>Processing noise</i>	<i>Leak</i>	<i>BIC</i>
Perfect Integration	26.2	1	33
Leaky Integration	28.9	0.86	23

pair of the sequence is overweighted, since it is presented last and unmasked. Thus, I introduced an extra free parameter w in both models and fitted again the average data based on the following equation:

$$P(t) = \lambda \cdot P(t-1) + w_t[V_A - V_B] + N(0, \sigma), \quad (3.3)$$

with $w_t = 1$ for $t < N$ and $w_t = w$ for $t = N$. As Figure 3.13(a) shows, both models are able to capture the increasing performance in the unbalanced trials. The addition of parameter w , allowed the perfect integration model to account for the main effect of primacy by overweighting the very last pair. However the model predicts that the temporal bias is flat and that it does not depend on the sequence length. A summary of the fitted parameters and the BIC values is given in table 3.3.

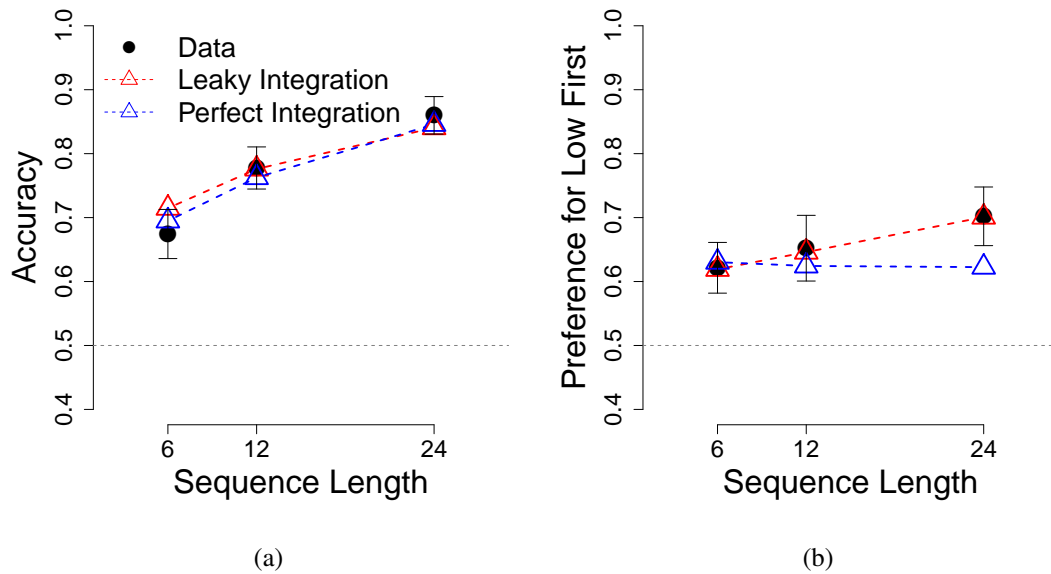


Figure 3.13: Data fits of perfect (blue) and leaky (red) integration with three parameters (Equation 3.3). a: Unbalanced trials, b: Balanced trials.

Thus, the only model capable of simultaneously capturing the increase of recency and the improvement in accuracy with longer trial durations, was the leaky integration. In

Table 3.3: Model parameters and BIC values for perfect and leaky integration).

<i>Model</i>	<i>Processing noise</i>	<i>Leak</i>	<i>Weight (w)</i>	<i>BIC</i>
Perfect Integration	5.3	1	8.1	20
Leaky Integration	14.1	0.94	3.4	12

order to demonstrate how the increasing recency pattern is accounted by leaky integration, I recast equation 3.2 in the form of a differential equation, assuming also that there is no internal noise and that the input of the accumulator is fixed to I :

$$\frac{dP}{dt} = (\lambda - 1) \cdot P + I.$$

The solution of the differential equation is:

$$P = \frac{I}{\lambda} (1 - e^{\lambda t - t}).$$

Therefore the weights assigned to the pieces of information decrease exponentially from the last to the first item. As a result the longer the sequences are the less the impact of the early items will be because the effective time constant of the integration is finite and equal to $1/\lambda$. Intuitively the leaky integration suggests that as time passes early items are forgotten more.

To conclude, leaky integration, a simple and parsimonious mechanism, can account for fast and automatic value-integration in tasks that require the online construction of preference to novel alternatives. Leaky-integration in combination with lateral (response) inhibition, is proposed to underlie the accumulation of evidence in perceptual choice (Usher & McClelland, 2001). However, unlike in perceptual choice where the temporal profile of evidence weighting indicates primacy (see Experimental Study 1) consistent with inhibition dominance, here the temporal profile indicates recency, consistent with leak-dominance (see also Hertwig et al., 2004 for a similar weighting profile in feedback-driven value decisions). This qualitative discrepancy can be explained by considering the differences in typical perceptual choice tasks and the fast value integration paradigm. While in the perceptual experiments of Study 1 the maximum trial length was 1750 ms, in the current study the maximum duration was 6000 ms (the fast group and 12000 ms in the slow group). One additional factor that might have suppressed primacy in the value integration experiment could be the fact that the last pair of numbers was presented unmasked and, inevitably, was overweighted. Despite the differences in the direction of the order effect found in the two studies, a fundamental

aspect characterizing both experiments is that information, either sensory or value related, is integrated in a continuous fashion with a limited effective time constant (i.e. subject to decay). This leads to differential information weighting (whose profile depends on factors such as task attributes or interpersonal differences) which violates the statistical optimality imposed by Bayesian inference and the SPRT. In the next section I recast the principle of optimality in decisions under uncertainty, showing that when the environment unpredictably changes the SPRT is a suboptimal decision strategy.

3.5 Optimality in Decisions under Uncertainty (Computational Study 3)

It is customary to interpret mechanistic models with evolutionary terms assuming that cognitive mechanisms that favour survival are more likely to have been maintained by natural selection. Across these lines, it is reasonable to assume that humans are equipped with decision mechanisms that generate fast and accurate decisions about the most likely causes of sensory experiences. As discussed earlier (see section 1.1), the view that humans make perceptual decisions in an optimal fashion is widespread, inspiring mechanistic models that are statistically optimal (e.g. the diffusion model which implements the SPRT Wald, 1947; Bogacz et al., 2006). However, as underscored in this chapter (sections 3.3 and 3.4) and also by a series of empirical findings on temporal biases (Hogarth & Einhorn, 1992; Usher & McClelland, 2001; Kiani et al., 2008; Gerstenberg et al., 2011), the order with which information is perceived causes distortions to its relative importance.

Order effects undermine the idea that humans, when confronted with noisy information, conform to the principles of Bayesian inference, weighing equally all pieces of evidence. A likely explanation for order effects is that the integration is continuous (i.e. all pieces of information are considered) but also subject to leakage or decay. If leakage is indeed an inherent aspect of information integration then one can superficially deduce that humans are suboptimal. Nevertheless, optimality is environment dependent and can be realized only in relation to the task at hand. Real world choice problems are often framed in volatile contexts, where the underlying structure of the environment can change without warning. It is thus more likely that natural evolution has promoted the flexibility and the ability to adapt in response to the environmental

demands rather than a mechanism that is optimal only in stationary environments (i.e. SPRT).

In unpredictable and non-stationary environments, observers need to discriminate whether unexpected events happen due to noise or due to a change in the state of the world (Yu & Dayan, 2005). For instance, an online order might be late, either because of typical delays (i.e. noise) or because the dispatch company closed (i.e. change of state). In order to assess the merits of leaky integration in such environments, I simulated a task where participants need to detect signals embedded in a continuous stream of noise. As I will demonstrate in the next section, perfect integration in an ever changing world proves to be a suboptimal decision strategy.

3.5.1 Methods

Each trial consisted of 400 time-steps. In half of the trials ($N = 10000$, labelled as “no signal trials”, Figure 3.14(a)) the input consisted of white noise only. In the other half ($N = 10000$, labelled as “signal trials”, Figure 3.14(b)), the input was constructed by embedding on top of white noise (i.e. $N(0, 1)$) a transient event (sampled from $N(1, 1)$) of 10 time-steps duration. The onset of the signal was sampled from a uniform distribution between 10 and 390 time-steps.

The input was integrated in one leaky accumulator according to the following equation:

$$\frac{dy}{dt} = -\lambda y + input(t). \quad (3.4)$$

For the special case where leakage is zero the integration is perfect and equivalent to the diffusion model in the 2AFC task. Once the activation y of the accumulator reaches a threshold (A) a response is initiated. If the response is made within the signal interval, in the signal trials, it is classified as a *hit*. On the other hand, if a response is made in the absence of signal (both in signal and no signal trials) then it is counted as a *false alarm*. If no response is generated in a signal trial it is a *miss*. Finally, if no response is made in a no signal trial it is considered to be a *correct rejection*.

3.5.2 Results and Discussion

The model in equation 3.4 was simulated for different values of the decay parameter and different criterion values ($A = 1 \dots 28$ with $step = 3$). For each instance of the model

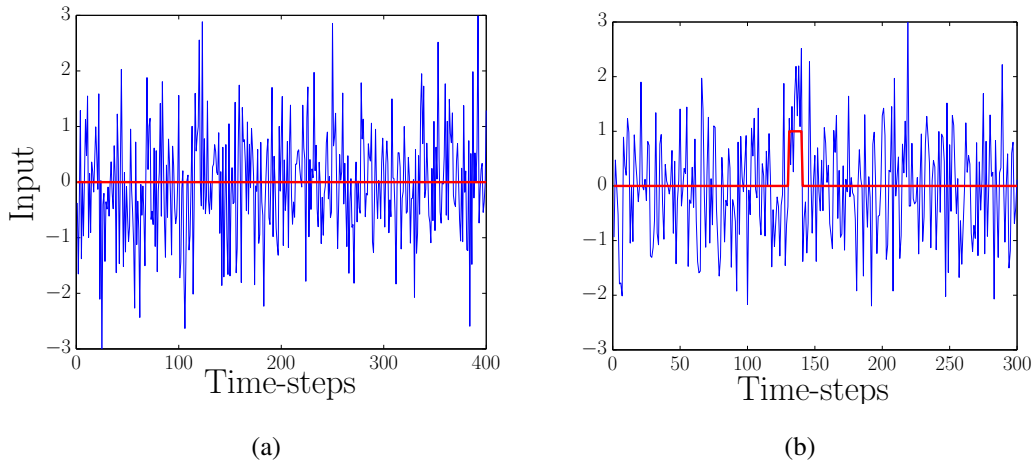


Figure 3.14: *a: Input in the no signal trials corresponds to white noise; b: Input in the signal trials corresponds to Gaussian signal of $N(1,1)$ superimposed on white noise for 10 time-steps, at a random point in the stream. The blue line indicates the actual input while the red curve shows the onset and duration of the signal.*

the True Positive Rate (TPR) was calculated as: $TPR = Hit / (Hit + Miss)$ while the False Positive Rate (FPR) as: $FPR = False\ Alarm / (False\ Alarm + Correct\ Rejection)$. The ROC curves in Figure (3.15) show that the perfect integration ($decay = 0$, black line) performs poorly being even below the line of “no discrimination” (black dashed line)². The reason why perfect integration fails in this task, is that noise is accumulated before the occurrence of the signal, resulting into a very high False Alarm rate. On the other hand, leakage enhances performance by limiting the accumulation of noise. From all the leakage (decay) values used, the optimal one was $\lambda = 0.1$. This value of leakage restricts the effective time constant to $1/\lambda = 10$ time-steps, which coincides with the duration of the signal.

In order to better understand the dependence of the optimal time constant on the length of the signal, I simulated the “signal trials” of the same task by varying this time the length of the inserted signal across four levels: 10, 20, 40 and 80 time-steps. For each leak value the criterion (A) that results in 20% False Alarm rate (in trials where there is only noise) was estimated. Subsequently, the stimulus strength was adjusted such that for each duration and leak value the Hit Rate is 80%. Figure 3.16 shows the signal

²An ROC curve below the “no discrimination” line implies that the model has negative predictive power. In that case the model obtains positive predictive power by reversing its decisions.

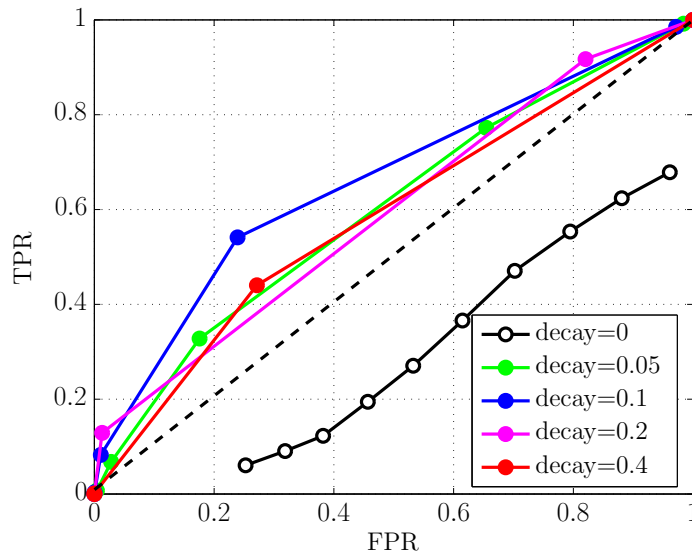


Figure 3.15: ROC detection curves for signal embedded in continuous noise for different levels of decay (λ in equation 3.4).

strength as a function of signal duration for each leak level.

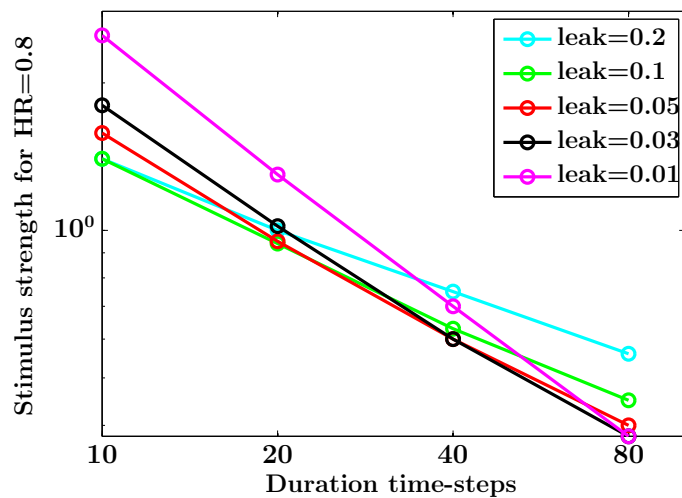


Figure 3.16: Signal strength as a function of the signal duration for each leak level in the log-log space.

For the shortest signal duration (10 time-steps) short time-constants (i.e. $\lambda = 0.2, 0.1$ and $\tau = 5, 10$ time-steps) work better, requiring the weakest signal to achieve a Hit Rate of 80%. However, short time-constants are outperformed by longer time-constants (i.e. $\lambda = 0.05, 0.03, 0.01$ and $\tau = 20, 33, 100$ time-steps) for the longest signal duration (80 time-steps, see Figure 3.17). Therefore, different leak values are ideal for different

signal durations with short time-constants being tailored to short durations but being less sensitive when the signal is longer and weaker (and vice-versa for long time-constants which are insensitive to the short signals).

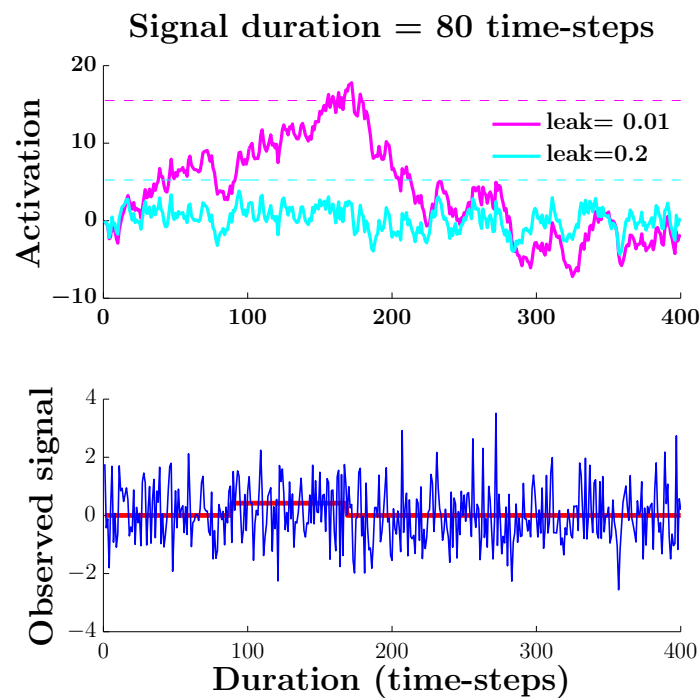


Figure 3.17: *Single trial activations (top) for long (magenta) and short (cyan) time constants, when the injected signal has duration 80 time-steps (bottom). The integrator with long time constant detects the signal. The strength of the signal is too low to cause a response in the short time-constant integrator, whose activation remains below threshold during the presence of the signal.*

To conclude, when there is uncertainty about the timing of the signal, perfect integration is sensitive to noise fluctuations and responds before the occurrence of the event. In these situations “forgetting” makes the decision process less susceptible to noise and improves the decision quality. However, when the duration of the signal varies there is no single forgetting rate (i.e. leak) for which the process is globally, for all signal durations, optimized. Whether a single optimal model exists, when the durations of the events are unpredictable is an open theoretical question. Additionally, future experimental work is needed to examine whether the time constant of the integration is adaptable to the statistics of the task (e.g. whether observers change their time constant in blocks where they anticipate mostly short signals compared to blocks where there are mostly long signals).

3.6 Summary and General Discussion

The present chapter addressed the differential weighting of information as a function of its temporal order. The dynamics of order effects were explored by questioning how deliberation length influences their magnitude and direction. In *Computational Study 1*, I showed a novel prediction generated by a neurally inspired process model of choice (LCA; Usher & McClelland, 2001): within the same parameter set, the information weighting automatically switches from primacy to recency as the length of the evaluated information increases. This prediction brings to light the possibility that order effects might be triggered by the bottom-up interaction of the input with a fixed decision algorithm (or strategy). It is important to note that this pattern emerges from a neurally plausible aspect of LCA, that activation cannot go below 0. And although the interaction of information processing with the input characteristics has been addressed in the past (Hogarth & Einhorn, 1992), it has never been done without assuming that model parameters are ad-hoc adjusted to reflect input or task contingencies.

In *Experimental Study 1*, I reanalyzed data from a perceptual experiment (Tsetsos, Gao, et al., 2011), aiming to detect a signature of the distinctive LCA prediction. Although the original scope of the experiment was to show the dependence of order effects on task demands (i.e. trial durations and response deadline), reviewing the data with the bias-duration interaction in mind, I detected this pattern in 2/7 participants. One limitation of this study was the task design which, encouraging primacy weighting, might have suppressed the interaction of the order effect with the trial length. This was not the case in *Experimental Study 2* where the interaction of length and weighting was explicitly looked for, using the fast value integration paradigm. There, recency was a universal pattern across all participants and it increased with the length sequence. The best fitting model for this experiment was a leaky integrator, which can be viewed as a leak-dominant instance of the perceptual LCA (which has additional parameters such as response competition and zero non-linearity applied on the preference states).

Comparing evidence (*Experimental Study 1*) to value integration (*Experimental Study 2*), one notices that mostly primacy is found in the first and recency dominates the latter. This discrepancy can be attributed to the different time scales used in the two different domains (much shorter trials in sensory decisions). This difference might underscore, in each case, different aspects of a single decision mechanism. Future experimental work needs to address whether the absolute duration of the decision drives

the main direction of the order effect, regardless of the decision domain. If so, support for a single, domain-independent decision mechanism will be provided. Meanwhile the theoretical analyses in this chapter (*Computational Studies 1 & 2*) are consistent with a single decision mechanism which enhances *response competition* in fast, sensory decisions and *leakage* in slower, value-based decisions. Intuitively, these different modes of information weighting can be understood in terms of response urgency. With a limited amount of samples, it makes sense to engage into comparative processing (which implies response inhibition) in order to quickly figure out which alternative is the best. On the other hand, when the time pressure is removed, it is likely that competition between the alternatives is relaxed and each of them is evaluated more independently in terms of its attributes only (i.e. a feed-forward integration).

The common factors underlying both the inhibition (i.e. perceptual decisions) and leak dominant (i.e. value-based decisions) computational accounts, is that choice is driven by the continuous integration of samples of information with the integration effective time-constant being finite. The latter, leads to imperfect integration which contradicts the optimal decision algorithm proposed by the SPRT for stationary environments. The suboptimality of leaky integration in stable, noisy environments, questions its evolutionary merit. Considered, however, within volatile environments, leaky integration is a robust and flexible decision mechanism as opposed to perfect integration which is vulnerable to noise fluctuations (*Computational Study 3*). Whether the amount of leakage that governs information integration in simple decisions is hard-wired or subject to adaptation to the environmental statistics is an interesting, open question which will, potentially, recast human optimality as the ability to modify choice strategies in response to the world demands.

Chapter 4

Context-dependent Weighting of Information

4.1 Overview

As discussed in the previous chapter, determining the goodness of an alternative is the result of the non-linear summation of its components, with the temporal order of the processing determining the relative importance of each component. Order biases, however, are not the only source of judgemental non-linearities. *Contextual biases* or the effect of the context ¹ on the subjective magnitude or goodness of an alternative, are widespread in cognition, from perception (Garner, 1953; Holland & Lockhead, 1968; Stevens, 1975; Luce, Nosofsky, Green, & Smith, 1982; Lockhead & King, 1983; Laming, 1997; Stewart, Brown, & Chater, 2005) to high-level judgement and decision-making (Tversky, 1972; Huber et al., 1982; Simonson, 1989; Birnbaum, 1992; Read & Loewenstein, 1995; Dhar & Glazer, 1996; Benartzi & Thaler, 1998; Dhar, Nowlis, & Sherman, 2000; Stewart et al., 2003; Pettibone & Wedell, 2007) and suggest that information is weighed differentially, depending on the context. One instance of contextual biases is when preferences between alternatives are reversed by the presence of decoy options (that are not chosen) or by the presence of other irrelevant options added to the choice set (see also section 1.2). This type of distortions, induced by the immediate context, beyond posing challenges for any theory of choice, can be

¹The context can refer either to the current choice set (i.e. immediate context) or to both the current environment and to items retrieved from memory (i.e. sampled context). Here I examine the influence of the immediate context.

revealing about the computational basis of preference construction. In this chapter, I will use contextual preference reversals (i.e., similarity, attraction and compromise effect) as a tool to elucidate the dynamics of evidence and value integration, in multi-alternative decisions. In the first section (*Experimental Study 3*), I use non-stationary, dynamic evidence in order to emulate multi-attribute, multi-alternative choice problems in a brightness discrimination task. There, from the three contextual effects, only a very strong similarity effect is found. The results are analysed within the sequential sampling framework of evidence integration (*Computational Study 4*) and reveal that context dependent weighting arises from the synergy of two mechanisms: the zero non-linearity of the preference states and the decisional (response) competition. Next, in the second section (*Experimental Study 4*), the presence of decoy effects are looked for in the fast value integration paradigm, with the results indicating robust attraction and similarity effects. These results, beyond their empirical significance², considered together with the findings of *Experimental Study 3*, indicate that evidence and value integration might be governed by distinct mechanisms.

4.2 Context Effects in Evidence Integration (*Experimental Study 3*)

The recent development of process models of multi-attribute decisions (Roe et al., 2001; Usher & McClelland, 2004; Tsetsos et al., 2010; Hotaling et al., 2010) has been based on the idea that, similar to evidence integration, value-based choice is driven by the integration of samples of values. Building upon earlier work by Tversky (1972), these theories assume that the samples of values are collected via a stochastic process which sequentially allocates attention to different choice aspects. One direct implication of this assumption is that preference states for different alternatives will have a temporal evolution that will reflect the similarity of the alternatives in the choice space. For example the preference state of two options that are similar, being advantageous and disadvantageous on the same dimensions, will be correlated rising and falling together. On the contrary, the consideration of two alternatives that are dissimilar will generate preference states that are anti-correlated across time.

²Attraction and similarity effects have never been reproduced within the same paradigm and within participants.

Although choice alternatives are often presented in different dimensions and value currencies, preference states are encoded in a single dimension, the firing rates of neural populations. Observing preference states, although possible using advanced neuroscience techniques, becomes more and more difficult when the decision problem is complex and the representation of the alternatives unknown. One alternative way to study preference dynamics is not to measure them directly but to invoke them by stimulating the cognitive system with dynamic information whose temporal profile is precisely controlled by the experimenter. That way it is possible to emulate the dynamics of preference that presumably occur in multi-attribute, multi-alternative problems, assuming of course that choice is driven by attentional switching across choice aspects. In order to test whether the attentional switching hypothesis is valid and whether stimulating the cognitive system can be a useful technique for the study of computational micro-mechanisms that are obscured by covert mental states involved in decision making, I created a psychophysical task where non-stationary, perceptual evidence is presented to the observer (see also section 2.2.2 for a more detailed description of the stimulus construction rationale). By manipulating the time course of the evidence, I created situations analogous to the attraction, similarity and compromise effects (see section 1.2 for a description of the effects) that are widespread in more complex, multi-attribute domains.

4.2.1 Method

Participants

Sixteen participants recruited from the University College London subject pool were tested over two sessions that took place in different days with a maximum difference of one week.

Stimuli

Each trial involved four alternatives of varying brightness, with mean brightness specified separately for each of the two phases. Thus, the brightness was non-stationary, based on a stochastic transition between two phases. In *phase 1*, the brightness of each patch (m) was sampled (at each time frame) from a normal distribution, $N(\mu_{m1}, \sigma_{in})$, while in *phase 2* it was sampled from $N(\mu_{m2}, \sigma_{in})$ ($\sigma_{in} = 0.1429$). One of the four patches (D) was so dim that it was virtually never chosen, with the effect that the ex-

periment effectively involves only three meaningful choice alternatives. The extra dim spot was added to balance the positions of the meaningful alternatives around the corners of an imaginary square. For the dim patch (D), the brightness fluctuation SD was only 0.01. The screen positions of the A, B, C, and D alternatives were randomized.

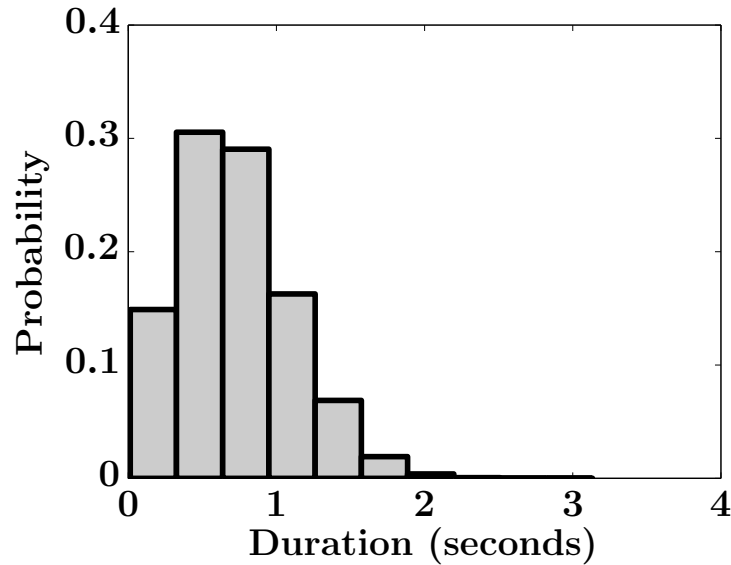


Figure 4.1: *Density distribution that determines the time of switching from one phase to the other.*

Each trial started randomly with either *phase 1* or *phase 2*. The transition times from the one phase to the other were selected from the distribution in Figure 4.1. Although the stimuli were presented on a monitor without applying a Gamma correction, the measurement of the monitor non-linearity with a photometer showed that the deviation from linearity was very small in the range (0.4-0.8). Gaussian noise added to the stimulus value could cause brightness to fall above 0.8, however, and the largest brightness value allowed was 1.0.

