THE EFFECT OF CONTEXT ON HUMAN DECISION-MAKING DURING INFORMATION INTEGRATION

SANTIAGO HERCE CASTAÑÓN

Lincoln College

Department of Oxford Neuroscience

University of Oxford

This thesis is submitted for the degree of Doctor of Philosophy in Neuroscience

June 2022

How I dream of a world where so much love, so much freedom and so much opportunity as I had, are not a privilege but a guarantee for everyone.

ABSTRACT

Perceiving and interacting with our world is an amazingly complicated computational problem. Our brains need to infer the best cause of action from a long chain of ambiguous and noisy signals reflecting properties from the environment. To solve this problem, our brains make use of multiple pieces of information available in a given context. In these thesis, we explore the influence of context on decision-making. First, we focus on decisions made in the context were three options are available to be chosen. Second, we focus on decisions made in the context of variable prior information about the probability of occurrence of a given event. A particular emphasis is placed on distinguishing mechanisms that optimize the accuracy of decisions from those that comply with axiomatic prescriptions of rational behaviour. We show that even as simple an assumption as having Gaussian noise distorting the representation of value is enough to drive irrational behaviour, calling into question the validity of the axioms of rationality as a yardstick for measuring human behaviour. We also show that uncertainty arising at different stages of the long chain of information processing, can drastically impact the nature of the inference problem that humans have to solve in different contexts. We show two mechanisms present when humans construct their value preferences between three available options. These mechanisms prioritize what information gets further processed in a way that minimizes the effects of noise arising at later stages of processing. Finally, we provide a novel computational explanation for the mismatch typically observed between near-optimal perceptual choices and suboptimal cognitive judgments. Namely, that people are metacognitive blind to noise arising when integrating of multiple pieces of evidence. Together, we show that the brain tends to make excellent use of contextual information to guide it decisions. Perhaps so much so, that we often trust our judgments blindly.

ACKNOWLEDGEMENTS

This thesis is the result of an incredibly long list of privileges I have had in my life thus far.

I feel incredibly privileged to have had the chance to spend the last five years in Oxford, living and working around a high concentration of passionate people and investigating the phenomena that puzzle me the most.

I was lucky enough to be awarded with the generous financial support from two great organizations. First, the Wellcome Trust, which I feel very proud to belong to and share the ethical commitment to science with. And Oxford University Press, from whom I have had access to some of my favourite pieces of reading. Thank you both because without your support the costs of studying here would have been prohibitive. I also want to thank Tessa and Guy for proofreading my application for the PhD, thank you for helping me put my passions in such beautiful words. Also to Margaret Esiri, thank you for your generous time on the day of my interview, you gave me all the encouragement I needed on that day. And Debbie and Andy, for always keeping an eye on me and on my PhD programme in general.

I want to thank all the great scientists I crossed paths with and collaborated during these five years. A lot of my PhD I owe to my supervisors. Starting with Chris, I cannot think of a supervisor with whom I could share greatest passion for the study of cognition. Thank you for allowing me the privilege of your supervision and of being a member of your lab. I am most grateful for the independence you gave me for focusing on the questions I liked the most, and for letting me choose the time and place at which I could work. I know it required supra-human patience to supervise me, but you were simply great. And to Masud, from the beginning you gave me all the freedom to spend most of my time at Chris' lab but opened the doors of your lab to me, I learnt a lot from coming to your lab meetings. It was a privilege to receive your wisdom and advice in both for academic and personal matters. There are so many great qualities from both of you that I hope I can one day possess.

My work in this thesis benefitted indirectly from many people in the department and elsewhere, thank you all from the Summerfield and the Husain labs. There were also many direct contributors to my work who, with endless hours of effort, allowed my PhD to be much better than I could have done it alone, thank you: Kostis, Sam, Alejo, Vickie, Jan, Howard, and Rani, your scientific input to all my projects was immense. Interacting with you made me much better investigator than I would have been otherwise.

I left four other people to mention separately, for my interactions with them were crucial for my PhD. Dan, who was a superb collaborator and spent many hours working with me on projects. I miss our fierce discussions over coffee(s) that would invariably end (often after several hours or days) with us realising we were actually defending the same point, or that we had already had that discussion previously. It really is a privilege to have a friend with whom to collaborate so productively. Hil, it was a privilege to traverse our PhDs alongside one-another stage by stage: I do not know if I would have been able to make it without your daily support (academic and personal) and without our epic squash matches. Thank you for being such a good friend and confidant. Jacqueline, you brought organization and joy to my PhD, you also made the collection of EEG data go from a "bit terrible" to a most enjoyable activity. I am lucky to have been the lab member who benefitted the most from your incredible empathy and brilliance, it is just a shame your table tennis was simply not up to standards. Andrés, in our discussions, you always pushed me to the limits of what I could explain. Discussing with you is one of my greatest privileges in science.

I also want to thank my thesis assessors, Matthew Rushworth and Nick Chater. Thank you for agreeing to read my thesis and hold my viva, I feel very privileged to have scientists of your quality generously contributing time and thoughts to my PhD.

Needless to say, my privileges date back way further than my time in the UK. I feel like the most privileged person in the world from having the luck of being born in a loving family: where the only obligation I ever had was that of pursuing my passions. I mainly have my mom to thank for that and for making sure I never lacked the most essential things in life. Thank you for showing me how to be an independent and strong person. I also want to thank Leonardo, from whom I learnt to nurture my intellect while taking school less seriously. Lety, for always making sure I had a second home. I know no one with so much love to give as you, and I am lucky to have been the most privileged recipient of it. Juan, I have always looked up to your intellect, humbleness and compassion, thank you for being such a good model to follow and for bringing Quillén into my life. María, my biggest fan and my greatest idol, from the time you arrived in my life you have been my best accomplice in life. Luis, for being my best teacher of how not to take life too seriously, and look at things with joy. And Pablito, my only oaxacan family member in the UK, it has been a great privilege to always have you by my side as we grew into adulthood. I also want to thank my friends, including the ones from Oaxaca, from UNAM, from my MSc, and from Oxford and the UK, some of my most important moments in life have been by your sides. A special mention must be given to all my MSc friends with whom we built very strong family bonds. Luki, Emmita, Annita, Ayesha and Sebas, I will always remember fondly the improvised dinners with sleepovers we had at the EPA, you became my Oxford family. I have received more support and love from Los Castañón, Los Pino, Los Herce, Los Somohano and all my friends, than any person can wish for. Thank you all for supporting me when around you or for being understanding of my long periods of absence from your lives.

My last mention of great privilege is to have had Ornela in my life for the last ten years. Thank you for your incommensurable contribution to my life as a PhD student and beyond, you are the most wonderful life companion imaginable.

To you all I owe everything.

Contents

1 INTRODUCTION 1

1.1 The problem of perception 1

1.2 Aim of the thesis 3

1.3 Combining information to resolve ambiguity 3

1.4 The brain as a Bayesian inference machine 5

1.5 Bayesian observers and generative models of the world 8

1.6 Beliefs about the world, decisions and metacognition 10

1.7 Sequential sampling and evidence accumulation 12

1.8 Rationality and Optimality 14

1.9 Taxonomy of uncertainty in decision-making 19

1.10 Brief summary of the following chapters 22

2 SIMULATING THE DISTRACTOR EFFECT IN MULTIALTERNATIVE DECISIONS 33

2.1 Introduction 33

2.2 Methods 36

2.3 Results 39

2.4 Discussion 46

3 DECISIONS WITH THREE OPTIONS IN A SEQUENTIAL INTEGRATION PARADIGM 50

3.1 Introduction 50

3.2 Methods 54

3.2.1 Task 54

3.2.2 Efficiency boxes 59

3.2.3 Regression models 59

3.2.4 Model fitting 61

3.2.5 Eye-tracking data collection and pre-processing 64

3.2.6 Pupillometric analyses 65

3.2.7 Electroencephalography data: collection and pre-processing 68

3.2.8 Single trial EEG decoding analyses 70

3.2.9 Cluster correction for multiple comparisons 71

3.2.10 Analysing protection against noise from models 71

3.3 Results 73

3.3.1 A sequential integration paradigm with three options available options. 73

3.3.2 Within-sample interactions of evidence across the three options................ 75
3.3.3 Modelling of behaviour. 77
3.3.4 Quantitative comparison of models.. 80
3.3.5 Evidence for Selective Integration from pupil dilation 86
3.3.6 Evidence for Selective Integration and Power transformation from EEG signals 88
3.3.7 Independent effects of Selective Integration and Power transformation 90
3.3.8 Can we look for these independent effects in Pupil dilation?....................... 91
3.3.9 Robustness against noise from Selective Integration and Power transformation 93
3.4 DISCUSSION 94
4 NOISE BLINDNESS AS THE SOURCE OF SUBOPTIMALITY 107
4.1 INTRODUCTION 107
4.2 METHODS 109
4.2.1 Participants and equipment.. 109
4.2.2 Short description of Task and stimuli... 110
4.2.3 Further details of task and stimuli 113
4.2.4 Computational model 118
4.2.5 Noise estimation 123
4.2.6 Drift diffusion model 124
4.2.7 Analyses 125
4.3 RESULTS.. 126
4.3.1 Matched costs to performance from encoding and integration noises 126
4.3.2 Are people optimal in using prior information under both sources of noise? 127
4.3.3 Are people metacognitively blind to integration noise?............................ 129
4.3.4 Noise blindness for an ideal observer model.. 132
4.3.5 Is noise blindness the simplest explanation to the observed suboptimal behaviour? 134
4.3.6 Neither a drift diffusion model or a variability blind model capture human behaviour.. 136
4.3.7 Could people simply be sampling a subset of the stimuli? 138
4.4 DISCUSSION 139
5 DISCUSSION.. 149

LIST OF TABLES

TABLE 1. PROPORTION OF CHOICES AND RESPONSE TIMES 103

TABLE 2. REGRESSION WITH VALUE OF DISTRACTOR. 103

TABLE 3. REGRESSION WITH NORMALIZED VALUE OF DISTRACTOR. 104

TABLE 4. REGRESSION WITH MOMENTARY INTERACTION PREDICTORS. 104

TABLE 5. SIGNIFICANCE (P-VALUES) FOR HUMAN AND MODEL REGRESSION COEFFICIENT ON CHOICES (WITH MOMENTARY INTERACTION PREDICTORS). 105

TABLE 6. SIGNIFICANCE (P-VALUES) FOR HUMAN AND MODEL REGRESSION COEFFICIENT ON CHOICES (WITH DISTANCE AND VARIABILITY PREDICTORS). 106

LIST OF FIGURES

FIGURE 1.1. AMBIGUITY IN SENSORY SIGNALS. 2

FIGURE 1.2. COMBINING INFORMATION TO SOLVE AMBIGUITY. 4

FIGURE 1.3. LIGHT FROM ABOVE PRIOR. 5

FIGURE 1.4. PREDICTIVE CODING. 8

FIGURE 1.5. SIGNAL DETECTION THEORY AND THE BRAIN. 9

FIGURE 1.6. DRIFT DIFFUSION MODEL. 13

FIGURE 2.1. DECOY EFFECTS. 34

FIGURE 2.2. VIOLATION OF RATIONALITY BY A SIMPLE NOISY AGENT. 40

FIGURE 2.3. WEAK VIOLATION FROM SIMPLE NOISY AGENT. 41

FIGURE 2.4. EFFICIENCY UNDER DIVISIVE NORMALIZATION. 42

FIGURE 2.5. STRONG VIOLATION FOR DIVISIVE NORMALIZATION. 42

FIGURE 2.6. A STRONG VIOLATION FROM SELECTIVE INTEGRATION. 44

FIGURE 2.7 STRONG VIOLATION OF RATIONALITY WHEN ASSUMING LAPSES. 45

FIGURE 3.1. TRIAL SEQUENCE. 56

FIGURE 3.2. AVERAGED TIME SERIES OF EYE TRACKING SIGNALS. 67

FIGURE 3.3. ELECTRODE LOCATIONS. 69

FIGURE 3.4. NO CLASSICAL DISTRACTER EFFECT IN HUMAN RESPONSES. 74

FIGURE 3.5. WITHIN SAMPLE INTERACTIONS OF EVIDENCE AFFECT CHOICES. 76

FIGURE 3.6. BEST FITTING PARAMETER SETS. 79

FIGURE 3.7. QUANTITATIVE MODEL COMPARISON. 81

FIGURE 3.8. QUALITATIVE PREDICTIONS FOR LOGISTIC REGRESSION. 83

FIGURE 3.9. INDEPENDENT EFFECTS FOR DISTANCE AND VARIABILITY. 85

FIGURE 3.10. DECODING RANK DEPENDENT VALUE FROM PUPIL SIZE. 87

FIGURE 3.11. REGRESSION COEFFICIENTS OF CONVOLUTION GLM. 88

FIGURE 3.12. DECODING OF RANK DEPENDENT VALUE FROM EEG SIGNALS. 89

FIGURE 3.13. NON-LINEAR DECODING OF DIGITS FROM EEG SIGNALS. 90

Figure 3.14. Competitive regression of rank dependent value and digit presence on EEG. 91

Figure 3.15. Competitive regression of rank dependent value and quantile presence on Pupil Dilation. 92

Figure 3.16. Protection against noise from Selective Integration and Power transformations. 94

Figure 4.1. Basic experimental paradigm. 111

Figure 4.2. Variations of the paradigm. 112

Figure 4.3. Distribution of orientations and noise. 118

Figure 4.4. Condition specific density functions. 119

Figure 4.5. Prior knowledge and blindness to noise. 120

Figure 4.6. Behavioural effects of stimulus contrast and variability on choices. 127

Figure 4.7. The influence of prior information. 128

Figure 4.8. Effects of stimulus contrast and variability on confidence. 130

Figure 4.9. Overconfidence and proportion of trials opted-in. 131

Figure 4.10. Comparison of model and human behaviour. 133

Figure 4.11. Behavioural results of experiments 4 and 5. 135

Figure 4.12. Drift Diffusion Model cannot account for data. 137

Figure 4.13. Participants are unlikely to engage in subsampling. 139

List of Appendices

Tables and appendices 103

1 INTRODUCTION

"Behold! human beings living in an underground cave […], here they have been from their childhood, and have their legs and necks chained so that they cannot move […]. Above and behind them a fire is blazing at a distance, and between the fire and the prisoners there is a raised way; and you will see, if you look, a low wall built along the way, like the screen which marionette players have in front of them, over which they show the puppets. [… To the prisoners], the truth would be literally nothing but the shadows of the images."

–Plato, Allegory of the cave, 360 B.C.E

1.1 The problem of perception

Perception comes to us with great ease. When we wake up, is reality not naturally there for us to perceive it? Is it not that sounds are there to be heard, light is there to be seen, objects are there to be touched, flavours are there to be tasted and scents are there to be smelled? Probably no other subject of scientific enquiry is as deceptively complex as perception. Our brains evolved to allow our bodies to interact with their environments, their realities, and they do. Paradoxically, brains accomplish such an excellent job at solving the problem of perception that we intuitively fail to see there was a problem in the first place. What is the problem of perception?

The only complete representation of reality is reality itself. Thus, the brain simply cannot represent its environment as it is. Instead, the brain relies on some cues from its environment to obtain information about it to interact with it. Through a process called sensory transduction, the brain obtains sensory signals by "translating" physical cues from the environment into neuronal firing (Kandel et al., 2000). In vision, for instance, information about light sources, such as the number of photons and their wavelength, are gathered and transduced into the firing of retinal cells. These sensory signals then serve as cues themselves to help the brain discriminate some aspects of the environment, such as the colour and the intensity of a light source.

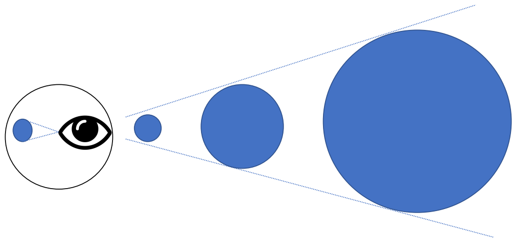

Figure 1.1. Ambiguity in sensory signals.

A single image projected to the retina at the back of the eye can be caused potentially by an infinity of world objects. In this example, the image of a blue ball, can be caused by different blue balls of different sizes at specific distances from the eye. Three such combinations of size and distance of balls are shown. Any balls that fit tightly in the same angle depicted by the blue lines would create the same image on the retina, i.e. an image of a ball that would subtend the same visual angle.

However, because sensory signals arise from incomplete information about the environment, perception is inherently a "guessing" process. Much in the way that the prisoners in Plato's "allegory of the cave" had to rely on the shadows projected at the bottom of the cave to make sense of their reality, the brain can do nothing more than to infer what causes the sensory signals evoked by the "images" that arrive to its retina.

The intrinsic ambiguity in the process of perception, however, has several sources. First, as suggested above, the incompleteness of the physical cues which are translated into sensory signals means that the same image that arrives to the retina can be caused by many different objects: for example, the image of a ball projected onto the retina can be caused by an infinity of different balls (Figure 1.1). Second, neuronal firing is a stochastic process, such that neuronal activity elicited by a unique image projected onto the retina, might differ from the activity evoked when presenting the exact same image in a different episode. Consequently, just as different objects can project the same image onto the retina, different images on the retina can evoke the same neuronal response, and those responses will evoke further responses (with their inherent ambiguity) down the chain of processing.

In Plato's "allegory of the cave", there was a long chain of events connecting the reality of the prisoners with the information they could access. Similarly, there is a long chain of events that connects our perception of the reality with the reality that generates it. Nevertheless, the brain eventually needs to come up with a good guess about which reality caused the current signals. This problem is an inverse inference problem: there is a causal chain of events going in one direction, where each event can be caused by several different events, and a chain of inferences going in the opposite direction (Helmholtz, 1925), where the aim is to infer the most likely ultimate cause of the events. So, the problem of perception is: how can we make so much sense of reality when all we have at hand is a long chain of ambiguous signals that arise from partial information about the environment?

1.2 Aim of the thesis

In this thesis, I explore how our brains are able to make sense of their reality by means of adapting to the context in which they receive their sensory information. I explored two main contexts in which humans and other animals often interact with their world. First, in cases where they need to combine information from multiple sources. Second, when the information they receive pertains to three discrete entities (options) that they have to choose from. Both types of context have been under-explored in studies on decision-making which typically have been concerned with one source of information and choice between two options. I studied these issues by concentrating on experimental paradigms that required participants to combine different pieces of information in order to make a decision. A special emphasis is placed throughout the thesis on describing uncertainty arising at different levels of the chain of processing. In the remaining sections of the Introduction, I will lay the conceptual groundwork for the rest of the thesis.

1.3 Combining information to resolve ambiguity

An important aspect for solving the inverse problem of perception is to combine multiple pieces of information. Why would the brain benefit from combining different pieces of information? Let us consider the example of the blue ball in Figure 1.1. The image that arrives to the retina of one eye is ambiguous as to the size and distance of the blue ball.

However, this ambiguity can be resolved, in principle, by simply combining the information that arrives to both eyes (Figure 1.2).

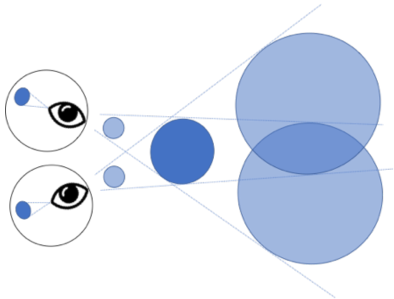

Figure 1.2. Combining information to solve ambiguity.

The size and distance of an object can be disambiguated by combining different pieces of information. For instance, by combining information about the object that arrives to the two eyes.

In order to disambiguate information, the brain often needs to combine not just two but several pieces of information. Take the example in Figure 1.2. In order to disambiguate the size and distance of the blue ball, the brain needs not only the information gathered from each eye; it also needs an estimate of the separation between the two eyes and the direction towards which each eye is looking at the time the information was gathered.

Sometimes it is not possible to disambiguate the cause of the sensory signals from the information gathered at a single point in time. Instead, the brain needs to combine information gathered at different points in time, or even at very different timescales, sometimes spanning the entire lifetime of the brain. For instance, when disambiguating which objects might have caused the images shown in Figure 1.3, our brains seem to consider one specific aspect of information, one which was acquired long before this image was presented, namely, that light sources tend to come from above (Sun and Perona, 1998). Using this *prior information* (prior, because it was acquired before), strongly biases the inference of the brain towards the cause of the image being two half spherical surfaces: i) a convex one on the left, a half sphere popping out of the surface being strongly lit at the top and casting shade on its bottom part, and ii) a concave one on the right, another half sphere but this time dipping into the surface, with the top part sheltered from the light source but the bottom part strongly lit.

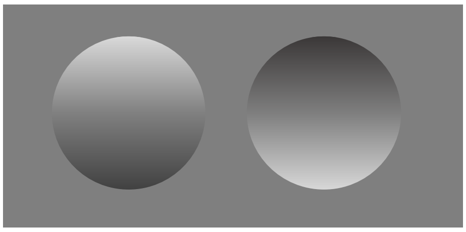

Figure 1.3. Light from above prior.

The image above is created by adding two circles that differ in the direction of the vertical gradient of grey shades, but are otherwise identical (i.e. light grey at the top and dark grey at the bottom or the other way around). When observing this image, however, most people tend to perceive a continuous grey surface with two half spherical surfaces, a convex one (popping out), on the left, and a concave one (dipping in) on the right. This inference seems to be guided by the prior acquisition of information about light sources usually coming from above (turning the page upside-down should flip your interpretation of the spheres).

We seem to come up with an almost instantaneous, albeit complex, interpretation of any image that arrives to our eyes, even when this requires the combination of several pieces of information. Think of how our visual interpretation of the world is typically three-dimensional even when images themselves have only two spatial dimensions. So, how does the brain construct these complex inferences? How does it know which sources of information to use at any time point and how much each can be trusted?

1.4 The brain as a Bayesian inference machine

The Bayesian inference framework offers answers to these questions by proposing that the brain uses probabilistic representations of information and that it combines different pieces of evidence consistent with the laws of probability. Central to this process is Bayes' theorem, which provides, in its most general conception, an optimal algorithm for combining probabilistic information: typically described in terms of obtaining a posterior probability distribution by combining a likelihood with a prior probability distribution. In what follows, I will provide a short review of the literature that supports the view of the brain as a Bayesian inference machine (Kersten et al., 2004; Knill and Richards, 1996; Ma et al., 2006).

Central to Bayesian inference, is the insight that, when combining different pieces of evidence, more weight should be given to evidence that is more reliable. In a canonical study (Ernst and Banks, 2002), an ingenious apparatus allowed the experimenters to control a visual stereographical stimulus (information independently projected to the two eyes making up a three-dimensional image of a bar) and a haptic (tactile) stimulus of a bar. Participants were instructed to simultaneously view and touch one “virtual bar” and then decide whether it was wider than a reference bar. Unbeknownst to participants, the tactile and visual information did not provide congruent information on all trials. They found that the judgments of height were strongly driven by the visual information when the visual height was clearly visible (highly reliable). In contrast, when the visibility of the bar was reduced, judgments were more strongly driven by the haptic information. Measuring the variability of within-modality discriminations (i.e. on trials where only visual or only haptic information was presented) allowed the experimenters to predict which modality should dominate the visuo-haptic judgement: in line with the predictions of Bayesian inference, the more reliable (less variable) modality contributed more strongly to the visuo-haptic estimate.

Such optimal weighting of information has been demonstrated in many other contexts, both in human and non-human animals (e.g. monkeys): for example, studies requiring subjects to combine visual and vestibular information (i.e. information about head movement) when judging the direction of travel (Fetsch et al., 2009; Gu et al., 2008; Morgan et al., 2008); studies that required subjects to make judgments about the horizontal displacement of an edge defined by two “visual textures” where the textures were made up of two dimensions (e.g. contrast and orientation) of varying reliability (Landy and Kojima, 2001); and studies that required subjects to judge the slant of three-dimensional objects based on texture cues and stereo cues (disparity between eyes) of varying reliability (Knill and Saunders, 2003; Saunders and Knill, 2001). In all studies, judgments were more influenced by the more reliable cues, a feature that has become a hallmark expected under Bayesian inference.

In the studies described above, judgments are made by combining different pieces of evidence presented simultaneously. However, as seen in the example of the ‘light-from-above’ prior (Figure 1.3), the brain can also combine evidence acquired at different timescales. Some experiments have tested whether trial-by-trial judgements about a stimulus feature are informed by the prior probability of that stimulus feature as learnt over many trials. For example, in one study, participants were asked to catch a virtual ball, having to trade off the time taken to process the ball’s trajectory with the time taken to generate a catching action towards the estimated landing (Faisal and Wolpert, 2009). Simply put, “spending more time” estimating the trajectories, results in more precise estimates, and "spending more time" generating movement trajectories leads to more precise movements. However, these processes have to be traded off against one another, because, as soon as an action was initiated, the ball was rendered invisible. The researchers showed that, as the task progressed, participants begun to initiate their catching actions at the point in time predicted by a Bayesian observer that optimally trades off sensory precision with motor precision. They also showed that individual differences in action-initiation times were explained by individual relationships between sensory precision and motor precision. Near-optimal combination of sensorimotor information has also been shown in other studies (Battaglia and Schrater, 2007; Körding and Wolpert, 2004; Körding et al., 2004).

Bayesian integration has also been observed even when perceptual information is used to make a non-perceptual judgment; think, for example, of predicting whether it will rain based on your perception of an overcast sky. In two studies, when guessing the landing point of a "Space-Invaders" ship (O'Reilly et al., 2013), or the position of a coin thrown into water (Vilares et al., 2012), people combine optimally the likelihood of perceptual information (the perceived stochastic trajectory of the falling invader, or the perceived stochastic splashes created by the coin) with the *a priori* (prior) distribution of landing points learnt across trials. In short, participants relied more on the perceptual likelihood when it was more reliable and more on the prior knowledge about the distribution of landing points when it was more reliable, as if they were using Bayes' rule.

Furthermore, evidence for thinking about the brain as a Bayesian observer does not come from analysis of behaviour alone. It has been proposed that neurons are perfect candidates for representing probabilistic quantities: their tuning curves can be conceptualized as probability density functions over the feature space that they respond to (Beck et al., 2008; Ma et al., 2006; Pouget et al., 2000; Zemel et al., 1998). And it has been theorized, under a framework called "predictive coding" (see Figure 1.4), that the brain constructs hierarchical beliefs about the world than can be continually updated (Friston and Kiebel, 2009; Kilner et al., 2007; Rao and Ballard, 1999).

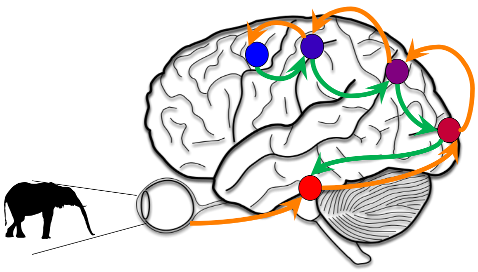

Figure 1.4. Predictive Coding.

A schematic of predictive coding, where the brain processes are seen as a hierarchy of inferential stages about the state of activity of the previous stage, in the cartoon, bluer stages represent higher parts of the hierarchy whereas lower ones are more red. Under this framework, two types of information flow are proposed, that together generate and update the expectations of the world: one being prediction errors between what was expected and what was "observed" that flow toward higher stages in the hierarchy (orange arrows), and the other one being the updated expectations flowing toward lower stages in the hierarchy (green arrows), to minimize future prediction errors.

Under this view, each stage of processing embodies an expectation (i.e. a belief) about the state of the stage immediately "below" i.e. (closer to the environment in the chain of processing), which is then contrasted against the actual observable state of that stage to generate a prediction error and update the initial prediction. The theory proposes that, neural activity flows in two directions: i) expectations or predictions flow from top layers down to lower ones (top-down) as they get updated; whereas ii) prediction errors flow from bottom layers to top layers to allow the updating of expectations (bottom-up).

There is increasing neural evidence supporting the existence of prediction error signals (Schultz et al., 1997), as well as neural correlates of expectations (Summerfield and Egner, 2009), and more generally the existence of probabilistic computations in the brain (van Bergen et al., 2015; O'Reilly et al., 2013; Vilares et al., 2012). Together, both from a theoretical perspective, and from behavioural and neural results, there is increasing support for the view of the brain as a Bayesian inference machine that works as an Ideal Observer. So, what exactly is then, an optimal or Bayesian Ideal Observer?

1.5 Bayesian observers and generative models of the world

The idea of a Bayesian Ideal Observer is a conceptual tool that serves as a benchmark against which we can compare human behaviour. As outlined in previous sections, the information that the brain uses to make decisions is inherently uncertain or ambiguous. This concept is formalized by assuming an agent that makes decisions about a latent state of the world based on a random variable (typically Gaussian) which provides some information about the state of the world (Green and Swets, 1974) (Figure 1.5). The decision is typically a categorical guess about which of two states of the world caused the observed random variable. A decision is made by thresholding the random variable; the threshold can be thought of as a criterion in the space of the possible values that the random variable can take. The criterion would be set so as to maximise a gain function (or minimise a loss function), such as the expected accuracy of the decision.

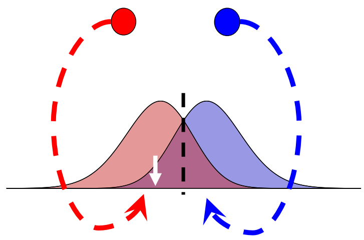

Figure 1.5. Signal detection theory and the brain.

Two possible states of the world (schematized by the red and blue balls), cause a random variable to take a value described by overlapping density functions (red and blue Gaussian curves). On a given “trial” an agent observes the random variable taking a particular value (schematized by the white arrow) and has to decide which what is the latent state of the world that caused it. A criterion exists (dashed black line) that maximizes the expected accuracy of the decision.

A Bayesian Ideal Observer is assumed to have a “generative model” of the world which consists of a set of beliefs about the probabilistic relationships between the states of the world and the noisy signals generated by those states (Yuille and Kersten, 2006). This generative model allows the observer to calculate the likelihood of observing a signal conditional on a state of the world. The probability density functions shown in Figure 1.5 reflect this relationship, with the curves describing the likelihood of observing a value of the random variable conditional on the red state or the blue state being true. Thus, when observing a signal on a given trial, the observer can “invert” the generative model (Friston, 2008; Yuille and Kersten, 2006) to estimate the probability that each state of the world is true given the observed signal (these posterior probabilities would be complementary when only two states are possible).

At first glance, this seems like a cumbersome way to make decisions. In the example above, probabilistic representations may not be needed and the agent could instead perform all inferences in stimulus space (Green and Swets, 1974). However, this formalization of the problem allows the agent to achieve several important behavioural features: 1) to have a continuous estimate of the certainty associated with a signal; 2) to use these estimates of certainty to combine independent pieces of evidence optimally by using Bayes' rule; 3) to update beliefs about the current state of the world, as well as the generative model itself, as experiences accumulate over a task.

1.6 Beliefs about the world, decisions and metacognition

What does it mean to hold beliefs about the world? Under the Bayesian framework, the brain holds continuous representations of the statistical relationship between states of the world given a specific brain signal. These statistical relationships are often termed *beliefs*. Central to this concept of belief is that these statistical representations are continuous measures of the knowledge the brain has about a state of the world. However, these continuous measures often need to be discretized in order to guide behaviour, as actions often are categorical. For example, when hearing a sound, an animal needs to decide whether to keep grazing, if it believes the sound was generated by the wind, or to run for its life if it believes there is a predator around. Indeed, many laboratory experiments require participants to discretize their beliefs to be able to reach decisions that need to be reported in a categorical way, typically by pressing one or two buttons with the fingers, or by directing their gaze to one of two target locations.

However, there are several scenarios where using the continuous distribution of belief is much preferred over using the simpler categorical judgment. For instance, during group decision-making, the explicit report of a more continuous belief, allows a group of people to combine their information in a near-optimal manner. A recent set of studies has shown that groups of people are able to make better decisions together than alone by means of sharing the confidence about their respective decisions (Bahrami et al., 2010; Bang et al., 2017). It has also been proposed that holding a continuous belief about the state of the world might be particularly useful during learning. For instance, receiving feedback about one's decision being incorrect should strongly advise us to revisit our model of the world when feedback is highly unexpected or in strong opposition with our current beliefs (Yu and Dayan, 2005). Humans and other animals have been shown to use their beliefs about the state of the world to adapt their behaviour optimally, even when making categorical decisions. In one study, for instance, when feedback was highly unexpected, human participants were more willing to update their belief about the underlying state of the world for subsequent decisions (Behrens et al., 2007). In another study, rats were shown to use their continuous belief about the state of the world to decide if it was worth it to wait for a potential reward or to forego the possibility but accelerate the start of a new trial. Rats were more willing to initiate a new trial when they were more likely to be wrong and thus would be unlikely to receive a reward (Kepecs et al., 2008).

Under the Bayesian framework, *metacognition* is a natural ingredient of cognition. There are several definitions of metacognition, but in a general sense, they all refer to knowledge we have about the quality of the inferences we make about the world, that seemingly goes "beyond" what is required for a specific task but that it is still accessible and reportable (i.e. ready to use). For example, how "certain" am I about the time it will take me to reach a destination, or how "confident" do I feel about my chances of winning a gamble? Certainty and confidence are core concepts of metacognition but what do they mean and how do they differ from one another? An emerging computational consensus (Pouget et al., 2016) is that certainty and confidence refer to distinct probabilistic quantities. Confidence is proposed to be reserved for an estimate of the probability that a choice or belief is correct, whereas certainty is proposed to be used for the representation of probability distributions over task-relevant variables which can be used to form a decision but which are represented in a decision-independent manner.

Regardless of whether we choose to call these two types of probabilistic representation as *confidence* or *certainty*, a conceptual distinction between them is that one refers to the distribution of probability and the other one to a summary-statistic measure of this distribution: we can hold a continuous representation of the distribution of probability for any particular state of the world, but we can also report a summary scalar of such a distribution (Meyniel et al., 2015). And the most popular views of confidence conceptualize it as a Bayesian distribution of probability about a state of the world, conditional on some other state (usually a choice or a premise).

At apparent odds with this idea of confidence as Bayesian probability, is the fact that confidence judgments seem to vary idiosyncratically across subjects, but consistently within subjects across tasks and time (Ais et al., 2016; Bang et al., 2017; Broihanne et al., 2014; Mann et al., 1998; Navajas et al., 2017; Niederle and Vesterlund, 2011). While it is true that these results are at odds with the idea that the brain is computing Bayesian probabilities under strict optimality, they are by no means at odds with the idea of the brain representing and dealing with probabilities in a Bayesian manner. After all, the

Bayesian framework is nothing but a conceptual tool to understand the challenges that the brain faces at a computational level, and any departures from optimality are crucial in constraining our understanding of the mechanisms that allow the brain to approximate a Bayesian solution.

1.7 Sequential sampling and evidence accumulation

One crucial aspect of cognition that I have neglected thus far in the introduction, is the fact that decisions take time to be made. Intuitively, it makes a lot of sense that harder decisions require more time for deliberation. But why should this be the case, and what does this tell us about how the brain works? When inferring the causes of our brain signals, when guessing “what is out there”, the brain reduces the ambiguity or uncertainty of its beliefs by combining different pieces of evidence. Delaying our decisions thus allows us to acquire more evidence to refine our beliefs, or to narrow down our guesses, such that our inferences are more likely to correspond with the reality that we are trying to infer.

The realization that response times offered a window into the mechanisms of human decision-making started around the decade of the 1950s (Chernikoff and Brogden, 1949; Gregg and Brogden, 1950). At about the same time, the mathematical idea that hypothesis-testing could be carried out in a sequential manner by adding discrete pieces of extra evidence one at a time, was formulated (Wald, 1945). This “Sequential Probability Ratio Test”, consists of adding with every step, or new piece of evidence, the ratio of two quantities: 1) the likelihood that some piece of evidence favours one hypothesis and 2) the likelihood that it favours another one. This strategy was shortly after deemed to provide the optimal solution to the speed-accuracy trade-off in inferential processes (Wald and Wolfowitz, 1948), maximising the expected accuracy for a fixed number of samples (pieces of evidence), or alternatively, achieve a target level of expected accuracy in the fewest number of steps.

It was not long before the ideas of “sequential probability ratio testing” became popular for describing the accuracy and response times in human decision making and memory retrieval (Laming, 1968; Ratcliff, 1978; Stone, 1960). The implementation of this algorithm in a more continuous way, is based on the idea of a “random walk” process, whereby a particle is believed to drift stochastically in time towards either of two boundaries, each representing one possible latent hypothesis, and the position of the drifting particle relative to the two boundaries at any point in time reflects the average amount of evidence gathered for one hypothesis (and against another hypothesis). This class of models has been termed Drift Diffusion Models(see Figure 1.6).

DDMs can be parameterized in terms of the boundary height (the separation between the two bounds), the drift noise (the amplitude of the stochasticity in time), the drift rate (the general tendency to drift towards the correct answer), and other parameters not discussed here. DDMs are able to provide exceptional fits to the response times and accuracy of humans in a variety of tasks and has been shown to conform well with optimality accounts (Bogacz et al., 2006).

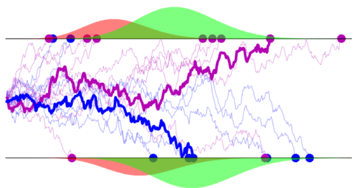

Figure 1.6. Drift Diffusion Model.

Under this model, a decision variable (particle) drifts stochastically in time (x-axis), akin to a random walk process. When the particle crosses either of two boundaries a decision is made. In the cartoon example above, there are two possible states of the world that must be inferred: a "purple" one and a "blue" one (the colour of each drift trace reflects the true state of the world). Similarly, there are two hypotheses (horizontal upper and lower boundaries), the upper one is the hypothesis that the state of the world is purple, the lower one that it is blue. When a particle crosses either boundary a decision is made that the hypothesis "crossed" is the most likely. This mechanism will determine the response time distribution (in this case separated by correct in green and incorrect in red responses) for each choice.

Formally, a quantity known as the log-likelihood ratio, is typically used for describing either 1) each step in the diffusion process or 2) the drift rate. The logarithm is a mathematical convenience for dealing with ratios such that a ratio of Y/X has the same magnitude but opposite sign as a ratio of X/Y, but the most important aspect of the algorithm is that the value of a given piece of evidence is determined by the relative support that it provides for one hypothesis over another, in this way, evidence *for* hypothesis X is necessarily evidence *against* hypothesis Y.

This algorithmic framework has been proposed to be well-implemented by neurons in Lateral Intraparietal cortex of monkeys (LIP) involved in the decision-making process (Gold and Shadlen, 2002; Roitman and Shadlen, 2002). The firing rate of these neurons ramps up in proportion to the amount of evidence presented that favours the neuron's preferred option (that which is in their receptive field) and up to a common level of excitation just before a decision is made. The same neurons would supress their firing rate where evidence is presented against their preferred option, and their suppression would be maximal also at the moment just before a decision is made. There has been recent controversy about whether the role of LIP (and analogue areas in the rodent brain) in evidence accumulation is causal or whether it reflects other processes related to the decision (Brunton et al., 2013; Churchland and Kiani, 2016; Katz et al., 2016). However, there is increasing evidence that DDM models with all their variants offer an algorithmic solution for optimal decision-making in the brain.

In short, the DDM has been extremely popular as an algorithmic implementation of decision-making in the brain because it offers an easy solution to the problem of adding extra evidence to an ongoing inferential process, and it turns out to do so in the optimal way: 1) the position of the particle relative to the boundaries dictates the best "guess" at that time given the evidence, 2) the decision will have maximal expected accuracy for any fixed amount of time or samples of evidence taken, or 3) for any target of expected accuracy it will reach a decision in the shortest amount of time. Similar algorithmic frameworks have been shown to be reducible to a DDM (Bogacz et al., 2006). Finally, while the DDM was originally proposed for making decisions between two options, there are slightly more complicated extensions of the model to deal with 3 or more options (Tajima et al., 2016).

1.8 Rationality and Optimality

The concepts of optimality and rationality are tightly interlinked. Rationality is a term that emerged earlier in the literature (compared to optimality), and is typically used to describe human judgment. Optimality, on the other hand, was originally coined in mathematics to refer to the best possible solution to a problem, and in psychology it is often used for perceptual decisions. However, the terms have been used interchangeably in the literature and their usage is increasingly interdisciplinary touching upon topics of economic decisions, psychological decisions and even morality. In this section I will explain some differences in their historical conception and usage and how they relate to important computational issues in current models of decision-making.

The challenge of making deliberate decisions under risk has been ubiquitous in human history: for instance, when predicting the weather for the day or when the best time to start planting seeds was. Moreover, there is archaeological evidence that "tallus bones" have been used as dice since prehistoric times (Hacking, 1975). However, it was not until the 1654 that the concept of probability was explicitly formulated (in epistolary interactions) by Blaise Pascal and Pierre de Fermat (Hacking, 1975). Interestingly, alongside probability theory came the idea of comparing "human behaviour" against what was deemed as good sense or "reasonable behaviour" under theoretical grounds. For instance, Pascal reasoned that the potential benefit of eternal life in heaven, even if the existence of heaven was highly improbable, demanded it was reasonable to believe in God. A similar comparison was made between human behaviour and the theoretically rational when analysing the "St. Petersburg paradox", but this time the two seemed at odds. In this paradox, a gamble is proposed to a player, where a fair coin is tossed repeatedly until it first lands on "heads", the number of consecutive tosses that the coin lands on "tails" dictates the amount of money the player gets as a reward, 2^k, where k is the number of tosses before the first "heads", such that the player would leave with £2, £4, £8, £16 if the first heads would be seen after 1, 2, 3 or 4 tails, respectively. While people are willing to pay between 5 and 10 pounds to play this game, the expected value of the game is theoretically infinite.

Daniel Bernoulli, the famous mathematician and physicist, proposed in the early 18th century to resolve the conflict between what people do and what is rational to do. He reasoned that the theory of rationality should include subjective aspects about the risk that each person was experiencing. Thus, he defined expected utility as an alternative to expected value using a logarithmic function to describe the relationship between objective value and subjective utility. Further, he allowed the incorporation, into the computation of utility, of other aspects that are intrinsic to each person, such as their personal endowment, reasoning that the same amount of money should have a lower utility for someone who already has a lot of money compared to someone who does not. More generally, Bernoulli resolved the mismatch between "human behaviour" and "the norm of what was rational" by re-formulating the norm, for instance by allowing for utility being a concave function of wealth (Bernoulli, 1954).

Perhaps because of the realization that rationality included a number of subjective components, mathematicians abandoned the scientific enquiry of this topic for the most part of the coming two centuries. Still, the sociological and moral aspects of rationality kept developing (Calhoun, 2012). Research of rationality in decision-making was only revived from 1950s onwards, with the advent of statistical methods in the economic and psychological sciences, and several theories of human decision-making were proposed contemporarily in the fields of biology, psychology and economics (Charnov, 1976; Gigerenzer and Selten, 2002; Kahneman and Tversky, 1979; Krebs, 1978; Luce, 1959; Sargent, 1993; Simon, 1956; Von Neumann and Morgenstern, 1945).

In the middle part of last century, the revival of the study of rationality started with the proposal of the choice axiom, four principles that defined rational behaviour: completeness, transitivity, independence and continuity (Luce, 1959; Von Neumann and Morgenstern, 1945). First, the **completeness** principle states that an individual presented with a choice set of option A and option B will either i) prefer A over B, ii) be indifferent between A and B or iii) prefer B over A. Second, the **transitivity** principle states that if an individual prefers option A over option B and option B over option C, then it *must* prefer option A over option C. Third, the **independence** principle states that the order of preference for options A and B, when presented in conjunction with option C, should remain the same as when options A and B are presented independently of option C. Lastly, the **continuity** principle states that if A is preferred over B and B is preferred over C, there is a mixture between A and C such that an individual is indifferent between this mixture and option B. As a summary, the axioms of rational choice, describe an agent whose preferences are dictated by a continuous utility function (which need not be a linear function of value). However, this utility function should to result in stable, logical and predictable preferences in diverse scenarios.

At about the same time, the concept of Sequential Probability Ratio Test (already described in the Sequential sampling and evidence accumulation section of the thesis) was developed as a tool for optimizing sequential sampling problems in terms of the speed and accuracy trade-off (see Wald, 1945). Soon after, problems of decision-making in humans were posed in terms of optimization, where a certain utility function should be maximized, or some cost function minimized (Edwards, 1954; Mosteller and Nogee, 1951). There was a general sense, that the rational choice axiom was still a useful description of human decision-making. However, by that point, there was an emerging acknowledgement that human decisions are not always consistent across time for a single individual.

Early scepticism arose against the choice axiom as the ultimate model for explaining human choices. The axiom assumes that rational individuals have complete knowledge of the choice set, and unlimited computation capacity, two assumptions that seem unlikely to hold for complex choice problems. Instead of *maximizing* some utility function, or *optimizing*, it was proposed that individuals *satisfice* their needs up to some desirable level, by using imperfect but simple mechanisms that exploit regularities in the environment (Simon, 1956). Interestingly, Herbert Simon proposed that this strategy of *satisficing*, could have some advantages over the complex mechanisms necessary for rational behaviour. In particular, it would be much more adaptable to changes in the goals or needs that an animal (individual) may experience. Small changes to a *satisficing* strategy could result in good-enough behaviour in each of the new scenarios (with different goals), whereas the challenge may be much greater for more complex but rational mechanisms when reusability is unknown (Simon, 1955, 1956).

After Herbert Simon's work, finding counter-examples to rational behaviour, became a popular avenue of research. Amos Tversky and Daniel Kahneman developed a (now extensive) list of cognitive illusions or biases that are at odds with the choice axiom, and culminated in their proposal of Prospect Theory (Kahneman and Tversky, 1979). One of the most important insights of this research is that individuals' choices cannot be described by a simple utility function of the person's wealth. Instead, choices are better described in terms of the ***local*** gains and losses for the particular choice problem. That is, local contextual information is providing a reference point according to which values will be re-framed and until then deemed as a gain or a loss.

To exemplify this, people in two different scenarios were asked to say how attractive a gamble was. In the first scenario, the gamble was: i) a 50% chance of winning £1,500, and ii) a 50% chance of losing £1000, which people didn't find very attractive. In the second scenario, the gamble was: i) to keep their own wealth, or ii) to have a 50% chance of having their wealth diminished by £1,000 and a 50% chance of having it augmented by £1,500. Since two scenarios are mathematically equivalent, a classic utility function would predict that people's preferences should not change between the two scenarios;

however, people actually found the proposal of the 50/50 gamble much more appealing in the second scenario (Kahneman and Tversky, 1979).

Across their experiments, Kahneman and Tversky showed an asymmetry in the way in which people treat losses and gains. Typically, gains need to be twice as big as losses before people are willing to engage in a 50/50 lottery. As described above, gains and losses themselves are not defined in terms of some objective function of the person's wealth but they are instead defined in terms of a flexible reference point. This anchor is in turn defined by information present or highlighted when the choice problem is stated but which should have been irrelevant according to the axiom of rational choice. Their theory proposed that human judgments often rely on intuitive thinking, or mechanisms that effortlessly result in a gist sensation, that disturbs or precludes rational thinking. Interestingly, the biases also affect people with extensive experience in a given domain (Kahneman, 2011; Kahneman and Tversky, 1979).

A parallel way of thinking about departures from optimality was also inspired by and more in line with Herbert Simon's work. Under this view, individuals making decisions have a toolbox, or a "bag of tricks", that allows them to achieve good-enough decisions in a variety of scenarios. Crucially, this toolbox frees individuals from having to perform the computations that optimize the particular choice problem (Gigerenzer and Selten, 2002). Gigerenzer and colleagues ascribe the discrepancies between normative and actual behaviour as a failure from experimenters and theorists to incorporate realistically the challenges of achieving such normative behaviour (Gigerenzer and Selten, 2002). This view differs radically from that proposed in Prospect Theory, where departures between normative and actual human behaviour, are deemed simple failures of rationality.

In a similar spirit, the theory of "optimal foraging" in biology assumes that the ultimate optimization occurs at the level of selecting mechanisms (through the process of natural selection), that trade-off the ability to find good answers to the problems with the challenges or costs of implementing the solutions (Charnov, 1976; Krebs, 1978). Under this view, any departures from optimality either i) have not taken into account some of the costs, or ii) explore behaviour in environments that are too alien to the ones where the mechanisms evolved.

A final view that I will discuss relating the terms of rationality and optimality, is one that resonates well with the theory of optimal foraging and has received recent attention in the perceptual decision-making literature (Tajima et al., 2016; Tsetsos et al., 2016). Under this view, individuals are seen as optimizing either accuracy (Tsetsos et al., 2016), or the frequency of reward over a given time window, that is, reward rate (Tajima et al., 2016). A crucial aspect of this view is that these optimizations are performed under constraints imposed by the computational tractability of the choice problem, and by the decision maker's availability of information and processing capacity. Under this view, direct violations of the transitivity and independence axioms of rationality can be explained as a natural consequence of the process of optimization of accuracy or reward rate under uncertainty.

In the remaining of the thesis, I will use "rationality" to describe choice made according to the axioms of Rational Choice Theory and "optimality" to describe choice made as a result of optimization with respect to some loss function (e.g. Bayesian inference). I will favour the view that the axioms of rationality do not serve as an adequate standard for prescribing optimal human behaviour; they neglect the real nature of the challenges that brains must deal with.

1.9 Taxonomy of uncertainty in decision-making

"The best material model of a cat is another, or preferably the same, cat."

–Arturo Rosenblueth (1900-1970) and Norbert Wiener (1894–1964)

The quote above is an undeniable truth, in the sense that the most accurate description of a system is given by incorporating all aspects of that system, however, for practical reasons, I could not agree less with it. Generating imperfect models of a system allows to construct abstractions (succinct descriptions of the system). Abstractions can in turn help us understand more aspects of that system and generalize that knowledge to other systems. We can capitalize on some regularities of a system, to create models that generate a distribution of predictions, such that even when we allow for "inexplicable variation" at some low levels of understanding of a system, we can still advance our understanding of the behaviour of that system at a higher level of understanding by assuming the emergence of new properties (e.g. meteorological forecasting).

My favourite example of such useful reductionism is *temperature*: a scalar summary measure of the dynamics of molecules in a system. The "correct model" of the dynamics would necessarily include information about the shape, mass and vector of velocity of each of the molecules of the system so as to be able to predict the exact position and vector of velocity of any molecule in the system at any future point in time. Just as we do not need to know anything about the mechanism behind radioactive decay to know that, on average, one gram of uranium-238 will give off 12600 alpha particles every second (Komura et al., 1990), we can *reduce* the dynamics of the molecules of water in a glass, to know that if it is at 100ºC and we put your hand in it will hurt.

Similarly, even though concepts such as entropy, noise, variability, ambiguity, risk, randomness and uncertainty might not reflect “real” aspects of the universe, they are essential ingredients of most models of neuronal firing, decision making and judgment. They probably are the most important ingredients when thinking about the brain as an inference machine, and have been paramount in our advancement of the understanding of brain function. However, these concepts encompass several different phenomena, each of which reflects different challenges for decision-making. In this section I will offer a taxonomy of uncertainty according to some of the most important criteria for the questions treated in this thesis.

The first distinction I want to make is between i) uncertainty in experimental measurements of a system and ii) stochasticity in how the system responds to experimental manipulations. For instance, when trying to explain the spiking patterns of a neuron, we, the experimenters, can build a tuning curve that models the neuron’s preference for a given stimulus, where the highest firing rate is achieved for the most preferred stimulus value. From the experimenter’s point of view, the neuron will not produce the exact same firing pattern for any two presentations of the same stimulus. While the experimenter can attribute this variability in firing patterns to the natural stochasticity of the subcellular components of a neuron, the variability can also be explained by incorporating more complex definitions of what the neuron’s preferred stimulus is. However, it is often convenient to model variability simply as being generic randomness that obeys some aggregate properties (as with Gaussian, or Poisson, noise). The distinction above is one that is often hard to make in reality, and the experimenter’s uncertainty haunts any field of scientific enquiry.

Perhaps the most intuitive classification of the uncertainty for a decision-maker, is that which separates: i) the uncertainty carried by the choice problem and ii) that which arises from the imperfect processing of information in the brain. For instance, take the choice between i) £10 if an unbiased coin lands on tails after being flipped or ii) £14 if an unbiased die lands on 5 or 6 after being thrown. The first source of uncertainty is reflected by the fact that the there is no sensible way in which the individual presented with the gamble can predict the event outcomes. In this sense, there is irreducible uncertainty about the outcome of those events, which can only be viewed probabilistically. The second source of uncertainty comes from the fact that when computing the value of the two gambles, the individual might not generate a good enough abstraction of the random process governing the toss, or might even commit errors of estimation (or arithmetic) when computing the value of her options. The variability in behaviour, observed by the experimenter, in different instantiations of otherwise equal gambles presented to an individual, is often modelled as noise. Using the brain as a reference point, we can classify noise as being “external” or “internal” when variability in behaviour is driven by uncertainty intrinsic to the environment or to brain processes, respectively (Howell and Burnett, 1978; Kahneman and Tversky, 1982).

Following the spirit of the last classification, but from a computational point of view (and with a rather cumbersome nomenclature), uncertainty can be classified as Thurstonian or Brunswikian. An almost direct mapping can be made between Thustonian and internal noise and between Brunswikian and external noise. A key aspect of Thurstonian noise is that it can be dissipated, for instance, when sensitivity for discriminating two stimuli increases with experience. On the other hand, a key aspect of Brunskwikian uncertainty is that it is irreducible, and it reflects the imperfect mapping between the cues that are available to the decision-maker and the latent states that she is trying to infer (Juslin and Olsson, 1997).

Interestingly, from a Bayesian point of view, the different sources of noise are all Brunswikian. At any stage in the processing of information, the inferences that need to be made about latent states of the world (or of the activity in earlier areas of processing) rely on cues whose states hold only a probabilistic relationship with the states of the latent variables, regardless of their origin. However, the temporal origin of internal noise in the inference process can play a crucial role on the eventual behaviour. For instance, distinguishing noise that occurs i) during sensory transduction or during early processing of individual sensory cues from noise that occurs ii) during the combination of different cues or afterwards has proven to be fundamental in justifying human behaviour that would otherwise be seen as suboptimal (Li et al., 2017; Spitzer et al., 2017; Tsetsos et al., 2016).

1.10 Brief summary of the following chapters

In the following three chapters, I present scientific work I carried out during my DPhil. In general, I was interested in understanding how decisions are affected by the context in which they are made. We tend to think of rational agents as having fixed decision policies that are translatable across contexts. For instance, it would seem weird if someone prefers apples over bananas when offered to choose from those two fruits but prefer bananas over apples when they are also offered pears. The remaining of the thesis contains three chapters of data. In the first two chapters, I explore the mechanisms that are at work when we make decisions in the context of three options. In the third chapter, I look at decisions in the context where explicit prior information is available to guide a perceptual judgment.

The **first chapter** is based on computer simulations of agents making decisions between three options. We focus on exploring systematically the sources of distractor effects: paradoxical changes in the sensitivity for choosing between two options (like bananas and apples) provoked by the value of a third option that is not chosen (like pears). While distractor effects have received recent attention in the literature, a systematic treatment has been lacking.

In the **second chapter**, we explore experimentally a crucial aspect of decision making with multiple options that has received little attention: how the value of the options is constructed before a decision can be made. We used a sequential integration paradigm, that allows us to control the stream of evidence (about the value of three options) that is available at every point in time as participants construct the value of the options. We focus particularly on how a third option in the choice-set can affect choices by means of truly distracting the process of value construction.

Finally, in the **third chapter**, we explore experimentally the sensitivity that people have about uncertainty arising at two different stages in the processing hierarchy: at early stages of processing, when feature values are encoded, and at later stages of processing, when multiple pieces of evidence are integrated to form a judgment. In particular, we wanted to see whether insensitivity to uncertainty at later stages of processing could be the cause of suboptimal behaviour observed in several cognitive tasks, including a tendency to overlook prior information or base rates.

Finally, the last chapter is a brief discussion of the findings presented in these data chapters.

References

Ais, J., Zylberberg, A., Barttfeld, P., and Sigman, M. (2016). Individual consistency in the accuracy and distribution of confidence judgments. Cognition *146*, 377–386.

Bahrami, B., Olsen, K., Latham, P.E., Roepstorff, A., Rees, G., and Frith, C.D. (2010). Optimally interacting minds. Science *329*, 1081–1085.

Bang, D., Aitchison, L., Moran, R., Castanon, S.H., Rafiee, B., Mahmoodi, A., Lau, J.Y., Latham, P.E., Bahrami, B., and Summerfield, C. (2017). Confidence matching in group decision-making. Nature Human Behaviour *1*, s41562–017.

Battaglia, P.W., and Schrater, P.R. (2007). Humans trade off viewing time and movement duration to improve visuomotor accuracy in a fast reaching task. Journal of Neuroscience *27*, 6984–6994.

Beck, J.M., Ma, W.J., Kiani, R., Hanks, T., Churchland, A.K., Roitman, J., Shadlen, M.N., Latham, P.E., and Pouget, A. (2008). Probabilistic population codes for Bayesian decision making. Neuron *60*, 1142–1152.

Behrens, T.E., Woolrich, M.W., Walton, M.E., and Rushworth, M.F. (2007). Learning the value of information in an uncertain world. Nature Neuroscience *10*, 1214–1221.

van Bergen, R.S., Ji Ma, W., Pratte, M.S., and Jehee, J.F.M. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. Nat Neurosci *18*, 1728–1730.

Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. Econometrica: Journal of the Econometric Society 23–36.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. Psychological Review *113*, 700–765.

Broihanne, M.-H., Merli, M., and Roger, P. (2014). Overconfidence, risk perception and the risk-taking behavior of finance professionals. Finance Research Letters *11*, 64–73.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98.

Calhoun, C. (2012). Classical sociological theory (John Wiley & Sons).

Charnov, E.L. (1976). Optimal foraging, the marginal value theorem. Theoretical Population Biology *9*, 129–136.

Chernikoff, R., and Brogden, W.J. (1949). The effect of response termination of the stimulus upon reaction time. Journal of Comparative and Physiological Psychology *42*, 357.

Churchland, A.K., and Kiani, R. (2016). Three challenges for connecting model to mechanism in decision-making. Current Opinion in Behavioral Sciences *11*, 74–80.

Daw, N.D., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. Nature

Neuroscience *8*, 1704–1711.

Dennett, D.C. (2004). Freedom evolves (Penguin UK).

Edwards, W. (1954). The theory of decision making. Psychological Bulletin *51*, 380.

Ernst, M.O., and Banks, M.S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. Nature *415*, 429–433.

Faisal, A.A., and Wolpert, D.M. (2009). Near optimal combination of sensory and motor uncertainty in time during a naturalistic perception-action task. Journal of Neurophysiology *101*, 1901–1912.

Fetsch, C.R., Turner, A.H., DeAngelis, G.C., and Angelaki, D.E. (2009). Dynamic reweighting of visual and vestibular cues during self-motion perception. Journal of Neuroscience *29*, 15601–15612.

Friston, K. (2008). Hierarchical models in the brain. PLoS Computational Biology *4*, e1000211.

Friston, K., and Kiebel, S. (2009). Predictive coding under the free-energy principle. Philosophical Transactions of the Royal Society of London B: Biological Sciences *364*, 1211–1221.

Gigerenzer, G., and Selten, R. (2002). Bounded rationality: The adaptive toolbox (MIT press).

Gold, J.I., and Shadlen, M.N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. Neuron *36*, 299–308.

Green, D.M., and Swets, J.A. (1974). Signal detection theory and psychophysics (Rev.

ed.). Huntington, NY: RF Krieger.

Gregg, L.W., and Brogden, W.J. (1950). The relation between duration and reaction time difference to fixed duration and response terminated stimuli. Journal of Comparative and Physiological Psychology *43*, 329.

Gu, Y., Angelaki, D.E., and DeAngelis, G.C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. Nature Neuroscience *11*, 1201–1210.

Hacking, I. (1975). The Emergence of Probability: A Philosophical Study of Early Ideas about Probability, Induction and Statistical Inference: A Philosophical Study of Early... Induction and Statistical Inference.

Heisenberg, W. (1949). The physical principles of the quantum theory (Courier Corporation).

Helmholtz, H. (1925). Physiological Optics, Vol. III: The Perceptions of Vision (JP Southall, Trans.). Optical Society of America, Rochester, NY.(Original Publication in 1910).

Howell, W.C., and Burnett, S.A. (1978). Uncertainty measurement: A cognitive taxonomy. Organizational Behavior and Human Performance *22*, 45–68.

Jowett, B. (1888). The republic of Plato (London; New York: Macmillan).

Juslin, P., and Olsson, H. (1997). Thurstonian and Brunswikian origins of uncertainty in judgment: a sampling model of confidence in sensory discrimination. Psychological Review *104*, 344.

Kahneman, D. (2011). Thinking, fast and slow (Macmillan).

Kahneman, D., and Tversky, A. (1979). Prospect theory: An analysis of decision under risk. Econometrica: Journal of the Econometric Society 263–291.

Kahneman, D., and Tversky, A. (1982). Variants of Uncertainty. ResearchGate *11*, 143–157.

Kandel, E.R., Schwartz, J.H., Jessell, T.M., Siegelbaum, S.A., Hudspeth, A.J., and others (2000). Principles of neural science (McGraw-hill New York).

Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. Nature *535*, 285–288.

Kepecs, A., Uchida, N., Zariwala, H.A., and Mainen, Z.F. (2008). Neural correlates, computation and behavioural impact of decision confidence. Nature *455*, 227–231.

Kersten, D., Mamassian, P., and Yuille, A. (2004). Object perception as Bayesian inference. Annu. Rev. Psychol. *55*, 271–304.

Kilner, J.M., Friston, K.J., and Frith, C.D. (2007). Predictive coding: an account of the mirror neuron system. Cognitive Processing *8*, 159–166.

Knill, D.C., and Richards, W. (1996). Perception as Bayesian inference (Cambridge University Press).

Knill, D.C., and Saunders, J.A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? Vision Research *43*, 2539–2558.

Komura, K., Yamamoto, M., and Ueno, K. (1990). Abundance of low-energy gamma rays in the decay of 238U, 234U, 230Th, 227Ac, 226Ra and 214Pb. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment *295*, 461–465.

Körding, K.P., and Wolpert, D.M. (2004). Bayesian integration in sensorimotor learning. Nature *427*, 244–247.

Körding, K.P., Ku, S., and Wolpert, D.M. (2004). Bayesian integration in force estimation. Journal of Neurophysiology *92*, 3161–3165.

Krebs, J.R. (1978). Optimal foraging: decision rules for predators. Behavioural Ecology: An Evolutionary Approach 23–63.

Laming, D.R.J. (1968). Information theory of choice-reaction times.

Landy, M.S., and Kojima, H. (2001). Ideal cue combination for localizing texture-defined edges. JOSA A *18*, 2307–2320.

Li, V.L., Castañón, S.H., Solomon, J.A., Vandormael, H., and Summerfield, C. (2017). Robust Averaging Protects Decisions from Noise in Neural Computations. BioRxiv 147744.

Luce, R.D. (1959). Individual Choice Behavior a Theoretical Analysis (John Wiley and sons).

Ma, W.J., Beck, J.M., Latham, P.E., and Pouget, A. (2006). Bayesian inference with probabilistic population codes. Nature Neuroscience *9*, 1432–1438.

Mann, L., Radford, M., Burnett, P., Ford, S., Bond, M., Leung, K., Nakamura, H., Vaughan, G., and Yang, K.-S. (1998). Cross-cultural differences in self-reported decision-making style and confidence. International Journal of Psychology *33*, 325–335.

Meyniel, F., Sigman, M., and Mainen, Z.F. (2015). Confidence as Bayesian Probability: From Neural Origins to Behavior. Neuron *88*, 78–92.

Morgan, M.L., DeAngelis, G.C., and Angelaki, D.E. (2008). Multisensory integration in macaque visual cortex depends on cue reliability. Neuron *59*, 662–673.