Conditions

A total of eight conditions were interleaved in the experiment. Each condition involved four alternatives of varying brightness, with mean brightness specified for each of the two phases. The critical conditions corresponded to the attraction, similarity and compromise effects. One fourth critical condition, labelled anti-compromise, was added. The non-critical or filler conditions were such that there was always one alternative with the highest integrated evidence (treated as the correct response and used to determine the participant feedback). The precise stimulus value in the critical and filler

conditions are given in Table 4.1. In total, for each condition, 50 trials were presented (25 at each session). At the critical conditions the integral of the evidence for two (or three) options had the same average value across the two phases. However, because the duration of the trials is limited, a small imbalance can occur, such that the alternative(s) that receive(s) more support at the beginning also receive(s) the most support on 65% of the trials.

Table 4.1: Experimental conditions. The mean values of the dim option D are omitted here since they were always the same: $\mu_1 = 0.1$ and $\mu_2 = 0.1$.

<i>Conditions</i>	Option A			Option B			Option C		
	μ_1	μ_2	Avg.	μ_1	μ_2	Avg.	μ_1	μ_2	Avg.
Attraction(C)	0.8	0.4	0.6	0.45	0.13	0.29	0.4	0.8	0.6
Similarity(C)	0.8	0.4	0.6	0.8	0.4	0.6	0.4	0.8	0.6
Compromise(C)	0.8	0.4	0.6	0.4	0.8	0.6	0.6	0.6	0.6
Anti-compromise(C)	0.8	0.4	0.6	0.8	0.4	0.6	0.6	0.6	0.6
Inconsistent-hard(NC)	0.55	0.3	0.425	0.3	0.5	0.4	0.3	0.4	0.35
Inconsistent-easy(NC)	0.7	0.6	0.65	0.4	0.7	0.55	0.4	0.2	0.3
Consistent-hard(NC)	0.6	0.7	0.65	0.55	0.4	0.475	0.2	0.4	0.3
Consistent-easy(NC)	0.8	0.6	0.7	0.3	0.5	0.4	0.4	0.2	0.3

For the critical conditions, the logic was to create temporal correlations in the evidence similar to attentional switches to dimensions that favour trade-off alternatives. In the attraction condition, options A and B were correlated, with option A always dominating B. Option C was anti-correlated to both A and C. In the similarity condition, options A and B were correlated and equal to each other, and anti-correlated with C (see Figure 4.2 for an example of a similarity trial). In the compromise condition, A and B were anti-correlated and equal on average, with option C being stationary and having always brightness equal to the average brightness of A and B. Finally, the anti-compromise condition was created in a similar way to the compromise condition with the difference that A and B are overall equal but correlated.

The filler or non-critical conditions, were labelled, inconsistenthard, inconsistenteasy, consistenthard, and consistenteasy, where consistent indicates that the evidence favours one of the alternatives at all time (consistent evidence), and inconsistent that the evidence favors different alternatives at different times.

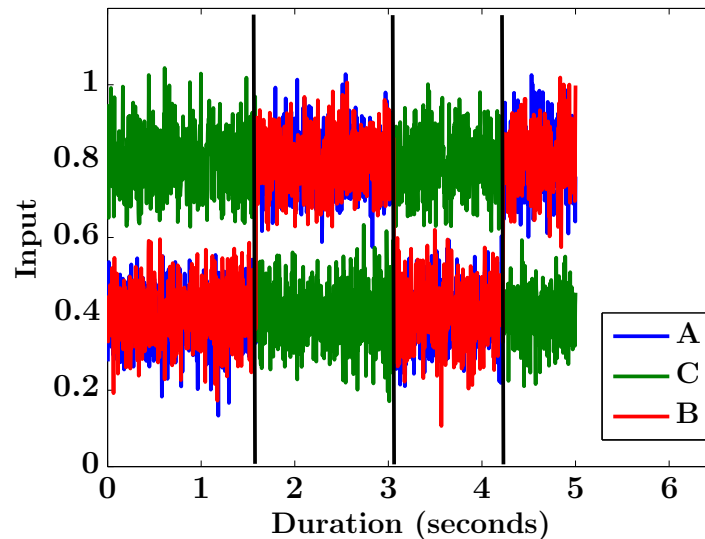


Figure 4.2: *The non-stationary evidence to three alternatives in a single trial of the similarity condition, with 5 s duration. The evidence includes Gaussian noise on top of a changing baseline. A-blue, B-green, C-red (phase switches are marked by vertical black lines).*

Procedure

The sessions were run on different days with a maximum of a week difference. Before the beginning of the experiment a brief explanation of the task was given and the participant was presented with 510 examples of the stimulus. The input values for these trials were randomly chosen. Immediately after the introductory trials, 2550 trials sampled from the experimental conditions were presented for practice (the introductory and practice trials were given in the first session only). The practice period ended when no error-beeps occurred for five consecutive trials (see below), but no earlier than 25 trials and no later than 50 trials. The main experiment had 200 trials per session (400 trials overall) and 8 conditions (50 trials for each condition). The 200 trials for each session were broken into 5 blocks (40 trials each). Trials within each session were randomized across all eight stimulus conditions. After each block participants were shown their accuracy score up to that point in the experiment and took a short break (15 min).

Each trial began with the presentation of a fixation cross. After 1 s, four patches appeared on the screen around the fixation cross, in a square formation. The brightness of each patch fluctuated across time (the brightness was updated every 13.3 ms, corre-

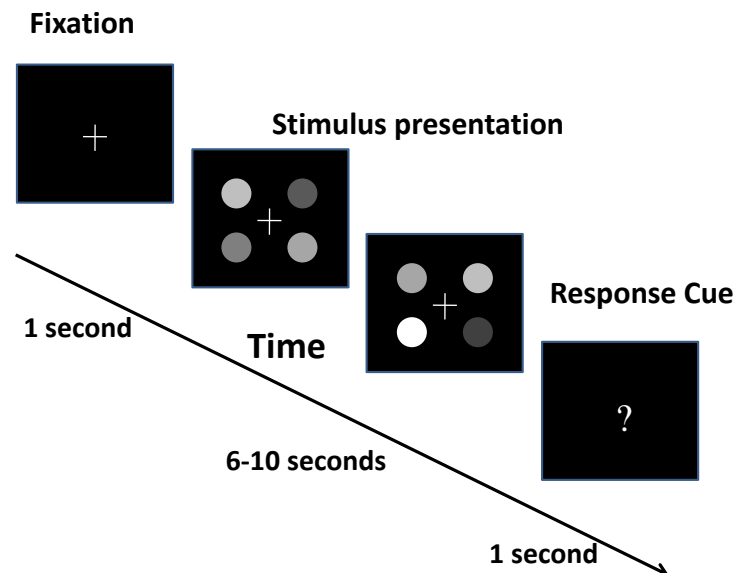


Figure 4.3: *The time course of an experimental trial.*

sponding to the frame rate of the monitor) and the participants had to select the patch that was the brightest overall (see Figure 4.3). The duration of the stimulus presentation was chosen randomly from a uniform distribution between 5 and 10 s. Upon termination of the stimulus presentation the participants had 1 s to make a response. If the participant failed to respond within this interval, a “Response deadline missed” screen was shown and the next trial started. For incorrect responses (in the non critical conditions, see Table 4.1) the participants received negative (error) feedback (beep sounds). For correct responses in these conditions and for trials in the correlation condition no feedback was given. The correct option in each trial was defined based on the average input brightnesses (average of μ_1 and μ_2 , in Table 4.1).

4.2.2 Results

The mean accuracy (averaged across the 16 participants) in the four non-critical (filler) conditions is shown in Figure 4.4, in terms of probability to choose the predominant (A) option. On average, participants chose the predominant option (A) more than 50% of the time in both inconsistent conditions (paired t-tests: $p < 0.001$ in both conditions); however, accuracy in both of these conditions was relatively low. For the consistent-hard and consistent-easy conditions, where the correct option dominated at all moments in a given trial, the subjects achieved very high accuracy. In particular, there was a big discrepancy between inconsistent-easy and consistent-hard, in favour

of the latter condition ($22 \pm 13\%$ *SD*; $t(15) = 6.46$; $p < 0.001$). This large difference in accuracy indicates that consistent information (i.e., evidence not reversing in time) has a positive impact on choice accuracy beyond what would be expected based simply on the integrated evidence advantage for the correct alternative; this advantage is 0.025, 0.1, 0.175, 0.3, in the four filler conditions, I-H, I-E, C-H, C-E, respectively.

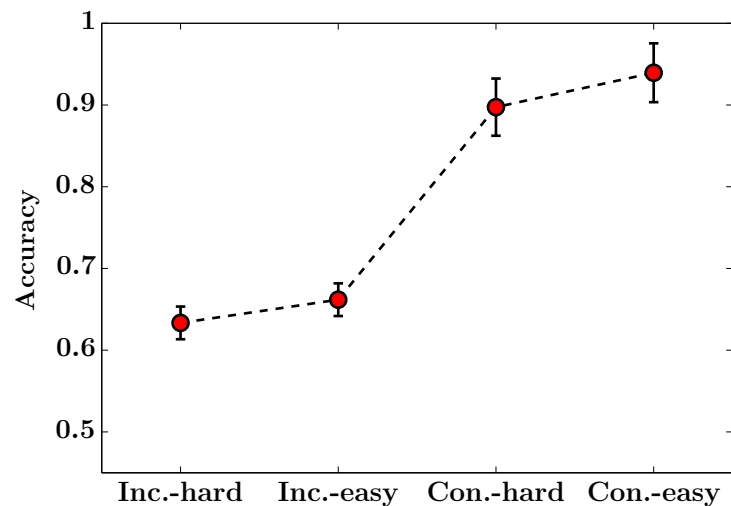


Figure 4.4: Mean accuracy in the non critical conditions (preference for option A). Error bars correspond to 1 SE.

Turning now to the critical trials, I first consider the attraction condition where the average brightness of options A and C are equal. The presence of option B (decoy), which is correlated but inferior to A, according to the attraction effect in multi-attribute choice (Huber et al., 1982), was expected to boost the preference for A (target). However, as Figure 4.5 shows the target option (A) does not benefit from the presence of a similar decoy. Instead, there is no difference between the mean preference for the target (A) and its competitor (C) (paired samples t-test: $t(15) = 1.18$; $p = 0.26$), suggesting that the placement of the decoy (B) has no effect in the choice outcome.

The situation is different in the similarity condition (Figure 4.6). There, all three alternatives have equal average brightness with options A and B being correlated to each other and anti-correlated to the dissimilar option C. If only the absolute brightness of each spot determined its probability to be chosen, one would expect the three alternatives to split their shares. Nevertheless, the preference for the dissimilar option C is much above the $1/3$ chance level ($t(15) = 5.09$; $p < 0.001$) suggesting that the fact that C rises and falls on its own, increases its probability of being perceived as the brightest.

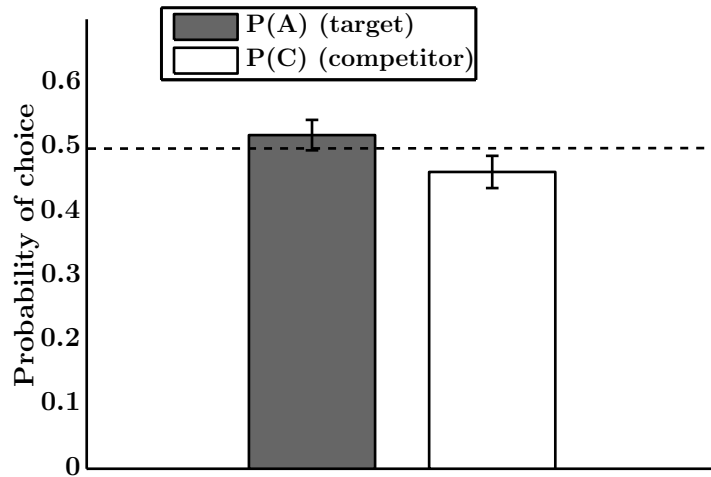


Figure 4.5: Mean choice for the target (A) and the competitor (C) in the attraction condition. Error bars correspond to 1 SE.

The increased preference for C is consistent with the similarity effect in multi-attribute choice (Tversky, 1972). As the individual data (blue symbols in Figure 4.6) indicate, the effect is robust but its magnitude varies a lot across the participants.

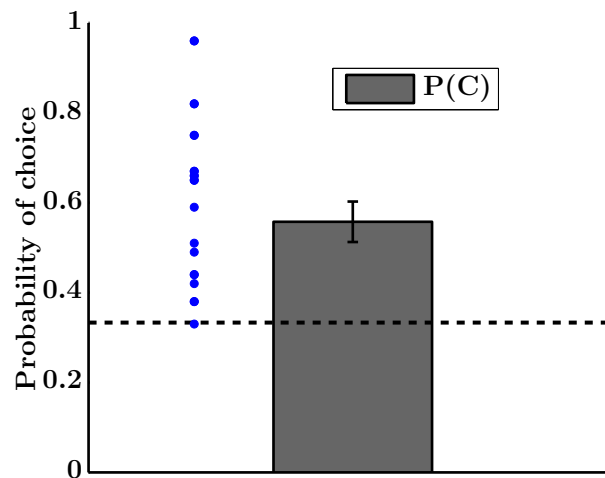


Figure 4.6: Mean choice for the target dissimilar option (C) in the similarity condition. Blue symbols show the $P(C)$ in each of the sixteen participants. Error bars correspond to 1 SE.

In both the attraction and similarity conditions, all three alternatives have non-stationary brightness. This is not the case in the two last conditions, the compromise and anti-

compromise, where the brightness of one option (C) is stationary. In the compromise condition, the stationary option competed for choice with two, anti-correlated to each other, alternatives. According to the multi-attribute choice literature, when confronted with a choice among two extreme options and an all-average, people are biased towards the compromise, all-average alternative (Simonson, 1989). Contrary to this phenomenon, in the current experiment the preference for the stationary option was much below the chance level (Figure 4.7 left, $t(15) = -6.7; p < 0.001$), suggesting that observers avoided it in favour of either of the two, anti-correlated extremes. This finding is surprising for one more reason; the representation of brightness is known to be logarithmically compressed which should set the two non-stationary options in a disadvantage. However, what seems to penalize the stationary option in the current experiment, is the fact that it is always dominated (i.e. always ranking second) by one of the two non-stationary alternatives. This conjecture is supported by the increased (but yet not above chance; $t(15) = -0.47; p = 0.65$) choice preference for the stationary option in the anti-compromise condition ($P(C, anti - compromise) > P(C, compromise), t(15) = 3.6; p < 0.01$; Figure 4.7, right). There, the fact that the two non-stationary alternatives are correlated, makes the stationary option salient (i.e. there are moments where it ranks first), consistent with the increased preference that the dissimilar alternative C received in the similarity condition.

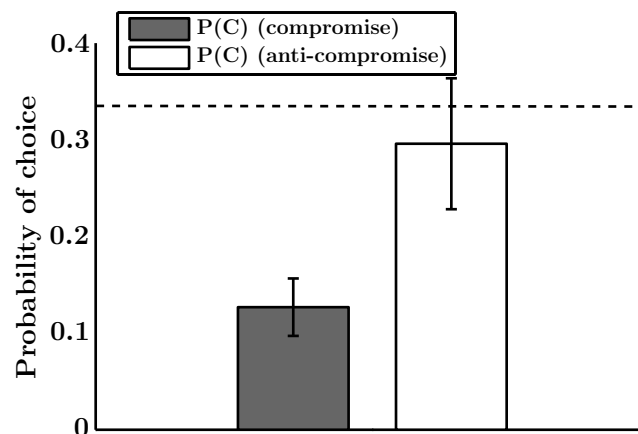


Figure 4.7: Mean choice for the stationary option (C) in the compromise and anti-compromise conditions. Error bars correspond to 1 SE.

To conclude, the attraction and compromise³ effects were not obtained, however, the

³Note that the way the compromise effect was measured in the current experiment is not completely valid and analogous to the compromise situation in multi-attribute choice. This is because the baseline

data in the other conditions provided support for a context-dependent integration mechanism. In particular, in the filler conditions the large performance improvement from the inconsistent to consistent trials, undermines a decision model where the brightness of each option is accumulated independently, in a feed-forward way. Additionally, the increased preference for the dissimilar option in the similarity condition and the boost that the stationary option received in the anti-compromise condition (compared to the compromise condition), indicate that the choice mechanism favours alternatives with peaks in their evidence that make them appear momentarily dominant (i.e. ranking first). In the next section these choice patterns will be further analysed by scrutinizing the individual data and using existing sequential sampling models of multi-alternative choice.

4.2.3 Computational Models of Multi-alternative Perceptual Choice (Computational Study 4)

As reviewed in Chapter 3, a series of process models have been proposed to characterize the integration of evidence in sensory decisions between two options (Ratcliff, 1978; ?, ?; Usher & McClelland, 2001) when the response time is externally controlled. In this section, I extend the *race*, *diffusion* and *LCA* models for choice involving more than two alternatives, in an attempt to determine the mechanisms that give rise to the context effects in *Experimental Study 3* and in particular the similarity effect result.

The first step in extending the models to multi-alternative choice, for the interrogation response protocol, is to assume a separate accumulator for each alternative (see Figure 4.8). Within the LCA or the race model, this extension is straightforward (?, ?; Usher & McClelland, 2001; Usher, Olami, & McClelland, 2002; Brown & Heathcote, 2008).

In the *n-choice* race model, each alternative is assigned to a separate accumulator, whose dynamics are governed by the following simple differential equation:

$$dx_m = I_m + N(0, \sigma). \quad (4.1)$$

Here the quantity dx_m represents the change in activation of accumulator m , I_m represents the external input, and $N(0, \sigma)$ represents processing noise thought to be intrinsic

preference for the stationary option when it is competing with one non-stationary target only, was not known.

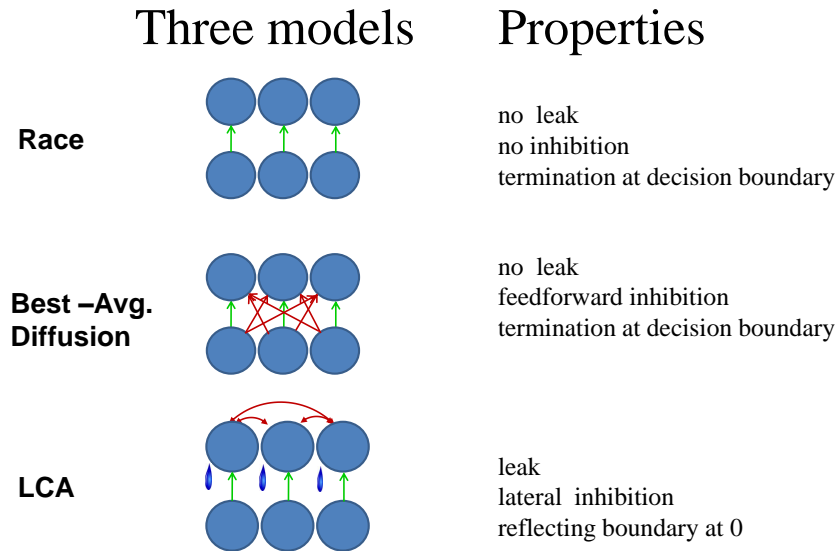


Figure 4.8: *Neural implementation of perceptual choice models for the interrogation paradigm (top row): Pure race model; (middle row) Niwa and Ditterich (2008) diffusion model; (bottom row) LCA model (Usher & McClelland, 2001)). Green arrows correspond to excitation and red to inhibition. Blue “tears” represent leakage in the LCA model.*

to the accumulators. This noise process, included in all the models, is assumed to be Gaussian, with 0 mean and standard deviation σ . In the n -choice LCA, each alternative is also assigned to a separate accumulator. The property of relative evidence integration is achieved through lateral inhibition, and the accumulators are also subject to leakage. The activation level of accumulator m is updated with each simulation time-step according to:

$$dx_m = I_m - kx_m - \beta \sum_{i \neq m}^n x_i + N(0, \sigma), \quad (4.2)$$

$$x_m(t+1) = \text{Max}(0, x_m(t) + dx_m).$$

Here k is the leak, β the inhibition, and the other terms are as before. The *Max* function in the second line of the equation implements a lower (reflecting) bound or floor imposed on the activations. The inclusion of the reflecting boundary was motivated by the fact that neural activity can never go below a minimum level (Usher & McClelland, 2001, p. 14 and Appendix-A; and also Bogacz et al., 2007). For the special case when $k = \beta = 0$, the LCA reduces to a classical race or pure accumulator model as long as all activations are greater than 0. When k and β are both non-zero but equal, the leak

and inhibition are said to be balanced, and the linearized 2-alternative version of this model is equivalent to the classical drift diffusion model (Bogacz et al., 2006).

It is less obvious how to extend the diffusion model to multi-alternative choice. One approach has been suggested by Niwa and Ditterich (2008) (see also Roe et al., 2001 for a similar scheme). For the case of 3 alternatives, 3 accumulators race towards a common decision criterion. The input to each accumulator, however, is the net evidence signal for that accumulator, defined as the momentary evidence for that alternative minus the evidence against it, which is in turn defined as the average of the evidence for the other two alternatives. Accordingly, the differential equation for the m -th accumulator⁴ is:

$$dx_m = I_m - \frac{\sum_{i \neq m} I_i}{2} + N(0, \sigma). \quad (4.3)$$

In both race and diffusion models, the choice was finalized even before the stimulus termination, if a decision criterion (A) was hit by any of the accumulators (i.e. absorbing boundary). In LCA, no decision criterion was assumed and the choice was made at the end of the stimulus presentation, on the basis of the accumulator with the highest activity. The models differ in the way with which they utilize stimulus information, with the race assuming independent accumulation and diffusion and LCA assuming competitive integration. In the following, I will focus on the individual differences obtained in the similarity condition and the boost that the dissimilar/ anti-correlated alternative received. I will attempt to explain this context-weighting pattern by utilizing the existing sequential sampling models described above and I will defer for the *Discussion* of this section a proposal for other mechanisms that could explain the similarity effect.

4.2.3.1 Method

Evidence alternation protocol

The transitions between the two phases of evidence (Table 4.1 and Figure 4.1) are simulated using a Markov process with a transition rate that increases at long intervals. In particular after staying to phase j for n time steps the probability of switching to phase k is $p(n) = 5 \cdot 10^{-4}n$. This transition formula resulted in the distribution of phase durations that is shown in Figure 4.1. Within each phase, for each alternative

⁴In the Niwa and Ditterich (2008) model, the noise variance is input dependent. Here we use a simpler variant of this model, with input-independent noise variance.

m , Gaussian noise with standard deviation σ_{in} (set to be 0.1429 - the value used in the experiment reported below corresponding to variability in evidence on a time scale faster than the characteristic Markov switch time), was added on top of the mean value of the evidence (designated μ_{m1} and μ_{m2} , see Figure 4.1, for an illustration of the input and Table 4.1 for the exact mean values that were used). The evidence values were restricted between 0 and 1, which correspond to minimal and maximal brightness values (in the RGB scale), in the experimental study.

Stimulus duration

Each simulation time-step corresponded to 13.3 ms (or 1 frame on a 75 HZ refresh rate monitor). The stimulus duration was uniformly chosen from the range 375-750 time-steps (or 5-10 seconds). Note that the duration of the last phase is truncated by the end of the trial, making the distribution of last phase durations different from the distribution shown in Figure 4.1.

Accumulator initialization and choice policy

In all three models, accumulators were initialized at 0 at the start of each simulated trial. For race and diffusion, if the bound was reached, the accumulator that reached the bound was chosen as the response on that trial. When the bound was not reached, or in the LCA where there is no bound, the alternative chosen is the one that is most active when the stimulus input is terminated.

Information integration in the three models.

Race. The race model involved three independent accumulators. Each of them (m) was updated according to Equation 4.1 above. Only two free parameters are needed in this model. The SD of the processing noise (σ), and the activation value corresponding to the upper absorbing bound, A . In accordance with the behavioural experiment, the inputs I_m vary in each time frame due to signal noise according to a Gaussian with mean $s\mu_{mi}$ (with μ_{mi} corresponding to the evidence for alternative m during *phase* $- i$ and with $SD = \sigma_{in}$: $I_m \sim N(\mu_{mi}, \sigma_{in})$).

Diffusion. The diffusion model was implemented using the same processing noise and absorbing bound parameters as in the race model. The activation state of each accumulator m was updated according to Equation 4.3.

LCA. In the following simulations, the LCA model was implemented using 4 free parameters including β , k and σ , which stand for the values of inhibition, leak and

processing noise (see Equation 4.2). The inputs I_m are computed as $N(\mu_{mi} + I_0, \sigma_{in})$, where I_0 is an additive input affecting all of the accumulators (set by default at 0.2 except from the simulation of Figure 4.11 where it was varied). This last parameter modulates the degree to which the model is affected by the reflecting boundary at 0; when the value of I_0 is large, activations tend to remain positive, avoiding the reflecting boundary.

4.2.3.2 Results

Temporal correlations in the 3 models

I start with an informal illustration of the models' choice pattern with two example stimuli chosen from the similarity-condition which is shown in Figure 4.2. Input parameters and simulations protocol is as described above. To keep the illustration simple, no processing noise is used ($\sigma = 0$) and the stimulus noise is reduced (from 0.1429 to 0.04). For this illustration only, I also constrain the total presentation time to be such that it gives an equal amount of time to the 2 phases of evidence. In Figure 4.9 (left panels) I show the response of the pure race model the model that simply accumulates incoming information. I consider two contrasting cases; in the first, the stimulus starts with evidence that favours alternatives A and B (the correlated options). In the second, the stimulus starts with evidence that favours C.

One can observe that, towards the end of the observation period, the activations of the accumulators converge, since all receive the same amount of input overall. At earlier integration times, however, one can see intervals where one of the correlated alternatives (A or B) dominates or where the uncorrelated alternative C dominates. If an absorbing bound is reached before the end of the observation period (as assumed by Kiani et al. (2008)), one finds that the likelihood of the dissimilar option to win is approximately 0.5, since the A and B activations (red and blue) are almost identical and therefore will be equally likely to cross the criterion at about the same time and thus split their wins. If extra noise (not correlated with the evidence) is introduced, then the likelihood to choose the dissimilar alternative decreases towards the chance level (0.33).

In the middle and right panels of Figure 4.9, I present the response of the diffusion model and the non-linear inhibition dominant LCA ($\beta = 0.019, k = 0.015$), using the same two example stimulus sequences that were used for the race model in the left

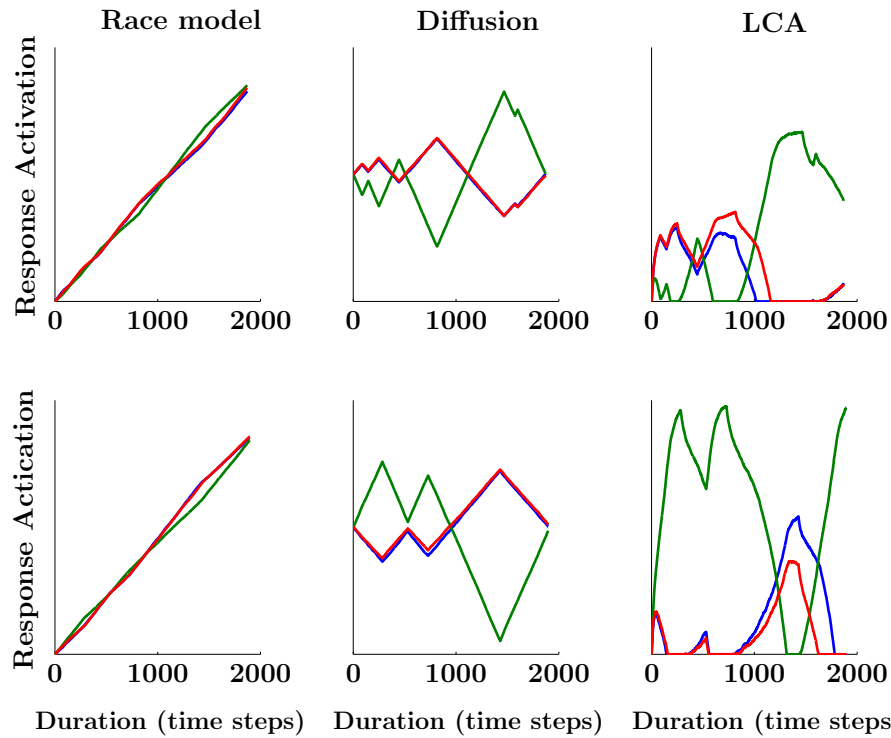


Figure 4.9: Single trial activations for race (left panels), diffusion (middle panels) and LCA (right panels) for initial evidence supporting A/B (top panels) and supporting C (bottom panels). The same two random streams (top/bottom panels) of evidence were used for all the models shown in this figure. Blue and red curves show the activations for the two correlated options (A-B) while the green curve shows the activation for the anti-correlated option (C). The processing noise σ is zero while the stimulus noise is reduced to 0.04.

panels. The activations for the diffusion model correspond to the differences between the activations of the accumulators in the race model. Looking directly at these differences, one can clearly observe moments in which either C or one of A or B dominates the choice. Again, since the total evidence to the 3 accumulators is equal, the 3 diffusion processes end up at the same level. If an absorbing bound is reached, this is likely to favour the alternative associated with the stimulus presented at the beginning of the trial; on average, then, C is likely to be chosen about 50% of the time. As before, with higher noise C may be chosen less than 50% of the time.

The situation is different for the non-linear LCA, when inhibition is larger than leak so that the process is inhibition-dominant. Here in the right panels we observe a clear advantage for the dissimilar option, C. Due to the non-linearity at zero-activation, the low evidence phases of the anti-correlated option C are not suppressed as much as they would be in the linear diffusion model or if their activation were allowed to go below 0. Also, since A and B are low when C is high, while A and B are both high together, the mutual inhibition causes A and B to suppress each other when they are high, while when C is high it receives no such suppression. This asymmetry allows the activation of C to rise more quickly than the activations of A and B, and tends to give C an advantage over A and B. As a result, within a particular range of parameter values, the LCA predicts a tendency to decide in favour of the dissimilar option more than 50% of the time, independent of whether the stimulus starts with A/B (as in the top panel) or with C (as in the bottom panel). As we shall see in more detail below, this phenomenon, an order-independent dissimilarity advantage, is not exhibited by either of the other models under consideration.

In order to demonstrate these differences in the conditions that are in force in the behavioural experiment, I present a second set of simulations. I ran simulations with stimuli of the type illustrated in Figure 4.2, driving the accumulators with inputs in accordance with the visual stimulation protocol used in the behavioural experiment. Note that the trials in the behavioural experiment differ from the single trial illustrations in Figure 4.9, where the total duration of the stimulus was set up to result in equal amount of time for the two phases. As noted in earlier (section 4.2.1), with a stimulus starting with one type of evidence, and then switching at random intervals, and with the trial ending at an independently chosen time, the evidence associated with the first event is more likely to be larger overall (this bias weakens and eventually disappears as the total length of the observation interval increases). For the protocol used in the

experiment, the proportion of trials that have C predominance in trials that start with C is 0.65. However, note that the degree of preponderance is moderate: the ratio between the integrated evidence corresponding to the 2 phases (A/B vs. C) only ranges in the interval 0.9-1.1.

I ran sets of 2000 simulation trials with such stimuli, for each of the 3 models (race, diffusion and LCA) with no processing noise ($\sigma = 0$; Figure 4.10, top panels) and high processing noise ($\sigma = 0.6$; Figure 4.10, bottom panels). For the race and the diffusion model, I examined the impact of an absorbing decision boundary A (Kiani et al., 2008); if the decision criterion is reached before stimulus termination, the evidence is not integrated after that time. I varied the boundary over a wide range to understand its effects. The fraction of C choices is shown as a function of decision boundary for the diffusion model (Figure 4.10, left) and the race model (Figure 4.10 middle). For the LCA, I plot the fraction of C-choices as a function of the ratio between leak (which was fixed at $k = 0.0457$) and inhibition, which varies in the range (0.00043, 0.08571) (Figure 4.10, right). For each model three curves are shown. The green curve corresponds to the trials where the initial evidence favours the dissimilar option C, the blue curve is obtained from the trials in which the early evidence favours the similar options A and B; the red curve is the average of the two other curves.

For low levels of processing noise, we observe that in most models, the total fraction of C-choices is at the 50% range for some range of parameters (red lines, top panels) while with higher processing noise the mean preference for C can go below 50% (red lines, bottom panels). In both the race and the diffusion models, we observe that the fraction of C-choices is above 50% when the evidence starts favouring C (green lines), and below 50% when the evidence starts favouring A and B (blue lines), which is consistent with the fraction of trials that have more A/B or more C evidence, overall. Note that, while true chance level is 33%, a 50% baseline is predicted by any model that decides on the basis of a random sample of momentary evidence, as the correlated alternatives are splitting their wins. On the other hand, since for the stimuli used here, the fraction of trials that have C predominance in trials that start with C is 0.65, a perfect integrator should converge to this choice value. Indeed this value is reached with high decision boundary values in both the race and diffusion models.

An important deviation from the primacy pattern shown by the race and diffusion models occurs in the non-linear LCA, where we see an order-independent advantage for the dissimilar alternative. With low noise, and when the inhibition-leak imbalance is small

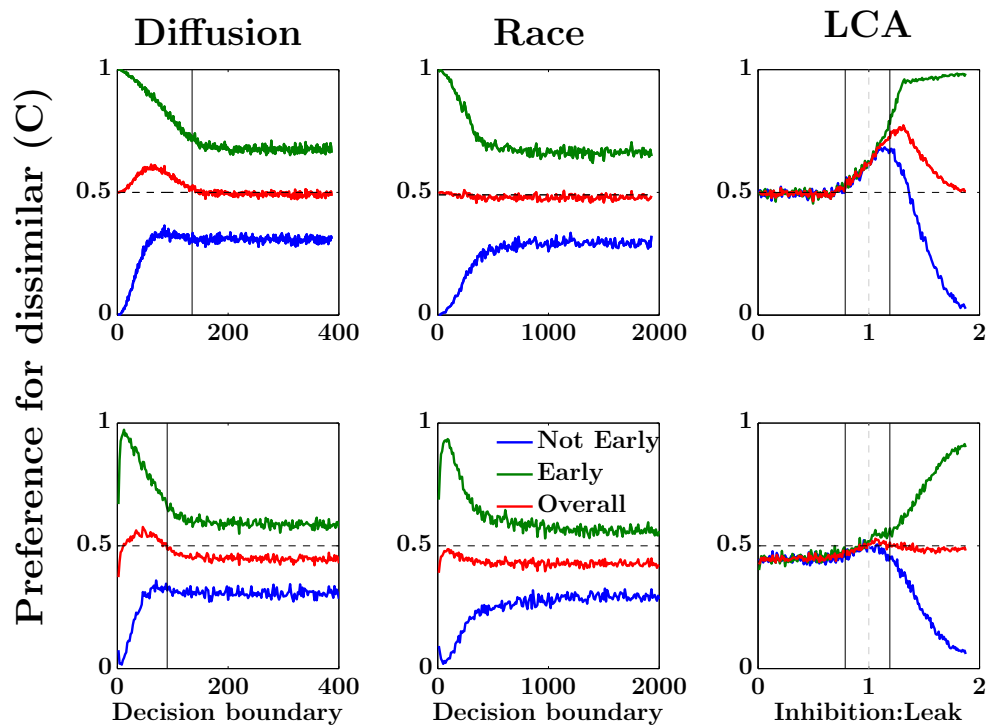


Figure 4.10: Predictions for the bounded diffusion and race models and the LCA model with zero ($\sigma = 0$, top) and high ($\sigma = 0.6$, bottom) levels of processing noise. The green curve shows the choice probability for C in trials when it is favoured at the beginning; blue shows the same when A and B are favoured at the beginning, and the red curve is the average of the other two. For the diffusion model (left panels) and the race model (middle panels) these probabilities are graphed as a function of the decision boundary position. For the LCA (right panels) they are shown as a function of the inhibition/leak ratio.

(Figure 4.10, top right panel, range between vertical black lines), the probability of choosing C is independent of whether the initial evidence favours C (green curve) or A and B (blue curve) and is higher than 50%. This arises from the advantage that the dissimilar option gains from the non-linear dynamic as previously discussed in relation to the single trial trajectories in Figure 4.9.

The area to the right of the second vertical line shows the probability of choosing the uncorrelated alternative, when inhibition becomes more than a little bit stronger than leak. Here the green (strong evidence for C at the beginning) and blue (weak evidence for C at the beginning) are initially both maintained above 50%, but start to progressively diverge as the relative strength of inhibition increases further. Eventually for inhibition much higher than leak, the LCA shows a strong primacy (large difference between green and blue lines), like in the diffusion/race models. For the LCA as well as the other models, the impact of an increase in processing noise is to push the fraction of C-choices down, towards the 33% chance level (Figure 4.10 bottom panels). In summary, we see that with low levels of processing noise, and in a particular range of the ratio between inhibition and leak, the LCA shows an advantage for the uncorrelated alternative over the correlated alternatives, even when the uncorrelated alternatives receive stronger activation at the beginning of the trial.

There is a situation within the diffusion model in which the C choice is made on more than 50% of trials. This occurs in the diffusion model for low decision boundary (left of the vertical black line, at $A = 42$, in the left panels of Figure 4.10). The low decision boundary strongly favours stimuli with larger initial support. It especially favours C, however, because the diffusion associated with the dissimilar option (Figure 4.10 left panels) raises with higher rate (green curve) and thus it is more likely to hit the decision boundary at the beginning of the trial than when the trial begins with greater support for the similar options A/B, which mutually suppress each other and thus have lower slopes. These differences produce the result that, averaging over trials where the evidence supports C first and those where it supports A and B first (red curves in left panels) the probability of choosing C can be greater than 50%. Crucially, though, the probability of choosing C is never above 50% in trials where the evidence associated with A and B is stronger at the beginning, so that the model never exhibits the order-independent advantage for C that we can observe in the LCA model. Thus a distinctive prediction of the non-linear LCA is that $P(C)$ can exceed 50%, both for the trials when C starts with stronger evidence, as well as for those when it starts with

weaker evidence. This prediction takes place for low additional noise (σ) and with inhibition moderately stronger than leak (close to the second vertical line in Figure 4.10, right panel).

To summarize, using the input protocol for 3-alternative choice in which the evidence is non-stationary and temporally modulated, I examined the effect of temporal correlations in the evidence for the various alternatives. I demonstrated that the LCA (with *inhibition* > *leak*) can predict an advantage beyond 50% for the dissimilar option, which is independent of evidence at the stimulus onset and is a result of inhibition dominance, combined with non-linear dynamics. This distinctive pattern the probability of choosing the dissimilar option more than 50%, independent of order of presentation distinguishes the LCA from the race and diffusion variants. Both of these patterns are examined in the individual data of the behavioural experiment (*Experimental Study 3*).

Individual differences in the similarity condition

In Figure 4.11 (upper-left) I report the C-choice pattern for each participant of the behavioural experiment, in a 2D plot, in which the x-axis corresponds to the preference for the dissimilar option, $P(C)$, in the trials where A/B received stronger input at the beginning of the trial, while the y-axis corresponds to $P(C)$, when it received stronger evidence at the beginning of stimulus presentation. Each o-symbol corresponds to the mean choice pattern of a participant and error-bars correspond to 90% confidence intervals. The red diagonal line in Figure 4.11 indicates the range of choice patterns expected if the choice mechanism is not sensitive to the initial evidence. Eight out of sixteen subjects conform to that pattern and for 5 of them in the top-right, $P(C)$ is significantly greater than 50% in both conditions. The other eight participants (in the upper left quarter) showed an increased preference for C when it received stronger input in the beginning. The magenta cross [at point (0.35,0.65)] indicates where the preference of a perfect integrator should lie since, given the limited duration of the trials, the options that receive strong evidence in the beginning will receive more total evidence 65% of the time. I next examine how the 3 choice models can account for these individual differences in the choice of the C alternative.

Model predictions (for the race, diffusion and the LCA, indicated by cyan dots on the figure) were generated by systematically varying the parameters in each model. For the diffusion/race models this involved varying the variance of the Gaussian noise, σ ,

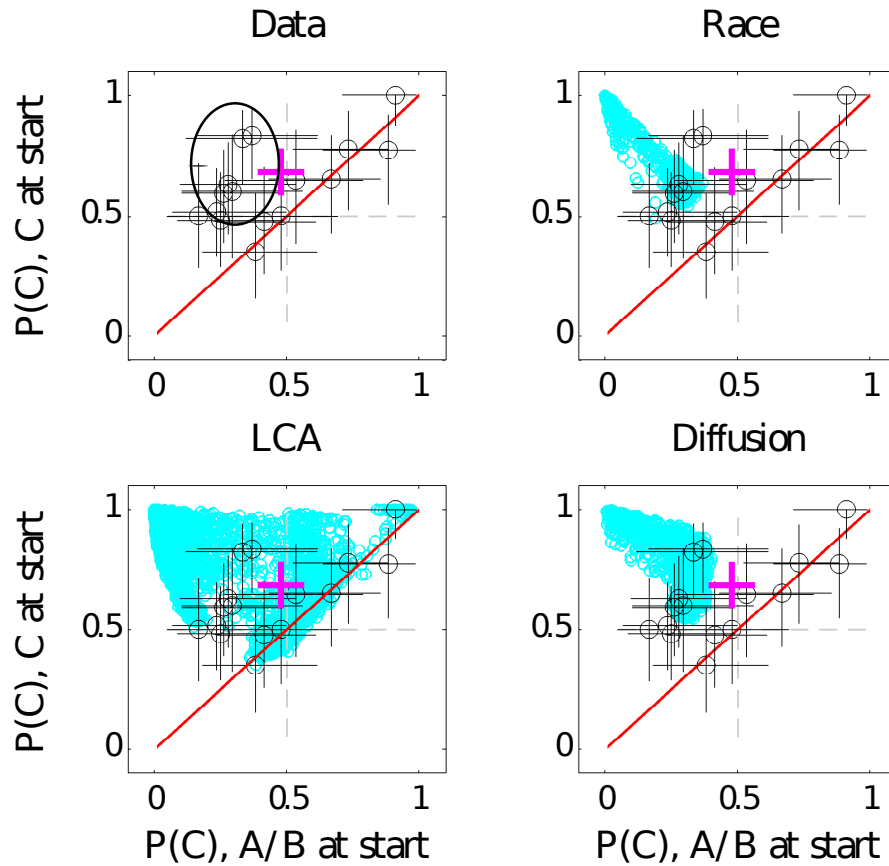


Figure 4.11: Individual choice for the dissimilar alternative, $P(C)$, in the similarity condition of the experiment (upper left) and in the models (others panels). Open-circles show the fraction of C -choices for each participant (error bars are 90% confidence intervals). The pink cross at $(0.35, 0.65)$, indicates where the preference of a perfect integrator should be, based on the input statistics. The model predictions are shown with cyan.

and the evidence value corresponding to the decision criterion on a 2-D grid. For these two models noise was varied in the interval (0.1, 4) with increments of 0.1 while the threshold (A) was varied in the interval (5, 400) with increments of 5 for the diffusion and in the interval (10, 1600) with increment 40 for race. Overall 3200 points were derived for each of these models. For LCA, the predictions in Figure 4.11 were derived using two sets of simulations. In the first set I varied four parameters (a 4-d grid): inhibition (0-0.384, step=0.024, leak (0-0.192, step=0.012, (0-2, step=0.5) and processing noise σ (0-3, step=0.5). In the second set of LCA simulations, I_0 was constant at 0.3, processing noise σ was set to zero and six levels (three low and three high) of leak were used (0.0076, 0.0051, 0.0038, 0.0305, 0.0457, 0.0610). For each leak level, inhibition started equal to leak and increased with a step of 0.00014 for 150 values. This set of parameters was chosen on the basis of the simulations reported above, as well as novel exploratory simulations, as they covered the relevant behaviours in the models. For example, the noise parameter did not exceed 3, so as to maintain accuracy levels in the range obtained in the experiment, and the value of the inhibition parameter in the LCA did not exceed 0.384; stronger inhibition would cause evidence early in the trial to predominate to the extent that it produces decisions that are too fast and of a too low level of accuracy.

Consistent with the simulations reported above, (Figure 4.10, right panels), we find that the non-linear LCA is the only model that is able to predict an order-independent advantage for the dissimilar alternative, as exhibited by the four participants whose choice pattern falls near the diagonal in the upper right quadrant of Figure 4.11 (it must be noted, however, that none of the models accounts for the extreme participant near the (1,1) corner). Data points on the upper-right portion of the main diagonal correspond to choice rates higher than 50% in favour of the dissimilar option, C, both when the evidence starts with a C-phase and when it does not. As previously discussed, this pattern is exhibited by the LCA with low noise, in the area of modest inhibition-dominance (left of the second vertical lines in Figure 4.10 right panels). As previously noted, a perfect integrator would choose C with a rate of 65%, when the trial begins with $C > (A/B)$ and with a rate of 35% when the trial begins with $(A/B) > C$. The ability of the LCA to predict data points on the upper diagonal implies that the model's choice (like the participants in the upper right quadrant) can be insensitive both to primacy and to the small differences in overall evidence. This is the case in the LCA with leak dominance (where early evidence has little weight), and for the LCA with

moderate inhibition dominance.

Additionally, LCA (with higher internal noise; see Figure 4.10, bottom-right panel) is the only model able to account for the C-choices of the other three participants near (0.5, 0.5), who show a preference for the dissimilar option of about 50%, but are still invariant to initial evidence. To account for the individual differences in C-choice probability among these participants, the LCA mainly varies the amount of processing noise in the simulation. This leads to a simple prediction. The five “low noise suspects” (participants with data in the upper right quadrant) should have a higher accuracy in the predominant trials, compared with the 3 “high-noise suspects” (those near the center of the figure, with $P(C)$ close to 50% regardless of the identity of the first stimulus). This prediction is confirmed: $83\% \pm 7\%$ vs. $73\% \pm 6\%$, for low-noise vs. high-noise suspects, respectively. As illustrated in Figure 4.11, the diffusion and the race models cannot account for the C-choices of the 8 participants on the diagonal. As Figure 4.10 suggests, diffusion and race both predict that when C initially receives stronger evidence it will be preferred more than when A/B receive stronger initial evidence. Therefore both the models are restricted to the upper-left quarter of Figure 4.11.

Finally the third group of subjects (in the upper-left quadrant) show a primacy pattern which can be explained qualitatively by all three models, with the race slightly worse for the two data points near ($x=0.4$, $y=0.8$). The LCA can encompass a wider range of patterns, spanning the participants whose performance falls near $y=0.5$ in Figure 4.11. The choice values for these participants are consistent with the LCA model with moderate noise and stronger inhibition dominance (inhibition right of the second black line in Figure 4.10, right-bottom panel).

4.2.4 Discussion

As the pattern of choices of the dissimilar alternative in the similarity condition was subject to considerable individual differences, I examined how well the models can capture this variance. First, I found that some of the participants (ellipse, in Figure 4.11 upper-left) showed a preference for the dissimilar option (C) that is larger for stimuli that start with evidence that favours that option than for stimuli where the initial evidence favours the A and B options. This pattern can be accounted for in all 3 models and it can also be accounted for by a perfect integrator, since the preponderance of evidence tends to favour the option that starts the trial. Second, I found that the pattern

of individual performance is better covered by the LCA. These participants showed little or no sensitivity to order effects (red-diagonal). This pattern is difficult to explain under the race and diffusion models (they can do so only if the overall proportion of C choices is very low, by assuming high noise levels), but is naturally explained by the LCA. Two properties of LCA model work together to produce a preference for the uncorrelated option with little or no primacy bias (Figure 4.10, upper-right, between the 2 vertical lines): moderate inhibition dominance and non-linear dynamics (preventing activation from going below 0). This mechanism is also in a position to explain why the preference for the compromise option was very low; the activation of the stationary alternative has a smaller slope compared to the two extremes and the zero non-linearity turns the high steepness into an advantage (by suppressing the “falls” or disadvantages). The situation improves for the stationary option in the anti-compromise condition where its activation has a higher slope due to the correlation among the two other, non-stationary alternatives. Note that these two mechanisms (inhibition and non-linearity) were also responsible for the unique LCA prediction regarding the interaction between the trial duration and the order effect (*Computational Study 1*).

Within this computational account, no attraction effect is predicted. For the attraction effect case, the presence of the decoy is negligible since the activation state of the inferior option will be soon suppressed to the zero boundary resulting in a binary competition between the two stronger alternatives. This, although complying with the data of the current experiment, stands in contrast to the multi-attribute domain findings where the decoy option has a pivotal role in shifting the choice preference. One reason why this effect was not obtained in the current experiment, might have been that the average brightness of the decoy was significantly lower compared to the stronger, anti-correlated alternatives (i.e. 0.29 compared to 0.6), which might have caused its elimination from the choice process. It is still conceivable that the use of a relatively better, but still inferior, decoy (e.g. with average brightness of 0.5) will bias the choice towards the correlated dominant option. In that case, the multi-alternative LCA will not be able to capture this pattern since the competition between the two similar alternatives (the decoy and the target) will boost the probabilities of choosing the competitor, as in the similarity condition.

Since the crucial mechanisms accounting for the similarity data are the response inhibition and the zero non-linearity, one can equip the, inherently competitive, diffusion

model with a similar non-linear mechanism and examine its predictions. Since the activation of the dissimilar option in the diffusion model (Equation 4.3) rises and falls faster, the zero non-linearity might enable it to encompass the data. Second, an alternative extension of the diffusion model to $n - alternatives$ has been suggested by McMillen and Holmes (2006) and is equivalent to the multi-hypothesis sequential ratio test (MSPRT; Bogacz, 2009). In this model, N accumulators integrate evidence independently and at each moment, the quantity L is computed, where L is the state of the accumulator with the maximum activity minus the activity of the next highest accumulator. When L exceeds a threshold a decision is made. This approach is asymptotically optimal but its neural realization is complex requiring the online computation of the max and the next-max functions. Unlike the diffusion model I focussed on here, this (max-next) diffusion model can account better for the tendency to choose the dissimilar option in the similarity condition. This is due to the fact that the decision criterion is applied to the two maximally activated alternatives, and this penalises alternatives that have correlated evidence (their support goes up together). The LCA can be seen as a natural biological approximation of this optimal choice model, without requiring a complex architecture or a complex computational algorithm. Indeed, inhibition among any number of alternatives can closely approximate the max-next computation. This happens since in LCA, all the choice units compete with each other, but the weak units drop out of the process due to the non-linearity at zero activation, leaving the ones that have the strongest evidence to compete at the end (Bogacz et al., 2007), and thus it does not require a change of weights with set-size.

The possibility that participants may shift their attention among the alternatives, either covertly or overtly via eye movements is not examined in the current experiment. It is conceivable that the attended alternative could exert a stronger influence on the corresponding accumulator (Krajbich & Rangel, 2011), and/or that shifts of attention could reset the integrators. If there were also a tendency to direct attention to the momentarily brightest alternative, these factors could potentially lead to a preference for the uncorrelated alternative; the dissimilar alternative will be more times the brightest (since it peaks alone) compared to the correlated options which will alternate, in their strong phase, in the first rank. Accordingly, the attraction effect is not predicted since the two competitors have the same amount of positive peaks, regardless of the presence or not of the decoy alternative. The low preference for the stationary option in the compromise condition can also be captured by the fact that it is always dominated by

one of the two non-stationary alternatives. This rank-dependent (where the alternative that is first in rank is further boosted) alternative model, captures simultaneously the competition induced by the LCA and the prioritization of the positive peaks caused by the zero non-linearity. Therefore, these two computational accounts are perfectly compatible and can be viewed as capturing the same function at different levels (i.e. neural versus more cognitive).

The emulation of multi-attribute choice problems using perceptual evidence bears limitations. Sensory evidence is processed in a much faster and parallel way compared to attribute values of multi-dimensional items whose representation is discrete. As a result of the continuous and noisy flow of brightness (i.e. noise is updated at each refresh frame which is every 13.3 ms), attention might focus only to the maximum and ignore the other options that rank lower (which can explain why the attraction effect did not occur). On the contrary, in the value integration paradigm where each sample has a symbolic representation, a full ranking of all alternatives could be achievable in each integration time-step. If so, due to the difference processing mode, one would expect different choice patterns to be observed. This possibility will be examined in the next section where context effects will be examined using the fast value integration paradigm.

4.3 Context Effects in Value Integration (Experimental Study 4)

Contextual reversals have been traditionally studied within value-based tasks (Tversky, 1972; Huber et al., 1982; Simonson, 1989). There, participants are typically confronted with a choice among alternatives that are statically presented and they are free to sample the presented information at will. The information sampling process, although crucial for the decision outcome, is covert to the experimenter. This difficulty to measure the information sampling process motivated *Experimental Study 3* where the evidence time course shares the same dynamic properties with the decisional input, when attention fluctuates to different choice aspects. In this first attempt to emulate multi-attribute choice problems in one dimension, only the similarity effect was obtained.

In this section, the fast value integration paradigm replaces the perceptual task of *Ex-*

perimental Study 3 while the logic of presenting non-stationary information is maintained. In a way, using numerical values to represent the input to the decision making process, is closer to preferential choice problems. This is because each information sample has a discrete, symbolic representation similar to the representation that attribute values might have in multi-attribute choice problems. Representing samples of information symbolically allows more complex computations to take place, such as the full ranking of all the alternatives at any given step, as opposed to perceptual tasks where evidence fluctuated at a much faster rate and where the mode of processing might have been more automatic. Because of the limitations in measuring the compromise effect without first obtaining a binary baseline (i.e., preference for the compromise versus the one extreme) and due to the stationary-option aversion that participants showed in the previous study (*Experimental Study 3*), *Experimental Study 4* focused only on the attraction and similarity effects.

4.3.1 Method

Participants

Participants were 20 adults (12 females; aged 21-44; mean age 28.4) recruited from UCL's subject pool that participated for payment.

Stimuli and Experimental Task

At each trial, participants saw 12 triples of numbers presented sequentially and in triangular arrangement (Figure 4.12), around a white fixation cross that stayed on screen throughout the trial. The numbers associated to the different options had different colours (orange for left, magenta for top and green for right) and were surrounded by a white frame. The background of the screen was gray. At the end of the presentation a white question mark appeared at the centre of the screen and the participant had to determine, within 3 seconds, which of the three sequences had the highest mean value. Error feedback (beep sound) was provided. Responses were indicated by the press of the left, right or top arrow on the QWERTY keypad of a standard PC.

Experimental Conditions

Overall there were 4 conditions. In each condition, each option was associated with two distributions labelled here as "blue" and "red", with standard deviations fixed at 7. The distributions are named after colours in order to facilitate the description of

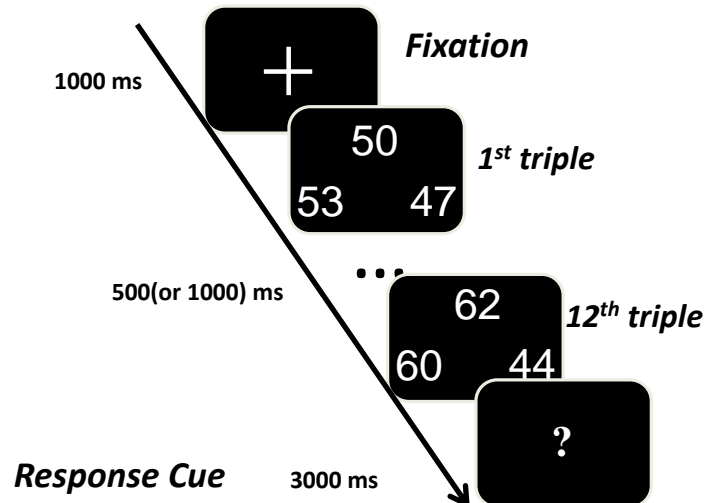


Figure 4.12: The time course of an experimental trial in the fast value integration among three alternatives.

the design. Participants were unaware that each alternative was associated to two distributions and what they perceived was the sequences of numbers without further information on the underlying structure. The positions of the alternatives were always randomized. At each trial, 6 triples were generated from the “blue” distributions while the other 6 triples were obtained from the “red” ones (the switching in the distribution type was covert to the participants and not indicated by any external cue). The triples were reshuffled and thus at each frame there was a 50% probability for all three values to be sampled from either the “blue” or the “red” distributions. Table 4.2 shows the means of the “blue”- “red” distributions for each option in each condition (Figure 4.13). In the attraction condition the values were sampled such that A-values were always greater or equal to B values. In the consistent condition the values were constrained such that $A > B > C$ at each frame, while in the inconsistent condition for the “blue” distribution frames $B > C > A$ and for the “red” distribution frames $A > C > B$.