Mosteller, F., and Nogee, P. (1951). An experimental measurement of utility. Journal of Political Economy *59*, 371–404.

Navajas, J., Hindocha, C., Foda, H., Keramati, M., Latham, P.E., and Bahrami, B. (2017). The idiosyncratic nature of confidence. BioRxiv 102269.

Niederle, M., and Vesterlund, L. (2011). Gender and competition. Annu. Rev. Econ. *3*, 601–630.

O'Reilly, J.X., Jbabdi, S., Rushworth, M.F., and Behrens, T.E. (2013). Brain systems for probabilistic and dynamic prediction: computational specificity and integration. PLoS Biology *11*.

Pouget, A., Dayan, P., and Zemel, R. (2000). Information processing with population codes. Nature Reviews Neuroscience *1*, 125–132.

Pouget, A., Drugowitsch, J., and Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. Nat Neurosci *19*, 366–374.

Rao, R.P., and Ballard, D.H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. Nature Neuroscience

2, 79–87.

Ratcliff, R. (1978). A theory of memory retrieval. Psychological Review *85*, 59.

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

Rosenblueth, A., and Wiener, N. (1945). The role of models in science. Philosophy of Science *12*, 316–321.

Sargent, T.J. (1993). Bounded rationality in macroeconomics: The Arne Ryde memorial lectures. OUP Catalogue.

Saunders, J.A., and Knill, D.C. (2001). Perception of 3D surface orientation from skew symmetry. Vision Research.

Schultz, W., Dayan, P., and Montague, P.R. (1997). A neural substrate of prediction and reward. Science *275*, 1593–1599.

Simon, H.A. (1955). A behavioral model of rational choice. The Quarterly Journal of Economics *69*, 99–118.

Simon, H.A. (1956). Rational choice and the structure of the environment. Psychological Review *63*, 129.

Spitzer, B., Waschke, L., and Summerfield, C. (2017). Selective overweighting of larger magnitudes during noisy numerical comparison. Nature Human Behaviour *1*, s41562–017.

Stone, M. (1960). Models for choice-reaction time. Psychometrika *25*, 251–260.

Summerfield, C., and Egner, T. (2009). Expectation (and attention) in visual cognition. Trends in Cognitive Sciences *13*, 403–409.

Sun, J., and Perona, P. (1998). Where is the sun? Nature Neuroscience *1*, 183–184.

Tajima, S., Drugowitsch, J., and Pouget, A. (2016). Optimal policy for value-based decision-making. Nature Communications *7*.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., and Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. Proceedings of the National Academy of Sciences 201519157.

Vilares, I., Howard, J.D., Fernandes, H.L., Gottfried, J.A., and Kording, K.P. (2012). Differential representations of prior and likelihood uncertainty in the human brain. Current Biology *22*, 1641–1648.

Von Neumann, J., and Morgenstern, O. (1945). Theory of games and economic behavior (Princeton University Press Princeton, NJ).

Wald, A. (1945). Sequential tests of statistical hypotheses. The Annals of Mathematical Statistics *16*, 117–186.

Wald, A., and Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. The Annals of Mathematical Statistics 326–339.

Yu, A., and Dayan, P. (2005). Uncertainty, neuromodulation, and attention. Neuron *46*, 681–692.

Yuille, A., and Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? Trends in Cognitive Sciences *10*, 301–308.

Zemel, R.S., Dayan, P., and Pouget, A. (1998). Probabilistic interpretation of population codes. Neural Computation *10*, 403–430.

2 SIMULATING THE DISTRACTOR EFFECT IN MULTIALTERNATIVE DECISIONS

2.1 Introduction

A tenet of rational choice behaviour establishes that the relative preference between any two specific options should not be affected by the addition of a third option to the choice-set. This principle of the *Independence of Irrelevant Alternatives* (IIA), is a cornerstone for describing rational behaviour in the context of choices with multiple alternatives (Luce, 1959; Von Neumann & Morgenstern, 1945). More specifically, IIA states that with the addition of an option to a choice-set, the number of times that this new option gets chosen should be "obtained" from the number of times the original options in the choice-set were chosen, in proportion with their original choice probability. However, humans (Huber, Payne, & Puto, 1982; Simonson, 1989; Tsetsos, Chater, & Usher, 2012; Tversky, 1972) and other animals (Bateson, 2002; Bateson, Healy, & Hurly, 2002; Hurly & Oseen, 1999; Shafir, Waite, & Smith, 2002; Waite, 2001) have been shown to systematically violate this tenet of rationality. Variations of the "irrational effects" have been named according to the properties of the choice-set used to describe them.

Historically, for instance, economists have described "decoy effects" when studying decisions between three options with two attributes each. They described violations of rationality in terms of the properties of the choice-set in the two-dimensional (two attribute) space of value. Thus, with decoy effects, an option (i.e. a decoy, which is not preferred over the other two), biases the proportion of times that two other options are chosen (relative to one-another), by virtue of its position with respect to the other two: generating the attraction, compromise and similarity effects (Figure 2.1).

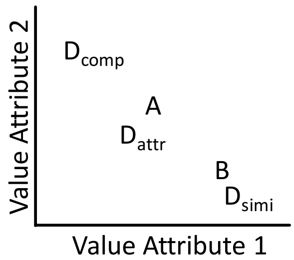

Figure 2.1. Decoy effects.

Consider two options, A and B, that are equally preferred when presented as the sole two options of a choice-set. The proportion of choices becomes biased towards option A with the addition of any one of the three different decoy options schematized (D_{attr}, D_{comp} and D_{simi}). In the attraction effect, the decoy, D_{attr}, is dominated in both dimensions by option A. In the compromise effect the decoy, D_{comp}, renders the value of option A on attribute 1 as no longer the lowest possible value, but the value of option B as clearly the worst on attribute 2, thus option A now appears as a good compromise between the two other options. Finally, in the similarity effect, the decoy, D_{simi}, competes with option B for having the highest value on attribute 1 whereas option A remains a clear winner for attribute 2.

More recently, neuroscientists have become interested in understanding the neural principles that may underlie violations of rationality. However, these violations of rationality have been described in a parallel frame of reference, where the one-dimensional (i.e. the overall, or average) value of a "distracter" option, D (which is either the option with lowest average value or one which is explicitly marked as unavailable), exerts an influence on people's accuracy for choosing the highest valued option, A, over the lower valued option, B. Two opposing effects have been described within this alternative frame for defining options in a choice-set, whereby the "efficiency" of choices (the probability of choosing option A over option B), can either: (i) decrease (negative effect; see (Hunt, Dolan, & Behrens, 2014; Louie, Khaw, & Glimcher, 2013); or (ii) increase (positive effect; see Chau, Kolling, Hunt, Walton, & Rushworth, 2014), as a function of increasing values of the distracter option, D.

These positive and negative distractor effects are explained in the different studies by means of separate neural mechanisms. The negative effect has been captured, in the original study, by means of divisive normalization. Under this view, the value of each option is encoded in relation to other values being encoded concurrently in a compressive way: higher overall values in a choice set lead to more overlapping values in the encoding of any pair of options (Louie et al., 2013). Chau and colleagues (2014), propose that the positive effect, on the other hand, depends on a model of competition, through lateral inhibition, between the pools of neurons that represent the value of each of the three options (Wang, 2012). In this model, higher overall values in a choice-set lead to stronger competition, and this in turn results in better dissolution of noise, leading to less overlapping values when encoding any pair of options.

A very different approach for understanding deviations from rationality has recently been proposed (Tsetsos et al., 2016, 2012), capitalizing on the fact that values are not readily available for individuals to encode. Under this view, subjective values for each option have to be constructed by means of accumulating values (for each of the available options) across a number of dimensions. During this accumulation period, decision-makers are proposed to prioritize the integration of values that rank highest on each dimension, leading to irrational preference reversals. However, little is known about the relationship that such a selective integration model could have on distracter effects.

In the original reports, the authors stress the fact that distractor effects are independent of whether the distractor option gets chosen. Even though higher distractor values are typically accompanied by a higher probability that the distracter gets chosen, the authors aim to control for this aspect of data by excluding trials where the distractor option gets chosen. However, there has been no systematic exploration of the effect that can result from excluding trials where the distracter gets chosen. In other words, there is no proof that the choices of a simple noisy observer are conditionally independent, such that the probability of one option over another is independent of the probability of choosing a third option.

Any departure from the independence of irrelevant alternatives axiom is at odds with the idea of a rational agent. However, we introduce a subtle but important distinction between weak and strong violations of the axiom. The weak violation refers to the fact that the increase in value of the distracter option results in a decrease in the proportion of choices for both other options, but without respecting the proportionality in the reduction of their original choice probabilities (which results in a change in the relative preference between the two). Instead, the strong violation, refers to an increase in the probability of choosing an option (whose value remains constant) that is accompanied by an increase in the value of a distracter option.

While a full characterization of the empirical choice probabilities would require prohibitive amounts of data, simulations can allow to characterise such choice probabilities for the mechanisms that are proposed to underlie choices. In this chapter, we provide a systematic exploration of choice probabilities under a variety of noise assumptions and decision-making models. Our simulations show that both weak and strong violations of the principle of irrelevant alternatives are achievable even under very simple assumptions of noise and even when decision-making agents have no explicit mechanisms of context dependent valuation. We also show that the experimenter's selection of a choice-set can allow them to recover positive or negative distracter effects even if a single mechanism governs choices.

2.2 Methods

We generated an arbitrary, but scale invariant, "choice-set" by systematically exploring values of three options A, B and D. First, we varied the value difference: the difference in value between two relevant options A, a high valued option, and B, a low valued option. We kept constant the mean value of the two options at an arbitrary value of 100 arbitrary units but varied the difference between them from 0 (no difference) to 100 (when the value of A is 150 and the value for B is 50), in one hundred linearly spaced steps. We orthogonally manipulated the value of D, the third, distracter option, to take values between 0 (when it is much lower than options A and B regardless of their value difference) and 200 (when it is much higher), also in one hundred linearly spaced steps.

First, we simulated a model whose choice variability is entirely driven by Gaussian noise in the representation of the value of the three options. In any "trial" (a particular triplet in the value of A B and D), the noisy representation of value for any one of the three options would be given by:

$$nV = V + \epsilon(0, \sigma^2) \qquad \text{Equation 1}$$

where V, is the true value of the option and ϵ is a scalar error value drawn from a zero mean Gaussian distribution with variance σ^2. For our simulations, we fixed the amplitude of the noise to have a standard deviation of $\sigma = 20$ arbitrary units.

We computed the unconditional probability of choosing an option, on each trial, as the probability that the noisy value of that option would be higher than that of the other two options as follows:

$p(Ch_{opt1}) = p((nV_1 > nV_2) \cap (nV_1 > nV_3))$ Equation 2

where nV_1, nV_2 and nV_3, are the noisy independent representations of the three options.

To obtain the efficiency for choosing between options A and B, for any of the models we simulated, we simply computed the ratio of probabilities for choosing the highest valued option, A over choosing either Option A or option B:

$$Eff = \frac{p(Ch_A)}{p(Ch_A)+p(Ch_B)} \qquad \text{Equation 3}$$

note that this is mathematically equivalent to computing the proportion of trials where option A was chosen over option B, after excluding cases where option D was chosen, for an infinite number of instantiations of Gaussian noise for each of the options.

The simulations we present could be approximated by relying on numerical simulations. With this method, a big enough number of random instantiations of noise is simulated for each trial and unconditional choice probabilities are estimated as a function of the proportion of the instantiations where each choice was selected. However, this method can be slow (requiring many instantiations of each "trial"), and the estimates of efficiency (i.e. the probability of choosing A over B, conditional on option D not being chosen) can be unreliable when the probability of choosing option D is very high. Instead, we computed the unconditional choice probabilities using an analytical solution:

$$p(Ch_{opt1}) = \int_{-\infty}^{x=\infty} \varphi(x|\, V_{1,}\sigma^2) \cdot \Phi(x|\, V_{2,}\sigma^2) \cdot \Phi(x|\, V_{3,}\sigma^2) \qquad \text{Equation 4}$$

where $V_{1,}$, $V_{2,}$,and V_3, represent the values of option 1, 2 and 3, respectively. Similarly, $\varphi(*)$ and $\Phi(*)$ represent the *probability density* and *cumulative distribution* functions, respectively, of a Gaussian process with the specified mean ($V_{1,}$, $V_{2,}$, or V_3) and variance equal to σ^2.

For the simulations under divisive normalization, the values themselves are transformed, prior to any noise, using the divisive normalization equation described originally (Louie et al., 2013), and noise is then applied to obtain the noisy value of any option:

$$nV_{\mathrm{i}} = \kappa \frac{V_{\mathrm{i}}}{\sigma H + \sum_{j=1}^{3} \omega \cdot V_j} + \epsilon(0,\sigma^2) \qquad \text{Equation 5}$$

where V_{i} is the actual value of the i^th^ option, and j is an index that goes through all three options, κ is a parameter to scale the value into a sensible firing rate range, σH is a semisaturation parameter and ω is a weight parameter that controls the degree of divisive normalization, when ω is zero the model has no divisive normalization. For our simulations we used the parameter values of the original report (Louie et al., 2013), and the option values we used are in a similar range to the original report, too. Thus, for our simulations $\kappa = 100$, $\sigma H = 50$, $\omega = 1$, and $\sigma = 8$. Note that σ, corresponds exactly with the "fixed noise" described in the original report and the value of σ was chosen to match the one used in their simulations.

Finally, for the simulations under selective integration, we assumed a non-linear transformation that gives a multiplicative boost to the value of the option that ranks highest. Importantly however, variability in the choices for this model is governed by noise that is decomposed in two parts: one that happens before the non-linearity and another one that happens after it. The process can be broken down as follows. First, the value of the three options is encoded in a noisy way (following Equation 1). For our simulations, we fixed the amplitude of the noise at this first stage to have a standard deviation of $\sigma = 10$. Second, we introduce a non-linearity: the maximum of the three noisy values is scaled by a parameter ω, which is typically assumed to be higher than 1, and for our simulations we assumed it to be $\omega = 2$. Thirdly, noise is again added to the three noisy values (i.e. to the two non-scaled values as well as to the one that was scaled). We controlled the amplitude of the noise at this stage to have a standard deviation of $\sigma = 20$. At the end, the highest of the three noisy values is selected for choice. Our decision of having a higher amplitude for late noise than for early noise is informed by previous reports about the main sources of human choice variability (Tsetsos et al., 2016, 2012; Wyart & Koechlin, 2016).

Note that selective integration is meant to occur when serially integrating the value across several attributes for the three options, where the parameter ω, controls the weight given to the maximum value for a particular attribute when it gets integrated into the overall value of the option that it belongs to. However, for direct comparison with the other models, and for simplicity, the simulations were carried out under the assumption that there is only one attribute to be integrated.

For estimating choice probabilities assuming the lapse rates, we again followed Equation 1. However, for any triplet combination, we split the probability of choosing the distractor option, equally among the other two options. Thus, bringing the probabilities of choosing both of the other options closer to 0.5.

2.3 Results

To obtain a benchmark, we first examined through simulations, the effect that the value of a distractor option, D, would have in the efficiency for discriminating two other options when variability in choices is merely driven by Gaussian noise corrupting independently each of the options' values. An agent is presented with a choice-set composed of options A (a high valued option), B (a low valued option) and D (a distracter option). This agent does not observe the values of the options directly, instead she observes a noisy version of each of them, nA, nB and nD, respectively. Where each noisy version of the options is drawn independently from a Gaussian distribution, with equal variance for the three options, σ^2, and whose mean is centred at the true value of each option. Efficiency is thus defined as the probability that nA would be higher than nB, conditional on nD not being the highest of the three (i.e. equivalent to computing the proportion of trials where nA was higher than nB after excluding trials where nD was the highest of the three noisy samples).

At first glance, there should be no reason for such a simple agent to exhibit a violation of rationality in her choices. To our surprise, we found that even this simplest of agents, whose choices are governed by nothing but Gaussian noise, revealed a violation of the principle of irrelevant alternatives in the form of a positive distracter effect: where higher values for a distracter option increase the efficiency for distinguishing the two other options (Figure 2.2).

Interestingly, the relationship between efficiency and the value of the distractor seems to be stronger in cases where the distractor takes values that are close to or above the value of the other two options. The violation of rationality observed from this agent, whose choice variability is governed simply by Gaussian noise, is counterintuitive because Gaussian noise is sampled independently to generate each option's noisy value.

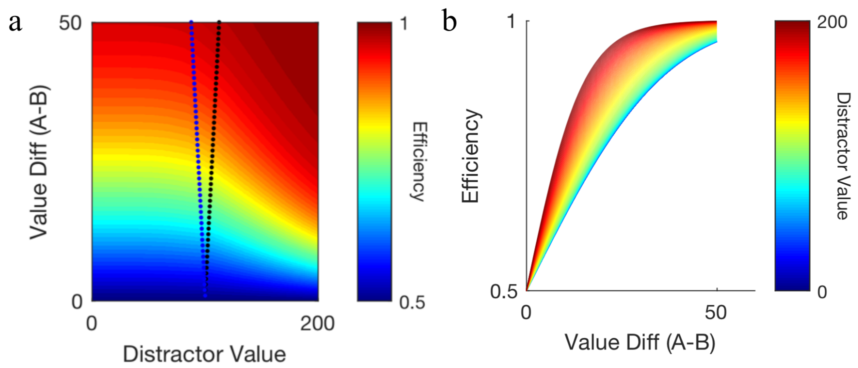

Figure 2.2. Violation of rationality by a simple noisy agent.

As expected, increasing value differences (between option A and option B) result in higher efficiency (i.e. higher probability of choosing the higher valued option over the lower valued one). Unexpectedly, however, efficiency also increases with the value of a distracter option. In panel a, the position on the abscissa of the black and blue dots represent the value of the higher- and lower-valued options for each value difference (position on the ordinate). In panel b, psychometric curves (efficiency as a function of value difference).

However, the effect arises because efficiency between options A and B depends on a third outcome: option D not being chosen. In other words, as the distracter option increases in value, the it will increase the probability that it gets chosen itself, but efficiency will not be driven by this probability but by the probability that either nA or nB end up having an even higher value that that of nD. Because B has a lower value than A, the probability that it will reach a higher (under Gaussian noise) decreases more rapidly than for A, driving thus an increase in efficiency.

Intuitively, this increase in efficiency for distinguishing between options A and B, can be present even if the unconditional probabilities of choosing both options decrease with higher values of the distractor. In other words, the violation of rationality could be weak, where the distracter option “steals” the probability that both options will get chosen but “steals” more from option B than from option A. Thus, we analysed how the unconditional choice probabilities for both options varied as a function of the distracter value. As we suspected, the change in efficiency for this simple agent, was the reflection of a weak violation of rationality: both probabilities decreased with increasing distracter values but they did not decrease in proportion to their unconditional choice probabilities (Figure 1.3). Finally, the relationship between efficiency for a fixed value difference and the probability that the distracter gets chosen, turns out to be monotonically increasing.

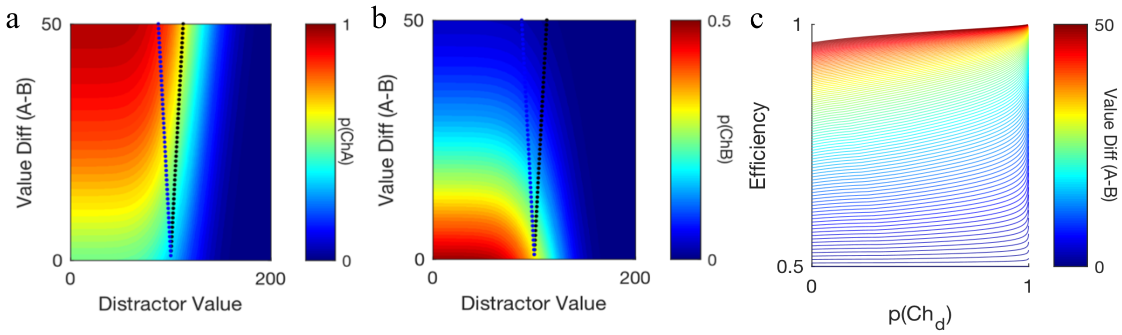

Figure 2.3. Weak violation from simple noisy agent.

The probability of choosing option A (panel a), but also option B (panel b), decreases as a function of higher distracter values. This reveals that the increase in efficiency is driven by a decrease that is not proportional to the original choice probabilities. Consequently, efficiency is a monotonically increasing function of the probability of choosing the distracter (panel c) for any fixed value difference.

Together, these simulations show that we need to rethink if the choice axioms, that we use as a yardstick to judge human behaviour, actually hold true as we incorporate realistic assumptions about the nature of human decision-making: if simply adding Gaussian noise to the process is a sufficient ingredient to violate rationality, there is no surprise if other mechanisms are also able to achieve the feat.

Divisive normalization has previously been shown to violate the principle of irrelevant alternatives (Louie et al., 2013). However, we examined systematically the effect that this mechanisms for relative encoding of value has on the efficiency for choosing between two options as a function of the value of a distracter option. Using the divisive normalization with the parameters reported before, we replicated the effect that efficiency drops with increasing values of a distracter option (Figure 2.4).

The intuition behind this effect is that divisive normalization occurs prior to any addition of noise, therefore, increasing values of the distracter option diminishes the values of all options making the effect of noise (of constant amplitude) on choices relatively larger, i.e. lowering the accuracy. A feature of this model that was previously reported but whose nature was not discussed, is that the change in efficiency reverts when the distractor approaches or surpasses the value of the other two options. Here, we replicate this reversal, and show that it a natural consequence of the addition of Gaussian noise that corrupts the values after being divisively normalized, thus driving the effect described in Figure 2.2.

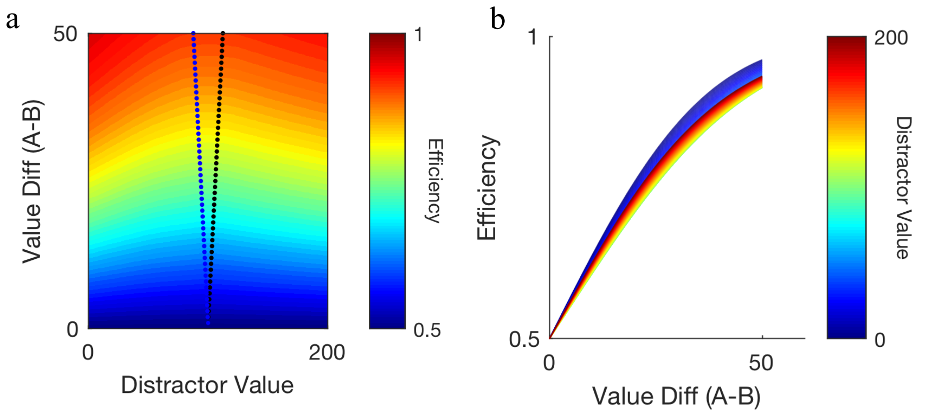

Figure 2.4. Efficiency under Divisive Normalization.

Under divisive normalization, efficiency for distinguishing between options A and B decrease due to the compression of value accompanied by higher distracter values. However, this increase in efficiency plateaus and starts to revert as the distractor value approaches or surpasses the value of the other two options. Axes and legends as in Figure 2.2.

We suspected that because divisive normalization is assumed to occur on the noiseless values of the available options, it would generate a strong violation of rationality. Indeed, while increasing distracter values always hinder the probability of choosing option A, their effect on the probability of choosing option B is non-monotonic: first boosting the probability of choosing option B, but dampening the probability as the distracter values approach the value of the other two options because the probability of the distracter being chosen increases. The relationship between efficiency and probability of choosing the distracter option, is thus non-monotonic for any given value difference (Figure 2.5).

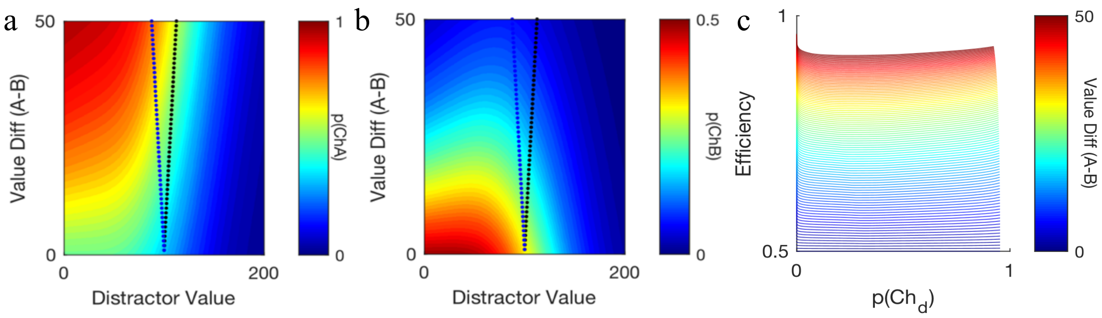

Figure 2.5. Strong violation for Divisive Normalization.

Legend as is Figure 2.3

Selective integration has previously been proposed as a mechanism driving violations of rationality (Tsetsos et al., 2016, 2012). Here, we explore systematically how selective integration relates specifically to the distracter effect. Typically, selective integration is implemented when there are several attributes to be considered for each option. Under selective integration, the highest value (across all options) in each of the attributes receives a multiplicative boost: which can be either conceived as a change in the representation of the value itself, or more typically, as an increased probability of being accumulated into the overall value of the option it belongs to.

In these simulations, for the sake of simplicity and for increased comparability with the other models, we assume selective integration to occur even when there is a single attribute for the three options. For the implementation of this model, we partition the total amount of noise that drives choices, into early noise (affecting the encoding of the values) and late noise (that which occurs after the selective boost is applied). More specifically, early noise refers to Gaussian noise conceived as corrupting the initial encoding of the values of the options (akin to the one described in Figure 2.2).

After the initial noisy encoding stage, a non-linearity occurs whereby a multiplicative boost is given to the noisy representation with the highest value of the three noisy options (on average the option with the highest value will end up being the one that receives the boost, but early noise can change the ranking of the noisy values). After this non-linearity, a further corruption of additive Gaussian noise occurs and the option with the final highest value is selected for choice. This latter stage of noise reflects noise that continues to build up as the decision evolves (in a multi-attribute paradigm, this noise could be mainly driven by imperfect computations during the accumulation of value across the attributes).

Akin to divisive normalization, selective integration drives both negative and positive distractor effects, with increasing values of the distracter first boosting efficiency but then decreasing it (Figure 2.6). As the distracter option increases in value, it is more likely to steal the multiplicative boost from option A, thus decreasing the distance in the representation between options A and B, and thus affecting their discriminability (i.e. decreasing efficiency).

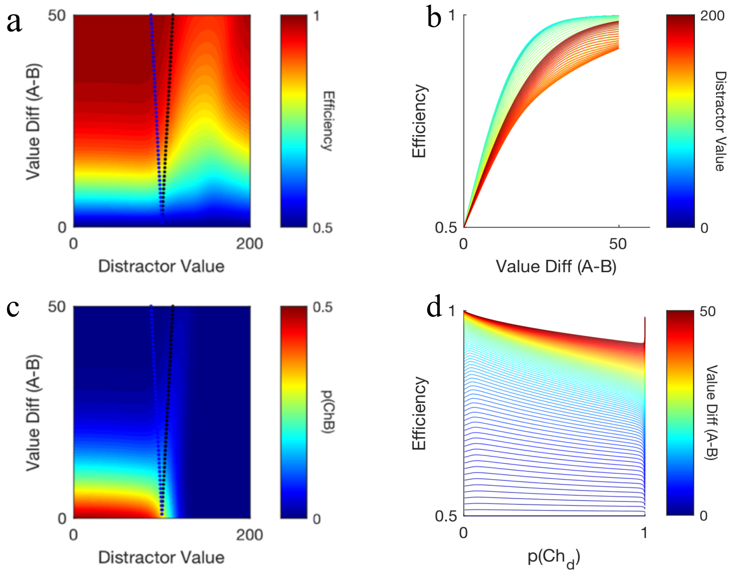

Figure 2.6. A strong violation from selective integration.

As the distracter option first increases in value, it reduces efficiency by means of stealing the boost from option A, but as its values increase further, efficiency starts increasing again because the its probability of being chosen increases (panels a and b). The violation from selective integration is strong, as the unconditional probability of choosing option B transiently increases with the value of the distractor (panel c). The relationship between the probability of choosing the distracter option and efficiency, is non-monotonic (panel d).

However, because a second stage of noise corrupts these representations, the option with the applied boost need not end up being the chosen one. Eventually, though, with much greater values of the distracter, its probability of being chosen increases much more, driving a positive distracter effect as described in Figure 2.2. For a fixed value difference between options A and B, the probability of choosing option A monotonically-decreases as the distractor increases its value (figure not shown). Conversely, increasing distracter values have a non-monotonic effect on the probability of choosing option B with a first transient increase (as option A loses its boost) and then an eventual decrease driven by the increase in the probability that the distractor is chosen itself.

An interim conclusion could be made that, while a weak violation of rationality could be achieved even by the simplest of agents, it requires an explicit mechanism of contextual value representation to drive a strong violation of rationality. However, we reasoned that other simple assumptions about the sources of noise governing decision could result in strong violations. For instance, if lapse rates increased as a function of the distractor option, then an option that is never preferred could increase it probability of being chosen as the distractor option increases.

As a proof of principle, to formalize this intuition, let us assume an experiment where the distractor is not only marked as a distractor, but it also cannot be chosen, such that the only accepted choices are one of the other two options. Let us also assume that an agent makes decisions on the basis of the noisy samples exactly in the way that is described in Figure 2.2 but with the following differences: when the distractor ends up having the highest noisy value out of the three options, then the agent “attempts” to choose that option only to realise that it is not allowed and resorting to choose randomly between the other two options. Note that this agent has two features that are taken to the extreme: (i) the agent that simply can’t block the value of the distractor option (a more realistic assumption would allow the distractor value to filter through on a small proportion of trials), and (ii) when the distractor would be chosen, the agent forgets the value of the other two options (a more realistic assumption would incorporate a shallower rate of forgetting). This agent exhibits a simple reduction in efficiency as the distractor increases in value (Figure 2.7).

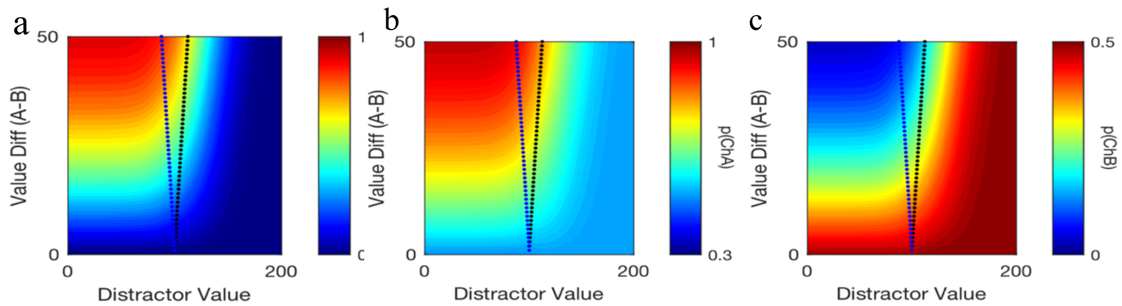

Figure 2.7 Strong violation of rationality when assuming lapses.

As expected, increasing the distractor value has a monotonic effect in decreasing efficiency (panel a). The increase in lapses that accompanies high distractor values, monotonically decreases the probability that option A will get chosen (panel b) but monotonically increases the probability that option B gets chosen (panel c).

This change in efficiency is driven both by the decrease in probability that option A gets chosen as well as the increase in the probability that option B gets chosen. We do not believe that the behaviour of this agent reflects an accurate depiction of human behaviour. Instead, we simply simulated it to show that if lapse rates increase with distractor values (a likely ingredient of a realistic agent), there is a tendency for strongly violating rationality even without explicit contextual distortion of value.

2.4 Discussion

Irrational decision-making has recently been described as positive and negative distracter effects, potentially originating from separate neural mechanisms (Chau et al., 2014; Louie et al., 2013). For both effects, participants are said to violate the principle of independence of irrelevant alternatives, whereby the value of a distracter option changes the relative preference between the other two options: with a positive effect referring to an increase of efficiency and a negative effect to decrease in efficiency.