Procedure Participants were assigned to two different groups (between subjects factor). The “slow” group performed the task at a presentation rate of 1 triple/ second while the “fast” group at a rate of 2 triples/ second. Overall there were 4 conditions, with each condition having 55 trials each. The trials were fully randomized and presented in 11 blocks (20 trials) each. After each block the participant could see her accuracy up to that point. Error-feedback was given only in the dominance conditions and accuracy scores were presented at the end of each block, after the exclusion of the

Table 4.2: The mean values of the sequences in the 4 experimental conditions.

Alternatives	Decoy Conditions				Dominance Conditions			
	Attraction		Similarity		Consistent		Inconsistent	
	Blue	Red	Blue	Red	Blue	Red	Blue	Red
A	70	40	70	40	60	65	75	40
B	65	35	70	40	55	55	55	55
C	40	70	40	70	40	60	40	60

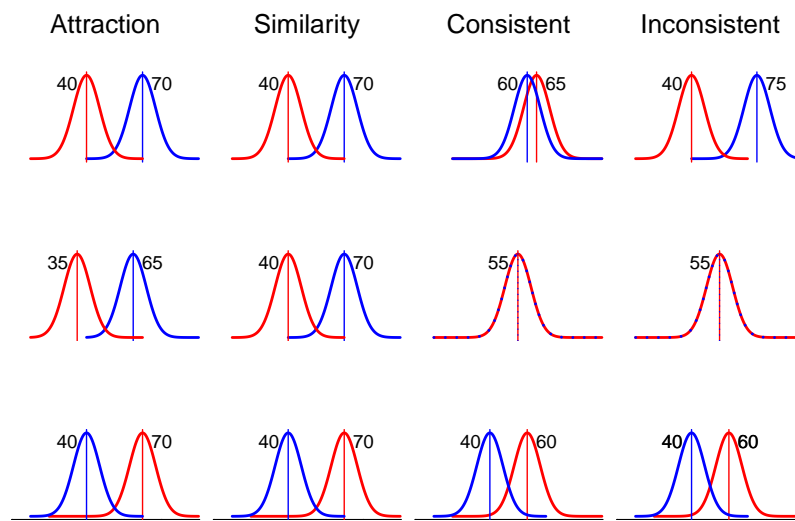


Figure 4.13: Each alternative is associated with two distributions, one red and one blue (colours used for illustration purposes only), and at each frame the values for all three alternatives are sampled from either the red or the blue Gaussian distributions (randomly determined, $p=0.5$). The mean values of the distributions for the 4 experimental conditions are shown for options A (top row), B (middle row) and C (bottom row).

decoy conditions trials.

4.3.2 Results and Discussion

The effect of the presentation rate was examined by comparing the choice patterns of the two groups in the 4 conditions, using independent-samples t-tests. The response mode had no effect in any of the conditions [DV was preference for A and $t(18)=0.52$,

$p=0.61$ for attraction; $t(18)=0.89$, DV was preference for C and $p=0.39$ for similarity; DV was preference for A and $t(18)=1.52$, $p=0.15$ for consistent; DV was preference for A and $t(18)=-0.22$, $p=0.82$ for inconsistent condition]. In the attraction condition the magnitude of the effect was quantified by comparing the relative preference for A to the one for C [i.e. $P(A)/(P(A) + P(C))$ vs. $P(C)/(P(A) + P(C))$] using a paired t-test, while in the similarity condition the effect was quantified by comparing $P(C)$ to the chance level preference (33%), using a one sample t-test. In the two dominance conditions the preference for option A was the accuracy measure and was compared with one sample t-tests against the 33% chance level⁵. Additionally the preference for the two dominated options was used for subsequent analysis of the error pattern [$P(B)$ vs. $P(C)$, paired t-test].

In the dominance conditions, participants successfully chose the highest value alternative ($t(19) = 27.77, p < 0.001$, Figure 4.14, right panels). Furthermore, they showed the predicted choice patterns corresponding to preference reversals, both in the attraction and in the similarity condition. They preferred the alternative A that dominates the decoy (B), at every time step, instead of the anti-correlated alternative (C) ($t(19) = 5.04, p < 0.001$; Fig.2D, left). In the similarity condition, where overall all three alternatives had equal net values, the observers preferred the anti-correlated alternative (C), compared to the two correlated ones (A and B) ($t(19) = 3.40, p < 0.005$; Figure 4.14, second from left).

The attraction and similarity effects provide constraints on the decision mechanism, ruling out a context-independent integration of values, according to which alternatives with the same net mean values should be equally preferred. In order to better understand how the decoys effect might arise in the fast value integration paradigm, I analysed the error pattern in the dominant-inconsistent condition (Figure 4.14 right). When failing to select the best option (A), the observers chose the worst overall option (C) significantly more than the second best (B) ($t(19) = 4.37, p < 0.001$). What makes C stand out is that in half of the frames its values are ranked first (red distribution) while option B is always ranked second.

This pattern indicates that momentary ranks play an important role in the integration mechanism that drives preference formation. However, the higher than chance accu-

⁵Note that A and B in this conditions are indistinguishable. One could measure the similarity effect by comparing the preference for C against the average preference to A and B. This is equivalent to one-sampled t-test of $P(C)$ against the chance level preference.

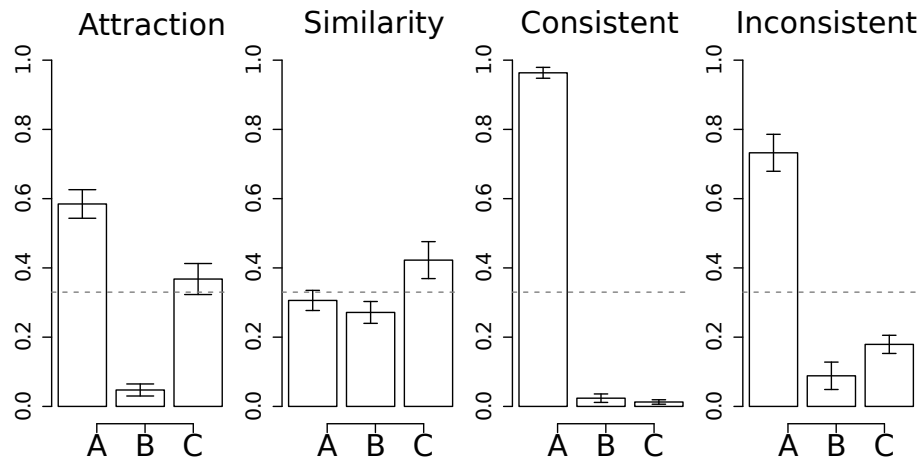


Figure 4.14: Results for the four conditions. From left to right the choice preference in the attraction, similarity, consistent and inconsistent conditions. Error bars correspond to 95% CI.

racy in the dominant-inconsistent condition, in which options A and C are identical in terms of ranks, rules out a model based solely on ordinal comparisons and indicates a decision mechanism that combines absolute magnitudes with momentary ranks (e.g. by weighing the values by their ranks). Accordingly, in the attraction condition, alternative A is preferred because it is overall higher ranked (1st and 2nd) compared to C (1st and 3rd). In the similarity condition, the preferred alternative C is ranked 1st (red) or 3rd (blue). By contrast, A and B alternate in the 1st /2nd (blue) and 2nd /3rd (red) positions and this shared advantage in the blue distributions, weakens their overall value.

The co-occurrence of attraction and similarity effects within participants and within the same study is an important empirical finding. The coexistence of these effects have challenged many theories of choice (with theories being able to predict the one or the other effect, within the same parameter set). In particular, the attraction effect rules out the Elimination by Aspects model that explains the similarity effect (Tversky, 1972). As proven in *Appendix A* in Roe et al. (2001), the context-dependent advantage model developed by Tversky and Simonson (1993) to account for the attraction and compromise effects, cannot explain the similarity effect. The similarity and attraction effects are maintained in DFT (Roe et al., 2001) and LCA (Usher & McClelland, 2004) for multi-attribute choice, however in both models the two effects are in tension (i.e., parameters that boost the similarity effect diminish the attraction effect and vice versa; see Tsetsos et al., 2010 for details). The data of the current experiment do not reveal

any significant negative correlation ($r = -0.303, n = 20, p = 0.195$) between the two effects.

Note that a rank-dependent weighting scheme could explain the data in the perceptual experiment (*Experimental Study 3*) but with one crucial difference. While in the perceptual paradigm only the maximum (brightest) is boosted, here, in the fast value integration task, the attraction effect indicates that participants are able to fully rank all the options and assign weights accordingly. In other words, the symbolic input in the current study, triggers a ranking computation that might be more efficient when the information is consistent but potentially leads to biases when information changes directions (i.e. inconsistent and decoy conditions).

4.4 Summary and General Discussion

In the current chapter, I examined the influence of the context on the weighting of information. By using non-stationary input, which allowed for reversals in the direction of evidence, I emulated multi-attribute choice problems with one dimensional stimuli and questioned the existence of well-known context effects such as the attraction, the similarity and the compromise effects. The potential benefit of looking for contextual effects using dynamic, one dimensional stimulus is twofold. First, it is empirically interesting to obtain preference reversal in a less conventional paradigm which relies on one dimensional stimulus. And second, the psychophysical nature of the experimental tasks, constrains significantly the number of mechanisms and models that could be tested to account for the behavioural patterns. Thus, obtaining the contextual effects within these paradigms can lead to the bottom-up construction of novel accounts for these effects and more generally to simpler theories of multi-attribute choice.

In *Experimental Study 3*, I emulated contextual effects (i.e. attraction, similarity and compromise) in a brightness discrimination task. The findings indicated a very strong similarity effect, no attraction effect and a negative compromise effect. These results were further analysed in *Computational Study 4* using existing models of multi-alternative perceptual choice (i.e. race, diffusion and LCA). From all the models, the non-linear LCA was able to provide the best account. In particular, the increased preference for the dissimilar option in the similarity condition, was accounted for by the interplay of the non-linearity and response competition; due to mutual inhibition at the

decision layer, the preference state of the anti-correlated option raises and falls faster and because preference states are bounded from below at zero, the dissimilar option accumulates overall more evidence. In practice, this mechanism favours positive peaks in the evidence and can be realized at the cognitive level as increased attention towards the momentarily brightest alternative. It is noteworthy that in the previous chapter (*Computational Study 1*, Chapter 3) the same synergy between inhibition and non-linearity generated a unique prediction regarding the interaction between the temporal bias and the stimulus length.

In *Experimental Study 4*, the attraction and similarity effects were obtained using the fast value integration paradigm with three alternative numerical sequences. Note that given the negative compromise effect that was obtained in the brightness discrimination task and the need for extra binary measurements (between the compromise option and the one extreme), I deferred the study of the compromise effect for future studies. The obtained attraction and similarity effects in combination with the results in two novel conditions (i.e. dominance conditions) indicated a simple mechanism in play, whereby the absolute numerical values are combined with the momentary ranks of the alternatives. A model based on this mechanism is further developed in the next chapter (Chapter 5).

Comparing the results of the two experimental studies, one can see clear differences. While only a strong similarity effect was obtained in the perceptual experiment, a weaker similarity effect together with an attraction effect occurred in the numerical integration task. The brightness task results were realized within existing mechanisms provided within LCA, however at a different level of analysis they could be captured within a rank-dependent model, similar to the one underlying the numerical task, where attention prioritizes the processing of the momentarily best option. Therefore evidence and value integration across multiple alternatives, can be realized within the same rank-dependent model and the only difference would be that in the sensory decision task the boost goes only to the momentarily best option whereas in the value integration task it is allocated to all options and is proportional to their momentary rank order. Hence, the differences in the behavioural patterns in the two decision domains can be attributed to the stimuli characteristics, with the faster dynamics of the perceptual task resulting in more automatic processing and a weaker influence of qualitative aspects of the stimulus (i.e. ranks) on the choice outcome. Future experimental work is needed in order to seek convergence between the two domains. One possibility would be to repeat the

brightness discrimination task but update the Gaussian noise that is superimposed on the brightness values much slower. That way the temporal correlation pattern would be more easily detectable by the observers and could encourage more cognitive processing, such as rank-based weighting, which would lead to an attraction effect bias.

Chapter 5

Rank-dependent Leaky Integration

5.1 Overview

The experimental studies so far revealed two main distortions in the information integration mechanism. The first distortion, induces differential weights on the pieces of information based on their temporal order, with this type of differential weighting being already part of existing frameworks of perceptual integration (e.g., Usher & McClelland, 2001). The second distortion was brought to light by the presence of context effects in both evidence (*Experimental Study 3*) and value integration (*Experimental Study 4*). This weighting, as it occurred in sensory integration in *Experimental Study 3* can be attributed to the synergy of low level mechanisms such as the zero non-linearity and the response inhibition (*Computational Study 4*). However, the richer behavioural patterns encountered in value integration in *Experimental Study 4*, indicated a mechanism which weighs absolute values by their momentary ranks in the decision context.

Interestingly, this mechanism can capture the results of the perceptual experiment (*Experimental Study 3*) under the assumption that only the maximum, and not the full rank-order, is used as an auxiliary cue in the predecisional, evidence distortion. This assumption, that in value integration distortions are conferred on all the options while in evidence integration only on the maximum, can be justified by taking into account special aspects of the stimulus in each case. In particular, the perceptual stimulus used in the brightness discrimination task changes rapidly (i.e. every 13.3 ms) and the calculation of the full ranking of all alternatives might be computationally demanding as opposed to the mere detection of the momentarily best option. Accordingly

the symbolic representation of the information in the value integration task and the slower presentation rate (i.e. changes occurred the fastest i.e. every 500 ms) renders the consideration of the full rank order feasible. Cross-validation of this account can be achieved by tweaking the stimuli in the two domains and examining whether the results converge (i.e. slowing down the rate of changes in the perceptual experiment or speeding up the presentation in the value integration task).

In this chapter, I first develop the rank-dependent account for information integration and present its mathematical implementation (*Computational Study 5*). Note that one non-crucial aspect of the model is that the integration at the response layer is leaky and not perfect. This assumption is motivated by the order effects found in *Experimental Study 2* and is accommodated to the model for completeness. Next, I demonstrate how this model accounts for the main behavioural patterns in *Experimental Study 3* and I fit it on the behavioural results of *Experimental Study 4*, discussing also the rank-dependency as a viable aspect of multi-attribute choice. Finally, I test the sensitivity of the model to second-order aspects of the information i.e. the variance, revealing an intriguing prediction of the model which is fully examined in the next chapter (Chapter 6): risk-seeking in the domain of gains (*Computational Study 6*).

5.2 Rank-dependent Leaky Integration Model (Computational Study 5)

5.2.1 Model Implementation

The intuition that underlies the proposed model is that information is weighed by its salience. Now the salience of a given sample is determined by the momentary ranking of the sample in the current decision context. The preference state P in favour of alternative i at moment t is given by the following equation:

$$P_i(t) = \lambda \cdot P(t-1) + [w(\text{rank}_i(t)) \cdot I_i(t)] + N(0, \sigma). \quad (5.1)$$

In the above equation, $I_i(t)$ is the magnitude of the sample in favour of alternative i at time t , λ corresponds to integration leakage (motivated by the order effects found in Chapter 3), and σ is the standard deviation of Gaussian, processing noise. The core

weighting mechanism is implemented in w which corresponds to a decreasing function, that assigns larger weights to high ranks (i.e. 1st) and smaller weights to low ranks (i.e. last). The momentary rank of a sample at time t for alternative i is denoted by $rank_i(t)$ and is always a positive integer value.

The most critical element of this model is the weighting function w which imposes a type of relativity and competition in the processing, similar to the lateral inhibition at the response layer posited by the perceptual LCA model (Usher & McClelland, 2001). A way to understand the practical role of this function is to assume that visual attention fluctuates from option to option and as a result some samples are lost and do not get to be integrated. A higher weight in the highly ranked alternatives means that the probability that their samples are encoded and processed is larger. The exact form of the weighting function is not assumed to be fixed but to rather change depending on the type of the task at hand. In the next subsections I will review what form this function needs to have in order to account for the data in *Experimental Study 3* and *Experimental Study 4*.

5.2.2 Rank-dependency in Perceptual Decisions

In the brightness discrimination study presented in Chapter 4 (*Experimental Study 3*) the flow of the sensory evidence was temporally manipulated so as to establish temporal correlations among the alternatives, analogous to shifts of attention to different choice aspects in trade-off, multi-attribute decisions. When two of the three alternatives were correlated to each other (i.e., raising and falling together) and anti-correlated to the third one, participants preferred above 50% the third alternative. On the contrary when the one of the two correlated options was always inferior to the other, respondents turned out to be indifferent between the inferior correlated and the anti-correlated one. Finally, between two anti-correlated alternatives and one with stable brightness in the middle of the range of brightness of the other two, people systematically avoided the stationary, mediocre one even if its brightness was overall equal to the brightness of the anti-correlated options.

In order to understand what weighting function can give rise to the above patterns, I simplify the experimental conditions by assuming that there are two alternative phases of equal overall length ¹. The mean brightness of each option in each phase is given

¹In the actual experiment the length of the phases was sampled from a distribution while the trial

symbolically in table 5.1, with H denoting high brightness, L low brightness and d a small constant that is smaller than L .

Table 5.1: Conditions in *Experimental Study 3*.

Alternatives	Conditions					
	Similarity		Attraction		Compromise	
	Ph_1	Ph_2	Ph_1	Ph_2	Ph_1	Ph_2
A	H	L	H	L	H	L
B	H	L	$H-d$	$L-d$	$(H+L)/2$	$(H+L)/2$
C	L	H	L	H	L	H

The above table is next transformed such that the absolute brightness values are converted into ranks (table 5.2). In cases where the brightness of two options is equal in a given phase, the noise (which was present in the actual experiment) will randomly push the one option above the other. For example, in the similarity condition in $phase_1$, A and B have both high brightness values and thus will randomly alternate in the first and second ranks. Now I assume that the first in rank option will receive a weight of a , the second b and the third c , with $a \geq b \geq c$. The overall goodness of an option will be determined by adding its ranked-weighted brightness in each phase, assuming of course that the two phases occur with equal likelihood (i.e. 50% each). In cases where there is a tie, as for example between A and B in $phase_1$ of the similarity condition, the ranks of an option alternate with equal probability (e.g. option A in $phase_1$ of the similarity condition will rank first 25% of the time and second the other 25%, summing up to 50% which is the probability of $phase_1$ to occur in a given trial).

In the similarity condition the increased preference for C that was found in the experiment indicates that the integrated, rank-weighted brightness V of option C is higher than those of A and B. Given that A and B are identical the following inequality should hold: $V_C > V_A = V_B$. Expanding the inequality gives: $0.5Ha + 0.5Lc > 0.25Ha + 0.25Hb + 0.25Lb + 0.25Lc$ or :

$$Ha + Lc > Hb + Lb \quad (5.2)$$

duration was independent to the phase switching process. It was therefore possible that the overall phase durations were unequal in a given trial.

Table 5.2: Rank ordering of the options of table 5.1.

<i>Alternatives</i>	<i>Conditions</i>					
	<i>Similarity</i>		<i>Attraction</i>		<i>Compromise</i>	
	<i>Ph₁</i>	<i>Ph₂</i>	<i>Ph₁</i>	<i>Ph₂</i>	<i>Ph₁</i>	<i>Ph₂</i>
A	1,2	2,3	1	2	1	3
B	1,2	2,3	2	3	2	2
C	3	1	3	1	3	1

The indifference between option A and C in the attraction condition is translated into $V_C = V_A$ which leads to $0.5Ha + 0.5Lc = 0.5Ha + 0.5Lb$ and finally to :

$$c = b. \quad (5.3)$$

In other words the weighting that the second and third ranked options receive is the same. In order to see the way this constraint affects the similarity condition, I plug equation 5.3 back to equation 5.2:

$$Ha > Hb \Leftrightarrow a > b. \quad (5.4)$$

Now, given equations 5.3 and 5.4, I turn into the compromise condition where either of the extremes (A, C) are systematically preferred over the stationary B. This leads to $V_A > V_B$ or $0.5Ha + 0.5Lc > 0.5b(H + L) \Leftrightarrow 0.5Ha + 0.5Lb > 0.5b(H + L)$ which turns into:

$$Ha > Hb \Leftrightarrow a > b, \quad (5.5)$$

which already holds from equation 5.4.

Overall the patterns in all three conditions are qualitatively predicted if the weights are constrained such that $a > b = c$. In practice this means that the momentarily brightest option attracts the attention while the ones that follow in the rank-order are not further amplified. As already mentioned, this is somehow natural given the nature of the brightness stimulus which changed every 13.3 ms, rendering as a result the computation of the full ranking at each frame demanding. Whether the form of the weighting

function changes by slowing down the updating rate of the stimulus is an open empirical question. In the next subsection I perform a similar analysis on the value integration task by quantitatively fitting the average data of *Experimental Study 4*.

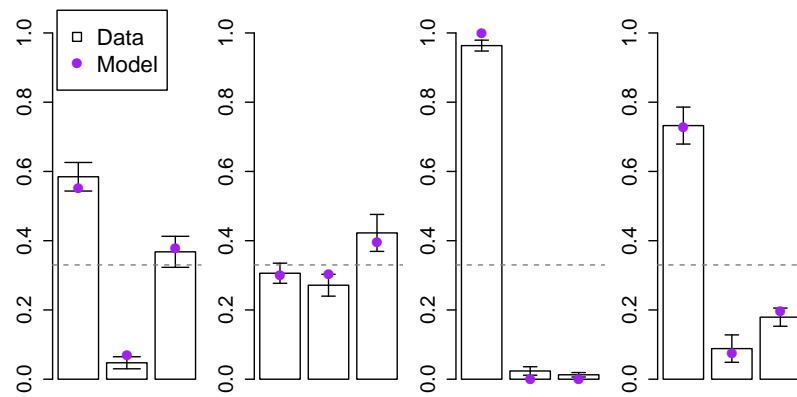
5.2.3 Rank-dependency in Value Integration

In *Experimental Study 4* in Chapter 4, I examined the presence of the attraction and similarity effects in value integration. As opposed to the perceptual experiment, both the effects were obtained, indicating differences in the underlying integration mechanism in the two domains. In order to examine what form the weighting function of equation 5.1 needs to have in order to account for both the effects simultaneously, I quantitatively fitted the data averaged across all participants. In table 5.3 the values associated with each option in each distribution (see table 4.2), are converted into ranks. In the stochastic simulations the input to the model was constrained such that these ranks were always maintained (e.g. in the attraction condition A values were forced to be always higher than B values), identical to the input the experimental participants observed.

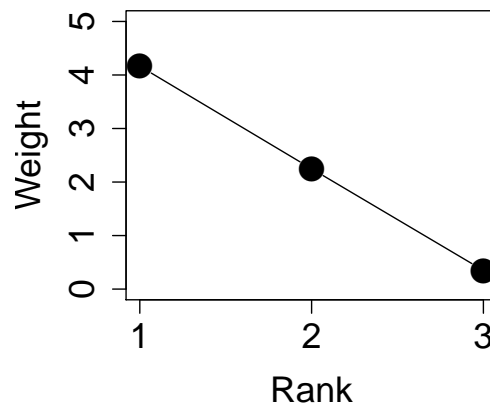
Table 5.3: The ranking of the sequences in each distribution in the the 4 experimental conditions of *Experimental Study 4*.

<i>Alternatives</i>	Decoy Conditions				Dominance Conditions			
	Attraction		Similarity		Consistent		Inconsistent	
	Blue	Red	Blue	Red	Blue	Red	Blue	Red
<i>A</i>	1	2	1,2	2,3	1	1	1	3
<i>B</i>	2	3	1,2	2,3	2	3	2	2
<i>C</i>	3	1	3	1	3	2	3	1

The model fits are given in Figure 5.1(a) and the optimized parameters in table 5.4. As also shown in Figure 5.1(b), the weighting function linearly decreases with the rank-order. This means that all three options in a given frame are taken into account at some extent, and the likelihood of each sample to be considered is linearly dependent on the local ranking. In practice though, the nearly zero weight of the third in rank sample (i.e. $c = 0.35$) implies that in most frames only the two best samples get processed.



(a)



(b)

Figure 5.1: *a*: Data fits of Experimental Study 4 for the attraction, similarity, consistent and inconsistent conditions (from left to right). The three bars in each condition correspond to the mean preference for alternative A, B and C respectively. Error bars correspond to 95% CI; *b*: The optimized rank-weighting function.

Table 5.4: Optimized parameters of the rank-dependent model for *Experimental Study 4*.

Parameters				
a	b	c	λ	σ
4.16	2.25	0.35	0.97	83.27

The fact that the weighting function is more continuous than the one of the perceptual experiment, accounts for the increased preference for option A in the attraction condition. This follows naturally from the fact that alternative A has higher ranks (i.e. 1st and 2nd) compared to its competitor C (i.e. 1st and 3rd) and that the 2nd in rank sample is more amplified compared to the 3rd in rank. The similarity effect is captured by the shared advantages/ disadvantages the similar options have. Applying the optimized weights on the actual values of the alternatives (4.2) gives the integrated value over the two distributions. For alternative A and B this will be $V_A = V_B = 0.25a \cdot 70 + 0.25b \cdot 70 + 0.25b \cdot 40 + 0.25c \cdot 40 = 138.175$. For alternative C the integrated value will be $V_C = 0.5a \cdot 70 + 0.5c \cdot 40 = 152.6$ and thus $V_A = V_B < V_C$.

Finally, the choice pattern in the dominance conditions shows signature of rank-dependence. First the higher accuracy in the consistent condition compared to the inconsistent one, although the overall value of option A was better in the first case (i.e. 125 against 115), can be attributed to the fact that in the inconsistent condition, in half of the frames alternative A was ranked last (see table 5.3) and consequently it was less often preferred. Second, in the inconsistent condition, the fact that when the respondents failed to select the best option (A) chose the worst overall (C) compared to the second best (B) was due to rank-dependent integration; option C in half of the frames was ranked first and thus further boosted as opposed to B which was always ranked second and never stood out.

5.2.4 Discussion

The rank-dependent weighting account posits that the decision input is distorted by its salience in the local decision context. With rapidly changing perceptual input (i.e. brightness) this distortion takes the form of overweighting the momentarily maximum option, while with slower and symbolic input (i.e. numerical sequences) it takes the

form of more continuous differential weighting, across the whole span of ranks in the immediate context. The simultaneous occurrence of the attraction and similarity effects in *Experimental Study 4* (through which the rank-dependent integration model was inspired), shows that the technique of controlling the sampling process is a good proxy to the underlying process of decisions in richer domains. A model of multi-attribute choice follows immediately from the combination of rank-dependent value integration and Tversky's proposal (Tversky, 1972), which was extended in Decision Field Theory (Roe et al., 2001), that people process multi-dimensional options by sequentially switching their focus from one choice aspect to another. An account whereby the decision maker's attention emphasizes the processing of attribute values that are highly ranked in a given dimension, provides a sufficient explanation of how people integrate values across attributes and why their preference is subject to reversal in multi-attribute problems.

One caveat of a rank-dependent multi-attribute account is that it cannot produce the attraction, similarity and compromise effects simultaneously. In particular, under the weighting function obtained for the value integration experiment (i.e. $a > b > c$) the mediocre compromise option will be avoided. For this effect to be obtained one needs to assume that the attribute values of the two best options in a given dimension are considered equally often (i.e. $a = b$) and that the last in rank option is ignored or has a very small weight (i.e. $a = b > c$, this weighting function accounts for the attraction effect but predicts a negative similarity). This mode of processing can be triggered other from different experimental material (i.e. different types of consumer products) or merely by the distribution of the choice alternatives in the choice space. In the latter case the presence of an all-average option, which is not outstanding in any dimension but whose overall value seems as good as those of the extreme options, could encourage a more deliberative and cautious strategy where the two best options in a given dimension are equally weighted. If the shape of the weighting function was shown to be flexible and dependent on the structure of the choice problem, then the simultaneous prediction of all three effects within the rank-dependent account would be feasible.