Using simulations, we show that Gaussian noise alone in the process of representing the value of the options, is an enough assumption to violate this principle of rationality, by means of increasing the probability of choosing the highest of two options when a distracter option increases in value (and in its probability of being chosen). This surprising finding, forces us to reconsider if violations of the axioms are to be taken as meaningful representations of irrationality when they cannot even accommodate the simplest (and most canonically accepted) assumptions about decision-making.

At first glance, the violation of rationality for this simple agent can be disqualified for being simply a weak violation. That is, the distractor option does change the relative preference between the other two options (A and B). However, it does so not by means of enhancing the unconditional probability that any will be chosen, but instead by reducing the probability that both will be chosen in a manner that does not respect the initial proportion of choices. On the other hand, we show that mechanisms which explicitly affect the representation of value in a contextual way (i.e. the representation of each value is not independent of the representation of the other values), can result in strong violations of rationality. These violations take the form of complex enhancements and reductions of the positive and negative distracter effects.

These contextual models (i.e. divisive normalization and selective integration), are able to generate strong violations of rationality, i.e. those that do not hinge simply upon an unbalanced reduction of the probabilities that the two available options are chosen. Instead, they are able to increase the unconditional probability of one of the two options to be chosen. However, we show that other simple assumptions about the sources of noise, even without a contextual representation of value, can result in strong violations of rationality.

Finally, by systematically exploring the role of that a distractor value can have on the probabilities of choosing two other available options, in an arbitrary but scale independent range of values, we were able to reveal complex (non-monotonic) relations. This finding shows that an experimenter's choice of experimental design (in particular the distribution of choice problems in the choice-set), can determine whether an overall positive or negative distractor effect is found. Together, our findings (i) question the value that has been previously placed on violations of the axioms of rational choice (particularly the independence of irrelevant alternatives) as a standard for revealing irrational human choices and (ii) raise caution on the selection of choice-sets that are used to investigate distractor effects.

Chapter References

Bateson, M. (2002). Context-dependent foraging choices in risk-sensitive starlings. *Animal Behaviour*, *64*(2), 251–260.

Bateson, M., Healy, S. D., & Hurly, T. A. (2002). Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, *63*(3), 587–596.

Chau, B. K., Kolling, N., Hunt, L. T., Walton, M. E., & Rushworth, M. F. (2014). A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nature Neuroscience*, *17*(3), 463–470.

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

Hunt, L. T., Dolan, R. J., & Behrens, T. E. (2014). Hierarchical competitions subserving multi-attribute choice. *Nature Neuroscience*, *17*(11), 1613–1622.

Hurly, T. A., & Oseen, M. D. (1999). Context-dependent, risk-sensitive foraging preferences in wild rufous hummingbirds. *Animal Behaviour*, *58*(1), 59–66.

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, *110*(15), 6139–6144.

Luce, R. D. (1959). Individual Choice Behavior a Theoretical Analysis. John Wiley and sons.

Shafir, S., Waite, T. A., & Smith, B. H. (2002). Context-dependent violations of rational choice in honeybees (Apis mellifera) and gray jays (Perisoreus canadensis). *Behavioral Ecology and Sociobiology*, *51*(2), 180–187.

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

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, *109*(24), 9659–9664.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, *113*(11), 3102–3107.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*(4), 281.

Von Neumann, J., & Morgenstern, O. (1945). *Theory of games and economic behavior*. Princeton University Press Princeton, NJ.

Waite, T. A. (2001). Background context and decision making in hoarding gray jays.

Behavioral Ecology, *12*(3), 318–324.

Wang, X.-J. (2012). Neural dynamics and circuit mechanisms of decision-making. *Current Opinion in Neurobiology*, *22*(6), 1039–1046.

Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*, *11*, 109–115.

3 DECISIONS WITH THREE OPTIONS IN A SEQUENTIAL INTEGRATION PARADIGM

3.1 Introduction

Decisions are typically investigated in contexts in which the decision-maker is offered two options to choose from. These binary choices have served as the basis for most models of decision-making. There has been, however, recent interest in exploring the mechanisms underlying decisions with multiple options. In these multi-alternative contexts, humans and other animals show counterintuitive deviations from optimality, for example violating the principle of independence of irrelevant alternatives (Bateson, 2002; Bateson, Healy, & Hurly, 2002; Chau, Kolling, Hunt, Walton, & Rushworth, 2014; Huber, Payne, & Puto, 1982; Hunt, Dolan, & Behrens, 2014; Hurly & Oseen, 1999; Louie, Khaw, & Glimcher, 2013; Shafir, Waite, & Smith, 2002; Simonson, 1989; Tsetsos, Chater, & Usher, 2012; Tversky, 1972; Waite, 2001).

Most conceptualizations of value incorporate a subjective component. In some studies, the true value of an option is determined by a combination of a probability of obtaining a reward and the amount of the possible reward, but the actual value as estimated by participants is guided by nonlinear subjective transformations of those quantities prior to their combination (e.g. Chau et al., 2014; Kahneman & Tversky, 1979). In some other studies, the conception of value is entirely subjective. For instance, one recent study estimated the subjective value of a snack item to be the proportional to the maximum bidding that participants were willing to pay in order to obtain the snack (Louie et al., 2013). Most studies of multi-alternative decision-making investigate deviations of rationality by means of exploring inconsistencies in the subjective values that participants confer to the options in different contexts: for example, when choosing between two or between three options.

Most accounts of multi-alternative decisions assume that variability in choices is arises from how value is represented in the brain, or how these representations change in different contexts. These models hinge their predictions on the dynamic interactions between the representations of the choices, for instance, by means of divisive normalization of the values (Louie et al., 2013). However, this type of approach assumes that the value of an option is readily available to be encoded: the true values of the options are given as inputs to the relevant mechanism for comparing values. However, little attention has been given to the process of construction of value itself, and the mechanisms that underlie this construction process. For instance, the value of many real-world options is not composed of a single attribute but often several. Any scalar estimate of value of an option must be constructed by combining several pieces of evidence. Furthermore, even in the simplest of choice tasks, where evidence is seemingly unidimensional, the most successful accounts of behaviour see this process as a sequential acquisition of samples of evidence (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006; Ratcliff, 1978; Shadlen & Shohamy, 2016), as if the value of an options was being retrieved from discrete events in memory.

Evidence acquisition requires resources and is costly. It takes more time to gather more evidence. Further, the mere act of adding new evidence to a current estimate can incur extra sources of errors (Wyart & Koechlin, 2016). Resource allocation is, thus, likely to play a vital role during the process of value construction of the available options. We have previously shown that selective allocation of processing resources can confer benefits in terms of accuracy even if at first glance it goes counter to what a rational agent would do (Li, Castañón, Solomon, Vandormael, & Summerfield, 2017; Tsetsos et al., 2016). Mechanisms that exploit statistical regularities of the environment can potentially maximize accuracy in the face of shortcomings (e.g. limited processing capacity or sources). However, these protective mechanisms can result in counterintuitive violations of rationality, apparent lack of consistency, or seemingly illogical behaviour; when analysing in detail a subset of the trials.

We believe that during multi-alternative decisions, human behaviour is likely to reveal such apparent departures from rational behaviour, likely arising as a consequence of the mechanisms employed to optimize behaviour under constraints. For instance, it has recently been shown that divisive normalization falls out as an ingredient necessary when optimizing the speed-accuracy trade-off in decisions with multiple options (Tajima, Drugowitsch, & Pouget, 2016). We believe there is a key further aspect that is likely to constraint the accuracy of choices in multi-alternative contexts. Namely, when integrating values for multiple available options across several of their attributes participants will face an information overload. As a result of this information overload, pressure will arise for prioritizing what information gets processed or the quality with which it is processed. This prioritization can occur simultaneously in two related but independent frames of reference: (i) in a global frame, by focusing on ranges of value (for any attribute) that are more predictive of the overall value of an option, and (ii) in a local frame, by focusing on pieces of simultaneously available evidence are more likely to be relevant for choices.

Here we propose that people will exploit the statistical regularities in the environment to solve the prioritization problem described above. First, people will bestow higher emphasis on processing values towards the higher part of the range of values (rather than towards the lower part of the range) of any single attribute (Li et al., 2017). Second, people will also give more weight to that evidence which ranks higher in value compared to other simultaneously available evidence. Based on previous reports, we predict that they will do so by means of prioritizing the integration of the highest-ranking evidence into their estimate of the overall value of each option (Tsetsos et al., 2016, 2012).

To explore these hypotheses experimentally, we use a sequential integration paradigm, akin to the one used in the weather prediction task (Knowlton, Squire, & Gluck, 1994), but expanded to explore decisions between three options. Thus, evidence for three options is simultaneously available for any single attribute, but evidence across different attributes is delivered sequentially at a high rate (2-4 attributes per second). The fast presentation rate allows us to emulate the information overload that is likely to haunt most real life multi-alternative decisions. The paradigm allows us to control the evidence that is available for the decision-maker at any point in time. And this in turn allows us to quantify if people are systematically prioritizing the processing of different pieces of information.

Crucially, we predict that the paradigm will allow us to uncover a novel effect that a distracter option can have on choices: by virtue of truly distracting the process of value formation (and subsequent comparison) of the two other options. Specifically, we predict that the number of attributes in which an option ranks highest, will have a strong positive effect on the probability of choosing that option. Even when the distracter option's overall value may be much lower than that of the other two options, its value in some attributes might be more similar or even higher, thus competing with the other two for being integrated, and decreasing the resources spent in constructing the value of the other two options.

Across four behavioural experiments using a sequential integration paradigm with three alternatives, we show: (i) that the local (within-attribute) ranking of information plays an important role in guiding human decision making; and (ii) that a model of selective integration, whereby the most valuable item at any point in time is integrated more strongly, is able to capture quantitatively and qualitatively the human pattern of results better than a divisive normalization model. We also collected an online estimate of the pupil size as well as electroencephalographic (EEG) signals while participants performed the task (Experiments 1 and 4, respectively). These physiological recordings have previously been used to decode the values that participants use to guide their decisions (Cheadle et al., 2014). We collected these signals with the hope that we could find evidence for the two types of mechanism of efficient allocation described above.

3.2 Methods

One hundred and thirty-one human volunteers gave informed consent to participate in the study across four experiments (Experiment 1, n = 23; Experiment 2, n = 59; Experiment 3, n = 28; Experiment 4, n = 21). All reported normal or corrected-to-normal vision and no history of neurological or psychiatric impairment. The study was conducted in accordance with the guidelines of the local ethics committee.

3.2.1 Task

Participants participated in one of four sequential integration tasks (Experiments 1 to 4) in which they had to choose their most preferred (highest valued) option out of three available ones. In all tasks, the overall value of each option was defined as the arithmetic sum of its values across all attributes (8 attributes for Experiments 1, 2 and 3, and only 6 attributes for Experiment 4). The values of the attributes were presented one at a time in a sequence of samples (8 or 6 depending on experiment). Each sample contained the values for all three options (presented simultaneously) for a specific attribute. In Experiments 1, 2 and 3, the values of each attribute were indicated by the length of bars presented within 3 horizontally adjacent boxes placed at the centre of the screen. In Experiment 4, the values of each attribute consisted of single digits (one to nine) presented within circles arranged in the vertices of an invisible equilateral triangular array: left, up and right from the fixation point. The circles' centres were equidistant with one another and equidistant to the fixation cross. For all experiments, participants were informed to wait until the end of the sequence to choose the option (left middle or right) with the highest average value. Because of the similarity of the tasks, we present their results side-by-side, except when results differ across tasks.

3.2.1.1 Experiment 1

Stimuli were presented on a standard cathode ray tube monitor set to a display refresh rate of 60 Hz with a resolution of 1600 x 1200 pixels, at a distance of 60 cm. The experiment was conducted whilst participants were seated in a darkened room. Participants were instructed to place their head on a chin rest to minimize head movements. The height of the chin rest relative to the screen was kept constant across all participants but they adjusted their seat height for comfort.

Throughout the experiment, the pupil size and gaze position of participants were monitored using an Eyelink 1000 eye-tracking system, recording monocularly from the left eye at a sampling rate of 1000 Hz. A standard recalibration and validation of the eye tracker occurred every 84 trials (5 times throughout the experiment). During these procedure participants had to foveate 13 evenly spaced dots on the screen. The stimuli were generated and presented using Psychtoolbox (Brainard, 1997) and custom scripts in MATLAB (MathWorks).

On each trial, participants viewed three simultaneously occurring streams of visual stimuli, presented side-by-side in the centre of the screen against a grey background. Each stream was composed of 8 vertical bars of varying height occurring at a presentation rate of 4 Hz. Each bar was presented centrally within a corresponding rectangular placeholder that was the same width as the bar (400 pixels) but surrounded by a frame (10 pixels wide). Pixels falling in the residual space at the top and bottom of the placeholder were filled with an isolumiant colour; bars and placeholder were blue or brown (counterbalanced within participants across blocks, see Figure 3.1). Participants performed 10 blocks of 42 trials.

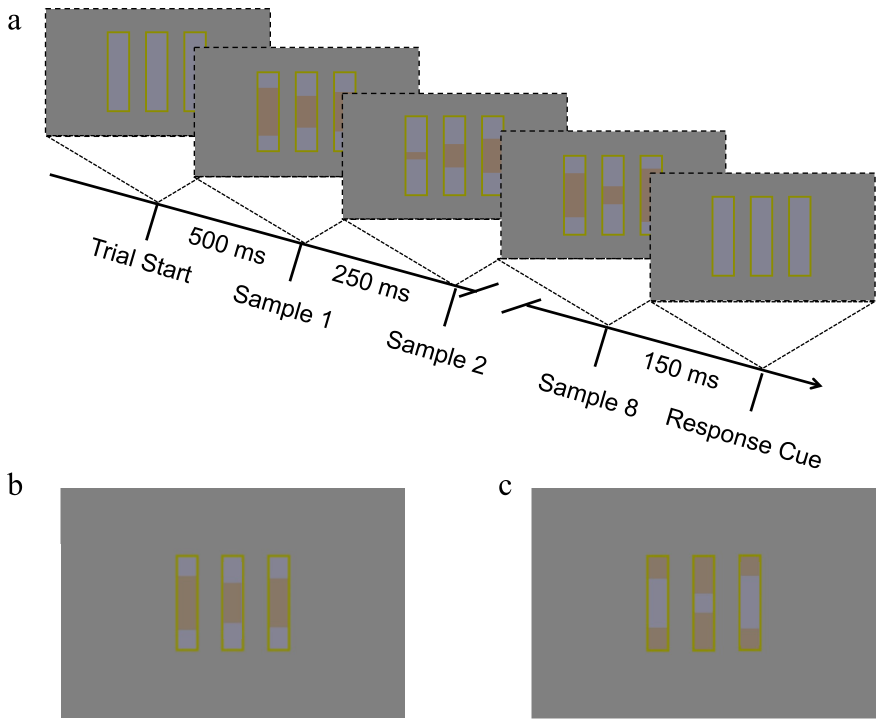

Figure 3.1. Trial sequence.

A schematic of the time course of a typical trial (panel a). The colours of the bars and box backgrounds were chosen to be close to isoluminant and were counterbalanced across blocks such that across the task, such that in half of the blocks (in random order) the background was blueish and the bars brownish (panel b) whereas in the other half of blocks the reverse was true (panel c), this manipulation guarantees that across the task the bars and the background are isoluminant.

After offset of the final bar, participants were asked to report their decision by pressing the left, middle or right key (on a standard QWERTY keyboard) to express which of the three streams (left middle or right, correspondingly) contained on average the longest bars. Following their response, the average bar length of the chosen stream was displayed as text at the centre of the placeholder. Participants were told that they would accumulate the points from their chosen option across the experiment, and that the number of points would translate into a monetary reward at the end of the experiment, with a maximum of £10 if they accumulated the highest possible number of points and £0 if they chose randomly, with a linear correspondence for the values in between.

On each trial, bar lengths for the low-valued, middle-valued and high-valued streams were sampled from Gaussian distributions with means μ_C, μ_B, μ_A, respectively, and a single standard deviation s_{all}. We began by drawing μ_B from a uniform distribution over pixels (range 152 to 248, in steps of 4). We then defined two levels of difference between the medium- and high-valued option (high+: $\mu_A = \mu_B +16$; high++: $\mu_A = \mu_B +32$) and two levels of difference between the medium and low-valued option (low-: $\mu_C = \mu_B -16$; low--: $\mu_C = \mu_B -32$). These were crossed in a 2 x 2 factorial design, yielding 4 trial types: high+/low-, high+/low--, high++/low- and high++/low--.

The standard deviation of the distributions was adjusted on a trial-to-trial basis (from an initial value of 80 pixels) using a staircasing algorithm: it was increased by 10% after two consecutive correct answers (when they chose the highest of the three boxes) and decreased by 10% after every incorrect answer (when they chose either of the lower valued boxes). For the purposes of staircasing (and monetary reward, see above) trials were defined as “correct” if participants chose the stream whose bars were on average longest, irrespective of the mean of the generative distribution from which they were sampled. The staircasing algorithm targeted a performance level of ~70%, in order to have a substantial number of error trials to increase the power of our logistic regression analyses (see below). The correspondence between spatial location (left, middle, right) and distribution (high-, medium-, low-valued) was randomised across trials. Experiments 2, 3 and 4 are very similar to Experiment 1, in the following section the differences are described.

3.2.1.2 Experiments 2 and 3

Participants were recruited using an online platform (Amazon Mechanical Turk). The conditions in which participants performed the task were not controlled by the experimenter. We modified the feedback in the task compared to Experiment 1, instead of showing them the amount of points from their chosen option, we gave them categorical feedback about whether that option was the best of all 3 or not. Thus, feedback was delivered by changing the colour of the frame for the chosen box into green, for correct, or into red for incorrect choices. To encourage engagement in the task and good performance, we also introduced a time penalty after incorrect responses compared to correct responses: if participants chose correctly the chosen box would turn green for half a second but if they chose incorrectly their chosen box would turn into red for 3 seconds. Participants were informed of this manipulation and that as a consequence the time taken to complete the task would depend on their performance.

In Experiment 2 the stimuli distributions were selected using the same algorithm described for Experiment 1. In Experiment 3, we generated the stimuli for option A and B the two highest valued options in the same way as described in Experiment 1. However, we used a different algorithm for generating the values of C, the lowest valued option, to allow it to reach values much lower in the range compared to the other two values. The distributions from which the values were drawn had a mean that was anchored to the value of option B, (the second highest valued option), such that the mean of the distribution from which values were drawn for option C was anything from 10% to 90% (in 20% steps) relative to that of option B.

3.2.1.3 Experiment 4

Stimuli were presented in controlled laboratory conditions, as in Experiment 1. However, in this task, the streams of information did not consist of bars contained within boxes. Instead, the streams were composed of a stream of 6 single digits (one to nine), contained in 3 circles (circular placeholders) arranged, close and equidistant to the fixation dot at the centre of the screen, in a triangular manner, the middle circle was positioned vertically above the fixation dot and the right and left at 120˚ clockwise and counter-clockwise, respectively. The stimuli were drawn from uniform distributions across the nine digits. Because of the extra demands of the task, the presentation rate was slowed, compared to Experiments 1-3 from 4Hz, down to 2Hz (i.e. two samples per second). Where the digits were shown inside the placeholders for 250 ms and the fixation alone remained on the screen for the following 250 ms when a new sample arrived.

3.2.2 Efficiency boxes

For visualization purposes, we present the average efficiency (i.e. the proportion of trials where option A was chosen over option B, after excluding trials where the distractor options was chosen) across participants for trials grouped according to two dimensions. The first one being the value difference between options A and B, which is the only variable that should affect efficiency. And the second one being any of a number of "task irrelevant" variables. Trials were split into 5 quantiles independently for each of the two dimensions. Average efficiencies were then computed for all trials in each of the 25 combinations of quantiles. The efficiency boxes plotted represent the average efficiencies across participants (with variability between participants not shown) in hues of red to blue, with red being perfect efficiency (i.e. 100% of choosing option A) and blue being chance (i.e. 50% of choosing either option). Because the quantiles were computed independently, not all combinations are warranted to have the same numbers of trials. We therefore added a layer of white tonalities (transparencies) to indicate the discrepancy in the number of trials: whiter boxes meaning fewer trials. The exact transparency was computed by dividing the number of trials in each box by the maximum number of trials across all boxes, and taking the logarithm of all these values.

3.2.3 Regression models

We used logistic regressions to predict the influence on the efficiency of choices: predicting whether option A would be chosen or not, after excluding all trials where the distracter option was chosen. First, we asked whether the overall value of the distracter option was able to exert any influence in the probability of choosing option A, the highest of the two other options, over option B, the lowest of the two. Our regression model included three predictors of choice: the overall value of option A, the overall value of option B and the overall value of option C:

$$p(Ch_A) = \boldsymbol{\phi}[\beta_0 + \beta_1 \cdot A + \beta_2 \cdot B + \beta_3 \cdot C] \qquad \text{Equation 6}$$

where $\boldsymbol{\phi}$, is a probit linking function. We also replicated the regression that was used in the original report of the positive distracter effect (Louie et al., 2013), using the overall value of option A, the overall value of option B and the normalized value of option the distracter option, nC (in line with the previous report, we computed the normalized value of the distractor by dividing its value over the sum of the other two values).

We ran another logistic regression model to investigate the effect of momentary (sample-to-sample) interactions of evidence onto choice. For this we included 5 predictors of choice (regressors). The first two being the overall values of options A and B, respectively. The momentary interactions were captured by three extra regressors. First, the number of samples for which option A ranked highest (out of the three) in the momentary evidence, $\#A_{1st}$. Second, the modulation of the momentary value of option A by virtue of the within-sample distance between this option and the distracter option, $A \cdot \delta_{AC}$. Similarly, for the third predictor, the modulation of the momentary value of option B by virtue of the within-sample distance between this option and the distracter option, $B \cdot \delta_{BC}$:

$$p(Ch_A) = \boldsymbol{\phi}\begin{bmatrix}\beta_0 + \beta_1 \cdot A + \beta_2 \cdot B + \beta_3 \cdot \#A_{1st} \\ + \beta_4 \cdot (A \cdot \delta_{AC}) + \beta_5 \cdot (B \cdot \delta_{BC})\end{bmatrix} \qquad \text{Equation 7}$$

Finally, we ran a third logistic regression. This time, with the aim of disentangling the effects of distance (described above) from the previously reported pro-variance effect (Tsetsos et al., 2012) whereby options whose values (across attributes) are more variable have an increased subjective value compared to options with less variability across their attribute values. In this regression, then, we included seven regressors: the overall value of options A and B, the modulation of their values by their average within-sample distance to the distractor, $A \cdot \delta_{AC}$ and $B \cdot \delta_{BC}$, and three extra regressors for the across attribute variability of each of the options VarA, VarB and VarC:

$$p(Ch_A) = \boldsymbol{\phi}\begin{bmatrix} \beta_0 + \beta_1 \cdot A + \beta_2 \cdot B + \beta_3 \cdot (A \cdot \delta_{AC}) + \beta_4 \cdot (B \cdot \delta_{BC}) \\ +\beta_5 \cdot (VarA) + \beta_6 \cdot (VarB) + \beta_7 \cdot (VarC) \end{bmatrix} \qquad \text{Equation 8}$$

An independent regression model was fit for each participant. The significance of the fitted weights (beta coefficients) were computed by computed a Student's t-test (two-tailed) on the weights for all participants on each experiment and each regressors independently. Because the significance of the regressors was either extremely high or really low, we did not compute any correction for multiple comparisons.

3.2.4 Model fitting

To explore what mechanisms allowed to capture the quantitative and qualitative pattern of responses observed from participants. We fitted 4 different models to our data by specifying in each model up to 4 free parameters: parameters that are assumed to take different values for different participants. We found the best fitting parameter set for each participant by using a minimization algorithm. Thus, we found the parameter set that would minimize the mismatch between the human and model responses across all trials (i.e. that would provide the maximum likelihood of observing the human responses given the model parameters) but independently for each participant. For this purpose, we used the Bayesian Information Criterion (BIC, also known as Schwarz criterion) as the measure of mismatch (Kass & Raftery, 1995):

$$BIC = -2\sum_{t=1}^{t=n} log\big(p(HCh_t|PS)\big) + log(n) \cdot \pi \qquad \text{Equation 9}$$

where the BIC is twice the negative sum across all trials (from trial 1 to the total number of trials, n) of the logarithm of the probability of observing the human choice on each trial, HCh_t, given a specific parameter set, PS. The last term penalizes models according to the number of parameters they have, where π is the number of parameters.

For performing all model fitting procedures, we used the genetic algorithm available in Matlab's optimization toolbox. We used the following algorithm parameters (irrespective of which model was being fitted): a population size of 100 individuals, a maximum generation time of 1000 generations, and a function tolerance of 0.01. The initial parameter ranges for all parameters were from zero to three, with an absolute lower bound of zero (as none of the parameters has sensible interpretation for non-positive values) but no absolute upper bound.

In order to find the best fitting parameter set (for each participant and each model, across all trials), we needed to be able to compute, for each trial, the probability of observing the same choice as the human participant given the parameter set of the model in question as shown in Equation 4. To find the best fitting parameters and perform the quantitative comparison between models, we made use of all trials and did not exclude trials where the distracter option was chosen.

For brevity, we first describe the common aspects of all the models before describing the specific features of them. The first common aspect of the models is that momentary evidence needs to be integrated across samples in order to reach a choice. However, informed in previous sequential integration paradigms (Cheadle et al., 2014; Tsetsos et al., 2012), we allowed for a free parameter of exponential decay, α, that allows earlier or later samples in the trial sequence (typically later ones) to be more influential on choice. The second common aspect to the models is that after the sum of momentary evidence across samples is carried out, a noise parameter controls the standard deviation of the Gaussian noise added to the final tally for each option. This noise parameter controls the level of variability or stochasticity in the choices.

Thus, irrespective of the model, the final noisy tally for any option, nT_i, will be determined by:

$$nT_i = \epsilon(0, N^2) + \sum_{s=1}^{ns}(MV_{is}) \cdot \left(\alpha^{(ns-s)}\right) \qquad \text{Equation 10}$$

where, $\epsilon(\cdot)$, is a number randomly drawn from a Gaussian distribution with zero mean and standard deviation, N; MV_{is}, represents the momentary value of option i at sample s, where $s = 1$ is the first sample in the trial sequence and ns is the total number of samples in the trial sequence. An exponential decay parameter, α, controlled the influence on choice of evidence according to its temporal position within a trial. To compute the probability of observing the same choice as the human participant in a specific trial, given a parameter set, we determined the probability that the noisy tally of the option chosen by the participant is the maximum of all three noisy tallies (as described in the previous chapter). We now describe the specific parameters for each model.

For the divisive normalization model, we used the equations described in the previous chapter, to compute the momentary values of the evidence, MV_{is}, leaving as model specific free parameters the semisaturation parameter, σH, and the normalization weight, ω. Note that instead of adding noise to the normalized values on each sample, we simply added noise after all samples had been integrated. This model has a total of 4 free parameters.

For the selective integration model, to obtain the momentary values of the evidence, MV_{is}, we simply scaled the maximum momentary evidence by a free parameter, W_1, and the lowest momentary evidence by a second free parameter, W_2. Again, in this model noise was added after all samples had been integrated. This model also has a total of 4 free parameters.

For the power model, we power transformed all momentary evidence by means of a free parameter P, controlling the exponent of the power transformation. Note that this model has a total of 3 free parameters only.

Finally, a power and selective integration model, incorporate some aspects of the power and of the selective integration models: an initial power transformation of the values was implemented by means of a free parameter P. Then, the maximum of the power transformed momentary evidence was scaled by a weight parameter W. Note that for this model only the maximum momentary evidence receives a multiplicative boost. Thus, this model also has 4 free parameters.

Quantitative model comparison, was performed using the difference in the Bayesian Information Criterion (BIC) of the best fitting parameter set for each participant and each model. So, for any pair of models we could compute the subject by subject as well as average BIC difference. For completeness, we also computed the Akaike Information Criterion (AIC) which penalizes less strongly models with a higher number of free parameters. For this, we simply replaced the last term in Equation 4, by the number of parameters multiplied by a factor of two (Burnham & Anderson, 2003; Kass & Raftery, 1995).

3.2.5 Eye-tracking data collection and pre-processing

Throughout Experiment 1, the pupil size and gaze position of participants were monitored using an Eyelink 1000 eye-tracking system, recording monocularly from the left eye at a sampling rate of 1000 Hz- With the use of a chin-rest, the height of the eyes relative to the centre of the screen was kept constant across all participants (but they adjusted their seat height for comfort). A standard recalibration and validation of the eye tracker occurred every 84 trials (5 times throughout the experiment). During this procedure, participants had to fixate their gaze (foveate) on thirteen evenly spaced dots on the screen appearing consecutively.

Blinks were determined either by missing data from the eye-tracking device or by outlying rates of change in the pupil size: for each time point (across the whole experiment for each subject), the difference between the previous and the next time points was taken. We then detected samples with momentary differences that were 3 standard deviations away from the zero or further (in either direction) and defined them as missing points. We ignored the 10 ms before a blink onset and the 10 ms after a blink offset, also defining them as missing points. A custom interpolation method was used during periods of missing data, to predict independently: (i) the pupil size, and (ii) the horizontal and

(iii) the vertical positions of the gaze. This interpolation procedure, was carried out using a shape-preserving piecewise cubic interpolation (using Matlab's *interp1* function specifying *pchip* as the interpolation method).

High-frequency components were removed through data smoothing using a sliding uniform window (using Matlab's *smooth* function specifying *moving* as the method for smoothing). The size of the smoothing window used was 51 ms in length. Only for the pupil size data, a slow drifting component computed as above but using a much longer window of 15 seconds, was subtracted from the time series. Finally, the data from the pupil size and the gaze position were down-sampled to 100 Hz, by taking the median value of the 10 ms of each window.

3.2.6 Pupillometric analyses

Two different analyses were carried out to describe the changes in pupil size as a function of task relevant factors, with particular interest to testing the hypothesis that the momentary value of the three options on the screen would be encoded differently as a function of their rank. The first analysis, following from Cheadle et al., 2014, does not assume any particular shape for the response to a stimulus. The second analysis is based on using the canonical pupil impulse response function and convolving it with a time series of ones and zeros were ones indicate the occurrence of a specific task relevant event (de Gee, Knapen, & Donner, 2014; Lempert, Chen, & Fleming, 2015; Wierda, van Rijn, Taatgen, & Martens, 2012). This analysis assumes a fixed shape of responding (only modulated in terms of its amplitude by task relevant factors) and is suited well for responses that overlap in time with one another.

For the first analysis, we fitted an independent multiple-linear regression to predict the pupil size at each of the time points (relative to sample onset). We ran a separate regression for each of the samples of each of the subjects, using the values of our task relevant factors as regressors. To find the main effects of our task relevant factors, we first computed the average beta weights across all samples for each of our subjects; we then found the mean and standard error of the mean for the sample-averaged beta weights across all subjects. This approach is equivalent to fitting an independent linear mixed-effects model (LME) for each time point, using our task relevant factors as fixed effects and using subjects and samples as random interacting effects, thus allowing each of the fixed effects to have a different contribution (slope and intercept) for each of the samples of each of the subjects.

We analysed the period spanning 6 seconds starting from 1 second before the onset of each sample. The multiple linear regression model fitted to the pupil size, P_{size}, at each time point for each sample for each subject, was as follows:

$$P_{size} = \begin{bmatrix} \beta_0 + \beta_1 \cdot (Max) + \beta_2 \cdot (Med) + \beta_3 \cdot (Min) \\ + \beta_4 \cdot (HDist) + \beta_5 \cdot (VDist) \end{bmatrix} \quad \text{Equation 11}$$

where Max, is the value of the bar with the maximum length (regardless of its spatial location on the screen), Med, is the value of the bar with the median length, and Min, is the value of the bar with the minimum length of the three bars on that sample. People fixated close to the centre of the screen throughout the sequence presentation (Figure 3.2). However, even when gaze position remained relatively quiescent throughout the sequence of evidence was shown, participants fixated on the option they chose (as it revealed the average length of the chosen bar). Thus, we included two additional regressors to control for gaze position: (i) the horizontal distance of gaze to the centre of the screen, $HDist$, and (ii) the vertical distance of gaze to the position of the centre of the screen, $VDist$.