5.3 Predictions for Binary Choice: Relativity of Value and Risk-Attitudes (Computational Study 6)

5.3.1 Relativity of Value

The rank-dependent analysis presented so far, concerned choice among three alternatives and revealed salience-based distortions in the integration of values. It is questionable whether this type of differential weighting is applicable also to binary choice problems which bear less computational complexity and where the utilization of the ranks might be less imperative. If indeed the local ranking distorts the absolute values in binary problems then the evaluation of an option will be always relative and context-dependent. For example the computed value of an alternative will be suppressed when it is paired with a better option relative to the case when it is compared against an inferior option (e.g., Louie, Grattan, & Glimcher, 2011).

Although the purpose of *Experimental Study 2* was the examination of order effects in value integration, two of the conditions (i.e. unbalanced) involved selection between two sequences, with the one sequence being always better. Crucially, in one of the two unbalanced conditions (i.e. Condition 2) the overall maximum value sample was placed in the sequence with the lowest mean. This was done in order to check whether respondents followed a heuristic rule whereby they chose according to the globally maximum sample only. Despite their original scope, the data in these two unbalanced conditions can be reviewed with a rank-dependent account in mind.

Comparing the accuracy scores in the two unbalanced conditions reveals a marginal difference in favour of Condition 1 (Figure 5.2, $F(1, 15) = 4.96, p = 0.042$). Note that Condition 2 trials are identical in difficulty to those in Condition 1². This difference in the accuracy scores in the two conditions is consistent with the rank-dependent weighting model. Accordingly, attention is driven towards the momentarily maximum values and in Condition 2 there will be certainly at least one pair in which the low sequence has the maximum sample. This signature of rank-dependency even in binary choices provides a potential explanation of relativity of value phenomena (Stewart et al., 2003; Vlaev, Seymour, Dolan, & Chater, 2009; Kurniawan et al., 2010).

²The same numerical values of Condition 1 were used in Condition 2 but modified in two pairs; in one pair I added a constant to the low sequence such that this value is globally maximum. From a second sample of the low sequence (randomly chosen) I subtracted the same constant. Therefore the integrated differences were equal in all trials of the two conditions.

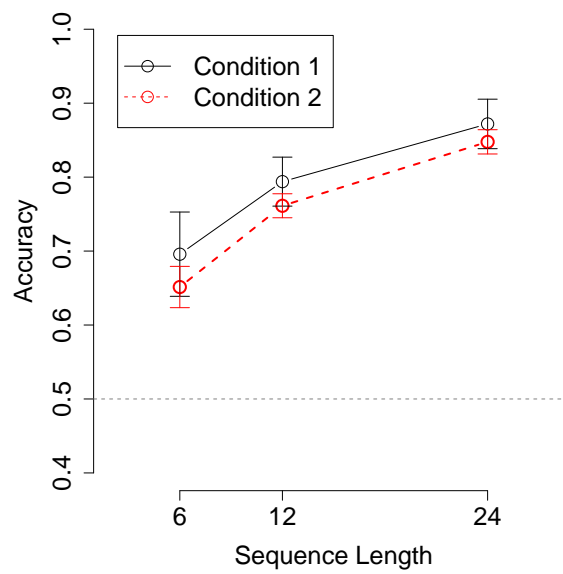


Figure 5.2: Performance in the two conditions of the unbalanced trials. Condition 2 differs from Condition 1 in that the maximum number appears always in the low-average sequence. Both conditions are matched in difficulty. Error bars correspond to 95% CI.

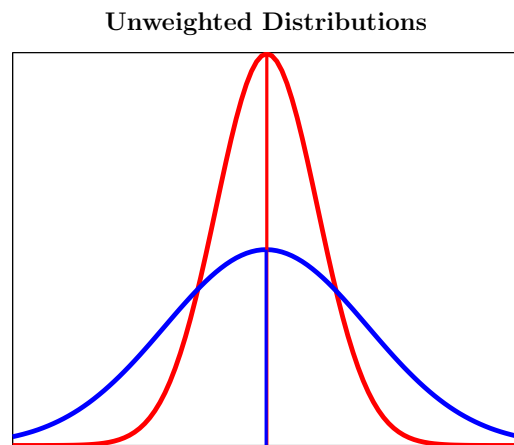
5.3.2 Risk Attitudes

If rank-dependency underlies binary choices, in cases where the alternatives have equal variances and different means, this strategy might facilitate the detection of the best sequence as, statistically, the high-mean option dominates the weak one, in most of the pairs. Therefore overweighting the maximum value in a given pair would amplify the accumulated differences between the two sequences, increasing the probability to detect the best one. On the other hand though, when the two alternatives have equal means but different variances, the focus of attention will be directed towards the extreme large values at the right tail of the high variance distribution. Consequently, the decision maker would be blind to the extreme low values generated from the left tail of the high variance distribution, developing a propensity to choose the riskier option that is associated to the broader Gaussian.

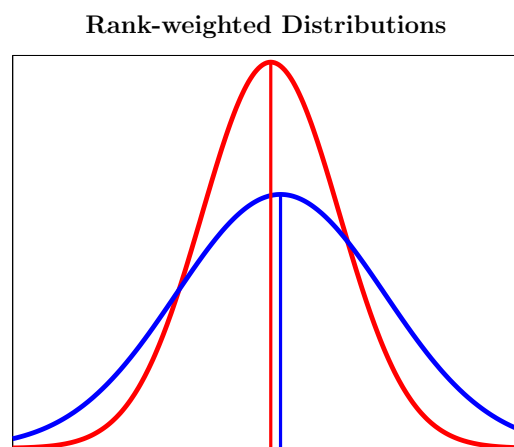
In order to demonstrate this one-sided pro-risk bias, I simulated two Gaussian sequences with mean at 50 and different standard deviations; the first distribution had a standard deviation of 10 while the second of 20. Ten thousand numbers were generated from each distribution and were subsequently rank-weighted as if they were presented paired together in a long sequence and assuming that the momentarily maximum number receives a higher weight (i.e., $a = 2$ for the maximum and $b = 1$ for the second in rank). As Figure 5.3 shows, indeed the transformation of the values according to the rank-dependent model shifts the mean of the broad distribution higher than the mean of the narrow one. This prediction of risk-seeking behaviour is quite surprising as it collides with findings from the mainstream line of research in risky choice, decisions by descriptions. There, people are provided with an explicit description of the probabilistic pay-offs of monetary gambles and typically exhibit risk-aversion in the domain of gains (Kahneman & Tversky, 1979).

5.3.3 Discussion

Despite that the rank-dependent aspect of value integration was launched for choice among three sequences, it is possible that it underlies also binary decisions. This conjecture is supported by existing data (see Figure 5.2) and provides an explanation for value relativity phenomena in other domains. Relying on the local rank can be auxiliary in contexts where the agent is unfamiliar with the absolute values (i.e., does not



(a)



(b)

Figure 5.3: *a: One narrow (red) and one broad distribution (blue) of equal means; b: The transformation of the distributions in a) according to the rank-dependent model.*

have a grasp of how good a value is on its own and considers also its relative ranking) or when the two options have different means but equal variance. In these cases the rank-weighting will accentuate the accumulated differences in favour of the best sequence, improving that way choice quality. One further prediction of this mechanism, however, is that decisions will not be sensitive to the strength of the values only but also to their variability.

When the sequences have equal means but different variances a salience-based value integration will be sensitive to the variance of the sequences (Figure 5.3), favouring riskier options in the domain of gains, in direct contrast to the prediction of risk aversion from expected utility theory and prospect theory (Kahneman & Tversky, 1979). However, this prediction is in the same direction with a recent finding in experience-based decisions (Ludvig & Spetch, 2011), where people learn about probabilistic outcomes through active sampling, and also with a qualitatively theory applied to scenario based decisions, the reason-based decision framework (E. Shafir, 1993; E. Shafir, Simonson, & Tversky, 1993). According to the latter framework, the decision mechanism is flexible and subject to the task framing with advantages looming larger in selection decisions and disadvantages looming larger in rejection decisions. Both the rank-dependent integration and the reason based framework weigh information by its salience and the relationship between these two accounts is an interesting open question.

5.4 Summary and General Discussion

In the present chapter I proposed an implementation of an integration model which weighs information by its salience. Based on the experimental studies presented so far, two factors seem to affect the prominence of information. First, the temporal order with which information is presented makes a piece of information less or more salient (i.e. temporal biases). Second, the local ranking of the options in the immediate decision context affects their perceived magnitude. The first factor was captured by assuming that the integration is subject to decay, consistent with the order effects in Chapter 3. The second factor was incorporated in the form of a weighting function which at each time frame boosts differentially samples of information, depending on their momentary ranking. Although the exact form of this function was not a priori specified I naturally assumed that high ranked items are more strongly weighted. This can be understood by

assuming that not all the pieces of information get processed but the most noticeable ones have higher probabilities to be encoded. This is mathematically equivalent with applying a differential multiplicative boost on the most prominent (i.e. highly ranked) items in a given context.

As shown in *Computational Study 5* the exact form of the weighting function differs across experiments. A steep, step-function accounts for the results in the perceptual experiment of Chapter 4 (*Experimental Study 3*) while a more continuous, decreasing function explains the data in the equivalent value integration study (*Experimental Study 4*). Whether the shape of the weighting function depends on the type of the stimuli (i.e., perceptual vs. symbolic) or other aspects of the experiment (i.e., how fast the frames are updated) is an open empirical question. Rank-based integration with a continuously decreasing weighting function combined with a dimension-wise processing approach (Tversky, 1972; Roe et al., 2001; Usher & McClelland, 2004), accounts for preference reversal in multi-attribute choice. In particular the attraction and the similarity effects are readily predicted, however, this function does not capture the compromise effect. For the latter a steeper weighting function is required and it is questionable whether different configurations of the alternatives in the choice space trigger different modes of processing and different forms of weighting.

While relying on auxiliary cues of the input, such as the ranks, is justifiable in more intensive decisions among three options, it is not immediately clear whether this is the case in the computationally simpler case of binary decisions. Examining the choice pattern in the value integration experiment on order effects (*Experimental Study 2*) revealed a clear signature of rank-dependent integration; in trials where the globally maximum value was placed in the low sequence alternative, performance deteriorated even if difficulty was controlled for. This salience-driven value integration in binary choices can explain relativity of value phenomena (e.g., the computed value of a target sequence will be lower when it is paired with a superior option relative to when it is evaluated against a worse option). The utilization of ranks can be justifiable since in cases where two alternatives have equal variances and different means, this strategy facilitates the detection of the best sequence. However, a direct consequence of this strategy is the sensitivity of the choice process to the variances of the sequences. As *Computational Study 6* showed, when two options have equal means but different variances the rank-dependent integration model predicts risk-seeking, i.e. higher preference for the broad distribution. This prediction is in sharp contrast with the results

in the standard paradigm of risky choice, decisions by description (Kahneman & Tversky, 1979) but agrees with recent findings in experience-based decisions (Ludvig & Spetch, 2011) and a qualitative theory applied on scenario-based decisions (E. Shafir et al., 1993). In the next chapter I will attempt to further pursue this intriguing prediction of the rank-dependent model, by examining risk-preferences in the fast value integration task.

Chapter 6

Value Integration and Risk-attitudes

6.1 Overview

As the experimental and computational studies in this thesis has shown so far, the integration mechanism is sensitive not only to the mean strength of the information but also to its variance. This sensitivity was incorporated in a simple accumulation model where the absolute magnitude of an alternative is weighted by its local rank in the immediate context. Relying on the ranks is an idea that has been proposed in other theories of choice (Parducci, 1965; Stewart et al., 2006) and in combination to the magnitude of the information might have an ameliorative effect when the ranks are correlated with overall goodness. For example when evaluating candidates that have graduated from different universities, the examiner might have little idea about what the absolute grades reflect and might use them in combination to the relative ranking of the candidate in her department, in order to make an informed decision. On the other hand when the options have different variances relying on the ranks will favour the riskier option that has the highest variance (*Computational Study 6*). For instance when evaluating two candidates on the basis of their grades in several courses, a candidate with exceptionally high but also extremely low marks will be favoured compared to an all-average candidate.

This risk-seeking prediction stemming from the rank-dependent integration model stands in sharp contrast with findings in decisions by description. There, people are given explicit information about monetary gambles and are typically risk-averse in the domain of gains and risk-seeking in the domain of losses. In this chapter, I will ex-

perimentally probe risk-attitudes in the value integration task. In *Experimental Study 5* I start by examining the risk preference with positive sequences and confirm the rank-dependent prediction, finding strong risk-seeking. In *Experimental Study 6* I turn into negative sequences, corresponding to losses, and surprisingly obtain again a risk-seeking pattern. However, when examining the preference for mixed sequences (gains and losses) participants exhibit a risk-aversion behaviour (*Experimental Study 7*). Risk-aversion with mixed gambles has been typically attributed to loss-aversion, the asymmetry between gains and losses. In order to test this hypothesis I present to participants mixed sequences; crucially in half of them a negative number in one sequence is always paired with a positive number in the other while in the rest of the trials the pairs are homogeneous (i.e., negative-negative/ positive-positive). The findings suggest that risk-aversion comes into play only in the non-homogeneous sequences where a loss in the one side is always compared against a gain on the other side. Hence, what seems to trigger risk-aversion is not the asymmetry between gains and losses but a change in the cognitive perspective; while with purely positive and purely negative sequences attention is attracted by the locally maximum sample, in the mixed sequences in a comparison between a loss and a gain, the loss will be more noticeable. This flexibility of the cognitive perspective leads to *Experimental Study 8* where people's risk-attitudes in gains flip from risk-seeking to risk-aversion, when instead of choosing the best option they are given the logically equivalent task of rejecting the worst. Finally, in *Experimental Studies 9-10*, I consider the relationship between the fast value integration task and experience-based decisions by inducing self-paced sampling and by presenting bimodally distributed values respectively.

6.2 Risk-seeking in Gains (Experimental Study 5)

In the study of risky choice an almost universally employed tool is hypothetical monetary gambles (Kahneman & Tversky, 1979; Brandstätter, Gigerenzer, & Hertwig, 2006; Birnbaum, 2008). There, a *reflection* effect has been revealed (Tversky & Kahneman, 1992); people are found to be risk-averse for gains and risk-seeking for losses, which is attributed to the s-shaped utility curve held within prospect theory. For example the subjective utility of gaining £1000 is less than twice as good as gaining £500, while the disutility of losing £1000 is less than twice as bad as losing £500. Therefore when offered a gamble of winning £500 for sure or winning £1000 with

50% probability or £0 otherwise, people prefer the safe £500 option (and vice versa for losses). This pattern clashes with the rank-dependent model of value integration which predicts merely risk-seeking when the sequences are positive (*Computational Study 6*). The risk-seeking prediction in the numerical integration task is paradoxical for a further reason; research in numerical cognition has revealed that numbers are logarithmically compressed (Nieder & Miller, 2003) and thus a linear summation across the sequences should yield avoidance for the high-variance alternative. In the current study I examine whether value integration is indeed characterized by risk-seeking in the domain of gains.

6.2.1 Method

Participants. Participants were 16 adults [8 females; aged 20-31, mean 25.3] recruited from UCL's subject pool, and were paid for their participation.

Stimuli and Experimental Task. At each trial, participants saw pairs of numbers presented sequentially and had to decide, within 1500 ms, which of the 2 sequences had the highest average value. Each trial started with a presentation of a white fixation cross for 1000 ms, which was positioned at the centre of a black background screen. Afterwards, sequences of pairs of white numbers were presented at a rate of 2 items per second. The presentation of the last pair of numbers was followed by a green question mark at the centre of the screen for 1500 ms, which prompted the participants to indicate their response (left or right sequence) by pressing the left or the right arrow on the QWERTY keypad of a standard PC. After the response of the participant a black screen stayed on for 250 ms and then the next trial started. For incorrect responses error-feedback (a beep sound) was provided for the "feedback" group while the "reward" group received an extra sample. Failure to respond within 1500 seconds after the response cue's appearance was followed by a "deadline missed" message and a beep sound.

Procedure. The response mode was manipulated between participants. Half of the participants ($N = 8$) had to choose between two sequences the one with the highest average. Error-feedback was provided after each trial ("feedback" group). The other half of the participants had to choose which sequence they would prefer to draw an extra sample from ("reward" group). This division was done in order to ensure that the results apply to preference and not only to judgements of magnitude. After their

response, participants of the 2nd group saw on screen an extra sample, generated from their preferred sequence. At the end of the experiment they received one of the trial-rewards (randomly determined), with experiment units corresponding to GB pences. The sequence length was fixed at 12 pairs and the presentation rate at 2 pairs/ sec. The positions of the options were randomized.

Experimental Conditions. Overall 150 trials were presented fully randomized across conditions, in 5 blocks of 30. There were 3 overall conditions as depicted at the top panels of Figure 6.1. One alternative, labelled as “broad”, was always associated to a Gaussian with standard deviation of 20 while the other alternative (“narrow”) was generated from a Gaussian with standard deviation of 10. Two of the conditions were “unbalanced” as one alternative had always the highest mean (M) with the other alternative having a mean value 8 units lower ($M - 8$). In the “balanced” condition, the two alternatives had equal means (M). The variable M was sampled from a uniform distribution in the 45-55 range, at each trial. For Condition 3 no error-feedback was given (applicable for the “feedback” group where participants received error feedback).

6.2.2 Results

The effect of the response mode was examined in the unbalanced (right-top panel in Figure 6.1) and balanced trials (left and middle top panels in Figure 6.1) separately. For the unbalanced conditions a mixed ANOVA was performed, with condition (broad-best or narrow-best) being the within subjects factor and response mode the between subjects factor. The effect of response mode was not significant ($F(1, 14) = 1.43, p = 0.25$). The preference for the broad distribution sequence was also examined in the balanced condition by performing an independent-samples t-test between the two response mode groups. Again the effect of response mode was not significant ($t(14) = 1.48, p = 0.21$). Therefore the data of the two groups were analysed together.

Participants were able to select the best alternatives, associate to the high distributions, in both the unbalanced conditions (Figure 6.1 left, $t(15) = 8.34, p < 0.001$; Figure 6.1 middle, $t(15) = 13.36, p < 0.001$). Furthermore, accuracy was higher when the broad distribution had the highest mean ($t(15) = 3.62, p < 0.005$) indicating a bias towards the high variance distribution. This was confirmed by the choice pattern in the critical condition, where participants showed a clear risk-seeking attitude, preferring the high variance alternative above chance (Figure 6.1 right $t(15) = 5.39, p < 0.001$).

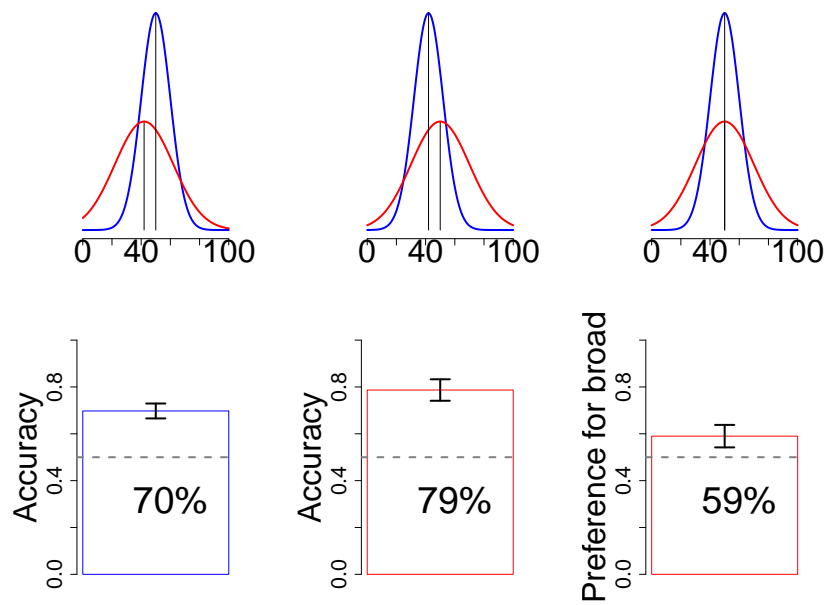


Figure 6.1: *Experimental Study 5 conditions (top) and results (bottom). Observers decided between two alternatives, each characterised by a sequence of 12 values, presented as pairs at a rate of 2/sec. Error bars correspond to 95% CI.*

6.2.3 Discussion

The results revealed a propensity towards the high-variance sequence, consistent with the predictions of the rank-dependent weighting model and *Computational Study 6*. This risk-seeking pattern quite surprising, colliding with findings from the mainstream line of research in risky choice, decisions by descriptions (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). However, this result is in the same direction with a recent finding in experience-based decisions (Ludvig & Spetch, 2011) where people experience probabilistic outcomes through active sampling, as opposed to the fast value integration task where sampling is passive and very rapid. Ludvig and Spetch (2011) found that when the risky option had two equally probable (50%) positive outcomes it is preferred over a safe option which always generates the same positive outcome. Moreover, when the outcomes were negative (losses) the pattern flipped to risk-aversion, yielding overall a reversal of the reflection effect with higher risk-seeking for gains than for losses, in contrast with the behavioural patterns in description-based decisions. Directly applying the rank-dependent model on negative sequences predicts the same risk-seeking pattern as the maximum (less negative) numbers will be more noticeable. However, it is not clear whether switching from gains to losses changes also the direction of the differential weighting, making the minimum (more negative) numbers more salient. Thus, in the next experiment I examine the risk-preference in negative numerical sequences.

6.3 Risk-seeking in Losses (Experimental Study 6)

6.3.1 Method

Participants. Participants were 9 adults [6 females; aged 20-33, mean 26.3] recruited from UCL's subject pool, and were paid for their participation £5.

Stimuli and Experimental Task. Identical to *Experimental Study 5* with the only difference that the presentation rate was set to 1 pair per 750 ms.

Procedure. All of the participants had to choose between two negative sequences, the one with the highest mean value (error-feedback in the unbalanced trials). At the end of each block participants could see their performance up to the point. The positions

of the options were randomized.

Experimental Conditions. The conditions and number of trials were identical to *Experimental Study 5*. The only difference was now that the domain was reversed from gains to losses and the sequences were always negative. The conditions are depicted in the top panels of Figure 6.2.

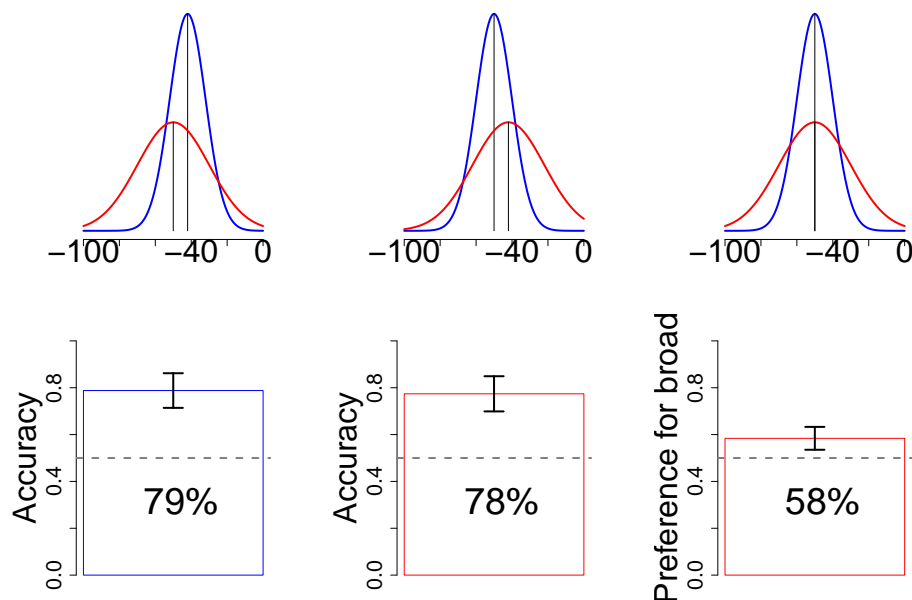


Figure 6.2: *Experimental Study 6 conditions (top) and results (bottom). Error bars correspond to 95% CI.*

6.3.2 Results

In the unbalanced conditions the participants were able to choose the sequence that was associated with the highest distribution above chance [narrow best (left panel in Figure 6.2) $t(9) = 7.63, p < 0.001$; broad best (middle panel in Figure 6.2) $t(9) = 7.15, p < 0.001$]. There was no signature of salience-based integration in the unbalanced trials since the difference in those two conditions was not significant ($t(9) = 0.45, p = 0.66$). However, in the critical condition where the two distributions had equal means, participants chose above chance the sequence associated with the broad distribution ($t(9) = 3.37, p < 0.001$), exhibiting risk-seeking behaviour.

6.3.3 Discussion

Although performance was invariant in the two unbalanced conditions, participants showed a clear and strong risk-seeking bias when the distributions had equal means. The lack of difference in the two unbalanced sequences might be attributed to a different mode of processing being in play due to the difficulty of processing negative sequences. In particular, because the representation of negative numbers is somehow less automatic (i.e., that is why the presentation rate was slightly reduced in this experiment), rank ordering the numbers at each frame might have been more difficult (e.g., ranking negative numbers might have been confusing for some participants). Nevertheless, in the balanced condition there was a clear indication of overweighting the locally maximum number although this tendency was somehow weaker compared to *Experimental Study 5* (i.e. by comparing the t-scores), consistent with the hypothesis that rank-based integration might be more difficult with negative numbers and fast presentation rate. In their study Ludvig and Spetch (2011) found an inverse reflection effect (i.e., risk-seeking in gains and risk-aversion in losses) and concluded that the representation of values might differ across domains e.g., between experience-based and description-based problems. Here, the same risk-attitude was found in both losses and gains, according to a rank-dependent mode of value integration. This mechanism is in a position to override any distortions possibly imposed at the representational level of values (e.g. s-shaped value function), revealing the important role of the micro-computations involved in choice.

6.4 Risk-aversion in Mixed Sequences (Experimental Study 7)

So far I examined the risk-attitudes in gains and losses using the fast value integration task. The results undermine explanations at the representational level of values (i.e. shape of value function) and favour the salience-driven integration account. However, there is a further behavioural pattern which has been attributed to the shape of the value function: risk-aversion with mixed gambles. In particular when a hypothetical gamble involves both gains and losses people are risk-averse (De Martino, Kumaran, Seymour, & Dolan, 2006; Tom, Fox, Trepel, & Poldrack, 2007). This tendency is assumed to arise because the value function is steeper in the negative domain and thus losses loom larger

than gains (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). In the next two experiments I examine risk-attitudes with mixed sequences.

6.4.1 Experiment 7a

6.4.1.1 Method

Participants. Participants were 14 adults [7 females; aged 20-31, mean 23.9] recruited from UCL's subject pool, and were paid for their participation £5 which could be reduced to £4 or increased to £6, depending on the choices they made during the task.

Stimuli and Experimental Task. Identical to *Experimental Study 5*.

Procedure. All of the participants had to choose between two sequences. The response mode was manipulated within participants. In half of the blocks the response question was to determine the sequence they would like to draw an extra sample from, while in the other half they had to choose the best sequence (error-feedback was given in these blocks). The response mode alternated from block to block, while the response mode in the initial block was counterbalanced between respondents. The sequences could be either purely positive or mixed, randomized across the experimental trials. After their response, respondents saw on screen an extra sample, generated from their preferred sequence. At the end of the experiment they received one of the trial reward/penalties (randomly determined), with experiment units corresponding to GB pences. The positions of the options were randomized. Overall there were 10 blocks and 300 experimental trials (30 trials per block).

Experimental Conditions. There were three conditions overall one unbalanced and two balanced (top panels in Figure 6.3). In the unbalanced condition the mean of the highest sequence was drawn from $U(0, 55)$. The other sequence was constructed with mean 8 units smaller. Both distributions had equal variances $\sigma = 10$. In the first balanced condition ("positive") the mean both distributions was equal and drawn from $U(45, 55)$, while one distribution was broad ($\sigma = 20$) and the other narrow ($\sigma = 10$). In the second balanced condition ("mixed") the mean of both distributions was set to $\mu = 0$ and again one distribution was broad and the other narrow as above. The trials were presented fully randomized and each condition had 100 trials.

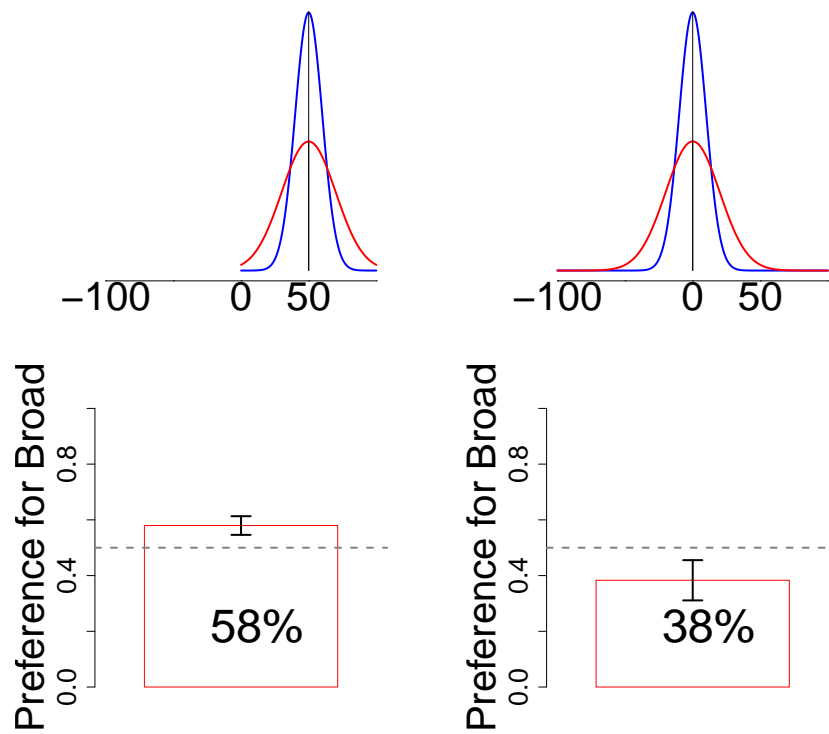


Figure 6.3: *Experimental Study 7a conditions (top) and results (bottom). The first condition (left) corresponds to strictly positive sequences while the second one (right) corresponds to mixed sequences (Gaussians with $\mu = 0$). Error bars correspond to 95% CI.*

6.4.1.2 Results

The effect of the response mode was not significant (similar to *Experimental Study 5*). In particular in the balanced condition the accuracy was invariant across the two modes ($t(13) = 0.94, p = 0.362$) as was the preference for the broad distribution in the “positive” and “mixed” conditions ($t(13) = -0.37, p = 0.719$ and $t(13) = 1.13, p = 0.278$, respectively). Therefore the results for each participant were averaged across response modes. In the unbalanced trials, participants had above chance accuracy ($M = 0.815, SD = 0.08$ and $t(13) = 14.1, p < 0.001$). In the positive-balanced trials participants showed above chance preference for the broad distribution ($t(13) = 4.6, p < 0.001$). However, this pattern reversed in the mixed-balanced where respondents were risk-averse, preferring less than chance the sequence associated to the broad distribution ($t(13) = -3.17, p < 0.001$). As also depicted in Figure 6.3, the difference in the preference for the risky option was highly significant in the mixed vs. the positive balanced trials ($t(13) = 4.7, p < 0.001$). In the next experiment examine whether this effect is driven by an asymmetry in the value function or by salience-driven integration.

6.4.2 Experiment 7b

6.4.2.1 Method

Participants. Participants were 16 adults [7 females; aged 19-36, mean 26.5] recruited from UCL’s subject pool, and were paid for their participation £5 which could be reduced to a minimum of £4, depending on the choices they made during the task.

Stimuli and Experimental Task. Identical to *Experimental Study 5*.

Procedure. All of the participants had to choose between two sequences the one they would like to draw an extra sample from. After their response, respondents saw on screen an extra sample, generated from their preferred sequence (could be positive, negative or zero). At the end of the experiment they received one of the trial reward/penalty (randomly determined), with experiment units corresponding to GB pennies. The positions of the options were randomized. Overall there were 8 blocks and 240 experimental trials (30 trials per block). Half of the blocks involved selection between homogeneous mixed sequences, i.e. a given pair would have strictly either neg-

ative or positive numbers (in both sequences). The second half of the blocks involved heterogeneous sequences, i.e. in each pair if one alternative had a negative value the other would necessarily have a positive one. Half of the participants ($N = 8$) did first four “homogeneous” blocks and then proceeded with the last four “heterogeneous” blocks (and vice versa for the other half of the participants).

Experimental Conditions. There were three conditions overall one unbalanced and one balanced. In the unbalanced condition the mean of the highest sequence was set to zero. The other sequence was constructed with mean 8 units smaller. In the “broad-best” unbalanced condition the highest distribution was broader ($\sigma = 20$) compared to the lowest distribution that was narrower ($\sigma = 10$) and conversely in the “narrow-best” condition. In the balanced condition, the mean of both distributions was zero but one sequence had highest standard deviation than the other (i.e., $\sigma = 20$ against $\sigma = 10$) and again one distribution was broad and the other narrow as above. In all trials each sequence had 6 positive and 6 negative numbers. The homogeneous and heterogeneous trials were identical in terms of the actual values and what differed was the order of presentation of the value samples (see Procedure). There were 60 balanced-homogeneous, 60 balanced-heterogeneous trials and 120 unbalanced trials (60 homogeneous and 60 heterogeneous, 30 for each unbalanced condition). Trial from the three conditions were randomized.

6.4.2.2 Results

The overall accuracy in the unbalanced trials was above chance ($t(15) = 11.09, p < 0.001$). A mixed ANOVA was conducted for the unbalanced trials with the type of the trial (broad-best vs. narrow-best) and the mode of presentation (homogeneous vs. heterogeneous) as repeated measures factors and the order of the blocks (homogeneous first vs. heterogeneous) as between subjects factor. The order with which the blocks were presented had no effect ($F(1, 14) = 1.54, p = 0.235$). The effect of the trial type was significant ($F(1, 14) = 8.80, p = 0.010$), showing in particular that accuracy was higher in the “narrow-best” trials ($M = 0.81, SD = 0.13$) compared to the “broad-best” ones ($M = 0.71, SD = 0.09$). Furthermore, the mode of presentation was significant ($F(1, 14) = 5.67, p = 0.032$) with performance in the homogeneous trials being higher ($M = 0.78, SD = 0.08$) compared to the heterogeneous trials ($M = 0.73, SD = 0.12$). No interaction was significant.

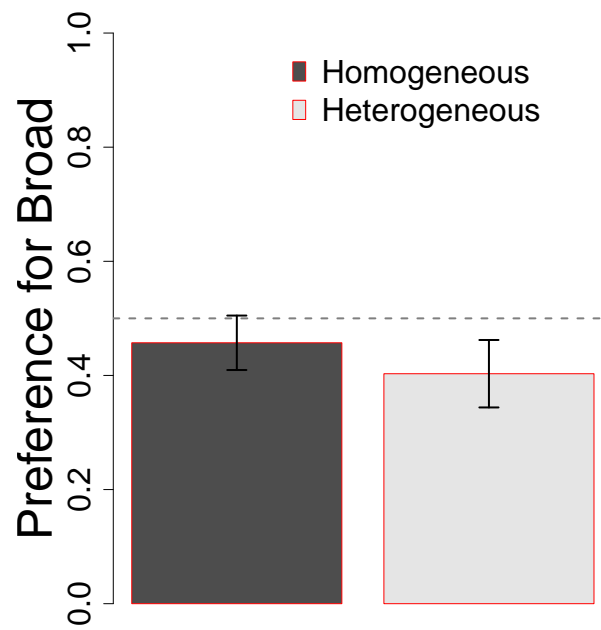


Figure 6.4: *Experimental Study 7b results*). The first condition (left) corresponds to homogeneous mixed sequences while the second one (right) corresponds to heterogeneous mixed sequences (Gaussians with $\mu = 0$). Error bars correspond to 95% CI.

The overall better accuracy in the “narrow-best” trials indicates risk-aversion (i.e., the low negative values of the broad distribution are penalized). Moreover, the higher accuracy in the homogeneous trials indicates a difference in the processing mode, depending on the type of presentation. In order to examine this hypothesis I turn now to the balanced trials. There, a mixed ANOVA with the type of presentation (homogeneous vs. heterogeneous) as repeated measures factor and the order of presentation as between subjects factor (homogeneous first vs. heterogeneous), was performed. The effect of the order of presentation was not significant ($F(1, 14) = 0.62, p = 0.443$) but the mode of presentation was ($F(1, 14) = 10.154, p < 0.001$). As Figure 6.4 shows, the preference for the broad sequence in the homogeneous trials, is not significantly different to the 50% chance level ($t(15) = -1.76, p = 0.099$). However, in the heterogeneous pairs it is below chance, indicating risk-aversion ($t(15) = -3.22, p < 0.001$). An interim conclusion is that in the mixed sequences the objective is to avoid the loss. Consequently, attention is driven towards the loss, in a pair where the one sequence has a negative sample and the other one a positive (heterogeneous trials).

6.4.3 Discussion

Risk-aversion in mixed gambles has been typically attributed to loss-aversion and the fact that losses loom larger than gains (Kahneman & Tversky, 1979). Whether this asymmetry in the value function is merely descriptive or reflects a hard-wired property of the brain is an open question (Tom et al., 2007). Recently, however, the loss-aversion hypothesis has been undermined in experience-based decisions (Hochman & Yechiam, 2011) where no behavioural risk-aversion was found with mixed outcomes. In the fast-value integration task, people exhibit risk-aversion (Experiment 7a) consistent with a gains/losses asymmetry assumption. However, as Experiment 7b showed, if negative numerical values were overweighted at the representational level then the risk-preference would be invariant between the homogeneous and heterogeneous trials where the numerical values were identical but their presentation order tweaked. The modulation of the risk-aversion by the presentation mode reveals that this pattern does not reflect any hard-wired value asymmetry but rather a differential attentional focus to losses: whenever offered a gain and a loss people will avoid the loss. At the microlevel of an experimental trial this is reflected in overweighting the negative value compared to the positive (this is why risk-aversion is boosted in the heterogeneous trials where evaluation of each pair penalizes further the loss, congruent with the trial objective

to avoid receiving a loss). This salience-driven account for risk-aversion (and consequently loss-aversion) is consistent with the view that loss-aversion reflects a change in the cognitive perspective (Ariely, Huber, & Wertenbroch, 2005) (i.e., when selling an item people focus more on what they have than on what they will gain from the transaction).

6.5 Risk-preferences and Task Framing (Experimental Study 8)

The experimental results so far provided ample support that salience-driven mechanisms underlie value integration. As *experiment 7b* revealed, what determines the salience of a sample is the long-term objective of the decision-maker. Accordingly, when presented with positive values only, the decision maker will try to maximize her gains and thus will overweight the maximum numbers that square well with her broader goal. In cases where a trial can lead to either a loss or a gain, people will try to avoid the loss which locally means that they will overweight the negative value in a pair that contains one negative and one positive sample. The top-down modulation of the salience of the items is reminiscent of findings in reason-based decisions (E. Shafir et al., 1993). In particular, when offered with a list of pros and cons for two alternatives, when asked to choose one, people prefer the option with the highest variability (with more cons and pros, similar to the result of *Experimental Study 5*). However, when they are given the logically equivalent task of rejecting one option, people prefer the less variable one (E. Shafir, 1993). In the next experiments I will test whether this flexibility obtained in reason-based decisions applies also to value integration and whether the salience of a value sample depends on the task framing.

6.5.1 Experiment 8a

6.5.1.1 Method

Participants. Participants were 16 adults [7 females; aged 21-36, mean 27.8] recruited from UCL's subject pool, and were paid for their participation £5.

Stimuli and Experimental Task. Identical to *Experimental Study 5*.

Procedure. Half of the respondents ($N = 8$) had to choose the best sequence between two throughout the whole experiment, while the other half had to choose the worst sequence (i.e. reject the worst). Error feedback was given in the unbalanced trials. The positions of the options were randomized. Overall there were 5 blocks and 150 experimental trials (30 trials per block).

Experimental Conditions. Conditions and number of trials were identical to *Experimental Study 5*.

6.5.1.2 Results

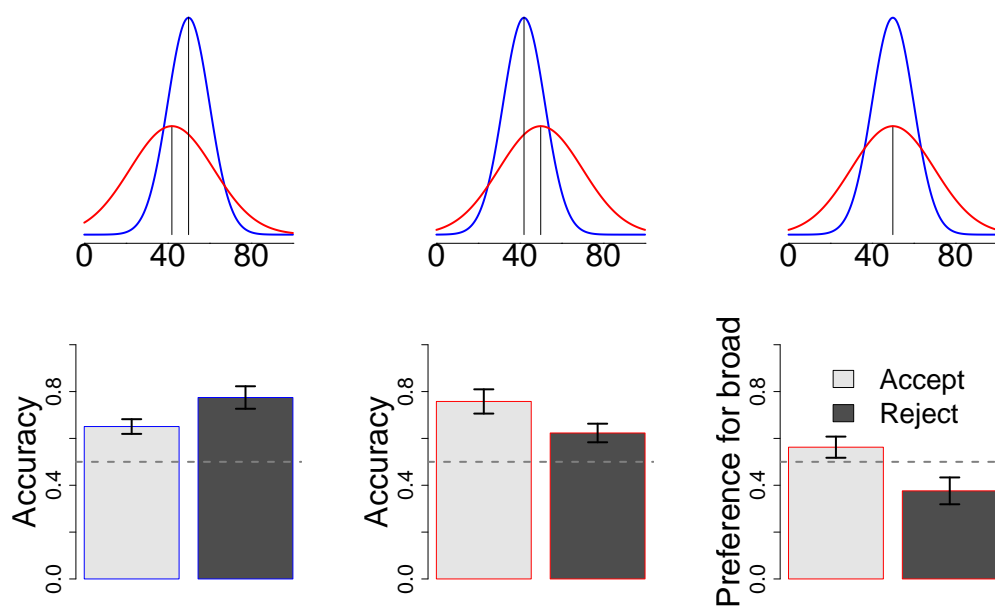


Figure 6.5: *Conditions and results in Experimental Study 8a.* Top: the two unbalanced (left, middle) and the balance (right) conditions. Bottom: accuracy in the unbalanced conditions (left, middle) and preference for the broad distribution in the balanced trials (right). Error bars correspond to 95% CI.

A mixed ANOVA was performed for the unbalanced trials (Figure 6.5 left and middle panels) with the type of the condition (broad-best against narrow-best) as repeated measures factor and the task framing (accept vs. reject) as within subjects factor. The type of the condition did not significantly affect performance ($F(1, 14) = 0.71, p = 0.415$) neither did the task framing ($F(1, 14) = 0.02, p = 0.882$). However, as shown in the bottom left and middle panels of Figure 6.5, there was a significant interac-

tion between the trial type (broad-best/ narrow-best) and the framing (accept/ reject) ($F(1, 14) = 24.01, p < 0.001$). Participants that chose the best option found easier the “broad-best” trials while participants that rejected the worst option, found easier the “narrow-best” trials. In other words the “selection” group showed a signature of risk-seeking while the “rejection” group an indication of risk-aversion (i.e., penalizing the low numbers). This was confirmed by comparing the preference for the broad option in the balanced trials between the two groups. As shown in the bottom-right panel of Figure 6.5, when selecting the best sequence people preferred the broad distribution but when rejecting the worst one they preferred the narrow distribution (independent samples t-test between the two groups: $t(14) = -3.54, p < 0.001$). This “reflection” effect confirms that the salience of a sample is determined by its congruency with the long term goal of the decision-maker, in accordance with reason-based theory and the results in E. Shafir (1993). In order to see how strong the effect of the task framing in modulating risk preferences is, in the next experiment I attempt to switch the risk-attitude within the same trial.

6.5.2 Experiment 8b

6.5.2.1 Method

Participants. Participants were 15 adults [7 females; aged 19-31, mean 24.1] recruited from UCL’s subject pool, and were paid for their participation £3 which they could increased up to £4, depending on their choices during the experiment.

Stimuli and Experimental Task. The background was gray and participants saw three sequences with each sequence number having a different colour (orange for left, magenta for top and green for right) and being surrounded by a frame. The position of the sequences was counterbalanced. The rest was identical to *Experimental Study 5*.

Procedure. The stimulus consisted of triples of numbers. Each trial consisted of two stages. During the first stage, participants saw 12 triples of numbers presented sequentially and in triangular arrangement, around a black fixation cross at the centre of the screen, at a rate of 750 ms. During the first stage the numbers were surrounded by black frames. At the end of the presentation the participant saw a black question mark, replacing the fixation cross, and had to indicate which option she wanted to eliminate (i.e., which option was the worst). Immediately after, the discarded option disappeared

from the screen and the two remaining options continued for another 12 frames at a presentation rate of 2 pairs/ frame. During the second stage the fixation cross and the frames around the numbers were white. At the end of the presentation a white question mark appeared at the centre and the participant had to choose which option was the best and from which she would like to draw an extra sample from. Upon selection, the extra sample from the selected option appeared on the screen indicating the reward at this particular trial. At the end of the experiment one of the extra samples (trial rewards) was randomly chosen and was given to the participant as monetary bonus (100 units corresponded to £1).

Experimental Condition. All three options were generated from Gaussians with the same mean value (at each trial the mean value was sampled from a uniform distribution in the 45-55 range). The two options had a standard deviation of 10 (narrow) while the third option had a standard deviation of 20 (broad). The positions of the options were randomized. In the trials where the broad option was eliminated at the 1st stage, beyond the participants' awareness, the distribution of one of the two remaining narrow options was turned from narrow ($\sigma = 10$) into broad ($\sigma = 20$) for the 12 frames of the second, selection stage. Overall there were 100 trials divided in 5 blocks of 20. None of the respondents detected the switch in the distributions when the broad option was rejected in the first stage.

6.5.2.2 Results

The results confirm the risk-attitude reversal with respondents rejecting the high risk alternative in the first stage more than the 33% chance level (Figure 6.6, $t(14) = 3.27, p < 0.001$). However, they subsequently showed risk-seeking by selecting the same alternative above chance (Figure 6.6, $t(14) = 4.81, p < 0.001$) at the second, selection stage of the trial. This preference reversal confirms the modulation of sample salience by the task framing. In order to see if there is a dependency between the first stage elimination and the second stage selection, I analysed the conditional probability of choosing the broad distribution given the result of the elimination stage [$P(\text{broad} - \text{chosen} | \text{broad} - \text{eliminated})$ vs. $P(\text{broad} - \text{chosen} | \text{narrow} - \text{eliminated})$]. The probability of choosing the broad distribution was independent of the option that was eliminated in the first stage ($t(14) = -0.69, p > 0.5$).

The flexibility of the decision mechanism to the task framing (bottom panel in Fig-

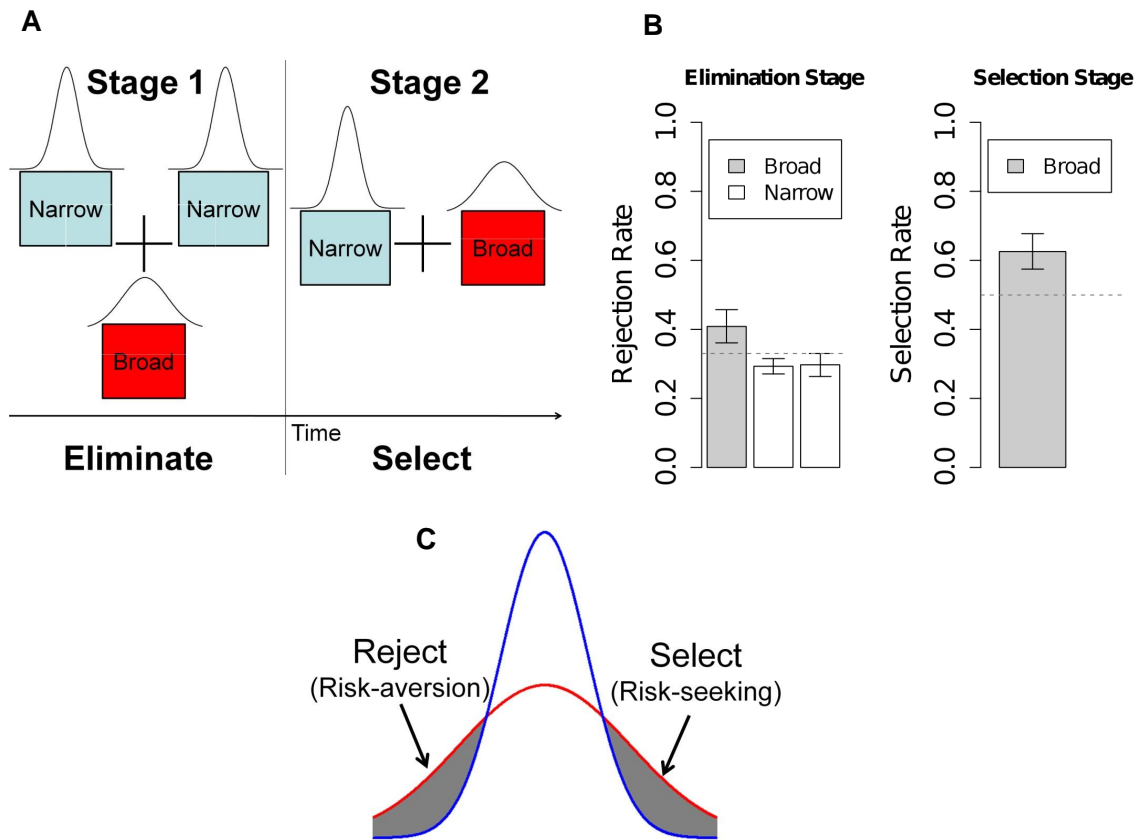


Figure 6.6: *Two-stage decision task and results in Experimental Study 8b.* Top left: Participants saw 12 triples presented at a rate of 750ms and were first asked to eliminate one of them (stage-1), and then to select one from the remaining two (stage-2), which were presented as a second sequence of 12 pairs at a rate of 500ms. Top right: In the first stage the rejection-rate of the risky-alternative was higher than chance (33%), while in the second stage the selection rate for it was also higher than chance (50%), consistent with an account that weighs different sides of the distribution depending on the task framing (bottom). Error bars correspond to 95% CI.

ure 6.6) can be understood in terms of a top-down mechanism, which modulates the salience of the samples depending on the framing. Assuming that some samples are ignored and not processed at all, in selection decisions, the maximum value in a given pair will be more noticeable and thus more often encoded while the same happens to the minimum values in rejection decisions. To mathematically capture this pattern, the values are weighed by their ranks and integrated in separate leaky accumulators:

$$P_i = \lambda \cdot P_i(t-1) + [V_i(t) \cdot w(\text{rank}_i(t))] + N(0, \sigma)$$

with $w(\text{maximum}) > 1$ and $w(\text{minimum}) = 1$ in selection and $w(\text{maximum}) < 1$ while $w(\text{minimum}) = 1$ in rejection decisions.

6.5.3 Discussion

The results of the two experiments presented in this section, violate the principle of invariance (Tversky & Kahneman, 1981) and are incompatible with theories of choice which assume that risk-attitudes are stable and task independent (e.g., Kahneman & Tversky, 1979). However, the findings are consistent with previous research in reason-based choice (E. Shafir, 1993) where people were found to reject the same option that they chose, depending on the question in the task. The relationship between the salience-driven value integration and the reason-based framework is apparent; information is weighted by its salience and the salience is top-down modulated by the long-term goals of the decision-maker. Such accounts undermine explanations of risky choice that attribute biases and reversals merely on the way values and probabilities are represented by internal functions. It is possible though that salience-driven accounts are more appropriate when information is presented sequentially rather than by description. In that case, the value integration paradigm, developed in this thesis, and experience-based decisions could be explained within a single framework. However, discrepancies in the two paradigms exist. In the next two experiments I will discuss and further examine these differences.

6.6 Risk-preferences and Self-paced Sampling (Experimental Study 9)

The decisions-by-experience paradigm (Hertwig et al., 2004) bears many similarities to the fast value integration task that was developed in this thesis. In particular, in both tasks, values are sampled sequentially and a decision needs to be made on the basis of the received samples. One procedural difference though, is that in the decisions-by-experience paradigm, the values are sampled actively through clicking boxes associated to alternatives. On the other hand, in the fast value integration task, values are perceived passively by the observer. Thus, in an experience-based task, the decision-maker is free to sample at will from all the alternatives, which induces also more complex exploration/ exploitation strategies (i.e., equally sample from both options or focus on one only). Although big discrepancies between experience-based decisions and the mainstream line of research in risky choice, decisions by description, have been revealed (Hertwig et al., 2004; Hertwig & Erev, 2009; Hadar & Fox, 2009; Ludvig & Spetch, 2011), it is not fully clear to what extent the experience-based protocol differs from the fast value integration task.

Going back to *Experimental Study 2* and order effects, the fast value-integration task was characterized by recency weighting and a large temporal span in play (i.e., accuracy kept improving with longer sequences). In experience-based task, in some studies there was an indication of recency weighting (Hertwig et al., 2004) which was not replicated though in another case (Ungemach, Chater, & Stewart, 2009). Furthermore, consensus has been reached that participants in experience-based decisions rely on a small subset of samples (Hau, Pleskac, Kiefer, & Hertwig, 2008), which is ruled out by the large time-constant and the steady improvement of accuracy with longer samples in *Experimental Study 2*. While the risk-seeking bias in gains of *Experimental Study 5* coincides with a recent finding in an experience-based task (Ludvig & Spetch, 2011), the risk pattern in losses differs (i.e., risk-seeking according to the result of *Experimental Study 6* and risk-aversion in Ludvig & Spetch, 2011). Additionally, the findings of *Experimental Study 7* indicated risk-aversion in mixed sequences which has not been obtained in decisions by experience (Hochman & Yechiam, 2011). In the next experiment I manipulate the type of sampling (active or passive/ fast or self-paced) in the fast value integration task, in order to examine at what extent the way that information is received affects risk preferences.

6.6.1 Method

Participants. Participants were 16 adults [9 females; aged 19-45, mean 26.6] recruited from UCL's subject pool, and were paid for their participation £3 which they could increase up to £4, depending on their choices during the experiment.

Stimuli and Experimental Task. Identical to *Experimental Study 5*.

Procedure. All respondents had to choose between two positive numerical sequences the one they would like to draw an extra sample from. In each trial, after their choice, the reward sample appeared on the screen. At the end of the experiment one of the trial rewards was randomly given to them as bonus, with experimental units corresponding to GB pennies. The experiment consisted of two blocks. In one block participants receive values passively at a rate of 2 pairs/ second. In the other block a pair appeared on the screen only after the participant press the space bar on the keyboard. Therefore, the sampling was self-paced and active. Half of the participants started ($N = 8$) with the self-paced block and proceeded with the passive block and vice-versa for the other half of the respondents. Each block contained 50 trials (100 trials overall).

Experimental Condition. Identical to the balanced condition in *Experimental Study 5*.

6.6.2 Results

A mixed Anova was performed with the DV being the preference for the broad sequence, the sampling mode (active or passive) the repeated measures factor and the order of presentation (active first or passive first) the between subjects factor. The sampling mode was not significant ($F(1, 14) = 0.10, p = 0.756$) while the order with which the blocks were performed was marginally significant ($F(1, 14) = 4.6, p = 0.05$). The effect of the order, as also shown in Figure 6.7, indicates that the risk-seeking pattern diminishes from the first block to the second (the sampling mode and order interaction was not significant; $F(1, 14) = 2.18, p = 0.162$), indicating a small learning effect. Finally, by averaging the preference for the broad distribution across the two sampling modes, a highly significant risk-seeking was found ($t(15) = 4.92, p < 0.001$).