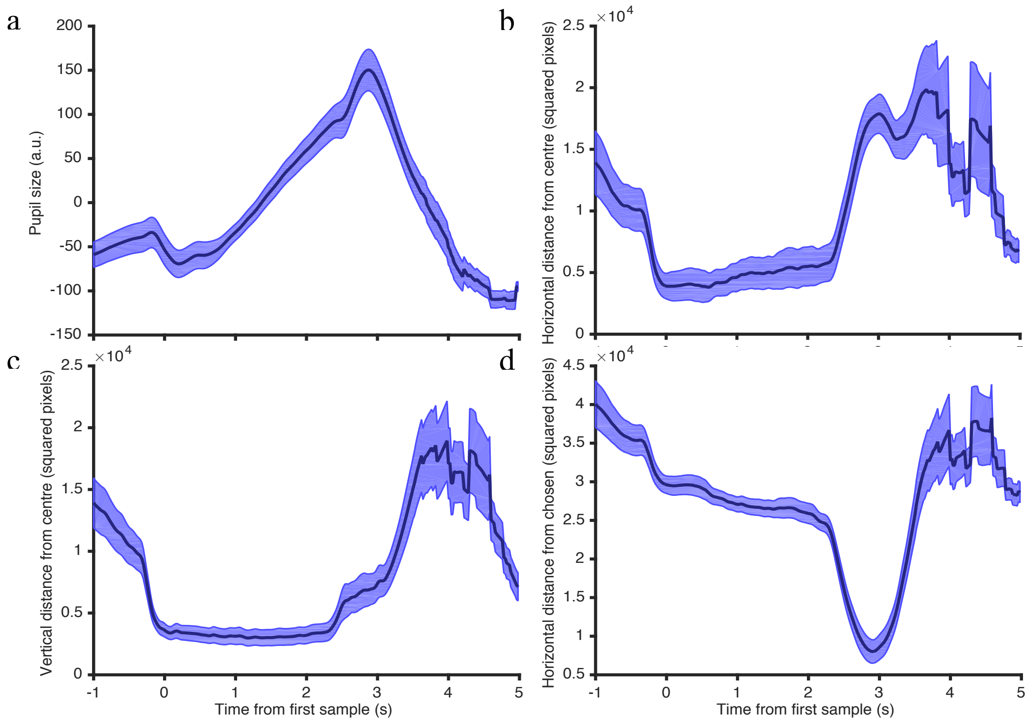

Figure 3.2. Averaged time series of eye tracking signals.

We plot the mean (dark line) ±SEM across subjects (light shadow) of 4 different measures derived from the eye tracker. All time series are locked to the onset of the first sample. The sequence of samples started at 0 seconds and finished 2 seconds later. Pupil size relative to participant across-task average (panel a). Distance of gaze on the horizontal axis from the centre of the screen (panel b). Distance of gaze on the vertical axis from the centre of the screen (panel c). Distance of gaze on the horizontal axis from the position of the chosen option (panel d).

For the second analysis, we defined trial epochs spanning 6 seconds, starting 1 second before the first sample onset of each trial. We concatenated the vector of pupil size for all our trial epochs into a single pupil vector. We then generated different predictor vectors. To construct each vector we generated predictor epochs filled with zeros except for the timings of occurrence of each sample. We generated 4 predictors, the first one was just the original vector denoting the mere occurrence of each sample, the other three predictors were modified versions of the original. For these other predictors we scaled the occurrence of each sample by: 1) the value of the maximum bar on that sample, 2) the value of the minimum bar and 3) the value of the minimum bar. The values of the three bars were normalized across the experiment before being added into the predictor vectors. We then convolved each predictor epoch with the canonical pupil impulse response function. And finally, we concatenated all the epochs of each predictor into 4 different predictor vectors. The impulse response function was computed using the equation below:

$$h(t) = s \cdot t^{w} \cdot e^{(-t \cdot w / t_{max})} \qquad \text{Equation 12}$$

where w controls the width and t_{max} controls the time (in milliseconds) to peak of the impulse response function, and s is just a scaling constant. For congruity with previous studies (de Gee et al., 2014; Lempert et al., 2015; Wierda et al., 2012), we set $w = 10.1$, $t_{max} = 930$ ms, and $s = 1/10^{27}$. The convolved predictor vectors were then used as predictors for the pupil vector (using Matlab's in-built *glmfit* function).

3.2.7 Electroencephalography data: collection and pre-processing

For Experiment 4. We recorded electroencephalographic (EEG) signals, while participants carried out the task. For this, we used a Neuroscan system. We used Synamps-2 amplifiers to amplify the EEG signals. We recorded from 60 Ag/AgCl electrodes in the locations shown in Figure 3.3. Additionally, four electrodes were placed in a bipolar montage: two as horizontal and vertical electro-oculograms (EOGs) and two electrodes (one in each mastoid) were used as reference. All electrode impedances were kept below a limit of 50 kΩ, or their signals discarded when this was not possible. We kept the impedances of the electrodes as low as possible and in general below 25 kΩ. The digitalization of the data was carried out at a frequency of 1KHz, with an online high-pass filter of 0.1Hz.

After collection, the data was pre-processed before analysis using custom scripts and the EEGLAB toolbox (Delorme & Makeig, 2004). First, the EEG data was further down-sampled to 250Hz. We subsequently used a band-pass filter to retain only frequencies between 1 and 30Hz.

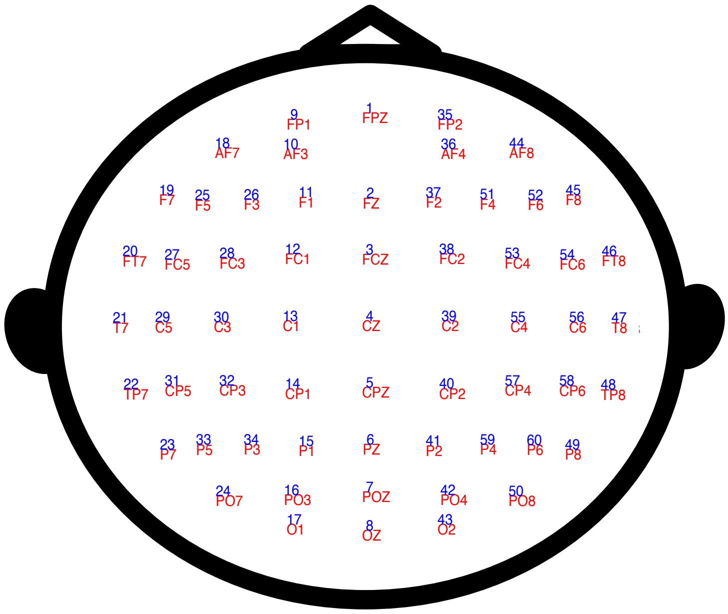

Figure 3.3. Electrode locations.

The approximate position of the 60 EEG electrodes is shown: each combination of electrode number (blue numbers) and electrode label (red text) demarks the approximate position of the electrodes relative to the head as seen from above (nose and ears are drawn for reference). An occipito-parietal region was defined as the group of eight most occipital (posterior) electrodes (electrode numbers: 7,8,16,17,24,42,43,50).

The data was cut into discrete chunks of time (epochs), spanning 6 seconds staring 1 second before the onset of the first sample of each trial. Visual inspection of all epochs was carried out to both (i) remove epochs with non-stereotypical artefacts from further analyses, and (ii) remove electrodes with electrical artefacts from further analyses. An automatic rejection of noisy data (from the EEGLAB toolbox) was carried out, for this, Independent Component Analysis (ICA) was run to detect and reject any abnormally noisy signals. We re-inspected the individual epochs to ensure no artefacts remained after the ICA rejection. The rejected channels were interpolated also using the standard parameters of the EEGLAB toolbox. Finally, all the channels were re-referenced, using the average of all channels as reference.

3.2.8 Single trial EEG decoding analyses

For analysing the EEG signals, we followed a similar decoding approach to that used for the pupillometric signal. We looked at how the EEG activity could be predicted as a function of the value of the digits shown on the screen by means of multiple-linear regression. We first defined a window of interest around the presentation of each evidence sample (a triplet of digits), the window spanned 700ms starting 200 ms before the onset of the sample and ending 500ms later. We then fitted an independent linear regression to predict each of the timepoints of each of the samples for each of the electrodes, by means of different groups of predictor variables (regressors).

The used three main groups of regressors. The first group consisted on three regressors with the momentary (local) value of the 3 options for the sample in question, and sorted according to their local ranks:

$$EEG_{amp} = \beta_0 + \beta_1 \cdot (Max) + \beta_2 \cdot (Med) + \beta_3 \cdot (Min) \qquad \text{Equation 13}$$

The second group or regressors consisted of nine dummy regressors signalling the presence or absence on the sample of each of the nine possible digits:

$$EEG_{amp} = \beta_0 + \beta_1 \cdot (D1) + \beta_2 \cdot (D2) + \cdots + \beta_9 \cdot (D9) \qquad \text{Equation 14}$$

where the regressors D1 to D9 represent dummy variables signalling the presence or absence of digits one to nine, respectively, within the sample of evidence at question. The last group of regressors consisted of the twelve regressors described above put together into a single regression.

We average the regression weights across samples (the sequence in each trial contained six samples), and across a group of Parietal and Occipital electrodes (O1 Oz O2 PO7 PO3 POz PO4 PO8).

3.2.9 Cluster correction for multiple comparisons

To estimate the significance of the resulting regression weights for the physiological signals (EEG and pupil size), we performed a cluster-based approach to correct for multiple comparisons (Maris & Oostenveld, 2007). Briefly, we generated a null distribution of regression coefficients. We did this by using the same regressors (thus preserving the statistics of the regressors) to predict the same timepoints of the physiological signals (thus preserving the temporal correlations in the timeseries) but for random pairings across trials for the predictor (regressors) vs the predicted (either the pupil size data or the anyone channel of EEG) variables.

We then averaged the timeseries of regressors across samples (eight samples in Experiment 1; six samples in Experiment 4), and in the case of EEG signals across the relevant channels. Afterwards, we computed a two-tailed Student's t-test on the average weight of each timepoint (across participants), to obtain a statistical result of how the average weights differed from zero. Next, we looked for "clusters" of significance by grouping all adjacent timepoints (for each of the timeseries of regressor coefficient) whose p-value was below a criterion of 0.01. We recorded which cluster had the highest sum of t-values, and constructed a null-distribution of the highest t-values expected under the null-hypothesis for each regressor.

Finally, we compared the sum of the t-values for all clusters in the regression using unshuffled data against the distribution of highest t-values obtain under the null-hypothesis. This allowed us to in which percentile of the null-distribution they each cluster of each regressor would fall. Only clusters that were above 95% of their null-distribution were marked as significant.

3.2.10 Analysing protection against noise from models

We created a hypothetical sequential integration experiment. We simulated 10,000 trials. In each trial, six samples of evidence, each containing evidence of the value for each of the three options (just as in our experimental paradigm), were presented sequentially. The evidence of in each trial was drawn randomly from a uniform distribution of integer values between 1 and 9 (just as the digits in our task), independently for each option and each sample. In each trial, we added an offset to all the evidence (equal to all samples of all options). This offset was different on every trial, the offset was selected randomly from a uniform distribution of integer values between 1 and 27. For convenience, we divided all of the randomly generated values by the maximum value obtained (so that all the values are constrained between zero and one).

We then modelled the choices (and thus accuracies) that would result under a variety of scenarios. We explored three different models: (i) Selective Integration, (ii) Power, and (iii) Power and Selective Integration. We explored the effect of adopting different parameters for the weight of the maximum option and of the power transformation (both varied between the values of 0.1 to 5 in 30 equidistant steps). We also explored the effect of post integration Gaussian noise, i.e. noise that occurs either (i) as the samples of evidence are tallied in time, or (ii) after the evidence is tallied. We explored Noise amplitudes with standard deviations spanning from 0.01 to .3 in 30 equidistant steps. The choice probabilities for each model were obtained using the equations described in the previous chapter of the thesis. For the combined Power and Selective Integration model we fixed the level of noise at a value of 0.1 and explored the effects of the other parameters in the same manner as described above.

The assumption motivating these simulations his analysis is that both Selective Integration and a Power transformation, are mechanisms that allow an observer to allocate their processing resources efficiently. This efficient allocation, creates some distortions in the representation of value, but can potentially prove beneficial in the face of post-integration noise. However, scaling the local maximum by a value greater than one, or transforming all evidence using a power exponent greater than one, can lead to trivial explanations for this protection against noise. Thus, we decided to control for this spurious inflation of evidence. For Selective Integration, we decided to keep constant the sum of the weights given to the three pieces of evidence in one sample by dividing each of the three weights by the sum of all weights. This ensures the total applied gain remains constant even when it is allocated differently according to the local ranks. Similarly, for the Power model, we divided all the transformed values by a gain factor:

$$g = 1/(1+P) \qquad \text{Equation 15}$$

where P is the exponent of the power transformation. This g factor captures the total scaling applied to the space in which the evidence values lie. Dividing by this factor ensures the total gain applied is constant, even if there is uneven allocation of that gain in the higher parts of the spectrum (with exponents greater than one) or in the lower parts of the spectrum (with exponents lower than one).

3.3 Results

3.3.1 A sequential integration paradigm with three options available options.

We labelled the three options in every trial according to their overall value (from highest to lowest) as A, B and C, with C being the lowest valued of the three option and referred henceforth as the distractor. The choice probabilities observed by our human participants averaged over experiments ranged between 70 and 72% for option A, between 20 and 26% for option B, and between 3 and 11% for option C. These choice probabilities are very similar to the ones recently reported in in other experiments (Chau et al., 2014; Hunt et al., 2014; Louie et al., 2013). We found response times to be on average the fastest when participants choose option A, then when they choose option B and the slowest when they choose option C. Average choice probabilities and response times with their standard error of the mean (SEM) are reported in

(in the Tables and Appendices section of this chapter).

We first looked at the effects of the overall value of the options on choice. With particular emphasis in investigating if we could recover a general distractor effect in our experiments (i.e. if the overall value of the distractor affected the preference for the other options). Not much evidence was found for either a positive or a negative distractor effect. As seen from visual inspection of the average efficiency (distinguishing between options A and B) when splitting trials according to five quantiles for two criteria: (i) the value difference (average value difference of options A and B) and (ii) average distractor value (Figure 3.4).

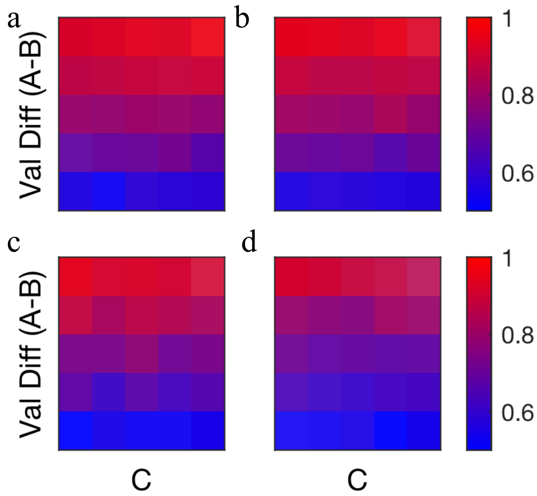

Figure 3.4. No classical distracter effect in human responses.

As expected the quantiles with greater value difference (y-axes) are associated with greater (towards red) levels of efficiency in choosing option A over option B. Trials where the distracter option gets chosen were excluded. However, differences in the value of the distractor option, C (x-axes), provoked little or no systematic change in efficiency. Panels a, b, c and d correspond to Experiments 1, 2, 3 and 4, respectively.

To formalize the analysis, we ran a logistic regression to quantity the effect that the overall value of A, the overall value of B and the overall value of C had on the probability of choosing option A (after excluding trials were the distractor got chosen). For all experiments, strong positive and negative effects were found for the overall values of option A and B, respectively (all t-values > 15, all p < 0.001, see Table 1. Proportion of choices and response times.

Table 2. Regression with value of distractor.

). With respect to the effect of the distractor value on choice, we found no significant effect for Experiments 1, 3 and 4 (all t-values < 0.6, all p > 0.5), and a significant but very small negative effect was found for Experiment 2 (t-value 2.15, $p < 0.05$, likely a false positive given weakness effect). A similar pattern of responses was observed when replacing the value of the distractor by its normalized value (see

Table 3. Regression with normalized value of distractor.).

3.3.2 Within-sample interactions of evidence across the three options.

A feature of our experimental design is that it allows us control over the evidence that is available to be evaluated at any point in time. We ran an additional regression model to explore the influence of the interaction of the momentary evidence onto choice. Thus, as well as including the overall values of options A and B, we included three extra regressors, one that captured the effect of the local (within sample) ranking of evidence, and two more that reflected the influence of the distractor as a function of its proximity to the other two options in the momentary values.

For the ranking regressor, we quantified the average ranking across samples (for each trial) for the momentary value of option A, r_A, with lower values reflecting more frequently ranking highest. We also reasoned that determining the local ranking of the values for the options would be harder when these values are more similar. We computed the average distance (unsigned difference) across samples between the distractor option and both option A, δ_{AC}, and option B, δ_{BC}. We then looked at how the distance between each of the options and the distractor was able to affect the value assigned to each of these options, using the following interactions: $A \cdot \delta_{AC}$ and $B \cdot \delta_{BC}$. Visual inspection of average efficiency (splitting trials according to their average value difference and each of the measures described above), revealed that the three seemingly irrelevant quantities were able to exert strong influences on choice (Figure 3.5).

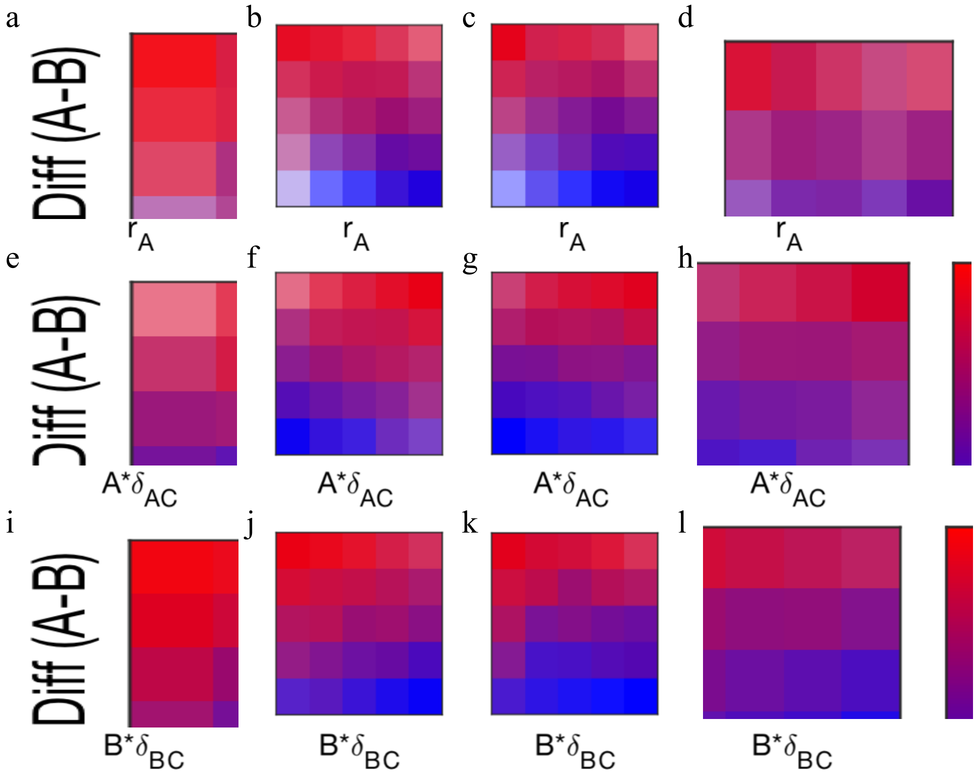

Figure 3.5. Within sample interactions of evidence affect choices.

Efficiency depended not solely on the value difference between options A and B (y-axes all panels). Three other factors (one for each row of panels) were also found to affect efficiency (x-axes). Each column of panels corresponds to an experiment. First, higher average rankings for option A (lower frequency of local winners) resulted in lower efficiency (panels a,b,c and d). Second, higher momentary distance between options A and C acted as a boost to the value of A, resulting in higher efficiency (panels e,f,g and h). Third, higher momentary distance between options B and C acted as a boost to the value of B, resulting in lower efficiency (panels i,j,k and l). Trials where the distracter option was chosen were excluded.

Indeed, the results from the logistic regression were able to confirm these predictions. We found a significantly negative effect of the average rank for option A: as the average ranking of option A increased (fewer times ranking highest) efficiency decreased for all experiments (r_A: all t-values > 3.17, all $p < 0.01$, see). We also found a significantly positive modulation in the value of option A by the distance between this option and the distractor ($A \cdot \delta_{AC}$: all t-values > 10.6, all $p < 0.001$), and between option B and the distractor ($B \cdot \delta_{BC}$: all t-values > 11.5, all $p < 0.001$). These distance effects can be seen as if the distractor competes on every sample against the values of the other two options, such that if it is close to either option, it steals value away from that option.

3.3.3 Modelling of behaviour.

3.3.3.1 Description of models and best fitting parameters

In order to see which aspects of our data can be accounted for by different mechanisms, we decided to fit 4 different models to our participants' choices. One model, Selective Integration, aims to capture the dynamic allocation of attention that we hypothesized to happen locally during the presentation of each sample, whereby the three values on the screen are given a differential weight determined by their ranks. The second model, Divisive Normalization, aims to capture the potential relative encoding of the values that are presented on the screen on each sample. Thirdly, the Power model aims to capture a global preference for attending big values compared to low ones. And fourthly, the Power and Selective Integration model incorporates features of the two models, with both a general bias for attending big values (globally) and dynamic allocation of attention towards each sample's highest value (locally).

We carried out the fitting of these models by finding the parameter set for each model that minimized the mismatch between the trial-to-trial choices between the model and each of the participants. The methods section of this chapter contains the full explanation of the models, but a brief explanation is provided next. All models have two common free parameters: (i) a noise parameter, N, controlling the amplitude of the stochasticity of choices after the integration of values across all samples (standard deviation of Gaussian noise), and (ii) an exponential decay parameter, α, controlling the degree of “leak” or forgetting when integrating the momentary value across samples within a trial. Each model has two extra model-specific parameters explained below.

The **Selective Integration** model has two key free parameters that control the gain given to pieces of available evidence as a function of its momentary ranks: a weight of W_1 is applied to the highest momentarily available evidence (henceforth local winner) and a weight of W_2 to the lowest (henceforth local loser). Note that a third weight is not needed, as the momentarily medium evidence is used as a reference point with a constant weight of one. The **Divisive Normalization** model also has two key free parameters related to the degree of normalization in the momentarily available evidence: a semi-saturation parameter σ_H, and a weight parameter W (which when taking values of zero indicates no normalization). The **Power Model** consists on a static non-linear transformation of the momentary evidence before being accumulated. This transformation is controlled by a single parameter P, a power exponent, which parameterizes whether a boost is given to higher values (when P >1) or to lower values (when P <1).

Finally, both the Power and the Selective Integration models, make predictions about high numbers being more influential on choice: one by means of a dynamic shift of attention on every sample (locally), the other by means of a fixed global preference for attending high values over smaller ones. We decided to construct a model that incorporated both of these ingredients. Thus, the **Power and Selective Integration Model** incorporates both an initial, power transformation that is applied globally to the evidence (controlled by a parameter P), then a multiplicative boost, controlled by a W parameter, is given solely to the evidence which ranks highest of the three (locally). A summary of the best fitting parameter sets for each experiment and each model (across all participants) is given in Figure 3.6.

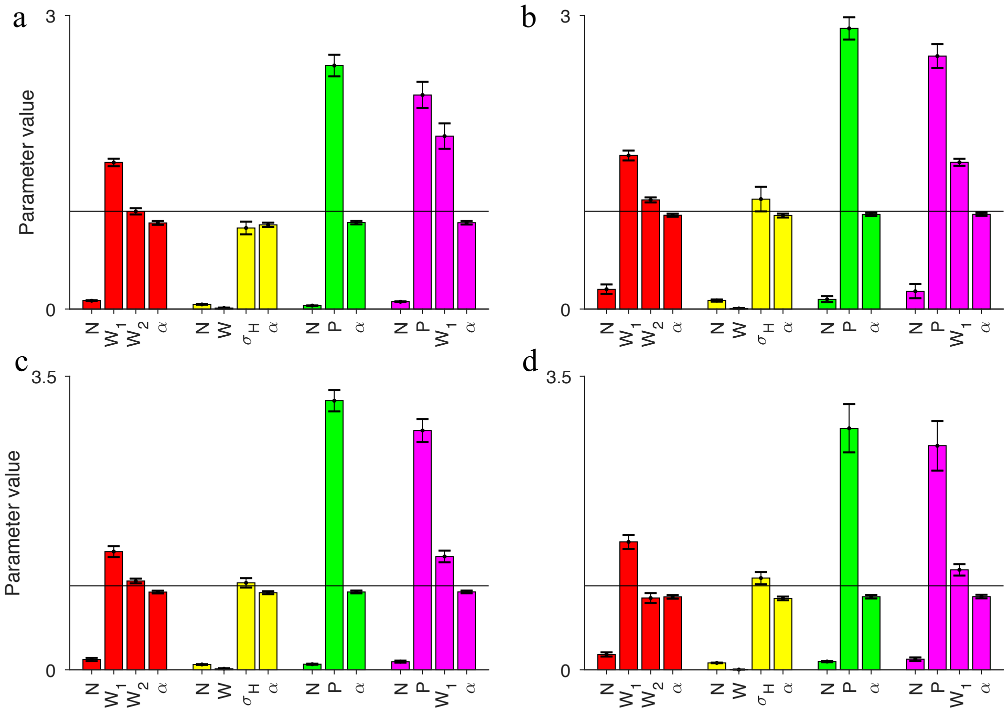

Figure 3.6. Best fitting parameter sets.

Mean and standard error (across participants) for the best fitting parameter sets for the "Selective Integration" (red bars), "Divisive Normalization" (yellow bars), Power (green bars) and for the Power and Selective Integration (pink bars) models. The fits were done independently for each participant and each model and for each of the four experiments: panel a, Experiment 1; panel b, Experiment 2; panel c, Experiment 3; and panel d, Experiment 4. The black horizontal line is set at the ordinate of 1 as a reference point. The parameters N and α, are common to all models and represent noise and the leak, respectively.

In general, for Selective Integration (red bars in Figure 3.6) we found that the weights for the local winner were significantly higher than one (W_1~ 1.5; all t-values > 6.4, all p < 0.001), and significantly greater than the weights of the local loser ($W_2 \sim 1$; all t-values for the difference > 4.5, all p < 0.001). These patterns of results are in line with the idea that local the local ranking of information determines the weight that is given to each piece of evidence.

For Divisive Normalization (yellow bars in Figure 3.6), however, we found the crucial parameter, W, to be close to zero (all W < 0.012, which is arbitrarily close to zero given that the fitting did not allow for negative values). While this finding is at odds with what would be expected from previous reports (Louie et al., 2013), it was unsurprising given the lack of a classical distractor effect reported above (Figure 3.4).

For the Power model (green bars in Figure 3.6), the parameter P, was significantly higher than one ($P \sim 2.8$; all t-values > 6.6, all p < 0.001). This result is in line with the idea that participants have a general tendency to focus on values that are high compared to values that are low, without the need to dynamically allocate their attention on every sample.

The Power and Selective Integration model (pink bars in Figure 3.6), we found both the P, and the W, parameters to be significantly higher than one (P ~ 2.5, all t-values > 5.7, all p < 0.001; W ~ 1.4; all t-values > 2.8, all p < 0.01). And in general, the fitted values are very close in value to their analogue parameters in the independent models (red and green bars). These results, suggest that participants' choices could be governed by both a fixed (local) and a dynamic (global) allocation of attention that favours higher numbers.

3.3.4 Quantitative comparison of models

We computed a pairwise quantitative model comparison based on the Bayesian Information Criterion (BIC) as well as Akaike Information Criterion (AIC). These two measures are tightly linked to each other, and differ only on the penalization they give for the number of free parameters for each model, with BIC having the strongest penalization. Thus, for any two models with the same number of free parameters their difference in BIC (ΔBIC) and AIC (ΔAIC) is exactly the same (Figure 3.7).

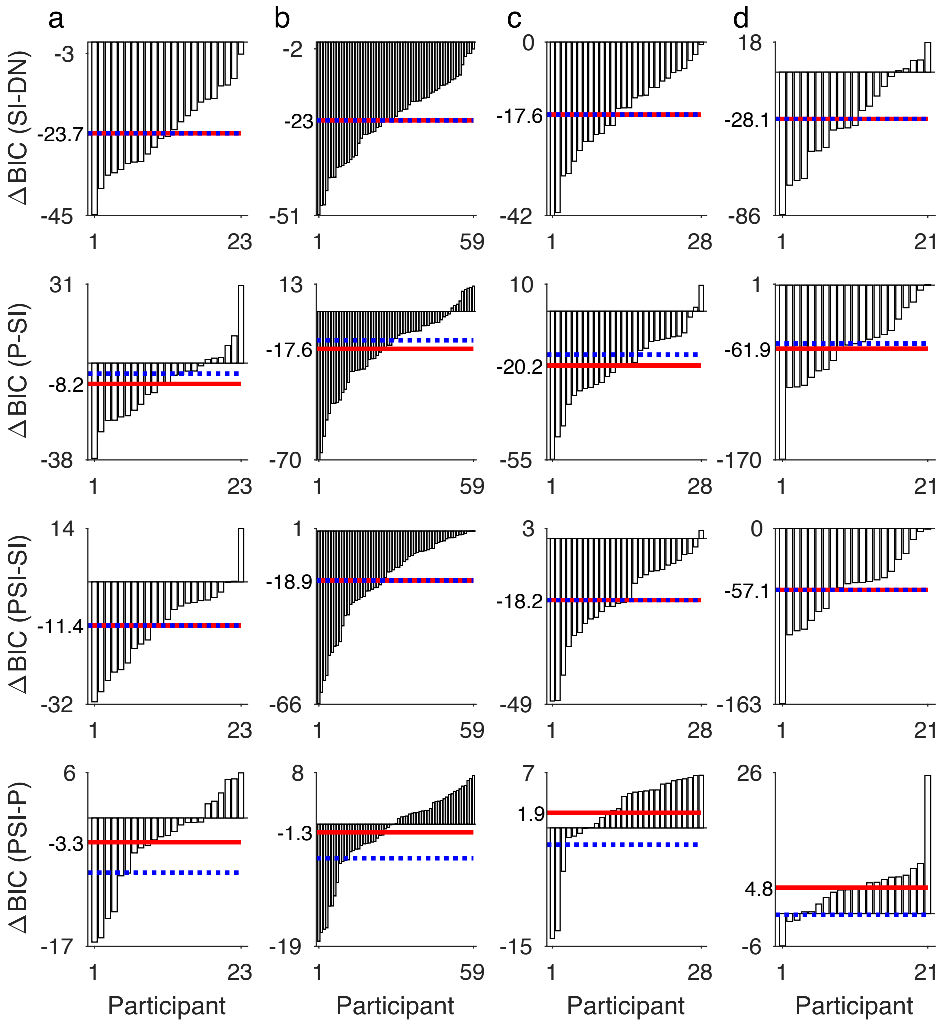

Figure 3.7. Quantitative model comparison.

Pairwise model Difference in the Bayesian Information Criterion (BIC) for each participant (bars, ordered from most negative to most positive) and averaged across participants (horizontal red lines), for comparison the average difference in the Akaike Information Criterion (AIC) is shown in dashed blue lines. Different experiments are presented in each column (panels a,b,c and d for Experiments 1,2,3 and 4, respectively). The different rows of figures present the pairwise comparisons between different pairs of models, from top to bottom: "Selective Integration" vs "Divisive Normalization", "Power" vs "Selective Integration"; "Power and Selective Integration" vs "Selective Integration"; and "Power and Selective Integration" vs "Power".