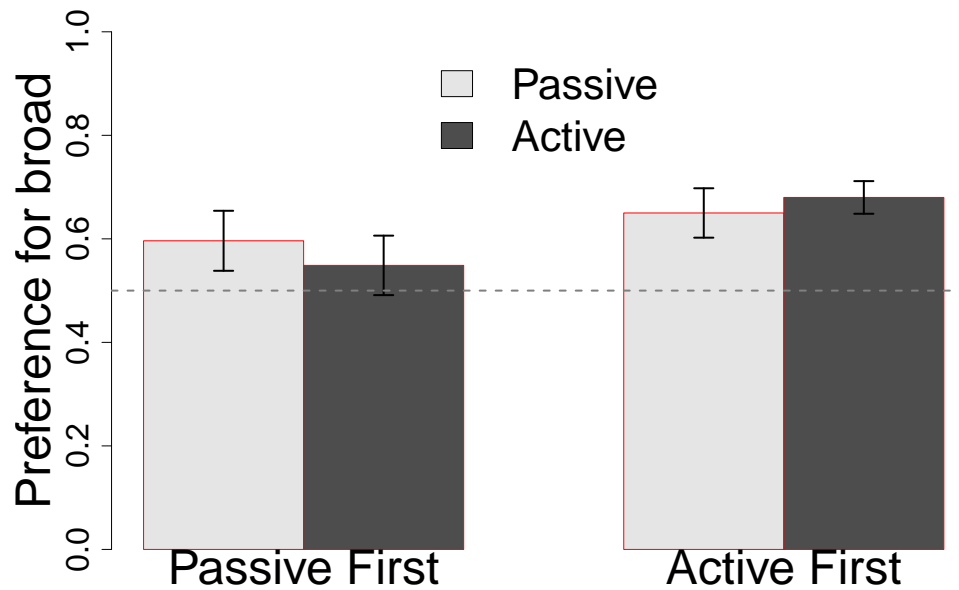


Figure 6.7: *Experimental Study 9 results. One half of the experiment involved active sampling (pressing a space bar to receive a sample) while the other half involved passive sampling. The order of the sampling mode was counterbalanced between participants (i.e., the one group sampled actively in the first half of the experiment and passively in the second and vice versa for the other group). Error bars correspond to 95% CI.*

6.6.3 Discussion

One of the distinctive characteristics of the fast value integration task is that values are sampled rapidly and in a passive mode. In this experiment I attempted to see whether the fast presentation mode caused or not one of the strongest patterns encountered up to the point: risk-seeking in the domain of gains. The results clearly reveal that when participants are free to receive the value samples at their own pace, the risk-seeking pattern persists. An interesting observation from the current study is that the risk-seeking pattern diminishes from the first to the second block (regardless of the sampling mode). Therefore it is possible that salience-driven integration is stronger at stages where people are unfamiliar with the task (and do not have a good grasp of the goodness of the absolute values) and rely more on auxiliary cues (i.e. local ranks), a tendency which dissipates with learning. In the next experiment I examine whether, similar to experience-based decisions, people underweight the rare events in the fast value integration task.

6.7 Risk-preferences and Rare Events (Experimental Study 10)

One tenet arising from the decisions-by-experience paradigm is that rare events are ignored or underweighted (Hertwig et al., 2004; Hertwig & Erev, 2009). This principle, combined with others such as the reliance on small samples, explains most of the patterns encountered in this paradigm. Considering the weight of the rare events in experience-based decisions is reasonable since, similar to descriptive monetary gambles, the alternatives are often characterized by bimodal distributions. For example people are often confronted with one risky and one safe option with equal expected values. The risky option might generate zeros 90% of the time and a large positive value the rest 10%, while the safe option will yield always the same moderate value. This stands in contrast with the fast value integration task, where the values are varying in a more continuous way, being generated by normal distributions. In this experiment, I use the fast value integration protocol and present to participants values generated by bimodal (risky option) or unimodal (safe option) distributions, similar to experience-based decisions.

6.7.1 Method

Participants. Participants were 9 adults [6 females; aged 22-29, mean 24.8] recruited from the University of Tel Aviv subject pool, and were paid for their participation £3 which they could increase up to £4 (the actual amount was converted into the corresponding Israeli currency), depending on their choices during the experiment.

Stimuli and Experimental Task. Identical to *Experimental Study 5*. The only difference was that between each pair, a blank black screen was interleaved for 100 ms. This was necessary since in some trials the same value sample was repeated sequentially within one alternative. Without this blank interval, the alternative which had always the same value would appear not to be updated at all, while the changing one would fluctuate attracting that way visual attention.

Procedure. All respondents had to choose between two positive numerical sequences, presented at a rate of 2 pairs/ second, the one they would like to draw an extra sample from. In each trial, after their choice, the reward sample appeared on the screen. At the end of the experiment one of the trial rewards was randomly given to them as bonus, with experimental units corresponding to GB pennies. The experiment consisted of two blocks. In one block participants receive values passively at a rate of 2 pairs/ second. The experiment consisted of 250 trials divided into 5 blocks.

Experimental Conditions. The task involved two types of conditions: balanced (150 trials) and unbalanced (100 trials). The balanced trials involved selection between two sequences, one risky and one safe, of overall equal expected value. First, the additive value of the sequences was sampled from $U(40, 60)$. In the safe sequence all twelve values were equal, summing up to the additive value (e.g. if the additive value was 60 then all twelve values of the safe sequence were equal to 5). For the risky sequence the probability of a non-zero value to occur in the risky option was manipulated across three levels (50 trials for each level). In the first condition the ten value samples were equal to zero while the two other samples had values that added up to the overall additive value (e.g. 10 values were zero and 2 values 30, $p(\text{non-zero}) = 2/12$). In the second condition, the probability of a non-zero value was $p = 6/12$ while in the third condition it was $p = 10/12$. Therefore, for $p = 2/12$ the rare events had high values while for $p = 10/12$ values of zero. The three balanced conditions were divided into two types: clear vs. perturbed sequences (75 trials each). The clear sequences were constructed as described above (i.e., with one discrete value for the safe option and

two discrete values- zero and a positive value- for the risky option). In the perturbed sequences a constant value sampled from $U(0,5)$ was added independently on each value sample. The unbalanced trials were constructed similarly (with one risky and one safe option, and the probability of a non-zero value to occur in the risky option manipulated across three levels), but the one option (for half of the trials the risky and for the other half the safe) dominated the other, having a higher additive utility (i.e., the difference set at 95 units). All conditions were fully randomized.

6.7.2 Results

An analysis on the filler trials was performed by dividing them on the basis of the dominating option. When the dominating option was the risky one, performance was not significantly above chance ($M = 0.52, SD = 0.27, t(8) = 0.22, p = 0.829$). On the other hand when the safe option (which always had non-zero values) was the best, performance was much above chance ($M = 0.83, SD = 0.13, t(8) = 7.46, p < 0.001$). The difference in the accuracy in the two types of trials was significant ($t(8) = -2.67, p = 0.027$). This difference in performance reveals a bias towards the safe sequence. In the balanced trials, the probability of the non-zero sample to occur in the risky option did not significantly affect the probability of choosing it ($F(2, 18) = 2.64, p = 0.102$). However, as Figure 6.8 shows, the manipulation of the probability of the non-zero event to occur, reduced the risk-aversion behaviour (the less the non-zero events in the risky option were, the more likely it was chosen). The lack of significance in this trend should be considered with respect to the relatively small amount of trials (25 trials for each data point in Figure 6.8). The effect of the type of sequences (clear vs. perturbed) was highly significant ($F(1, 8) = 18.82, p < 0.001$). As depicted in Figure 6.8, the perturbation of the sequences with a small positive constant (randomly chosen from $U(0,5)$) reduces risk-aversion. In particular, the preference for the risky option (averaged across the three conditions, i.e. $p = 2/12, 6/12, 10/12$) in the clear sequences was significantly below 50%, indicating risk-aversion ($M = 0.28, SD = 0.20, t(8) = -3.23, p = 0.012$) while in the perturbed sequences it did not differ from chance ($M = 0.42, SD = 0.23, t(8) = -0.99, p = 0.36$).

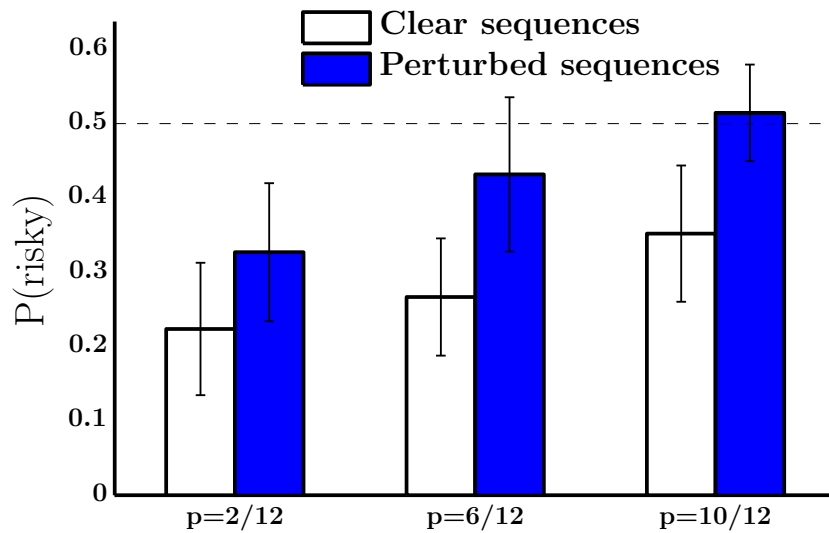


Figure 6.8: *Experimental Study 10 results. The probability of zero (in clear sequences) or a very small number (perturbed) appearing in the “risky” sequence was manipulated within participants. The other sequence (“safe”) had always moderate values (the same value in clear sequences and slightly differentiated in the perturbed sequences) spread evenly in all of its samples. The two sequences had equal integrated values. Error bars correspond to 1 SE.*

6.7.3 Discussion

The fact that participants avoided the risky option in the $p = 2/12$ condition, where that option had rarely non-zero positive values, is consistent with the ignoring rare events hypothesis (i.e. the positive rare values are ignored). If ignoring the rare events was in play, this pattern should be reversed in the $p = 10/12$ condition where the risky option had rarely zero values (which should be ignored). However, in the latter condition people did not show any propensity to choose the risky option above chance. Moreover, the significant effect of perturbing the value samples on the probability to choose the risky option, indicates the important role of the zero value-samples and supports a salience-driven integration account. In particular, it is reasonable to assume that in this task, the long-term objective of the respondents is to avoid getting at the end a zero sample as reward. If this was the case then zero values should be further penalized when paired with a small positive value (similar to penalizing the losses in *Experimental Study 7*) and that might explain why the risky option is avoided in the clear sequences. On the other hand, in the perturbed sequences, where the zeros are replaced by small positive values, the risk-aversion attitude disappears. It is interesting to repeat this experiment by perturbing the sequences with even higher values and see if, and at what magnitude of perturbation, risk-seeking is obtained. To conclude, the results of the current study suggest that the “ignoring the rare events” hypothesis does not apply to the value integration task. This finding in combination with the study by Ungemach et al. (2009), where people accurately (i.e., without underweighting) reported the frequencies of the encountered events in an experience-based task, opens the possibility for alternative interpretations (e.g. salience-based integration) of Hertwig et al. (2004) original results.

6.8 Summary and General Discussion

Risk attitudes have been traditionally elicited using hypothetical monetary gambles which provide explicit information about probabilistic outcomes (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992; Stewart et al., 2005; De Martino et al., 2006; Tom et al., 2007). Findings from this line of research (i.e. decisions-by-description), revealing risk-aversion with gains and risk-seeking with losses, have been attributed to non-linearities in the way values or probabilities are represented by internal functions.

In this chapter, I aimed to examine people's risk preferences using an alternative experiment tool, the fast value integration task. In the fast value integration paradigm, similar to experience-based decisions (Hertwig et al., 2004), people experience directly value samples. The value samples are normally distributed and therefore the risk associated with an option coincides with the variance of the numerical sequence. For example, choosing a sample from a sequence associated with a broad distribution can yield an extremely low or an extremely high value, while choosing a low-variance sequence will return a more mediocre value. Probing risky choice using this task can reveal the micromechanisms that lead to risk-biases, which potentially override explanations at the representational level.

In *Experimental Study 5*, I confirmed the rank-dependent prediction of risk-seeking in the domain of gains (*Computational Study 6* in Chapter 5). This finding, although clashing with results in description-based decisions (Tversky & Kahneman, 1992), coincides with a similar pattern that was revealed in experience-based decisions by Ludvig and Spetch (2011) who found an inverse reflection effect: risk-seeking in gains and risk-aversion in losses. The latter finding was not replicated in the fast value integration task, where participants showed again, as in the domain of gains, risk-seeking (*Experimental Study 6*). However, in accordance with decisions by description (De Martino et al., 2006) and contrary to experience-based decisions (Hochman & Yechiam, 2011), participants in the fast value integration task, exhibited risk aversion with mixed sequences that were normally distributed around zero (*Experimental Study 7a*). As *Experimental Study 7b* showed, what drives risk-aversion is not the magnitude asymmetry between losses and gains (as factored in by hard-wired loss-aversion in value-based accounts), but the change of cognitive perspective (e.g., Ariely et al., 2005) and the increased salience of losses: the goal of the decision-maker to avoid a loss at the end of the trial, is reflected at the microlevel by an increased penalty or attentional salience applied on the loss samples. The modulation of the samples salience by the long-term objective of the respondent was further confirmed in *Experimental Studies 8a* and *8b*. There, asking participants to reject the worst option instead of selecting the best one, resulted in a switch in the risk-bias from risk-seeking to risk-aversion in the domain of gains. This switch is consistent with data (E. Shafir, 1993) and theories (E. Shafir et al., 1993) of high-level, reason-based decisions and also with a salience-driven account of value integration, with salience being dependent on the goals of the agent.

It is questionable whether and how a salience-driven account is applicable to decisions-by-description with hypothetical gambles. What is missing in order to apply the micro-structure insights obtained here to description-based problems, is a detailed understanding of the deliberative algorithm that people undergo when faced with hypothetical gambles. Recent eye-tracking studies (e.g., Glöckner & Herbold, 2011) on risky choice promise to bridge this gap and provide a process-based description of the information sampling regularities that characterize choice among monetary gambles. Contrary to decisions-by-description, where the deliberative computations that people employ are covert, a salience-driven integration account is directly applicable to experience-based decisions, where information is actively and sequentially sampled. Yet, however, discrepancies between the fast value integration and the experience-based domains exist at the empirical front (i.e., risk-attitude in losses i.e., Ludvig & Spetch, 2011 and with mixed sequences i.e., Hochman & Yechiam, 2011). In order to further explore the basis of these discrepancies, I attempted to minimize the procedural differences between those two domains by first inducing active sampling in the fast-value integration task. The risk preferences in *Experimental Study 9*, were invariant to the nature of the sampling (passive against active). Next (*Experimental Study 10*), I changed the type of the sequence-distributions from normal to more discrete unimodal or bimodal (similar to the standard practice in the experience-based paradigm e.g., Hertwig et al., 2004) and examined how this change affects risk preferences. The results ruled out the standard explanation held within the experience-based literature, that rare events are ignored or underweighted (Hertwig & Erev, 2009), supporting a salience-driven account where the occurrence of zero outcomes (or extremely low ones) is more noticeable and negatively weighted. The latter finding opens the possibility of re-interpreting the results in Hertwig et al. (2004) under the rank-dependent computational framework that was developed throughout this thesis.

Chapter 7

Summary and Conclusions

The aim of this thesis was to examine the way that decision-relevant information is integrated across time in both perceptual and value-based choice. The two specific research questions I pursued, concerned how the temporal order and the decision context distort this integration process. In a series of experimental studies, by tightly controlling the time-course of the decision-relevant information, I obtained order, risk and task-framing biases as well as contextual preference reversal effects. These phenomena were explained by a simple mechanism based on the integration of information, weighted by its salience. The salience of a sampled piece of information depended on its temporal order and local rank in the decision context, while the direction of the weighting was determined by the task framing. That way, I demonstrated that choice regularities, that are traditionally attributed to the way information is represented, arise from the microstructure of the information integration process. Moreover, the salience-driven integration model I developed here, promises to establish a common theoretical framework between evidence-based (e.g., integrating perceptual or reward information) and goal-directed decisions, focussed on multiple goals with trade-offs (e.g. choice among cars or flats). In this last chapter, I will summarize the main findings of each chapter, discuss the implications of these findings and allude to outstanding issues and future directions.

7.1 Summary of Findings

In Chapter 3, I examined the presence and basis of order effects in both perceptual and value-based choice. I started by presenting a computational exploration (*Computational Study 1*) of the LCA choice model (Usher & McClelland, 2001), showing that under a single set of parameters it predicts primacy for short stimulus durations which switches into recency for longer durations. Evidence for this unique prediction, emerging from a neurally plausible aspect of the model (that preference states correspond to neuronal firing rates and thus cannot go below 0), was provided in a perceptual experiment (*Experimental Study 1*). A similar pattern was obtained in the value integration task (*Experimental Study 2* and *Computational Study 2*), in the form of recency bias which increased with sequence length. These results led to the conclusion that the temporal span of information integration is limited. The implications of this limitation for choice optimality were explored in *Computational Study 3*, showing that decay-based information integration is ameliorative in volatile and ever changing environments.

In Chapter 4, I examined the effect of the context on the integration process, using choice among three alternatives. By manipulating the exact time-course of the decision input, I induced temporal correlations among the alternatives in an attempt to create analogs of the deliberative process that underlies multi-attribute choice problems. In a brightness discrimination task (*Experimental Study 2*), a strong similarity effect (Tversky, 1972) was obtained and was explained within the LCA model of perceptual choice due to its zero non-linearity (*Computational Study 4*). In the corresponding value-integration experiment (*Experimental Study 3*), both the attraction (Huber et al., 1982) and the similarity effects were obtained, indicating that the fast value integration psychophysical paradigm is a valid proxy to the underlying process of goal-directed decisions.

Although the similarity effect could be accounted within the perceptual LCA, this was not the case with the attraction effect that was obtained in *Experimental Study 3*. This led to the development of an alternative computational model in Chapter 5, which accounted for preference reversal effects by assuming that pieces of information are boosted proportionally to their local/ momentary ranks (*Computational Study 5*). The exact form of this differential boost was specified for *Experimental Studies 3* and *4*. Finally, novel predictions from the rank-dependent model, involving risk-seeking in gains, were derived (*Computational Study 6*).

These predictions, led to Chapter 6 and the examination of risky choice using the fast value integration paradigm. First, contrary to the standard result obtained within the description-based literature (e.g. Tversky & Kahneman, 1992), respondents in the fast value integration task demonstrated risk-seeking with positive sequences (*Experimental Study 5*). This finding is consistent with a recent study in experience-based decisions which revealed a similar pattern (Ludvig & Spetch, 2011). However, contrary to the latter study in which losses were characterized by risk-aversion, participants in the value integration task, were still risk-seeking even when the sequences corresponded to losses (*Experimental Study 6*). Their attitude reversed to risk-aversion only when the sequences were mixed (*Experimental Study 7*), involving both positive and negative values, and when the task framing switched from selection (i.e. choose the best sequence) to rejection (i.e. discard the worst sequence) (*Experimental Study 8*). The reversal of risk preferences in mixed sequences and rejection decisions, revealed that the direction of the differential weighting applied on value samples, depends on top-down factors and is congruent with the decision-maker's global objective. In the final two experiments (*Experimental Studies 9, 10*), I explored the relationship between the fast value integration task and experience-based decisions with respect to risk preferences. The results ruled out the prominent view held within the experience-based literature, that rare events are underweighted, opening the possibility of re-interpreting data in this field (e.g. Hertwig et al., 2004) under a salience-driven computational framework.

7.2 Implications

Value Psychophysics: A Window on Motivation-based Choice

Value-integration is an essential process embedded on dynamic models of preference (Roe et al., 2001). There, the cognitive system integrates subjective values (rather than, say, pieces of perceptual evidence), which depend on how each alternative matches the decision maker's goals (Roe et al., 2001; Usher & McClelland, 2004). In particular, when alternatives are characterized by different attributes (e.g. price and quality of a product), preference is shaped through shifting attention across these attributes, assessing an item's subjective value on each attribute, integrating these values across time, and finally making a choice when some threshold is reached. A detailed understanding of these computations might explain the systematic anomalies observed in motivation-based decisions.

Studying the microstructure of motivation-based choice has been difficult to pursue because classical laboratory preference tasks provide little control of the moment-by-moment processes of value sampling and integration. Instead, choice alternatives are statically presented and the decision maker reports her preference after freely sampling information related to the alternatives. Although tracking the regularities in the information sampling process is feasible (Krajbich & Rangel, 2011; Glöckner & Herbold, 2011) internal shifts of attention between attributes cannot be measured and hence these tasks are not properly constrained to test the underlying cognitive operations. This stands in contrast with psychophysical paradigms for studying evidence-based perceptual choice where the flow of sensory evidence is fully controlled by the experimenter (Britten et al., 1993). To obtain more precise control on the decision input I introduced a novel experimental paradigm at the interface of psychophysics and motivation-based decisions. Participants simultaneously viewed two or three rapidly varying sequences of numerical values. Controlling the flow of the input values allowed to directly probe how people attend to and integrate values.

How close is this paradigm to the actual deliberative process that people employ when faced with complex trade-offs? By manipulating the temporal distributions of the value samples and by inducing temporal correlations among the alternatives, I obtained a number of choice patterns and paradoxes, typically encountered in motivation-based choice problems. In particular I obtained order (e.g. Furnham, 1986) and risk (e.g. Kahneman & Tversky, 1979) biases, violations of the invariance due to task framing (E. Shafir, 1993) and contextual preference reversal (Tversky, 1972; Huber et al., 1982; Maylor & Roberts, 2007). These effects link the simple, psychophysical task of value integration to high-order, multi-dimensional decisions showing that the technique of controlling the sampling process is a good proxy to the study of decisions in richer domains.

The use of psychophysics as a valid abstraction of complex decision-making problems has several advantages. First, it facilitates the construction of simple and parsimonious models, that are able to account primarily for the experimental data and capture, as a consequence, the core mechanisms that underlie more complex cognitive functions. Second, it allows the “masking” of experimental manipulations (e.g. collapse trade-off problems into one-dimensional input with covert temporal correlations) and prevents the corruption of the results from attitudes and knowledge held by the participants independent of the experiment. Furthermore, it allows to take repeated measures

of the same problem from the same person without the responses being affected by memory or a tendency for responding consistently across trials. Finally, the “value psychophysics” success in approximating the underlying process of goal-directed decisions, opens the possibility for the development of similar simplified paradigms for the induction of higher-level mental states in other domains (e.g. reasoning).

Microcomputations vs. Value Representations

The making of a decision arises from a series of computations over a set of mental representations related to the available choice options. Traditional theories of choice have focused especially on the representational level and attributed choice anomalies to distortions in the functions that represent decision relevant quantities such as values and probabilities. For example, the representation of utility in risky choices, following the normative theory (Von Neuman & Morgenstern, 1947), is assumed to be logarithmically compressed. Risk-biases have led to the refinement of this value function and the development of similar non-linear functions for the representation of probabilities (Kahneman & Tversky, 1979). A theoretical prerequisite arising from this approach, holds that these functions are relatively stable and their shape hard-wired. However, the existence of context effects (Huber et al., 1982; Tversky & Simonson, 1993; Pettibone & Wedell, 2007) and value relativity phenomena (Stewart et al., 2003; Vlaev et al., 2009; Kurniawan et al., 2010) has challenged, but not completely ruled out (since these accounts can incorporate the notion of reference points that shift around these functions without fundamentally changing their shapes), this claim.

Throughout this thesis, I focused on the algorithmic details that govern the choice process. The obtained effects were accounted merely by the nature of the computations towards a decision, regardless of the exact representational details of the decision-relevant quantities. For instance, it is well established that the representation of numerical values is logarithmic (Nieder & Miller, 2003). A direct implication, assuming that numbers are integrated via linear summation, would be that when selecting from two streams of numbers the one with the highest sum, people would avoid the one with the highest variance (risk-aversion). Despite that, the results of this thesis revealed that participants performing the fast value integration task exhibited risk-seeking behaviour. In other words, the microcomputations performed with values (however these values are represented) might exert a stronger influence on the choice outcome and override, in some cases, biases induced by non-linearities in the representation of values.

Process theories that emphasize the role of computations, rather than the role of the representations, have been recently put forward (e.g. Stewart et al., 2006; Roe et al., 2001), successfully accounting for the online construction of value and its dependency on environmental contingencies (Stewart, 2009; Ungemach, Stewart, & Reimers, 2011). I believe that micro-level accounts like Decision by Sampling (Stewart et al., 2006) or the salience-driven one developed here, are not necessarily at odds with the notion that there are some regularities in the representation of values at the neural level (Trepel, Fox, & Poldrack, 2005; Padoa-Schioppa & Assad, 2007; Tom et al., 2007; Levy, Snell, Nelson, Rustichini, & Glimcher, 2010). And in fact some behavioural patterns might be explained by non-linearities in these regularities alone. In most of the cases, however, decision-relevant quantities are expected to be subject to further important transformations by the decision algorithm.

Sensitivity to Higher Moments of the Decision Input and Choice Optimality

Research in evidence-based (e.g. perceptual or reward) choice has traditionally focused on the way people respond to the mean strength of the stimulus. Accordingly, the goal of the decision-maker is to average out the noise embedded in stationary distributions. The results of this thesis clearly show that people are also sensitive to the variance of the decision input. This sensitivity is captured by the salience-based integration framework, which further enhances peaks in the stimulus, and squares well with findings in other domains such as perceptual categorization (Summerfield et al., 2011) and judgemental forecasting (Harvey, 1995; Reimers & Harvey, 2011). It appears thus, that people are not processing the decision input in a merely bottom-up and automatic way, but rather engage into additional computations (e.g. ranking) in order to better infer the causal and probabilistic structure of the problem (i.e. model-based reasoning rather than pavlovian responding).

In stationary environments, picking up second-order regularities could be thought of as redundant. Why are then people sensitive to aspects of the stimuli that are seemingly irrelevant? One plausible answer is that the use of auxiliary cues can facilitate choice even under stable conditions. For example utilizing the rank order of the alternatives (Parducci, 1965; Stewart et al., 2006) in combination with their absolute values (i.e. the rank-dependent framework developed here) might be a more robust algorithm, speeding up spotting the best alternative (small differences are amplified) or guiding decisions in novel situations where the stimulus scale is meaningless to the decision-maker (e.g. evaluating candidates whose absolute grades mean nothing to the

examiner). Furthermore, in volatile environments and since boosting values by ranks magnifies small differences, changes in the underlying structure would be more readily detected.

Of course a strategy like this one is a double-edged sword; the increased sensitivity to differences (or changes) might lead to “overfitting” of the world’s structure and a costly false alarm rate. Therefore the optimality question is whether people are able to adjust the way they rely on secondary cues of the stimulus, in response to the structure of the environment (Summerfield et al., 2011). As I demonstrated in this thesis (e.g. order effects due to decay-based information, section 3.5), biases might be the price that humans pay for being employed with mechanisms that enable them to be optimized on an ever changing world. Being employed with such mechanisms should improve, on average, choice quality and choice paradoxes might just reflect exceptional cases where these mechanisms fail due to oddities or wittingly introduced twists in the structure of the choice problem.

Perceptual vs. Value-based Choice

Recent research on the psychology and neuroscience of simple, evidence-based choices (e.g., integrating perceptual or reward information) has led to an impressive progress in capturing the underlying mental processes as optimal mechanisms that make the fastest decision for a specified accuracy (Shadlen & Newsome, 2001; Ratcliff et al., 2004; Rorie et al., 2010). The idea that decision-making is an optimal process stands in contrast with findings in more complex, motivation-based decisions, focussed on multiple goals with trade-offs (e.g., choice among cars or flats). Here, a number of paradoxical and puzzling choice behaviours (Tversky, 1972; Huber et al., 1982; Tversky & Simonson, 1993) have been revealed, posing a serious challenge to the development of a unified theory of choice. Can a common theoretical framework between evidence-based and motivation-based decisions be established?

So far the two fields of perceptual and value-based decisions have been studied independently with different methods and techniques and only a set of recent theoretical models (Roe et al., 2001; Usher & McClelland, 2004; Stewart et al., 2006; Tsetsos et al., 2010) have attempted to bring those two fields closer, by theorizing goal-directed behaviour under sequential sampling models. Drawing on these models, I attempted in this thesis to create an experimental protocol where values are presented sample by sample, in a way similar to sensory stimuli. Despite that the perceptual and value-based

experiments differed only in the type of stimulus (sensory vs. numerical values), several behavioural discrepancies were obtained. In particular, concerning order effects, the perceptual experiment was characterized mostly by primacy whereas the numerical one by strong recency. Thus, different time constants seem to underlie each paradigm which can be presumably attributed to the longer trial durations (in absolute terms) employed in the value-based task. Regarding context effects, different effects (and of different magnitude) were obtained from each paradigm. Despite that, a common principle that emerged from both experiments was that choice is driven by peaks in the stimulus.

This led to the development of a salience-driven (rank-dependent) integration model whereby magnitudes are weighted by their local ranks. The exact form of this rank-dependency differed across the two experiments (sections 5.2.2 and 5.2.3). In particular, in the perceptual experiment, only the momentarily maximum option was subject to extra boosting while in the value-based experiment all options were boosted, proportionally to their rankings. This difference can be attributed to the nature of the stimulus; quickly rank-ordering the options is much easier with symbolic, numerical quantities as opposed to noisy sensory stimulus. Although numerical values and sensory stimulus are processed by different neural circuits, what this thesis showed is that this processing might be governed by similar abstract principles (i.e. leaky, salience-driven integration); and this is conceivable because these simple principles are in a position to explain phenomena from simple, brightness discrimination tasks to multi-attribute choice problems (i.e. preference reversal) and risky choice.

7.3 Future Directions

Visual Attention and Decision-making

One central aspect of the computational framework developed in this thesis was that magnitudes are weighted by their salience (i.e. ranks). A signature of this differential weighting could be sought in systematic fluctuations of visual attention, with visual gazes being directed towards the momentarily better options. An alternative possibility is that the differential weighting is not applied on the input values, due to biased visual attention, but at a later stage of the decision process as an internal top-down boost (e.g. scoring high in the ranks might be an extra hint that this option is the best). A detailed

understanding of the nature of this differential weighting and the involvement of visual attention is expected to further refine the salience-driven model, leading also to novel predictions. Furthermore, the role of attention is under-explored in sequential sampling tasks with humans. There, participants are assumed to fixate their gazes at the centre of the screen (which is not always controlled for) and to sample equally from all available options. However, it is likely that, during the course of a decision, especially with more than two alternatives, preference states feed-back towards the input stage, biasing the sampling process and facilitating a “winner takes all” effect. As Krajbich and Rangel (2011) showed, incorporating attentional parameters into classical sequential sampling models can provide a more complete account of decision behaviour.

Emergence of Ranking Structure in Dynamic Stimulus

One important difference found in this thesis between perceptual and value-based decisions among three alternatives, concerned the shape of the rank-dependent weighting function. While in the value-based task all three options were differentially weighted, in the perceptual experiment only the momentarily maximum option was boosted while the second and third in rank were indistinguishable (in terms of weighting) from each other. Two alternative hypothesis can explain this difference. First, the symbolic representation of numerical values might have just facilitated the utilization of the rank ordering of all three alternatives. On the contrary, sensory stimulus (i.e. brightness) might not be appropriate (by nature due to different circuits that process it) for an immediate and rapid rank ordering. Alternatively, the use of the full rank-order of perceptual stimulus might have been hindered due to external reasons and in particular the very fast updating of the noise that was superimposed on top of the signal (in the experiments presented here that was every 13.3 ms). In order to disentangle these two hypotheses, one could attempt to slow down the noise fluctuations in a brightness discrimination task and probably make the length of the phases longer (so as to make the rank-order structure more apparent). If, after these changes, an analog of the attraction effect (which can be explained only by the use of the full ranking) is obtained then the second hypothesis would be supported. In that case, specifying under exact what conditions (i.e. how long does the stimuli need to be stable in order to create a symbolic representation of the full rank order of all options) secondary cues are used (ranks) would be the next step.

Microcomputations vs. Perceptual Representations and Task Framing

As shown in this thesis, the computational details at the decisional level of the choice process might induce distortions that override the representational non-linearities associated with numerical values. In particular, people, among two numerical sequences of equal mean, chose more often the one associated with higher variance. This might be the case with sequentially presented physical quantities, such as brightness, which are known to be also logarithmically compressed. For example, between two spots with equal normally distributed brightness, observers might be biased towards the one with the highest noise because it has more positive peaks. This was found to be the case in one of the conditions in the brightness discrimination task (i.e. compromise condition in section 4.2), although this experiment was conducted with three alternatives and temporal correlations among the options. The visual attention in such an experiment can be controlled not to freely fluctuate so as to test whether a higher preference for the more noisy option is due to biased attention or because of non-linearities at the accumulation level (e.g. the zero non-linearity of LCA naturally promotes alternatives with higher drift rate/ variance without assuming that attention is fluctuating between the options). If a variance-seeking bias is obtained in a brightness discrimination task, it would be interesting to see whether this pattern is reversed when observers are given the, logically equivalent, task of rejecting the spot with the lowest brightness, similar to the task-framing bias obtained in the value integration paradigm.

Experience-based Decisions

Experience-based decisions share many common properties with the fast value integration task. In both cases, value samples are experienced (presented directly or actively sampled) and a decision is made on the basis of the whole stream of these received samples. Therefore, the probabilities of risky prospects are not given explicitly but are experienced implicitly. One of the most prominent result in experience-based decisions, is that people tend to underweight small probabilities (or rare events) opposite to what happens in description-based risky choices (i.e. overweighting small probabilities; Hertwig & Erev, 2009). However, a recent study by Ungemach et al. (2009), although replicated most of the choice patterns encountered in the original study by Hertwig et al. (2004), failed to obtain an underestimation of the small probabilities. This finding in combination with the results of this thesis (section 6.7), opens the possibility to re-interpret the original results in Hertwig et al. (2004) under a different computational framework where there is no explicit underestimation of rare events but

rather a salience-driven accumulation of samples.

7.4 Conclusion

In this thesis, in a series of computational and experimental studies, I examined the cognitive processes that underlie the integration of decision-relevant information across time. The findings demonstrated that information integration is distorted by differential weighting applied on the more salient samples. Two factors were found to affect the salience of the samples: i) their temporal order (primacy or recency, depending on task contingencies) and ii) their local ordering in the decision context (with the direction of the latter source of distortion being modulated by the task framing). Furthermore, the findings revealed that a simple salience-based integration model accounts for classical decision paradoxes (temporal, risk and task framing biases as well as preference reversal) and characterizes the deliberative process employed in richer domains, such as multi-attribute choice problems with trade-offs or decisions under risk. The explanatory success of this simple, salience-driven computational account, underscores the possibility that the roots of several decision anomalies lie on the algorithmic details of the choice process, rather than on non-linearities at the representational level of decision-relevant quantities (e.g. values and probabilities). Overall, the results of this thesis conferred new insights into the microstructure of complex decision making, clarifying how people weigh evidence, reverse their preferences and deal with risk, and providing a unifying neurocognitive foundation for decision making.

References

- Akaike, H. (1974). A new look at the statistical model identification. *Automatic Control, IEEE Transactions on*, 19(6), 716–723.
- Ariely, D., Huber, J., & Wertenbroch, K. (2005). When do losses loom larger than gains? *Journal of Marketing Research*, 134–138.
- Armel, K., Beaumel, A., & Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, 3(5), 396–403.
- Audley, R. (1960). A stochastic model for individual choice behavior. *Psychological Review*, 67(1), 1.
- Barnard, G. (1946). Sequential tests in industrial statistics. *Supplement to the Journal of the Royal Statistical Society*, 8(1), 1–26.
- Barth, H., La Mont, K., Lipton, J., Dehaene, S., Kanwisher, N., & Spelke, E. (2006). Non-symbolic arithmetic in adults and young children. *Cognition*, 98(3), 199–222.
- Bechara, A., Damasio, H., Tranel, D., & Damasio, A. (2005). The iowa gambling task and the somatic marker hypothesis: some questions and answers. *Trends in Cognitive Sciences*, 9(4), 159–162.
- Benartzi, S., & Thaler, R. (1998). Illusory diversification and retirement savings. *Unpublished manuscript. University of Chicago and UCLA.*
- Birnbaum, M. (1992). Violations of monotonicity and contextual effects in choice-based certainty equivalents. *Psychological Science*, 3(5), 310.
- Birnbaum, M. (2008). New paradoxes of risky decision making. *Psychological Review*, 115(2), 463.
- Blavatsky, P. (2007). Stochastic expected utility theory. *Journal of Risk and Uncertainty*, 34(3), 259–286.
- Bogacz, R. (2007). Optimal decision-making theories: linking neurobiology with behaviour. *Trends in Cognitive Sciences*, 11(3), 118–125.

- Bogacz, R. (2009). Optimal decision-making theories. In J. Dreher & L. Tremblay (Eds.), *Handbook of reward and decision making* (pp. 375–397). Academic Press Orlando, FL.
- Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700.
- Bogacz, R., & Cohen, J. (2004). Parameterization of connectionist models. *Behavior Research Methods*, *36*(4), 732–741.
- Bogacz, R., & Gurney, K. (2007). The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural computation*, *19*(2), 442–477.
- Bogacz, R., Usher, M., Zhang, J., & McClelland, J. (2007). Extending a biologically inspired model of choice: multi-alternatives, nonlinearity and value-based multidimensional choice. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1485), 1655.
- Brandstätter, E., Gigerenzer, G., & Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review*, *113*(2), 409.
- Britten, K., Shadlen, M., Newsome, W., & Movshon, J. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, *12*(12), 4745.
- Britten, K., Shadlen, M., Newsome, W., & Movshon, J. (1993). Responses of neurons in macaque mt to stochastic motion signals. *Visual Neuroscience*, *10*, 1157–1157.
- Brown, S., & Heathcote, A. (2005). Practice increases the efficiency of evidence accumulation in perceptual choice. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(2), 289.
- Brown, S., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178.
- Busemeyer, J., & Johnson, J. (2004). Computational models of decision making. *Handbook of judgment and decision making*, 133–154.
- Busemeyer, J., & Townsend, J. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, *100*(3), 432–459.
- Caspi, A., Beutter, B., & Eckstein, M. (2004). The time course of visual information accrual guiding eye movement decisions. *Proceedings of the National Academy*

- of Sciences of the United States of America*, 101(35), 13086.
- Chater, N. (2001). How smart can simple heuristics be? *Behavioral and Brain Sciences*, 23(05), 745–746.
- Choplin, J., & Hummel, J. (2005). Comparison-induced decoy effects. *Memory & Cognition*, 33(2), 332–343.
- Churchland, A., Kiani, R., & Shadlen, M. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6), 693.
- Debreu, G. (1960). *Individual choice behavior: A theoretical analysis* (Vol. 50) (No. 1). JSTOR.
- De Martino, B., Kumaran, D., Seymour, B., & Dolan, R. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, 313(5787), 684.
- Dhar, R., & Glazer, R. (1996). Similarity in context: Cognitive representation and violation of preference and perceptual invariance in consumer choice. *Organizational Behavior and Human Decision Processes*, 67, 280–293.
- Dhar, R., Nowlis, S., & Sherman, S. (2000). Trying hard or hardly trying: An analysis of context effects in choice. *Journal of Consumer Psychology*, 9(4), 189–200.
- Ditterich, J., Mazurek, M., & Shadlen, M. (2003). Microstimulation of visual cortex affects the speed of perceptual decisions. *Nature Neuroscience*, 6(8), 891–898.
- Donner, T., Siegel, M., Fries, P., & Engel, A. (2009). Buildup of choice-predictive activity in human motor cortex during perceptual decision making. *Current Biology*, 19(18), 1581–1585.
- Furnham, A. (1986). The robustness of the recency effect: Studies using legal evidence. *The Journal of General Psychology*, 113(4), 351–357.
- Garner, W. (1953). An informational analysis of absolute judgments of loudness. *Journal of Experimental Psychology*, 46(5), 373.
- Gerstenberg, T., Lagnado, D. A., Speekenbrink, M., & Cheung, C. (2011). Rational order effects in responsibility attributions. In C. Carlson, C. Hölscher, & T. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*. (pp. 1715–1720). Austin, TX: Cognitive Science Society.
- Gigerenzer, G. (1991). How to make cognitive illusions disappear: Beyond heuristics and biases. *European Review of Social Psychology*, 2(1), 83–115.
- Gigerenzer, G. (2006). Bounded and rational. *Contemporary Debates in Cognitive Science*, 115–133.
- Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102(4), 684.

- Gilovich, T., Griffin, D., & Kahneman, D. (2002). *Heuristics and biases: The psychology of intuitive judgement*. Cambridge Univ Press.
- Glöckner, A., & Herbold, A. (2011). An eye-tracking study on information processing in risky decisions: Evidence for compensatory strategies based on automatic processes. *Journal of Behavioral Decision Making*, *24*(1), 71–98.
- Gold, J., & Shadlen, M. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences*, *5*(1), 10–16.
- Gold, J., & Shadlen, M. (2002). Banburismus and the brain:: Decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, *36*(2), 299–308.
- Gold, J., & Shadlen, M. (2003). The influence of behavioral context on the representation of a perceptual decision in developing oculomotor commands. *The Journal of neuroscience*, *23*(2), 632.
- Gold, J., & Shadlen, M. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.*, *30*, 535–574.
- Guo, F., & Holyoak, K. (2002). Understanding similarity in choice behavior: A connectionist model. In *Proceedings of the twenty-fourth annual conference of the cognitive science society* (pp. 393–398).
- Hadar, L., & Fox, C. (2009). Information asymmetry in decision from description versus decision from experience. *Judgment and Decision Making*, *4*(4), 317–325.
- Hanks, T., Ditterich, J., & Shadlen, M. (2006). Microstimulation of macaque area lip affects decision-making in a motion discrimination task. *Nature neuroscience*, *9*(5), 682–689.
- Harvey, N. (1995). Why are judgments less consistent in less predictable task situations? *Organizational Behavior and Human Decision Processes*, *63*(3), 247–263.
- Harvey, N. (2007). Use of heuristics: Insights from forecasting research. *Thinking & reasoning*, *13*(1), 5–24.
- Hau, R., Pleskac, T., Kiefer, J., & Hertwig, R. (2008). The description–experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*(5), 493–518.
- Hertwig, R., Barron, G., Weber, E., & Erev, I. (2004). Decisions from experience and the effect of rare events in risky choice. *Psychological Science*, *15*(8), 534.
- Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

- Hochman, G., & Yechiam, E. (2011). Loss aversion in the eye and in the heart: The autonomic nervous system's responses to losses. *Journal of Behavioral Decision Making*, 24(2), 140–156.
- Hogarth, R., & Einhorn, H. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive Psychology*, 24(1), 1–55.
- Holland, M., & Lockhead, G. (1968). Sequential effects in absolute judgments of loudness. *Attention, Perception, & Psychophysics*, 3(6), 409–414.
- Holyoak, K., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General*, 128(1), 3.
- Hotaling, J. M., Busemeyer, J. R., & Li, J. (2010). Theoretical developments in decision field theory: Comment on Tsetsos, Usher, and Chater (2010). *Psychological review*, 117(4), 1294–1298.
- Huber, J., Payne, J., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, 90–98.
- Huk, A., & Shadlen, M. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *The Journal of neuroscience*, 25(45), 10420.
- Hurly, T., & Oseen, M. (1999). Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal behaviour*, 58(1), 59–66.
- Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision Making Under Risk. *Econometrica*, 47(2), 263–291.
- Kahneman, D., & Tversky, A. (2000). *Choices, values, and frames*. Cambridge University Press.
- Kiani, R., Hanks, T., & Shadlen, M. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *The Journal of Neuroscience*, 28(12), 3017.
- Knetsch, J. (1989). The endowment effect and evidence of nonreversible indifference curves. *The American Economic Review*, 79(5), 1277–1284.
- Krajbich, I., & Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences*, 108(33), 13852–13857.
- Kurniawan, I., Seymour, B., Vlaev, I., Trommershäuser, J., Dolan, R., & Chater, N. (2010). Pain relativity in motor control. *Psychological science*, 21(6), 840.

- Lagnado, D., & Channon, S. (2008). Judgments of cause and blame: The effects of intentionality and foreseeability. *Cognition*, *108*(3), 754–770.
- Laming, D. (1968). *Information theory of choice-reaction times*. New York: Wiley.
- Laming, D. (1997). *The measurement of sensation* (Vol. 30). Oxford University Press, USA.
- Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America*, *49*(2), 467–477.
- Levy, I., Snell, J., Nelson, A., Rustichini, A., & Glimcher, P. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of neurophysiology*, *103*(2), 1036.
- Link, S., & Heath, R. (1975). A sequential theory of psychological discrimination. *Psychometrika*, *40*(1), 77–105.
- Lockhead, G., & King, M. (1983). A memory model of sequential effects in scaling tasks. *Journal of Experimental Psychology: Human Perception and Performance*, *9*(3), 461.
- Louie, K., Grattan, L., & Glimcher, P. (2011). Reward value-based gain control: divisive normalization in parietal cortex. *The Journal of Neuroscience*, *31*(29), 10627.
- Luce, R., Nosofsky, R., Green, D., & Smith, A. (1982). The bow and sequential effects in absolute identification. *Attention, Perception, & Psychophysics*, *32*(5), 397–408.
- Ludvig, E., & Spetch, M. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PloS one*, *6*(6), e20262.
- Ludwig, C., Gilchrist, I., McSorley, E., & Baddeley, R. (2005). The temporal impulse response underlying saccadic decisions. *The Journal of neuroscience*, *25*(43), 9907.
- Massaro, D., & Anderson, N. (1971). Judgmental model of the ebbinghaus illusion. *Journal of Experimental Psychology*, *89*(1), 147.
- Maylor, E., & Roberts, M. (2007). Similarity and attraction effects in episodic memory judgments. *Cognition*, *105*(3), 715–723.
- McKenzie, C., Lee, S., & Chen, K. (2002). When negative evidence increases confidence: Change in belief after hearing two sides of a dispute. *Journal of Behavioral Decision Making*, *15*(1), 1–18.
- McMillen, T., & Holmes, P. (2006). The dynamics of choice among multiple alternatives. *Journal of Mathematical Psychology*, *50*(1), 30–57.

- Newell, B., Wong, K., Cheung, J., & Rakow, T. (2009). Think, blink or sleep on it? the impact of modes of thought on complex decision making. *The Quarterly Journal of Experimental Psychology*, 62(4), 707–732.
- Nieder, A., & Miller, E. (2003). Coding of cognitive magnitude:: Compressed scaling of numerical information in the primate prefrontal cortex. *Neuron*, 37(1), 149–157.
- Niwa, M., & Ditterich, J. (2008). Perceptual decisions between multiple directions of visual motion. *The Journal of Neuroscience*, 28(17), 4435.
- Oaksford, M., & Chater, N. (1995). Theories of reasoning and the computational explanation of everyday inference. *Thinking and Reasoning*, 1, 121–152.
- Oaksford, M., & Chater, N. (1998). *Rationality in an uncertain world: Essays on the cognitive science of human reasoning*. Hove, Sussex: Psychology Press.
- Padoa-Schioppa, C., & Assad, J. (2007). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nature neuroscience*, 11(1), 95–102.
- Parducci, A. (1965). Category judgment: A range-frequency model. *Psychological Review*, 72(6), 407.
- Pennington, N., & Hastie, R. (1993). Reasoning in explanation-based decision making. *Cognition*, 49(1-2), 123–163.
- Pettibone, J., & Wedell, D. (2000). Examining models of nondominated decoy effects across judgment and choice. *Organizational Behavior and Human Decision Processes*, 81(2), 300–328.
- Pettibone, J., & Wedell, D. (2007). Testing alternative explanations of phantom decoy effects. *Journal of Behavioral Decision Making*, 20(3), 323–341.
- Pietsch, A., & Vickers, D. (1997). Memory capacity and intelligence: Novel techniques for evaluating rival models of a fundamental information-processing mechanism. *The Journal of general psychology*, 124(3), 229–339.
- Pratkanis, A., & Farquhar, P. (1992). A brief history of research on phantom alternatives: Evidence for seven empirical generalizations about phantoms. *Basic and Applied Social Psychology*, 13(1), 103–122.
- Raab, D. (1962). Statistical facilitation of simple reaction times. *Transactions of the New York Academy of Sciences*, 24, 574.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59.
- Ratcliff, R. (2006). Modeling response signal and response time data. *Cognitive Psychology*, 53(3), 195–237.

- Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology*, *90*(3), 1392.
- Ratcliff, R., Gomez, P., & McKoon, G. (2004). A diffusion model account of the lexical decision task. *Psychological Review*, *111*(1), 159.
- Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: Theory and data for two-choice decision tasks. *Neural computation*, *20*(4), 873–922.
- Ratcliff, R., & Rouder, J. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347.
- Ratcliff, R., & Rouder, J. (2000). A diffusion model account of masking in two-choice letter identification. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(1), 127.
- Ratcliff, R., & Smith, P. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological review*, *111*(2), 333.
- Rayner, K. (1978). Eye movements in reading and information processing. *Psychological Bulletin*, *85*(3), 618.
- Read, D., & Loewenstein, G. (1995). Diversification bias: Explaining the discrepancy in variety seeking between combined and separated choices. *Journal of Experimental Psychology: Applied*, *1*(1), 34.
- Reimers, S., & Harvey, N. (2011). Sensitivity to autocorrelation in judgmental time series forecasting. *International Journal of Forecasting*, *27*(4), 1196–1214.
- Reutskaja, E., Nagel, R., Camerer, C., & Rangel, A. (2011). Search dynamics in consumer choice under time pressure: An eye-tracking study. *The American Economic Review*, *101*(2), 900–926.
- Roe, R., Busemeyer, J., & Townsend, J. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, *108*(2), 370–392.
- Roitman, J., & Shadlen, M. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of neuroscience*, *22*(21), 9475.
- Rorie, A., Gao, J., McClelland, J., & Newsome, W. (2010). Integration of sensory and reward information during perceptual decision-making in lateral intraparietal cortex (lip) of the macaque monkey. *PloS one*, *5*(2), e9308.
- Russo, J., & Rosen, L. (1975). An eye fixation analysis of multialternative choice. *Memory & Cognition*, *3*(3), 267–276.

- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 461–464.
- Shadlen, M., & Newsome, W. (2001). Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of Neurophysiology*, 86(4), 1916.
- Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory & Cognition*, 21(4), 546–556.
- Shafir, E., Simonson, I., & Tversky, A. (1993). Reason-based choice. *Cognition*, 49(1-2), 11–36.
- Shafir, S., Waite, T., & Smith, B. (2002). Context-dependent violations of rational choice in honeybees (*apis mellifera*) and gray jays (*perisoreus canadensis*). *Behavioral Ecology and Sociobiology*, 51(2), 180–187.
- Shanks, D., Tunney, R., & McCarthy, J. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, 15(3), 233–250.
- Simon, H. (1982). *Models of bounded rationality. vol. 2, behavioral economics and business organization*. MIT Press.
- Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of consumer research*, 158–174.
- Smith, P., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, 27(3), 161–168.
- Stevens, S. (1975). *Psychophysics: Introduction to its perceptual, neural, and social prospects*. Transaction Publishers.
- Stewart, N. (2009). Decision by sampling: The role of the decision environment in risky choice. *The Quarterly Journal of Experimental Psychology*, 62(6), 1041–1062.
- Stewart, N., Brown, G., & Chater, N. (2005). Absolute identification by relative judgment. *Psychological Review*, 112(4), 881.
- Stewart, N., Chater, N., & Brown, G. (2006). Decision by sampling. *Cognitive Psychology*, 53(1), 1–26.
- Stewart, N., Chater, N., Stott, H., & Reimers, S. (2003). Prospect relativity: How choice options influence decision under risk. *Journal of Experimental Psychology: General*, 132(1), 23.
- Stewart, N., & Simpson, K. (2008). A decision-by-sampling account of decision under risk. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for*

- bayesian cognitive science* (pp. 261–276). Oxford University Press, Oxford, UK.
- Stone, M. (1960). Models for choice-reaction time. *Psychometrika*, 25(3), 251–260.
- Summerfield, C., Behrens, T., & Koechlin, E. (2011). Perceptual classification in a rapidly changing environment. *Neuron*, 71(4), 725–736.
- Summerfield, C., & Koechlin, E. (2010). Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. *Frontiers in Human Neuroscience*, 4.
- Swensson, R. (1972). The elusive tradeoff: Speed vs accuracy in visual discrimination tasks. *Attention, Perception, & Psychophysics*, 12(1), 16–32.
- Todd, P., & Gigerenzer, G. (2000). Précis of simple heuristics that make us smart. *Behavioral and brain sciences*, 23(5), 727–741.
- Tom, S., Fox, C., Trepel, C., & Poldrack, R. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, 315(5811), 515.
- Townsend, J., & Nozawa, G. (1995). Spatio-temporal properties of elementary perception: An investigation of parallel, serial, and coactive theories. *Journal of Mathematical Psychology*, 39(4), 321–359.
- Trepel, C., Fox, C., & Poldrack, R. (2005). Prospect theory on the brain? toward a cognitive neuroscience of decision under risk. *Cognitive Brain Research*, 23(1), 34–50.
- Trueblood, J., & Busemeyer, J. (2010). A comparison of the belief-adjustment model and the quantum inference model as explanations of order effects in human inference. In *Cogsci 2010 the annual meeting of the cognitive science society* (pp. 1166–1171).
- Tsetsos, K., Gao, J., Usher, M., & McClelland, J. (2011). *Using time-varying evidence to probe decision dynamics*. (In preparation, for submission to the Frontiers in Cognitive Science and Frontiers in Neuroscience Special Issue: “Dynamics of decision making: from evidence to preference and belief”.)
- Tsetsos, K., Usher, M., & Chater, N. (2010). Preference reversal in multiattribute choice. *Psychological review*, 117(4), 1275–1291.
- Tsetsos, K., Usher, M., & McClelland, J. (2011). Testing multi-alternative decision models with non-stationary evidence. *Frontiers in neuroscience*, 5.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological review*, 79(4), 281.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology

- of choice. *Science*, 211(4481), 453.
- Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of business*, 251–278.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4), 297–323.
- Tversky, A., & Simonson, I. (1993). Context-dependent preferences. *Management Science*, 1179–1189.
- Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, 20(4), 473.
- Ungemach, C., Stewart, N., & Reimers, S. (2011). How incidental values from the environment affect decisions about money, risk, and delay. *Psychological Science*, 22(2), 253.
- Usher, M., Elhalal, A., & McClelland, J. (2008). The neurodynamics of choice, value-based decisions, and preference reversal. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for bayesian cognitive science* (pp. 277–302). Oxford University Press, Oxford, UK.
- Usher, M., & McClelland, J. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, 108(3), 550.
- Usher, M., & McClelland, J. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. *Psychological Review*, 111(3), 757–769.
- Usher, M., Olami, Z., & McClelland, J. (2002). Hick's law in a stochastic race model with speed-accuracy tradeoff. *Journal of Mathematical Psychology*, 46(6), 704–715.
- Usher, M., Tsetsos, K., & Chater, N. (2010). Postscript: Contrasting predictions for preference reversal. *Psychological review*, 117(4), 1291–1293.
- Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, 13(1), 37–58.
- Vickers, D. (1979). *Decision processes in visual perception*. Academic Press.
- Vlaev, I., Chater, N., Stewart, N., & Brown, G. D. (2011). Does the brain calculate value? *Trends in Cognitive Sciences*, 15(11), 546 - 554.
- Vlaev, I., Seymour, B., Dolan, R., & Chater, N. (2009). The price of pain and the value of suffering. *Psychological science*, 20(3), 309.
- Von Neuman, J., & Morgenstern, O. (1947). *Theory of games and economic behavior*. Princeton University Press.

- Wald, A. (1947). *Sequential analysis*. Dover Publications.
- Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics*, 326–339.
- Walker, L., Thibaut, J., & Andreoli, V. (1972). Order of presentation at trial. *The Yale Law Journal*, 82(2), 216–226.
- Wong, K., Huk, A., Shadlen, M., & Wang, X. (2007). Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Frontiers in Computational Neuroscience*, 1.
- Yu, A., & Dayan, P. (2005). Uncertainty, neuromodulation, and attention. *Neuron*, 46(4), 681–692.