Briefly, we found a very strong support for Selective Integration over Divisive Normalization as candidate mechanisms for explaining our participants' choices (for all experiments, ΔBIC=ΔAIC < -17.6). We also found the Power Model to be a better candidate mechanism over the Selective Integration Model (for all experiments, ΔBIC < -8.2, ΔAIC < -4.2). Still, the Power and Selective Integration Model (a model that incorporates ingredients from both models) was able to provide an even better fit than the Power Model alone. However, these results are not as conclusive as with the other comparisons. In general, we found an advantage for the Power and Selective Integration Model, particularly when using AIC which penalizes less strongly the extra free parameter, but this advantage is less evident for the fourth experiment.

3.3.4.1 Qualitative predictions for regressions

The quantitative model comparison, suggests that both a dynamic (local) and a static (global) allocation of attention for bigger values guides behaviour. However, we wanted to find if the different models made different qualitative predictions about behaviour that we could put at test, to confirm the independent contribution of fixed and dynamic allocation of attention. For this, we decided to generate simulated choices from the best fitting models (for each participant and each experiment), and use those choices to run the same logistic regression models that we ran for the actual human behaviour reported in Figure 3.5; with particular emphasis on the last three regressors: (i) the average ranking of the local values of option A, r_A, (ii) the modulation of the value of Option A by its local distance to the distractor, $A \cdot \delta_{AC}$, and (iii) the modulation of the value of Option B by its local distance to the distractor, $B \cdot \delta_{BC}$.We found an interesting dissociation between the mechanisms that can account for the first and the last two regressors (see Figure 3.8 and Table 5).

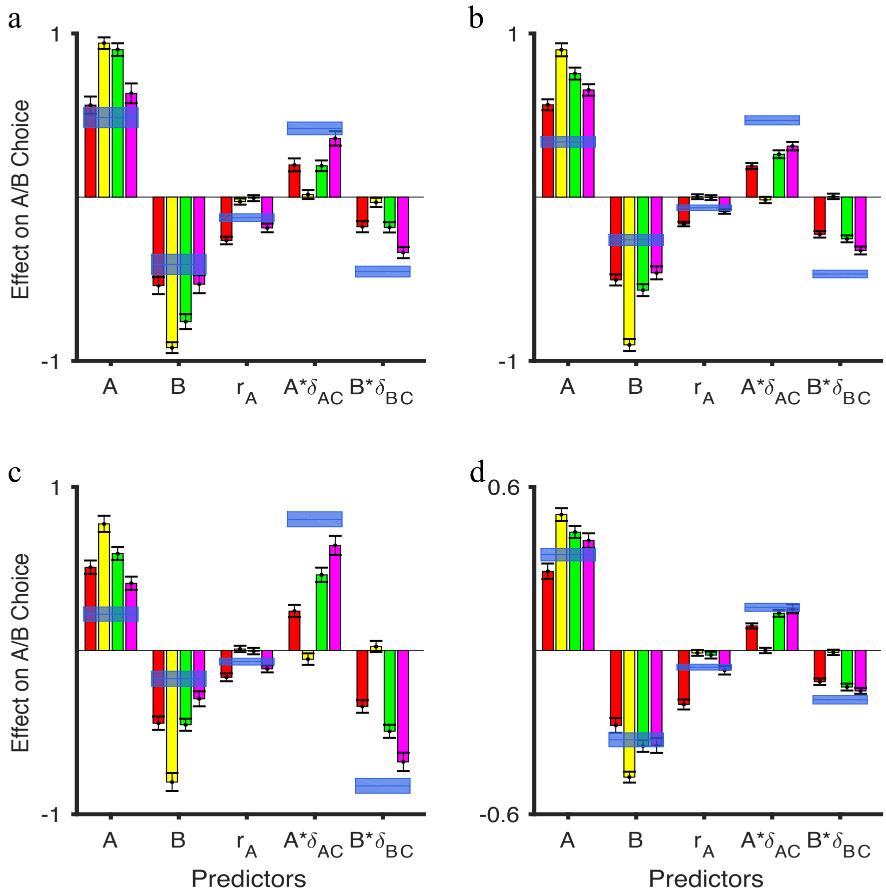

Figure 3.8. Qualitative predictions for logistic regression.

Mean and standard error (across participants) for the weights on choice for participants (blue horizontal shades), and the models: "Selective Integration" (red bars), "Divisive Normalization" (yellow bars), Power (green bars) and for the Power and Selective Integration (pink bars). Panels a,b,c and d, represent the regression results for Experiments 1,2,3 and 4, respectively.

In short, we found Selective Integration (i.e. a dynamic allocation of attention) to be a crucial ingredient for explaining the ranking effects, r_A. Only the "Selective Integration" and the "Power and Selective Integration" models were able to predict the negative effect on human choice resulting from higher average ranks of Option A (with higher values of r_A reflecting a less frequently having the local highest value). "Selective Integration" alone was able to also predict the modulations of value determined by the local distance to the distractor (i.e. regressors $A \cdot \delta_{AC}$ and $B \cdot \delta_{BC}$) in the same direction as human choices, however the fit was not perfect. The closest fits were provided by incorporating a power transformation (either with the "Power" model or the "Power and Selective Integration"). This shows that a fixed allocation of attention for general high values was a crucial ingredient to drive the strength of the last two regression effects.

An important aspect that has been previously emphasized about Selective Integration is its ability to predict a pro-variance effect (Tsetsos et al., 2012). This effect is reflected in people's increased preferences for an option that has high variability in its values across samples or attributes, above and beyond the average value of the option. Here, we wanted to ask whether the distance effects we found, were a simple reflection of the pro-variance effect: with higher distances expected in trials where the Option in question has higher variability across its attributes. We ran a regression including the average values of options A and B, as well as both distance related regressors described in Figure 3.8, and added three more regressors one for each of the options' variabilities (standard deviation) across samples within each trial (VarA, VarB and VarC).

Interestingly, we found the distance regressors to have an independent contribution for explaining participants' choices than the variability regressors (blue horizontal lines in Figure 3.9). The positive weights for the variability of option A reflect participants' increased probability of choosing option A when its values are more variable. Similarly, the negative weights for the variability of option B, reflect an increased probability of choosing option B when its values are more variable. Because trials when the distractor are excluded from the analyses, the variability of option C has little or no influence on choice (see Table 6 for a summary of the significance testing).

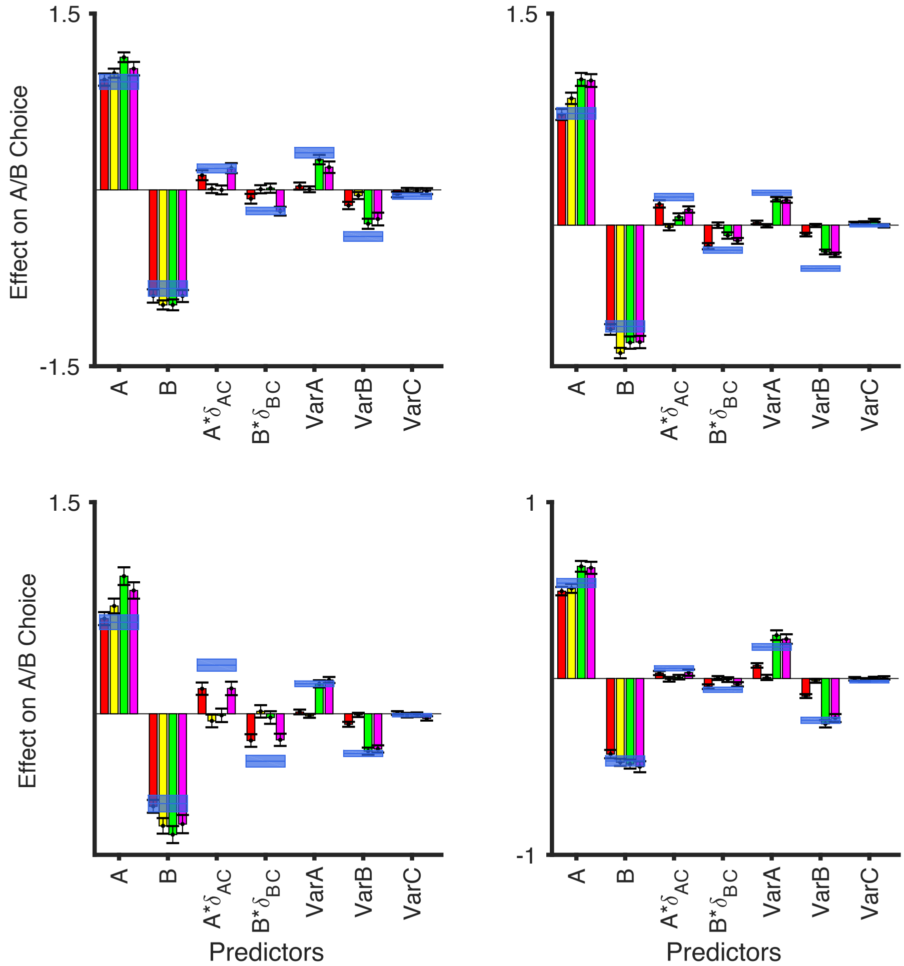

Figure 3.9. Independent effects for distance and variability.
Labels and colour schemes as in Figure 3.8.

Crucially, running the same regression model but on the simulated choices of the best fitting models, revealed that the pro-variance effect is more likely attributable to a Power transformation of the values, that is, a fixed allocation of attention to higher values (Figure 3.9). Choices from the Selective Integration model (red bars) show little or no consistent effect of variance (VarA, VarB and VarC) on choice; while the models that incorporate a Power transformation do (green and pink bars). Intuitively, this is explained because the power transformation results in greater "boost of high values" than “shrinkage of low values”. Instead, the distance effects are more likely attributable to a dynamic allocation of attention to the locally highest available piece of evidence (i.e. selective integration). The proximity of the local values of the distractor to the values of the other options increases the competition against the other options, thus decreasing their subjective value. This time, only the models that incorporate a Selective Integration component (red and pink bars) are able to show consistent distance effects.

3.3.5 Evidence for Selective Integration from pupil dilation

As participants carried out the task in Experiment 1, we recorded their pupil sizes. Based on recent work (Cheadle et al., 2014; de Gee et al., 2014; Lempert et al., 2015), we looked to see to see if we could decode the value of the samples on the screen from the size of participant’s pupil. In particular, we explored if the local ranking of information exerted an influence in the strength of encoding of the local values. For this decoding approach, we ran a series of independent regressions to predict pupil size at each timepoint (relative to the onset of each stimulus sample) as a function of the three values of evidence in the sample at question.

Crucially, we ordered the three pieces of evidence (from each sample in the sequence of each trial) according to their local ranks (i.e. maximum, median and minimum) before we entered the values as predictors in the regression model. In line with Selective Integration, we found that the steepness of encoding of value was the greatest for the evidence that ranked highest locally (Figure 3.10).

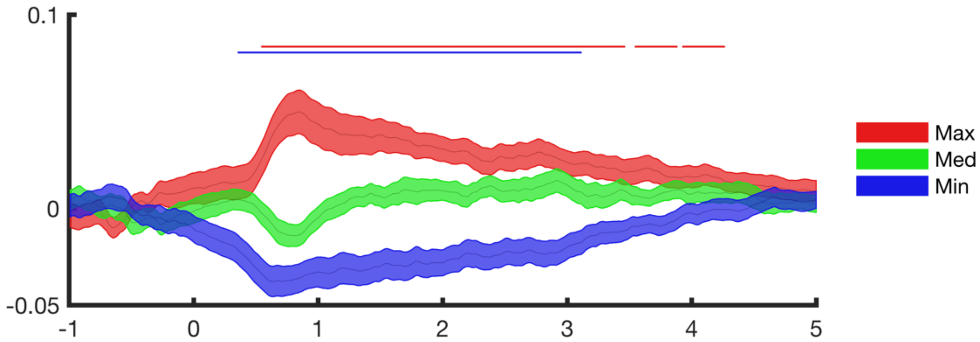

Figure 3.10. Decoding rank dependent value from pupil size.

Regression coefficients (y-axis) for predicting the size of the pupil (a.u.) at each timepoint (x-axis seconds relative to sample onset) using the value of the maximum, median and minimum pieces of evidence on the screen as predictors (using Equation 6). Gaze position was relatively quiescent during the sample sequence (see Figure 3.2). Still, the (momentary) vertical and horizontal gaze positions were included as regressors (weights not shown). Mean regression across subjects (dark lines) ±SEM (shadowed areas) shows a higher encoding of the information that ranks maximum of the three bars on a given sample (red) compared to the median (green) and the minimum (blue). The dashed lines at the top of the figure represent timepoints for which the regressors differed statistically from zero (p<0.05 after a cluster correction for multiple comparisons).

The peak of decoding occurs just before one second after stimulus onset. This spectrum of time is well in line with the timecourse of previously described decoded signals (de Gee et al., 2014; Lempert et al., 2015). However, the effect we recover lingers on for about 3 more seconds. This extended timecourse is likely the consequence of the fast presentation rate of the sequence, and slow dilation of the pupil in response to each sample, resulting in a high degree of overlap in the responses that was left un-modelled in our analysis. We thus decided to use a convolution approach to analyse the pupil timeseries, as it was originally designed to work with slow overlapping signals (such as in functional Magnetic Resonance Imaging). This analysis delivered similar results: with a higher influence on pupil size given by the value of whichever piece of evidence was higher on the screen (the Maximum, Max) compared to the Median and Minimum, Med and Min (Max>Med: $t_{(22)} = 2.43$, $p < 0.05$; Max>Min: $t_{(22)} = 2.96$, $p < 0.01$; Med>Min: $p > 0.1$); see Figure 3.11.

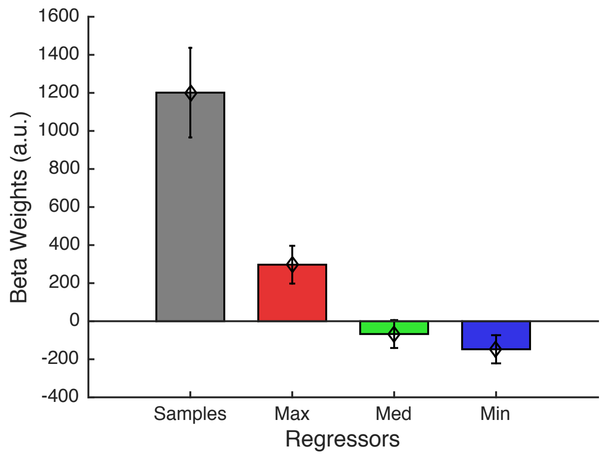

Figure 3.11. Regression coefficients of convolution GLM.

Regression coefficients for predicting the size of the pupil (a.u.) using Equation 10 the IRF convolution analysis. Bars denote the mean beta weights across subjects (dark lines) ±SEM (whiskers). The mere occurrence of the samples drives the pupil to increase (grey bar). However, we recovered a difference in the way the pupil diameter encodes information as a function its momentary rank. The maximum of the three bars on a given sample has a higher weight on pupil size (red bar) than the median (green bar) and minimum (blue bar) bars.

3.3.6 Evidence for Selective Integration and Power transformation from EEG signals

A similar decoding approach to that used for the pupil dilation signals (of Experiment 1) was used for the electroencephalographic (EEG) signals recorded in Experiment 4. EEG responses are much faster than those of the pupil dilation. We therefore looked for encoding of value on the EEG signals in a much shorter window of time, of 700ms, starting 200ms before the onset of each sample. In line with what we expected from Selective Integration, we found that the activity (averaging the eight most occipital electrodes, see methods), had a strong encoding of the value of the highest digit present on the screen, but not of the other two (Figure 3.12).

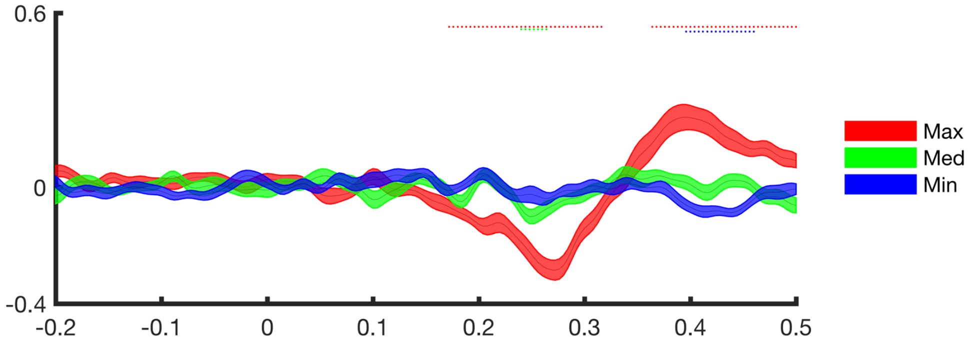

Figure 3.12. Decoding of rank dependent value from EEG signals.

Regression coefficients on EEG signal (averaged for occipital and parietal electrodes) for Experiment 4. Average regression coefficients ±SEM (shadowed areas), for predicting the EEG signal (a.u.) as a function of the value of the maximum (red line), median (green line) and minimum (blue line) samples on the screen (the dotted points signal the timepoints that are significantly different from zero after correction for multiple comparisons, see methods). The x-axis represents the time in seconds relative to the onset of each sample.

The discrete nature of the evidence (single digits from one to nine), in this experiment allowed us to explore how linear the encoding of value was. To do so, we ran a separate regression model whereby nine independent regressors signalled the presence or absence of each of the nine digits (for each of the samples of the sequence of each trial). We were able to decode the presence of the digits in a graded manner in two discrete time-windows (~200-300ms and ~300-400ms after stimulus onset), defined as all timepoints where at least two of the regression coefficient were statistically significant after cluster correction (Figure 3.13). In line with a Power transformation of value, we found that the higher digits had a much stronger influence on the EEG signal than the lower ones, when looking at the average weights for each digit within the defined windows.

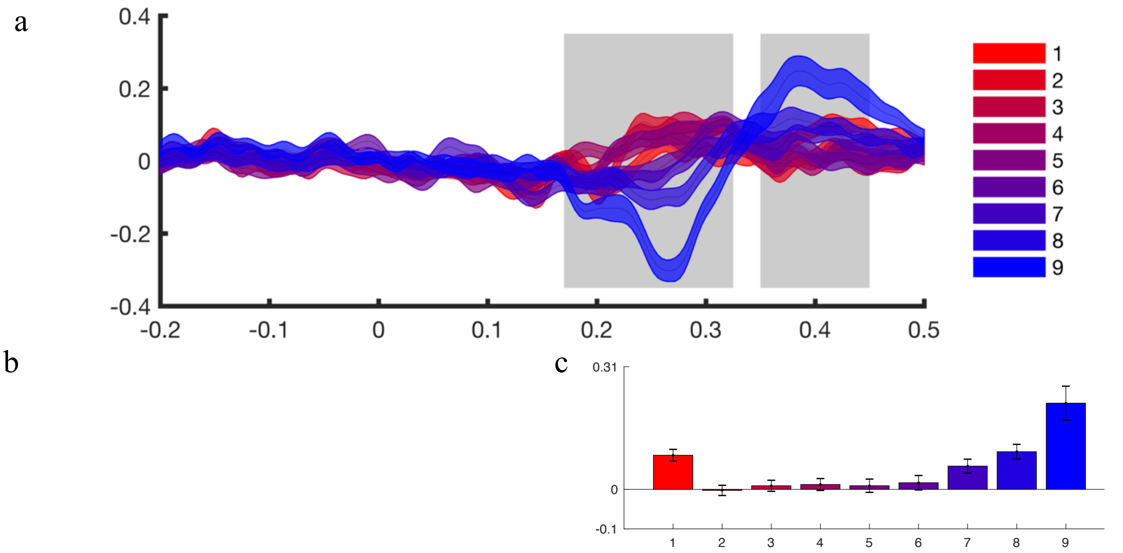

Figure 3.13. Non-linear decoding of digits from EEG signals.

Regression coefficients on EEG signal (averaged for the eight most occipital electrodes) for Experiment 4. Panel a, each coloured trace represents the encoding (average regression coefficients ±SEM shadowed areas), of each of the digits (the red-most being digit one and the blue-most being digit nine) on the EEG signals (a.u.), the two sections shadowed (marked in grey) are timepoints where at least two of the traces are significant after multiple comparisons correction. Panel b shows the average traces during the first of the shadowed areas for each of the digits. Panel c shows the same as panel b but for the second shadowed area. For panel a, the x-axis represents the time in seconds relative to the onset of each sample, for panels b and c it represents the digit in question.

3.3.7 Independent effects of Selective Integration and Power transformation

Because higher digits are more likely to be the maximum digit within their triplet (sample), the two effects described above could be the reflection of a single process, particularly when the timecourse of the signals is so similar. We decided to run a regression including all 12 regressors (the three rank regressors as well as the nine digit regressors) competing to explain the variance of the EEG signals. While both effects reduced in strength (likely as a consequence of the high degree of correlation between them), we were still able to recover similar effects as when the two groups of predictors were run in isolation (Figure 3.14). Interestingly, we found that the earlier part of the signal (~200-300ms) was primarily related to the (non-linear) representation of the digits, whereas the second part of the decoding of the signal (~300-400ms) was rank dependent.

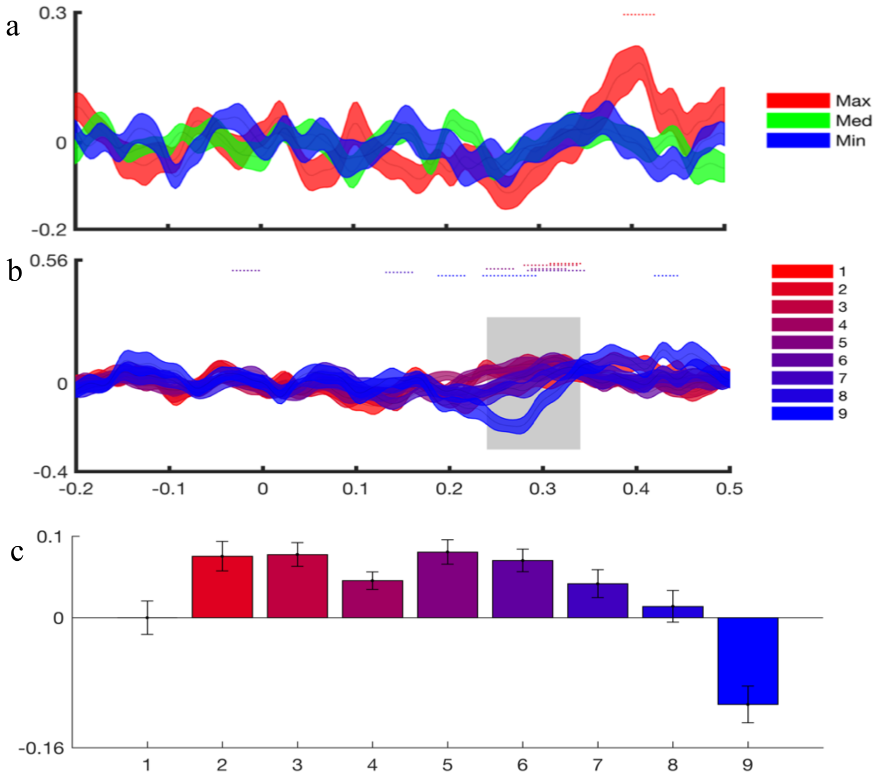

Figure 3.14. Competitive regression of rank dependent value and digit presence on EEG.

The rank regressors and presence of each digit, were included competitively as predictors of the EEG signals. The regression coefficients are shown separately for each of these groups of predictors: panel a for the rank dependent regressors and panel b for the presence of each digit. Dashed lines on top of each panel reflect the timepoints for which the significance was below 0.05 after multiple comparison correction (see methods). The shadowed area in panel b, represents the timespan where at least two of the digit-presence regressors were significantly different from zero.

3.3.8 Can we look for these independent effects in Pupil dilation?

Looking for evidence for non-linear encoding of value in EEG was easy, as the attribute values in Experiment 4 are composed of nine discrete stimuli (the digits), sampled uniformly. For our Experiment 1, however, the stimuli are much more continuous and sampled from several Gaussians. Still, to try and ask if we could find independent evidence for Selective Integration and a Power transformation from the pupil dilation signal, we decided to discretize the stimuli according to nine quantiles (across the entire distribution of stimuli in the task). We then replicated the analysis described in Figure 3.14, but this time, for the pupil dilation signal.

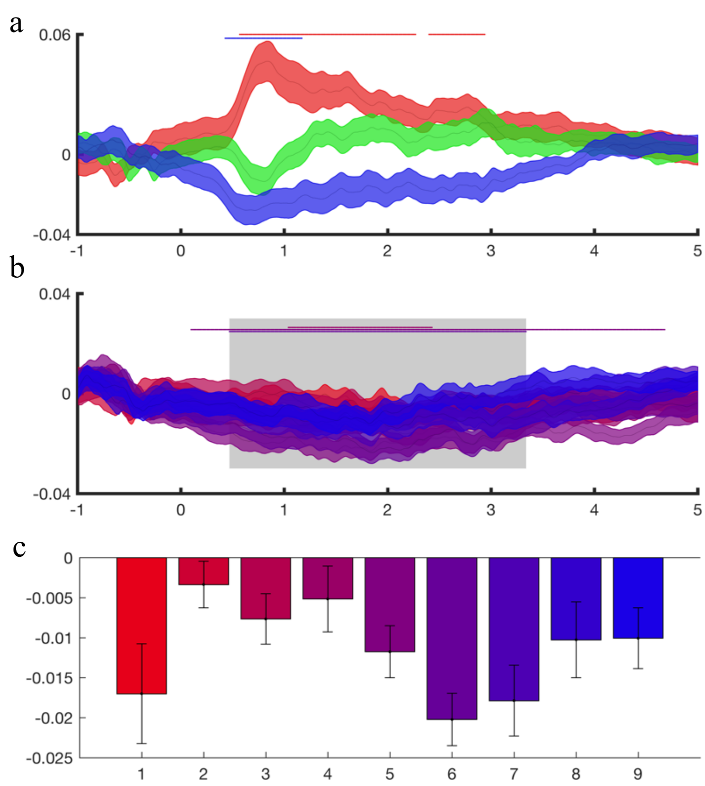

Figure 3.15. Competitive regression of rank dependent value and quantile presence on Pupil Dilation.

Figure legend as in Figure 1.14 but predicting the pupil dilation of participants in Experiment 1 (instead of the EEG signals in Experiment 4). We used nine regressors signalling the presence or absence of stimuli within each of nine quantiles across the entire distribution of stimulus values in the task.

While we found a strong effect of the rank dependent (local) prioritization as encoded on the pupil dilation signals, we found almost no evidence for encoding of Power transformed stimulus values (Figure 3.16). These results could suggest, that the signal decoded from the pupil dilation is not encoding value itself but a signal that encodes the gain given to the local winner (and this gain signal may be value dependent).

3.3.9 Robustness against noise from Selective Integration and Power transformation

We introduced Selective Integration and Power transformation as sensible efficient mechanisms for exploiting the regularities of the environment. Because resources for processing information are limited, allocating them efficiently should be crucial. Both Selective Integration and Power transformation allow for this efficient allocation of resources.

Selective integration allocates more resources (gives higher gain) to the piece of evidence on the screen with the maximum value. This is a sensible strategy when resources do not suffice to encode all available evidence at the best possible quality. Because the maximum evidence on the screen is likely to belong to the overall highest valued option, Selective Integration is a sensible way to encode with higher gain (better quality) the values of the options that are more likely to be the best. Similarly, a Power transformation allocates more resources to higher values compared to lower ones. This prioritization in the encoding of higher values is a sensible strategy because in general, higher values are more indicative of which option will be the best option.

The benefits of these two mechanisms are particularly evident when considering that noise is likely to corrupt the process of accumulation of evidence. Prioritizing the strength of encoding towards more relevant pieces of evidence (as in Selective Integration), or towards more informative feature values (as with a Power transformation), should allow for noise to exert a relatively lesser effect in corrupting choices. To test these ideas, we decided to simulate choices (from models incorporating either or both of these mechanisms) within a “hypothetical” three-option, sequential integration paradigm.

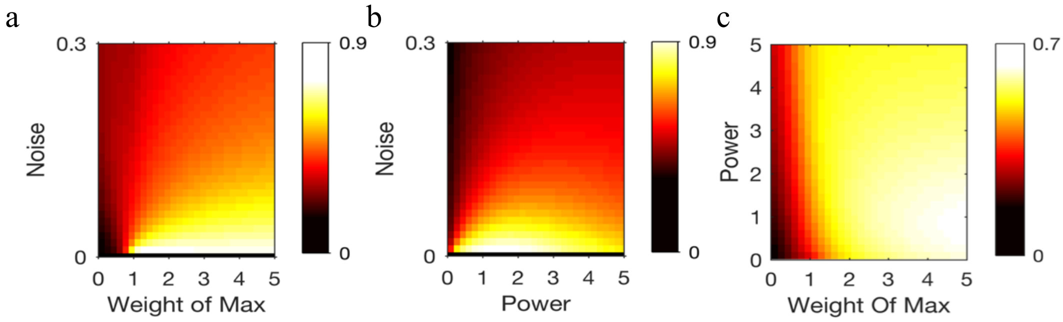

Figure 3.16. Protection against noise from Selective Integration and Power transformations.

Simulated accuracy under varying levels of noise and under different models. For all panels simulated accuracy is shown in hues of color–higher temperature hues meaning higher accuracy (proportion correct), see colorbars for each panel. Accuracy for a Selective Integration model is shown in panel a, with different weights for the momentary maximum evidence (x-axis), and varying levels of noise (y-axis). Accuracy for a Power model is shown in panel b, with different power values for transforming evidence (x-axis), and varying levels of noise (y-axis). Finally, accuracy for a Power and Selective Integration model is shown in panel c, with different power values for transforming evidence (y-axis) and different weights for the maximum momentary evidence, and under a fixed level of noise.

Indeed, we found that both Selective Integration and a Power transformation are able diminish the effects of noise onto choices (Figure 3.16). As noise increases, the best accuracy is achieved by means of allocating higher resources to the local maximum compared to the median and minimum. An alternative solution is to allocate more resources to the processing higher numbers compared to lower ones. Unsurprisingly, the two strategies trade one another off: if you allocate a lot of resources to processing the maximum value available on the screen, you may not need to allocate as many resources to processing high numbers in general.

3.4 Discussion

Decisions with multiple alternatives have recently been used to explore departures from rational behaviour (Chau et al., 2014, 2014; Louie et al., 2013). These studies have proposed mechanisms to explain the departures from optimality. These mechanisms invoke context dependent changes in the way the value of the options is represented. However, this view assumes that the value of the options is readily available to be represented, even when the options are composed of multiple attributes. Instead, we focus on how the construction of value for each option takes place while integrating their values across several attributes.

We designed an experimental paradigm where the value of three options has to be constructed as evidence for the value of the three options is delivered simultaneously for the three options but sequentially (one attribute at a time). In particular, we explored how two mechanisms of prioritization of information processing are likely to exert strong influences on choice, in the context of three alternatives. One mechanism acts at a local timescale, whereas the other acts at a global timescale.

The first mechanism (local), Selective Integration, prioritizes the integration of the highest of the simultaneously available pieces of evidence (Tsetsos et al., 2012). The second mechanism (global), a Power (exponential) transformation of value, prioritizes the processing of higher parts of the value ranges (Li et al., 2017). In a series of four experiments (with variations of the same paradigm) we were able to show independent evidence for these two mechanisms from: (i) the behavioural choice patterns, (ii) as well as from physiological signals recorded as participants carried out the experiments (i.e. pupil size for one experiment and EEG for another one). We were also able to show through simulations that the two mechanisms are able to offer an efficient allocation of resources that has the potential of making the representation of the overall value of the best option (and thus the accuracy of choices) more robust against noise.

In line with Selective Integration, we found that the local (within attribute/within sample) ranking of evidence was a strong predictor of choice. Using logistic regression, we found that the probability of choosing the highest valued option was highly related to the number of attributes for which it possessed the highest value (above and beyond the overall value it had across all attributes). Interestingly, the attribute-by-attribute proximity in value of the distracter option (that with the lowest overall value) to the value of either of the other two options, diminished the probability that that other option would be chosen. This proximity effect could reflect the fact that it is harder to establish the ranks of the options when their values are very similar.

It has previously been reported that humans have a tendency to choose an option whose values are more variable across its dimensions (Tsetsos et al., 2016, 2012). This pro-variance effect, has previously been attributed to Selective Integration, the intuition behind it being that options with higher variability in their values are more likely to receive a boost (of Selective Integration) in their values when they are high, and miss the boost when their values are not. Particularly with three or more options, the option with the highest variance is more likely to have values that surpass the value of the other options (thus receiving the boost more frequently than its rivals).

Instead, we found that the pro-variance effect in our four experiments is better explained not by Selective Integration (a local prioritization of high values) but by a Power transformation of the values (a global prioritization of higher valued evidence). Quantitative model comparison was able to show that both the local and global prioritization of high values was needed to best explain the responses of our participants. Interestingly, we were able to obtain specific signatures for the local and global prioritization of information both from behaviour as well as from EEG and pupil dilation signals.

Finally, using simulations, we were able to show how both of these mechanisms are potentially beneficial for the accuracy of choices in the face of noise. Without noise, any mechanisms that departs from linear representation of evidence (globally) or uniform integration of evidence (locally) will result in loss of information and thus it will affect the accuracy of choices. However, the process of accumulating evidence in time is noisy and costly (Wyart & Koechlin, 2016). Similarly, the quality of encoding of feature values has been shown to be controlled by a resource that can be dynamically allocated (Bays & Husain, 2008). We found that both the local and global prioritization of high values are protective against noise. While demonstrating that these mechanisms are optimal is a challenging task, we could at least show that they are sensible strategies to adopt in the face of resource limits and noise.

Our results raise the question: under which scenarios is a local (Selective Integration) strategy of prioritization of information preferred over a global one (Power transformation), and vice versa? While the current data does not allow us to answer this question, some aspects of the data allowed us to conjecture about these issues. For instance, the best fitting model for the different experiments incorporated variable degrees of the two strategies of prioritization: Experiment 1 has a notable contribution of the local prioritization (Selective Integration) whereas Experiment 4 has a notable contribution of the global prioritization.

There are a few differences in the design of the two experiments. In Experiment 1 the attribute values consist of the length of bars, whereas they are digits for Experiment 4. The rate of presentation of the different samples in the sequence was 4Hz for Experiment 1 and 2Hz (slower) for Experiment 4, as processing digits requires more time. A crucial difference in the design, however, is that for Experiment 1, the attribute values for each trial included an off-set that varied across trials, whereas the values where simply drawn from uniform distributions in Experiment 4. This means that for Experiment 1, all the attribute values in one trial could lie on the lower part of the value range, whereas they could be on the higher part of the value range for the next trial, making a purely global prioritization inefficient in all trials where values lie in the lower part of the value range.

Interestingly, we found evidence for the local prioritization of information processing of high values, even in Experiment 4 where the values were drawn from uniform distributions. This could be driven by the fact that integrating information is costly: either (i) not all the information in a sample can be integrated, or at least (ii) the quality of the integrated evidence is not equal for all evidence within a sample. However, we did not have conditions within an experiment where this cost was manipulated. An idea to test in future experiments is the degree to which we can manipulate experimentally the relative pressures for performing global vs local prioritization and see if we can recover the expected changes in the actual prioritization performed by participants. For instance, the pressure for local prioritization could be increased by means of increasing the presentation rate of the sequence in a subset of trials within an experiment.

In summary, across a series of four experiments, making use of behavioural analyses, computational modelling and simulations as well as analyses of physiological data, we were able to show evidence for two efficient mechanisms that are likely to exert influence on choices when confronted with multiple options. These mechanisms, while at first glance irrational, confer robustness against noise to the choices of individuals by means of exploiting the regularities in the environment where locally and globally high values are highly predictive of which option is the best one to choose.

Chapter References

Bateson, M. (2002). Context-dependent foraging choices in risk-sensitive starlings. *Animal Behaviour*, *64*(2), 251–260.

Bateson, M., Healy, S. D., & Hurly, T. A. (2002). Irrational choices in hummingbird foraging behaviour. *Animal Behaviour*, *63*(3), 587–596.

Bays, P. M., & Husain, M. (2008). Dynamic Shifts of Limited Working Memory Resources in Human Vision. *Science*, *321*(5890), 851–854.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Burnham, K. P., & Anderson, D. R. (2003). *Model selection and multimodel inference: a practical information-theoretic approach*. Springer Science & Business Media.

Chau, B. K., Kolling, N., Hunt, L. T., Walton, M. E., & Rushworth, M. F. (2014). A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nature Neuroscience*, *17*(3), 463–470.

Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., De Gardelle, V., Castañón, S. H., & Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron*, *81*(6), 1429–1441.

de Gee, J. W., Knapen, T., & Donner, T. H. (2014). Decision-related pupil dilation reflects upcoming choice and individual bias. *Proceedings of the National Academy of Sciences*, *111*(5), E618–E625.

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21.

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

Hunt, L. T., Dolan, R. J., & Behrens, T. E. (2014). Hierarchical competitions subserving multi-attribute choice. *Nature Neuroscience*, *17*(11), 1613–1622.

Hurly, T. A., & Oseen, M. D. (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 under risk. *Econometrica: Journal of the Econometric Society*, 263–291.

Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, *90*(430), 773–795.

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification learning in amnesia. *Learning & Memory*, *1*(2), 106–120.

Lempert, K. M., Chen, Y. L., & Fleming, S. M. (2015). Relating Pupil Dilation and Metacognitive Confidence during Auditory Decision-Making. Retrieved from http://dx.plos.org/10.1371/journal.pone.0126588

Li, V. L., Castañón, S. H., Solomon, J. A., Vandormael, H., & Summerfield, C. (2017). Robust Averaging Protects Decisions from Noise in Neural Computations. *BioRxiv*, 147744.

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, *110*(15), 6139–6144.

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG-and MEGdata. *Journal of Neuroscience Methods*, *164*(1), 177–190.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59.

Shadlen, M. N., & Shohamy, D. (2016). Decision making and sequential sampling from memory. *Neuron*, *90*(5), 927–939.

Shafir, S., Waite, T. A., & Smith, B. H. (2002). Context-dependent violations of rational choice in honeybees (Apis mellifera) and gray jays (Perisoreus canadensis). *Behavioral Ecology and Sociobiology*, *51*(2), 180–187.

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

Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decision-making. *Nature Communications*, *7*.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, *109*(24), 9659–9664.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, 201519157.

Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, *79*(4), 281.

Waite, T. A. (2001). Background context and decision making in hoarding gray jays. *Behavioral Ecology*, *12*(3), 318–324.

Wierda, S. M., van Rijn, H., Taatgen, N. A., & Martens, S. (2012). Pupil dilation deconvolution reveals the dynamics of attention at high temporal resolution. *Proceedings of the National Academy of Sciences*, *109*(22), 8456–8460.

Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*, *11*, 109–115.

TABLES AND APPENDICES

Table 1. Proportion of choices and response times.

	Option	A	B	C
Exp 1 (n=23)	p(Ch)±SEM	0.718±0.001	0.207±0.002	0.072±0.002
	RT(ms)±SEM	537.7±21.4	638.0±29.2	704.4±31.4
Exp 2 (n=59)	p(Ch)±SEM	0.720±0.01	0.204±0.006	0.066±0.006
	RT(ms)±SEM	514.7±20.4	592.1±22.5	651.0±26.8
Exp 3 (n=28)	p(Ch)±SEM	0.709±0.009	0.245±0.006	0.039±0.003
	RT(ms)±SEM	556.1±20.0	652.2±22.0	747.4±27.9
Exp 4 (n=21)	p(Ch)±SEM	0.623±0.01	0.261±0.005	0.11±0.006
	RT(ms)±SEM	570.0±29.4	594.5±30.7	649.4±45.2

Table 2. Regression with value of distractor.

		A	B	C
	β–weights	0.9430	-0.8984	0.0181
Exp 1 (n=23)	*t*-value	20.8601	-17.7046	0.6146
	p	**5.51667E-16**	**1.67472E-14**	0.5451
	β–weights	0.8479	-0.7954	-0.0330
Exp 2 (n=59)	*t*-value	28.8594	-27.2246	-2.1596
	p	**4.23996E-36**	**9.92909E-35**	**0.0350**
	β–weights	0.7146	-0.7383	-0.0108
Exp 3 (n=28)	*t*-value	16.8285	-15.3262	-0.6470
	p	**7.70635E-16**	**7.63628E-15**	0.5231
	β–weights	0.4914	-0.4712	-0.0036
Exp 4 (n=21)	*t*-value	21.4695	-16.8957	-0.2691
	p	**2.7781E-15**	**2.63031E-13**	0.7906

Table 3. Regression with normalized value of distractor.

		A	B	nC
Exp 1 (n=23)	β–weights	0.9512	-0.8923	0.0084
	t-value	21.0672	-18.7775	0.5054
	p	**4.48319E-16**	**4.95771E-15**	0.6183
Exp 2 (n=59)	β–weights	0.8362	-0.8061	-0.0235
	t-value	29.0939	-28.5810	-2.4487
	p	**2.732E-36**	**7.17374E-36**	**0.0174**
Exp 3 (n=28)	β–weights	0.7127	-0.7402	-0.0110
	t-value	16.8982	-15.2598	-0.6767
	p	**6.95798E-16**	**8.48691E-15**	0.5043
Exp 4 (n=21)	β–weights	0.4888	-0.4717	-0.0058
	t-value	21.2368	-17.2002	-0.4929
	p	**3.42468E-15**	**1.88114E-13**	0.6274

Table 4. Regression with momentary interaction predictors.

		A	B	r_A	$A \cdot \delta_{AC}$	$B \cdot \delta_{BC}$
Exp 1 (n=23)	β–weights	0.4864	-0.4107	-0.1256	0.4200	-0.4551
	t-value	7.7309	-6.5170	-5.3046	10.6039	-13.4444
	p	**1.03695E-07**	**1.48293E-06**	**2.52792E-05**	**4.10994E-10**	**4.35092E-12**
Exp 2 (n=59)	β–weights	0.3376	-0.2613	-0.0647	0.4692	-0.4700
	t-value	10.2439	-7.6924	-3.9463	16.7516	-19.8064
	p	**1.26141E-14**	**2.01029E-10**	**0.0002**	**6.84699E-24**	**1.71917E-27**
Exp 3 (n=28)	β–weights	0.2232	-0.1721	-0.0672	0.8021	-0.8268
	t-value	5.0812	-3.6412	-3.1734	17.1897	-17.8798
	p	**2.44872E-05**	**0.0011**	**0.0037**	**4.55373E-16**	**1.70934E-16**
Exp 4 (n=21)	β–weights	0.3514	-0.3272	-0.0605	0.1584	-0.1804
	t-value	15.2140	-12.5109	-5.5018	10.6272	-11.5330
	p	**1.85277E-12**	**6.48501E-11**	**2.19344E-05**	**1.12663E-09**	**2.73268E-10**

Table 5. Significance (p-values) for human and model regression coefficient on choices (with momentary interaction predictors).

		A	B	r_A	$A \cdot \delta_{AC}$	$B \cdot \delta_{BC}$
Humans	**Exp 1 (n=23)**	**1.03695E-07**	**1.48293E-06**	**2.52792E-05**	**4.10994E-10**	**4.35092E-12**
	Exp 2 (n=59)	**1.26141E-14**	**2.01029E-10**	**0.0002**	**6.84699E-24**	**1.71917E-27**
	Exp 3 (n=28)	**2.44872E-05**	**0.0011**	**0.0037**	**4.55373E-16**	**1.70934E-16**
	Exp 4 (n=21)	**1.85277E-12**	**6.48501E-11**	**2.19344E-05**	**1.12663E-09**	**2.73268E-10**
SI Model	**Exp 1 (n=23)**	**3.37E-10**	**5.55E-10**	**1.25E-10**	**5.89E-05**	**2.41E-05**
	Exp 2 (n=59)	**5.00E-24**	**1.45E-21**	**1.47918E-16**	**1.84E-14**	**5.41E-16**
	Exp 3 (n=28)	**1.24E-12**	**2.83784E-11**	**2.7811E-07**	**6.12E-07**	**1.51E-09**
	Exp 4 (n=21)	**2.31E-09**	**2.53E-09**	**8.77E-10**	**6.44E-09**	**1.01E-08**
DN Model	**Exp 1 (n=23)**	**3.23E-18**	**1.37E-18**	0.1587	0.5608	0.2663
	Exp 2 (n=59)	**1.81E-30**	**3.25E-32**	0.8674	0.3781	0.7777
	Exp 3 (n=28)	**7.27E-15**	**2.44803E-14**	0.6242	0.1519	0.4731
	Exp 4 (n=21)	**3.13E-15**	**5.68E-16**	0.4132	0.9913	0.5247
P Model	**Exp 1 (n=23)**	**6.05E-17**	**4.23E-14**	0.7478	**5.91E-06**	**7.40E-06**
	Exp 2 (n=59)	**2.46E-28**	**1.69E-22**	0.8580	**8.16E-16**	**1.93E-17**
	Exp 3 (n=28)	**1.40E-14**	**1.83884E-12**	0.8640	**6.11E-11**	**1.46E-12**
	Exp 4 (n=21)	**1.38E-14**	**6.26E-13**	0.2124	**1.12E-09**	**1.40E-09**
P+SI Model	**Exp 1 (n=23)**	**5.23E-10**	**2.88E-09**	**3.96E-07**	**4.47E-08**	**1.94E-09**
	Exp 2 (n=59)	**5.61E-26**	**6.81E-17**	**7.04026E-08**	**2.22E-17**	**1.87E-19**
	Exp 3 (n=28)	**9.77E-11**	**6.62225E-07**	**1.81885E-05**	**1.76E-11**	**2.24E-12**
	Exp 4 (n=21)	**2.07E-12**	**5.07E-11**	**0.0001**	**1.01E-08**	**1.10E-11**

Table 6. Significance (p-values) for human and model regression coefficient on choices (with distance and variability predictors).

		A	B	$A \cdot \delta_{AC}$	$B \cdot \delta_{BC}$	VarA	VarB	VarC
Humans	**Exp 1 (n=23)**	**1.53184E-12**	**1.6366E-11**	**0.0001**	**1.32306E-05**	**6.68024E-07**	**2.882E-09**	0.0564
	Exp 2 (n=59)	**8.139E-27**	**2.92231E-25**	**2.463E-10**	**4.86601E-11**	**2.14401E-18**	**6.32621E-23**	0.8335
	Exp 3 (n=28)	**1.66686E-12**	**3.34478E-11**	**1.74783E-08**	**1.56365E-08**	**8.18442E-11**	**4.60097E-13**	0.4175
	Exp 4 (n=21)	**5.73031E-15**	**8.60583E-13**	**0.0008**	**0.0004**	**1.09919E-08**	**2.43342E-11**	0.0342
SI Model	**Exp 1 (n=23)**	**1.04462E-14**	**1.43679E-13**	**7.41E-03**	0.0831	0.2927	**0.0003**	0.1079
	Exp 2 (n=59)	**9.18313E-28**	**2.59609E-27**	**1.99E-07**	**7.23E-09**	0.1397	**2.51E-06**	0.1399
	Exp 3 (n=28)	**2.45119E-14**	**3.16187E-14**	**0.0004**	**0.0002**	0.4468	**0.0002**	0.7143
	Exp 4 (n=21)	**4.22617E-15**	**6.4558E-14**	**0.0154**	**0.0023**	**1.09E-05**	**1.81E-08**	0.8844
DN Model	**Exp 1 (n=23)**	**5.99848E-18**	**9.26625E-18**	0.7805	0.8753	0.8402	0.1189	0.9551
	Exp 2 (n=59)	**3.88305E-30**	**2.77999E-32**	0.5256	0.9929	0.7259	0.8272	0.0974
	Exp 3 (n=28)	**2.55066E-14**	**3.3137E-14**	0.2724	0.7013	0.2929	0.5919	0.5545
	Exp 4 (n=21)	**1.03804E-14**	**5.17525E-16**	0.7418	0.8187	0.6842	0.2208	0.4582
P Model	**Exp 1 (n=23)**	**2.1034E-18**	**3.79931E-16**	0.9962	0.7132	**1.93827E-06**	**9.62977E-07**	0.8723
	Exp 2 (n=59)	**2.54292E-30**	**9.01246E-27**	**0.0185**	**0.0019**	**7.18305E-17**	**1.86951E-16**	**0.0023**
	Exp 3 (n=28)	**5.37534E-15**	**4.62665E-14**	0.8194	0.5542	**5.16298E-09**	**4.5044E-11**	0.6813
	Exp 4 (n=21)	**4.71765E-15**	**3.0264E-14**	0.5579	0.6424	**2.412E-08**	**6.26259E-11**	0.7424
P+SI Model	**Exp 1 (n=23)**	**1.29356E-14**	**1.69774E-14**	**0.0003**	**7.20E-05**	**0.0007**	**0.0001**	0.6796
	Exp 2 (n=59)	**1.95591E-30**	**1.22017E-26**	**7.35E-06**	**1.16E-06**	**3.10184E-13**	**8.5358E-20**	0.7143
	Exp 3 (n=28)	**1.187E-14**	**2.12001E-12**	**0.0008**	**0.0002**	**1.30177E-11**	**2.0898E-10**	**0.0398**
	Exp 4 (n=21)	**5.44758E-14**	**7.18978E-13**	**0.0761**	**0.0140**	**6.44978E-08**	**3.88111E-09**	0.7104

4 NOISE BLINDNESS AS THE SOURCE OF SUBOPTIMALITY

4.1 Introduction

An interesting mismatch exists in the reports of optimality between the perceptual and economic literatures. In the perceptual domain, it has been emphasized that humans and other primates make near optimal inferences. Under this view, information sources are weighted by their reliability when combined to make a decision (Ernst & Banks, 2002; Knill, Kersten, & Yuille, 1996; Körding & Wolpert, 2006; Ma, Beck, Latham, & Pouget, 2006; Mamassian, Landy, & Maloney, 2002; Trommershäuser, Maloney, & Landy, 2008). For example, in sensorimotor tasks, humans rely more heavily on knowledge of stimulus base rates (prior probability of occurrence) when stimuli are corrupted by noise and thus harder to perceive (Kersten, Mamassian, & Yuille, 2004; Körding & Wolpert, 2004; Sun & Perona, 1998). However, psychologists and behavioural economists have highlighted that in cognitive tasks, or when making real-world judgments, humans show systematic deviations from optimality that impair judgment and fail to maximise reward (Tversky & Kahneman, 1974). For example, when asked to guess an individual's occupation, humans neglect the prior probability of occurrence of different professions and rely instead solely on the descriptions of the individual provided by the experimenter (Kahneman & Tversky, 1973). Current theories attribute cognitive suboptimalities to an insufficiency of relevant past experience (Hertwig & Erev, 2009), limited stakes in laboratory experiments (Levitt & List, 2007), the format in which problems are posed (Jarvstad, Hahn, Rushton, & Warren, 2013), distortions in representations of values and probabilities (Ackermann & Landy, 2015), or to a reluctance to employ computationally costly executive processes in more complex decision settings (Gershman, Horvitz, & Tenenbaum, 2015; Kahneman, 2011). However, a computationally grounded account of human decisions that explains both perceptual optimality and cognitive suboptimality has yet to emerge (Summerfield & Tsetsos, 2015).

Human decisions are limited by uncertainty (noise) that is intrinsic to the outside world (e.g. gambles are inherently probabilistic), or to internal processes in the brain: sensory processing (e.g. imperfect transduction of sensory stimulation), general randomness in neural computation (Hunt, 2014; Juslin & Olsson, 1997; Ma & Jazayeri, 2014) and systematically biased processes of accumulation of information (Juslin & Olsson, 1997; Kahneman & Tversky, 1982; Renart & Machens, 2014). Noise within the brain can originate at different stages of neural processing and be subdivided accordingly. Performance can be limited during early stages of processing by the quality of encoding of sensory information (e.g. the cellular mechanisms that transduce sensory stimuli are stochastic). However, additional sources of loss may occur at later stages of processing. When accruing sufficient evidence to make a choice, for instance, the limited time-constants of integration of neurons contribute to a "leaky" accumulation process (Brunton, Botvinick, & Brody, 2013; Cheadle et al., 2014; Tsetsos et al., 2016; Usher & McClelland, 2001; Wang, 2002). However, many classic decision models conflate these potentially distinct sources of noise into a single error term without clearly specifying the multiple sources of uncertainty that may corrupt decisions (Beck, Ma, Pitkow, Latham, & Pouget, 2012; Wyart & Koechlin, 2016).

The different types of uncertainty limit performance to different degrees in different tasks. In most perceptual tasks, performance is mainly limited by encoding noise arising during sensory transduction or by the poor quality of a stimulus, such as when an object is hard to identify under conditions of low visibility ("encoding noise"). By contrast, in cognitive or economic decisions performance is largely limited by integration noise occurring when multiple signals must be combined to make a decision, such as when different stimuli evoke competing actions (Botvinick, Braver, Barch, Carter, & Cohen, 2001), or attributes of a consumer product provide discordant information about its monetary value (Scott, Constantinople, Erlich, Tank, & Brody, 2015; Wyart & Koechlin, 2016). Integration of diverse samples of information incurs additional computational costs ("integration noise") which typically lead to biased or suboptimal behaviour in cognitive tasks (including those involving economic choices). In most real-world settings, a major limitation for performance comes from lack of knowledge about the decision relevant factors (e.g. what are key things that make a car valuable to me?), however, we believe that this type of noise alone should not drive suboptimality or bias behaviour in a systematic manner. Instead, we propose that the "gap" in optimality between perceptual and cognitive decisions arises because of differential metacognitive sensitivity to the different types of noise that typically affect these two types of decision: humans perform optimally on perceptual tasks because they are sensitive to encoding noise, but suboptimally on cognitive tasks because they are "blind" to integration noise.

Here, over 6 experiments, we tested this hypothesis directly by manipulating encoding noise (stimulus visibility) and integration noise (stimulus variability) concurrently and orthogonally in a single behavioural paradigm. Our paradigm allowed us to estimate independently (i) the degree to which encoding and integration noise limited performance, (ii) how these two sources of noise predicted the suboptimality of inferences, and (iii) how this suboptimality related to metacognitive sensitivity to the two types of noise. We found strong support for the view that humans are sensitive to encoding noise, but blind to integration noise, and argue that this latter metacognitive failure is a major driver of suboptimal decision-making.

4.2 Methods

4.2.1 Participants and equipment

One hundred and eight healthy human participants with normal or corrected to normal vision were recruited to participate in five experiments. Three participants were excluded from all presented data because their performance fell below 55% correct on neutral trials. Analyses and data presented were carried out on the remaining 105 participants (72 females, 8 left-handed, 25.02 ± 4.25 years old). Participants performed the task while seated in a darkened room at ~60 cm from the computer monitor (model: Dell U2312HMT; refresh rate: 60Hz). Task and stimuli were generated using custom scripts for Psychtoolbox (Brainard, 1997). All research was conducted in accordance with local ethical guidelines.

4.2.2 Short description of Task and stimuli

The task details are described in more detail in the supplementary materials. Participants were required to judge on each trial if the average orientation of a circular array of 8 noisy Gabor patches (see Figure 4.1) was tilted clockwise (CW) or counter-clockwise (CCW) relative to horizontal, by pressing the right (CW) or left (CCW) arrow-keys of a standard "QWERTY" keyboard. The task consisted of 1296 trials, distributed in 36 blocks of 36 trials each. Participants could take a short self-timed break between blocks. On each trial the stimulus array was presented for 150 ms, followed by a 3000 ms response period. For Exp1 and 2, we manipulated the prior probability of occurrence of each category by presenting participants with either of three cues (letters L, N or R), for 700ms before each the appearance of the array of stimuli and until a response was registered, the cue font was Calibri of 16 points.

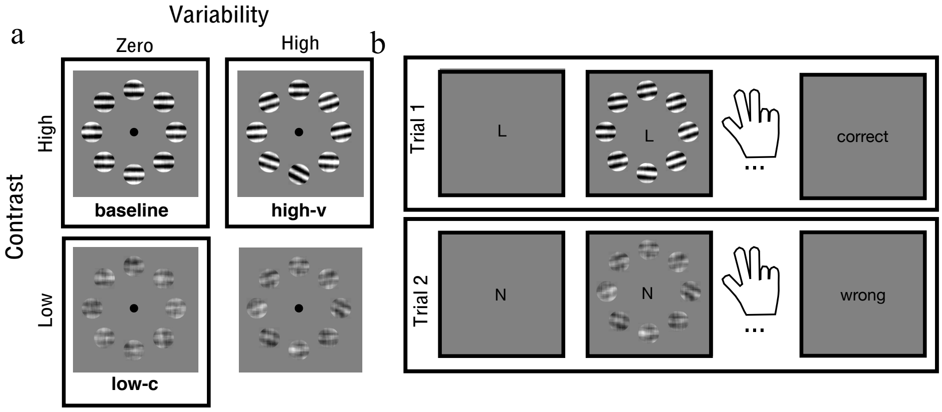

Figure 4.1. Basic experimental paradigm.

We manipulated stimuli contrast and orientation variability independently across trials. Panel a, shows example stimuli for two levels of contrast (root mean square contrast { 0.15, 0.6}) and two levels of variability (standard deviation of orientations {0°, 10°}). Participants categorized the average orientation of the 8 gratings as clockwise (CW) or counter-clockwise (CCW) from horizontal. In experiments 1 and 2, a fixation dot appeared at the centre of the screen for 300 ms (not shown) to announce the start of the trial, then the prior cue was shown for 700ms before stimulus appearance and until a response was made. This cue determined the prior probability of occurrence of each category (L=75% likely to be CCW, N = 50% and R= 75% CW). Stimulus duration was 150 ms, and participants responded within 3000 ms. For experiments 1 and 3, feedback was delivered for 500 ms after response (panel b).

When the cue was "N" it indicated that either category was equally likely to occur, when the cue was "R" it indicated a 75% probability that the category of the forthcoming stimuli was CW (and thus 25% CCW), similarly when the cue was "L" there was a 75% probability that the category of the forthcoming stimuli was CCW. For half of the blocks (neutral blocks), the cue in all trials was "N", for the other half of blocks (biased blocks), the cue on each trial could be either "R" or "L", in random but counterbalanced order. Block order was also randomized for each participant. Feedback about the correctness of each decision was given at the end of each trial by showing the word "CORRECT" or "WRONG" (Calibri, 16 points) at the centre of the screen for a duration of 0.5 seconds, after which a new trial immediately followed. Participants had a deadline of 3 seconds to report their decision, if no response was registered before this deadline, the word "LATE" appeared at the centre of the screen.

Five variations of this paradigm were implemented in experiments 2, 3, 4 ,5 and 6 (see Figure 4.2). For **experiment 2**, we asked participants to report how confident they felt about their response in each trial, participants were required to express their confidence by adjusting a horizontal marker up and down a scale (from 50% to 100% in steps of 1%) using a computer mouse. This confidence report occurred after their choice was made on each trial and was subsequently followed by categorical feedback about their choice.

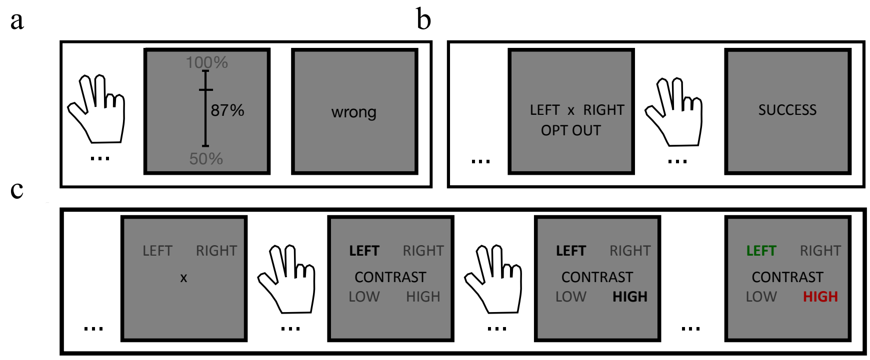

Figure 4.2. Variations of the paradigm.

For experiment 2, an interim stage occurred between response and feedback, requiring participants to report their confidence by sliding a horizontal marker up and down a scale (using a computer mouse). The confidence scale ranged from 50% to 100% in 1% increments, the value of the marker was displayed in real-time at the centre of the screen, participants pressed the spacebar to commit their confidence and proceed to the feedback screen (panel a). For experiment 3, after the stimulus disappeared a screen showed the options available ("LEFT" and "RIGHT" always and also "OPT OUT" in the opt out blocks), panel d. For experiment 4, two consecutive decisions had to be made about the stimulus before feedback was revealed: the first decision was always about the average orientation of the gratings, the second decision was to categorize either the contrast or the variability of the stimulus (randomly and unpredictably on each trial) as high or low, the chosen option at each stage was highlighted before the correctness of both choices was revealed by a change in the colour (green for correct and red for wrong), panel e.

In **experiment 3**, to probe the participants implicit sense of confidence, we replaced the "biased trials" of experiment 1 (those were the prior probability was different from 50%), by trials where a third option was available. This "opt-out" option, allowed participants to abort the current trial in exchange of having a 75% probability of obtaining positive feedback for the trial. Intuitively, participants should abort the trials more often when their confidence is lower.

For **experiment 4**, we required participants to make 2 subsequent decisions about the stimuli they saw on each trial. The first decision, as with all other experiments, was a categorical decision about the average orientation of the Gratings. The second decision, was to report (categorically) either: (i) if the contrast of the stimuli was low or high, or (ii) if the variability of the stimuli was low or high. This second decision allowed us to probe the sensitivity of participants to the two dimensions linked to the two sources of noise. Which dimension (contrast or variability) was enquired in each trial was determined randomly. After the decision for the two stages was made feedback about the correctness of each choice was delivered.

Experiment 5, was identical to experiment 4, however, instead of requiring participants to judge the level of contrast or variability they were in, we decided to keep at a constant for the duration of a block of trials, either: (i) level the contrast (half of the blocks low, half of the blocks high) or (ii) the variability (half of the blocks low, half of the blocks high).

Finally, in **experiment 6**, the task was the same as in experiment 2 (see more details below), with a crucial manipulation that on half of the trials (randomly interspersed) the stimulus array was composed of eight items but on the other half of trials the stimulus was composed of only four.

4.2.3 Further details of task and stimuli

The centre of each grating was located at a distance of ~4.3 degrees of visual angle (400 pixels) from the centre of the screen. A Gabor patch was used to create each of the gratings using the following parameter values: diameter of ~1.07 degrees of visual angle (100 pixels); spatial frequency of ~5 cycles per degree of visual angle (0.05 cycles per pixel); random phase. For each trial, all gratings had the same level of root mean square contrast (rmc), but the rmc of each trial was either 0 (no signal), 0.15 (low contrast) or 0.60 (high contrast). For experiment 4 (see description below), only the two higher levels of contrast were used. The average orientation of all gratings in each trial (henceforth trial mean) was drawn randomly from a Gaussian distribution with mean 3º either side of horizontal and standard deviation of 8º. The variability of orientations in each trial (stimuli variability) was also randomly drawn from a Gaussian distribution with mean 0º and standard deviation of either 0º (zero variability), 4º (medium variability) or 10º (high variability). For experiment 4 (see below), only the two most extreme variability levels were used. In each trial, we subtracted the mean deviation from 0 from the stimuli variability to ensure the originally drawn trial mean remained unchanged after adding stimuli variability. This allowed us to have orthogonal manipulations of stimuli contrast, stimuli variability and trial mean across the task. The correct category on each trial was determined by the actual average orientation of the 8 gratings and not by the distribution from which they were drawn.

To manipulate the quality of the sensory information, we generated low level random noise with similar spatial properties to the gratings (Wyart, Nobre, & Summerfield, 2012). To do this, we convolved a patch of white noise (unique for each grating across the task) with a two-dimensional Gaussian with zero mean and standard deviation of ~0.21 degrees of visual angle (20 pixels). The amplitude of the low-level noise was 10% of the maximum possible. We superimposed (additively) each grating with its low-level noise. Finally, we convolved the noisy patch with a 2-dimensional Gaussian envelope peaking at the centre of the grating and decaying with a standard deviation of ~0.21 degrees of visual angle (20 pixels).

For twenty of our participants (Exp2), we used a modified version of the task (the confidence task). The confidence task was identical to that in Exp1 with the following exceptions. The decision on each trial was reported using the left hand to press the “X” (clockwise) and “Z” (counter-clockwise) keys of the keyboard. Furthermore, we collected a judgment of confidence about the decision of each trial. Participants reported the confidence they had on their decision being correct by adjusting the vertical position of a marker (sliding the computer mouse with their right hand) on a continuous scale from 50% to 100% in 1% increments. During the adjustment period, the percentage corresponding to the marker's position was updated in real-time and shown at the centre of the screen. After adjustment, participants committed their confidence judgment by clicking the left button of the computer mouse. There was no deadline for the confidence judgment. Categorical feedback about the correctness of their decision occurred immediately after the confidence judgment was made.

We gave performance related monetary reward to our participants to incentivize good performance, while still keeping an average payment rate of ~£10/hour. For the task in Exp1 and Exp3, participants earned £10 for completing the experiment with a bonus of £1 for every 2% increase in their overall performance in excess of 60%. For the confidence task (Exp2) participants earned £5 for completing the task and a reward bonus of ~£12, based on their confidence judgments. For this bonus payment we used a quadratic scoring rule that incentivizes participants to: 1) maximize their proportion of correct responses and 2) reveal as accurately as possible their true belief about the probability that their decision was correct (Sonnemans, Theo Offerman, & others, 2001).

For Exp3 (n = 20), we substituted the biased blocks of trials with blocks of trials where a third option was available for choice. This "opt out" option, was rewarded on 75% of trials when chosen, if participants instead decided to judge the stimuli as titled CW or CCW, then they were rewarded according to whether they were correct or not, thus allowing participants to opt-out of trials when they felt more uncertain about the stimulus category membership. Participants were informed of that left meant CCW and right meant CW. To remind participants of the available options on each trial (even when the available options remained constant within a block of trials), the words "LEFT" and "RIGHT" appeared to the left and right of the fixation cross respectively, after the stimulus disappeared; on trials where the "opt out" option was available, "OPT OUT" appeared below the fixation cross signalling that the downwards key could be chosen to opt out of making a choice on that trial. Also, to make the feedback from choices with the uncertainty of the feedback when choosing the “opt out” option, we exchanged the feedback words given after every choice: instead of “CORRECT” and “WRONG” we used the words “SUCCESS” and “FAILURE”.

Experiment 4 (n= 24) was identical to Exp1 with the following exceptions. There were 32 blocks of 40 trials each. Both the medium-variability condition and the zero-contrast condition were removed from the set, leaving only the zero and high-variability and the low and high-contrast conditions manipulated in a factorial design. After categorizing the average orientation of the stimuli as CW or CCW from horizontal (by pressing the “X” and “Z” buttons of the keyboard using their left hand), the selected option was highlighted on the screen but without revealing the correctness of such decision. Participants were then required to judge if the contrast of the stimuli (on half of the trials) was high or low, or if the variability of orientations in the stimuli (on the other half of trials) was high or low. The dimension to be judged as high or low (contrast or variability) was randomly selected on each trial and revealed to the participants only after the average orientation decision was made. Either the word “CONTRAST” or “VARIABILITY” appeared at the centre of the screen replacing the fixation point, and the words “HIGH” and “LOW” appeared equidistantly to the right and left, respectively. Once participants judged the enquired dimension (contrast or variability) as high or low, using their right hand to press the right and left arrows of the keyboard, the selected option was highlighted for 0.3 seconds and then the correctness of the two decisions was revealed by changing the colour of the selected options to green when they were correct and to red when they were wrong, this feedback remained shown for 0.7 seconds before proceeding to the next trial.

Experiment 5 (n= 24) was identical to Exp4 with the following exceptions. Instead of prompting participants to discriminate the contrast or variability, we decided to block on each of the 32 blocks of trials one of the two dimensions (contrast or variability) at either of the two levels (high or low), while leaving the other dimension free to vary trial-by-trial. The 4 types of blocks occurred 8 times each in random intermingled order for each participant. Feedback for this task followed each choice as in Exp1. Crucially, our manipulations in this task relief participants from having to attend to the two critical dimensions (contrast and variability) simultaneously; and it allows us the experimenters to assess if the blindness is a consequence of participants' inability to switch, on a trial-to-trial basis, the dimension they condition their evidence on.

Experiment 6 (n=21) was identical to Exp4 with the following exceptions. Participants were not asked to categorize the contrast or variability. Instead, the crucial manipulation in this experiment was that, in half of the trials (randomly interleaved), the stimulus array was composed of 8 items (as in all previous experiments) and in the other half of trials the array was composed of only 4 items. The positions of the 4 items were contiguous and fixed within a block (randomized across blocks). Thus, there was no uncertainty about where the 4 items would occur in each block. However, whether the stimulus array would be composed of 4 or 8 items was unpredictable. As with all previous experiments the stimuli were chosen such that all the crucial manipulations (i.e. contrast, variability, prior cue, set-size, average orientation) were independent and uncorrelated. In particular, we used a correction to make sure the measured standard deviation was equal across the different set-sizes (Gurland & Tripathi, 1971). This set-size manipulation was implemented to probe if participants were integrating most of the items in the stimulus array, as "sub-sampling" would lead to performance costs in the 8 item condition compared to the 4 item one.

All participants from all experiments received a 10-minute presentation with detailed instructions of the relevant aspects of the task, such as: the types of stimuli used, the response contingencies, the different sources of difficulty in the task, how prior probabilities were manipulated, and the rules behind the payment bonus. Also, participants completed a short practice version of the task (the length of 1 block) to ensure they were familiar with it before they started the full-length version of the task.

4.2.4 Computational model

A noisy observer categorizing the average orientation of 8 gratings would not have access to the true distribution of average orientations in the task (Figure 4.3).

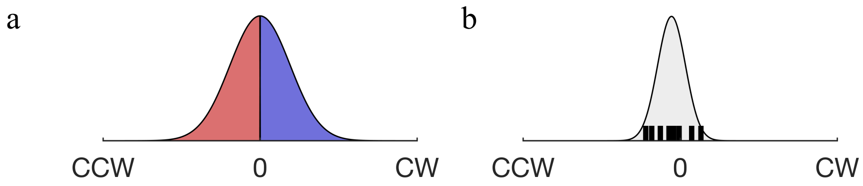

Figure 4.3. Distribution of orientations and noise.

The true category-conditioned distribution of average orientations (blue for CW; red for CCW). A noiseless observer could achieve 100% correct performance, panel a. Variability in choices is modelled as internally generated Gaussian noise, the position on the x axis of the black vertical lines represents the orientation of each of the 8 stimuli in an example trial (panel b). The white curve represents the likelihood of estimating, after the influence of noise, the average orientation of this example trial at all possible orientations.

In a trial, the observer's noisy estimate x, of the true average orientation μ, will be drawn from a Gaussian distribution with mean μ, and variance σ^2:

$$x = \epsilon(\mu, \sigma^2) \qquad \text{Equation 16}$$

Thus, a noisy optimal observer will construct a generative model of the task using its noisy estimates of orientations, conditioning them by the category that they belong to. This generative model of the task is composed of a pair of functions that describe the probability of observing a given orientation, x, conditional on each category:

$$PDF_{cat} = \mathrm{P}(x|cat) \qquad \text{Equation 17}$$

where cat is one of the two categories (CW or CCW).

The observer would have one PDF for the CW category and one for the CCW one. We constructed these category-conditioned density functions, PDF_{cat}, by convolving the true distribution of average orientations in the task, with a zero centred Gaussian distribution with variance σ^2 (Figure 4.3, panel b). For simplicity, we assume an optimal observer that would readily have knowledge of this structure.

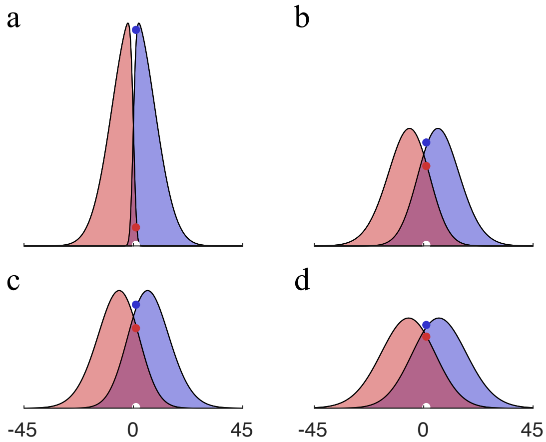

Figure 4.4. Condition specific density functions.

The generative model of the noisy observer for each experimental condition, is composed of a pair of category-conditioned density functions describing the likelihood of observing a noisy average estimate contingent on each category (blue for CW; red for CCW), the white dot on the x axis denotes the mean of an example trial, the blue and red dots represent the likelihood that such an estimate belongs to the CW and CCW categories, respectively. The four panels represent the density function for each of the four conditions: panel a, high-contrast and low-variability (lowest noise); panel b, high-contrast and high variability (intermediate noise); panel c, low-contrast and low-variability (intermediate noise); and panel d, low-contrast and high-variability (highest noise).

In different trials the prior probability of occurrence is different for each category, in some trials the prior cue favours the CW category, in others it favours the CCW category and in others it favours neither. The prior probability of occurrence of each category $\mathrm{p}(cat)$, needs to be incorporated into a posterior probability of the category given the noisy evidence and the prior cue. An optimal observer, would scale the density function of each category by the its prior probability of occurrence. See FigS1.D for an example of the scaled density functions when the prior cue favoured the CW category, and thus $\mathrm{p}(cw) = 0.75$ and $\mathrm{p}(ccw) = 0.25$.

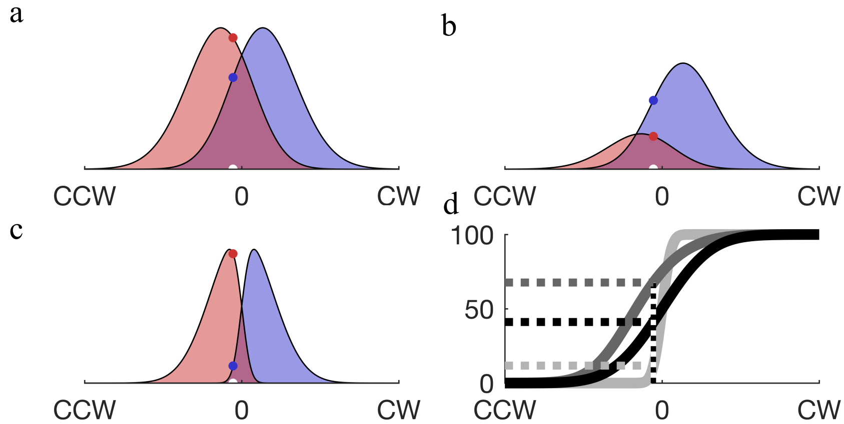

Figure 4.5. Prior knowledge and blindness to noise.

The category-conditioned density functions for a single condition under two different prior cues (panel a and b). Knowledge about the prior probabilities is incorporated by weighting each category-conditioned density function by its prior probability 50%CW (and thus 50% CCW) in panel a, and 75% CW (and 25%CCW) in panel b. The generative model of the "noise-blind" model, has less overlapping density functions because it only incorporates encoding noise in its generative model but ignores its integration noise (panel c). The posterior belief about average orientation being CW (panel d) after accounting for prior information (solid curves). The three solid curves show the posterior belief as a function of the true average orientation for the cases shown in panel a (optimal model in neutral trials), panel b (optimal model in trials where the prior favours CW with 75% probability) and panel c ("noise-blind" model in neutral trials). Using the same example trial, we show the posterior belief that the stimulus is CW (horizontal dotted lines) for each of the examples shown in panels a, b and c. Note that for the same noisy estimate, the noise-blind model in a neutral trial (light grey line) holds a different belief that the average orientation is CW compared to the optimal model in a neutral trial (black line) and from the optimal model in a trial cued as 75% likely to be CW (dark grey line). Also, the steepness of the posterior function for the noise-blind model is much greater than that of the optimal model, denoting the noise-blind model's tendency to hold stronger (further from 50%) beliefs about the category membership of the average orientation of the stimulus.

Then, the posterior probability of the category being CW given the noisy evidence and the prior cue will be computed as follows:

$$\mathrm{p}(cw|\ cue, x) = \frac{\mathrm{P}(\mathcal{X}|CW)\cdot \mathrm{p}(cw)}{(\mathrm{P}(\mathcal{X}|CW)\cdot \mathrm{p}(cw)+\mathrm{P}(\mathcal{X}|CCW)\cdot \mathrm{p}(ccw))} \qquad \text{Equation 18}$$

To make a decision about the category of the average orientation, the observer selects the category with the highest posterior. In a given trial (CW or CCW), the observer compares the posterior probability that the category is CW against a 0.5 threshold, if $\mathrm{p}(cw|\ cue, x)$ is above the threshold it chooses cw, otherwise it chooses ccw.

After the observer makes a choice ch (cw or ccw), its confidence about having made the correct choice will be:

$$p(ch|\ cue, x) \qquad \text{Equation 19}$$

Because the amount of noise in a trial depends on the condition that the trial belongs to (e.g. conditions with low contrast stimuli have higher noise), an optimal observer would construct a separate pair of category-conditioned density functions for each of the conditions in the task in order to compute a posterior belief about the category of the average orientation of the stimulus. When making its choices, this observer is able to know which of the six conditions it is in on the current trial. Having a separate pair of category-conditioned density functions for each condition, thus, allows the observer to i) optimally use the information provided by the prior cue when making its choices, ii) estimate a truthful probability of its choices being correct, and iii) be able to optimally opt-out of more difficult trials.

For instance, in a condition with high noise (a difficult condition), the density functions for the CW and CCW categories will overlap much more than in an easy condition, this difference in overlap of the density functions means that the same noisy estimate in a more difficult trial will result in a more uncertain posterior belief. Thus, in more difficult conditions the observer should, compared to an easier condition: i) display in its choices a higher influence of the prior cue, ii) report lower confidence about the probability of being correct, and finally iii) have a higher proportion of trials opted out of.

To support our hypothesis that people are sensitive to encoding noise but blind to their integration noise, we created two modified versions of the ideal observer model: the "noise blind" and the "variability-blind" models. The "noise-blind" model, builds the category-conditioned density functions after the evidence is affected by encoding noise (see below) but it is blind to the fact that it further incurs in error while integrating disparate evidence from the 8 samples (i.e. it ignores it has integration noise). Therefore, in conditions of high integration noise (high stimulus variability), its generative model of the task will consist of density functions that are less overlapping than they should be (see Figure 4.5, panel c). Essentially, this “noise-blind” model has only 2 distinct pairs of category-conditioned density functions, one pair for each of the levels of encoding noise (i.e. the two contrast levels). These two pairs of density functions are equivalent to the density functions of the optimal model in the two zero variability conditions (zero-variability-high-contrast and zero-variability-low-contrast). In other words, this model behaves as if there was zero variability in the orientations of the stimulus array but is otherwise optimal. Because this “noise-blind” model ignores it has integration noise, it will tend to hold systematically stronger beliefs about category membership than the optimal model (Figure 4.5, panel d). Thus, behaviour in the conditions with higher integration noise (higher variability conditions), will be characterized by i) a lower influence of the prior cue on choices, ii) overconfidence in its belief about making a correct choice, and iii) a reduced proportion of trials opted out of.

As a control, the “variability-blind” model is not blind to its integration noise but it is just incapable of determining which condition of integration noise it is on in a particular trial (i.e. what the stimulus variability condition is for each trial). Thus, this model too, has 2 pairs of category-conditioned density functions, one for each of the two contrast levels, however these density functions are built as a mixture of the distributions across all levels of integration noise (all variability conditions). In other words, because this model is not blind to integration noise, but simply cannot pin down which condition of late noise it is on, it just treats all trials as coming from either of the contrast conditions (therefore two pairs of density functions) but from a single integration noise condition made up of the mixture of the densities across all integration noise conditions. Thus, the two pairs of category-conditioned density functions for this model are equivalent to the functions of the optimal model in the two contrast conditions averaged across the variability conditions. Behaviourally, this “variability-blind” model makes very similar predictions to the “noise-blind” model. However, a key difference is that this model predicts underconfidence in conditions with below average integration noise and overconfidence in conditions with above average integration noise, whereas the “noise-blind” model predicts overconfidence only in conditions with higher integration noise.

Finally, all the models presented (optimal, noise-blind and variability-blind) make the same predictions about choice probabilities in neutral trials and are only distinguishable in terms of their second order judgments (confidence and opting out) or in their overall behaviour in biased trials (choices, confidence and opting out).

4.2.5 Noise estimation

Both the level of contrast and the orientation variability of the stimuli determined the amount of noise in human choices. In an independent procedure from the modelling described above, we estimated, for each participant, the amount of noise incurred by each of our experimental manipulations (2 contrast levels and 3 variability levels).

Four parameters allowed us to characterize the amount of noise in each of the six conditions. Two parameters described the amount of encoding noise for the different contrast conditions: one for low contrast and one for high contrast. Two further parameters allowed us to characterise the amount of integration noise arising from the conditions with medium and high orientation variability. For each condition, the total amount of noise in a given condition, σ_{cond}, would be a function of the amount of encoding noise and integration noise in that condition:

$$\sigma_{cond} = \sqrt{(\varepsilon\sigma_{cond}^2) + (\iota\sigma_{cond}^2)} \qquad \text{Equation 20}$$

where $\varepsilon\sigma_{cond}$, is the amount of encoding noise in the particular condition, and $\iota\sigma_{cond}$, is the amount of integration noise in the particular condition.

The four noise estimators for each participant, were obtained by maximizing the likelihood of the participant's choice in neutral trials (the likelihood for each trial was calculated based on the total noise in its condition). Finding the best parameter set, was carried out using a genetic algorithm with the following parameters: a population size of 1000 individuals and a maximum generation time of 1000 generations.

Note that, because the estimation of noise is done on neutral trials only, and is independent of the modelling process described in the previous section, the amount of noise estimated for each participant in each condition was then validly used to generate the behaviour of all models in the biased trials.

4.2.6 Drift diffusion model

Intuitively, slower response times for higher variability trials are incompatible with the idea of having a single threshold towards which evidence is accumulated. However, to illustrate this point, we used a simple fitting procedure to be able to capture the cost to performance of decreasing contrast and increasing variability. Three free parameters allowed us to do this: the first parameter is a scaling factor for the drift rate that allowed low contrast trials to have reduced drift rates compared to high contrast trials, the second parameter is a baseline amount of noise in the drift process, and the third parameter controls the degree to which increases variability trials have increased noise in the drift process compared to zero variability trials. To find the best fitting parameter set for each participant we used a genetic algorithm with a population size of 1000 individuals and a maximum generation time of 1000 generations, to minimize the sum of squared error in the average accuracy for each of the 6 conditions between human responses and model responses.

This simple drift diffusion model allowed to capture well the pattern of accuracies across the 6 conditions, with lower accuracy in the model both for decreased stimulus contrast ($F(1,39) = 1946.8$, $p < 0.001$) and increased orientation variability ($F(1.1,42.9) = 58.2$, $p < 0.001$). However, and as predicted intuitively, the response times became faster with increased variability ($F(1.0,39.5) = 30.1$, $p < 0.001$) a pattern that is in strict opposition with what we observed from humans.

4.2.7 Analyses

All the statistics reported are computed at the group level. Trials with no signal (rmc=0) in Exp1-3, allowed us to infer participants' baseline level of cue usage but because the variability manipulation could not make a difference in such trials, we thus excluded them from most analyses. Experiments 4 and 5 no longer included the manipulation trials with zero contrast or medium variability. We performed within-subject analyses of variance (ANOVA) to compare accuracy, cue usage, confidence and opting out as a function of the level of stimuli contrast (low and high) and variability (zero, medium and high in Exp1-3 and only zero and high in Exp4). Subjective confidence reports were only available for experiment 2 (n=20) and the possibility to opt-out only in Exp3 (n=20). We carried out the analyses of accuracy and confidence on all trials from the neutral blocks, whereas for the analysis of cue usage we used trials from the biased blocks.

To compute cue usage, we separated all biased trials by the type of cue (R or L) and calculated the decision criterion, *c*, which provides a signed estimate of the effect of the cue on biasing responses independent of participants' sensitivity (Macmillan & Creelman, 2004; Stanislaw & Todorov, 1999). We then compared the difference in values of *c* when cued clockwise (C_{CW}) and criteria when cued counter-clockwise (C_{CCW}) as our measure of cue usage (bias index) for each stimuli contrast and variability condition (Prior Usage = C_{CW}-C_{CCW}). Values closer to zero indicate lower usage of cue whereas stronger departures from zero (positive or negative) indicate stronger usage of cue. To further explore the effect of variability on confidence and opt out, we performed multiple linear regressions to predict the confidence judgments (n=20) and the opt-out responses (logistic regression, n=20). The model description and the details on noise estimation for the modelling are given in the supplementary information.

4.3 Results

4.3.1 Matched costs to performance from encoding and integration noises

In experiments 1 and 2 (n = 20 each), healthy humans discriminated the average orientation of 8 tilted gratings presented in a circular array (see Figure 4.1 in the methods section), as oriented clockwise (CW) or counter-clockwise (CCW) from the horizontal axis. We manipulated encoding noise by degrading the average contrast of the gratings (root mean square contrast {0.15, 0.6}) and integration noise by increasing the orientation variability of the gratings in the array (standard deviation of orientations {0°, 4°, 10°}). Participants received deterministic trial-to-trial feedback following every choice thus ensuring that only internal noise was responsible for errors in the task. As expected, accuracy decreased with lower contrast (Exp1: $F(1,19) = 15.54$, $p < 0.001$; Exp2: $F(1,19) = 41.08$, $p < 0.001$; collapsed: $F(1,39) = 49.3$, $p < 0.001$) and higher variability (Exp1: $F(1.3,24.7) = 8.51$, $p < 0.01$; Exp2: $F(1.6,32.2) = 26.0$, $p < 0.001$; collapsed: $F(1.4,57.3) = 30.61$, $p < 0.001$).

Our analysis focusses on (i) a baseline condition where overall noise is lowest (high contrast, 0.6; zero variability, 0º), and two increased-noise conditions that were matched for performance: (ii) a "low-c" condition with lower contrast (increased encoding noise) compared to baseline (low contrast, 0.15; zero variability, 0º), and (iii) a "high-v" condition with higher variability (increased integration noise) compared to baseline (high contrast, 0.6; high variability, 10º). We first compared these conditions in a "neutral" context, where prior probability of occurrence of CW and CCW stimuli was equated. This analysis confirmed that relative to baseline, accuracy was reduced on both low-c and high-v conditions by ~12% compared to baseline (baseline>low-c: $t(39) = 9.24$, $p < 0.001$; baseline>high-v: $t(39) = 9.70$, $p < 0.001$). Importantly, the manipulations of contrast and variability reduced performance to an equivalent extent (Exp1, high-v>low-c: $t(19) = 0.36, p > 0.7$; Exp2, high-v>low-c: $t(19) = 0.11, p > 0.9$; collapsed, high-v>low-c: $t(39) = 0.34, p > 0.7$; see Figure 4.6).

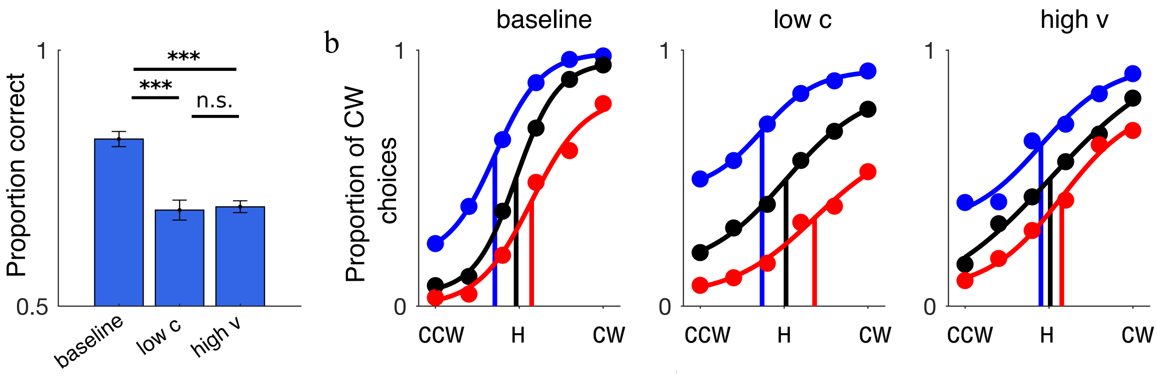

Figure 4.6. Behavioural effects of stimulus contrast and variability on choices.

Accuracy for the baseline, reduced contrast (low-c) and increased variability (high-v) conditions across participants (panel a, blue bars, n = 40), statistical comparison between these conditions was carried out using a Student's t-test (* *p < 0.001*). Panel b, psychometric curves for neutral trials (black curves) showing shallower slopes (increased noise) for the low-c and high-v conditions. Shifts towards the cued category in biased trials (blue lines, 75% CW; red lines 25% CW) are stronger in the low-c condition and weaker in the high-v condition (vertical lines show inflexion points for visual purposes only).**

4.3.2 Are people optimal in using prior information under both sources of noise?

This matched performance allowed us to compare the optimality of inferences, by measuring how participants incorporated prior probabilistic information into their perceptual decisions. To this end, we introduced on the remaining half of trials a "biased" context, in which a cue occurring before the stimulus provided probabilistic information about the correct category of the forthcoming stimulus (CW vs. CCW) with 75% fidelity. An ideal observer should be more influenced by the cue when sensory evidence is less reliable, i.e. when encoding or integration noise are higher. First, we used decision-theoretic tools (Macmillan & Creelman, 2004; Stanislaw & Todorov, 1999) to compute a bias index to quantify the adjustment of the decision criterion under different prior cues, with greater influence of prior information reflected by larger values (Figure 4.7).

Consistent with previous research showing humans weight information by its reliability (Ernst & Banks, 2002; Körding & Wolpert, 2004), participants were more swayed by prior information in the low-c than in the baseline condition ($t(39) = 4.89$, $p < 0.001$).

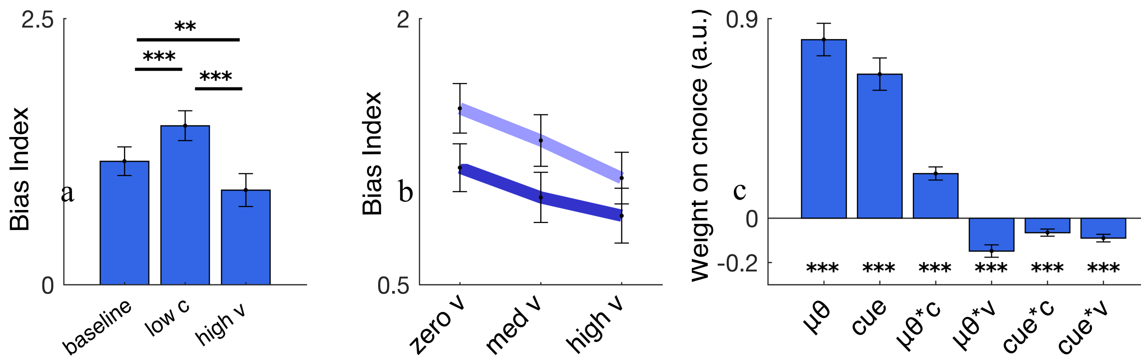

Figure 4.7. The influence of prior information.

Panel a, the influence of the prior cue increases with higher encoding noise (low-c) and lower integration noise (high-v) compared to baseline. Panel b, an overall increase in bias index for lower contrast but a decrease for higher variability, seen through all combinations of contrast (high contrast, dark blue; low contrast, pale blue) and variability (zero, medium and high variability). Panel c, the influence of prior information on trial-to-trial choices, measured through a logistic regression, decreases with higher variability (cue*v). The remaining predictors: average orientation (μθ), prior cue (cue), and the interaction between each of these predictors and the stimulus contrast (c) and variability (v), show an effect in the expected direction. Error bars represent standard error of the mean (SEM) across participants. We assessed statistical significance, for the full factorial design by two-way analysis of variance (ANOVA) correcting degrees-of-freedom for violations of sphericity when appropriate, for pairwise comparisons by Student's t-tests, and the significance of the regression coefficients by a one-sample t-tests; (*$p < 0.05$, ** $p < 0.01$, *** $p < 0.001$).

Strikingly, and contrary to expectations from ideal-observer frameworks, participants were less swayed by prior information in the high-v than in the baseline condition ($t(39) = 2.85$, $p < 0.01$; see shifts in psychometric curves in Figure 4.6). The full factorial analysis of the bias index confirmed a reliable main effect of contrast in the expected direction ($F(1,39) = 24.02$, $p < 0.001$), and of variability in the opposite direction ($F(1.9,37.1) = 9.9$, $p < 0.001$; see panel b of Figure 4.7). Using the entire dataset, we then relied on logistic regression to test how contrast (c) and variability (v) mediated the influence of the cues and sensory evidence (μθ) on trial-to-trial choices (μθ, cue, μθ*c, μθ*v, cue*c, cue*v; see panel c in Figure 4.7). The cue influenced choices more under low than high contrast ($t(39) = 4.05$, $p < 0.001$) but also under low than high variability ($t(39) = 5.21$, $p < 0.001$). Together, these findings suggest that participants systematically neglect the fact that stimuli of higher variability are harder to categorize, as if they were blind to the uncertainty arising while integrating more variable pieces of evidence.

4.3.3 Are people metacognitively blind to integration noise?

The metacognitive account described above proposes that participants are suboptimal when information is variable because they are blind to the loss incurred during integration of discordant decision information. In Experiment 2, thus, we tested directly whether participants' certainty was reduced with lower contrast but not higher variability, by asking participants to report their confidence about being correct on a continuous percentage scale (**Fig. 1C**). We incentivised participants both to make correct decisions and report their confidence faithfully by implementing a proper scoring rule (Sonnemans et al., 2001). Analysis of the full factorial design indicated that, as predicted, confidence decreased with increased contrast ($F(1,19) = 32.97$, $p < 0.001$) but not variability ($F(1.2,22.5) = 0.73$, $p > 0.4$).

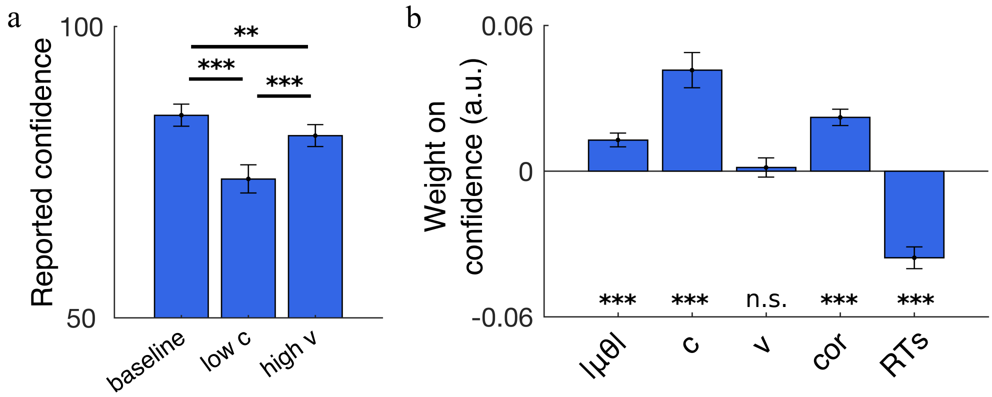

Figure 4.8. Effects of stimulus contrast and variability on confidence.

Mean reported confidence for the baseline, low-c and high-v conditions across participants (blue bars), showing significantly greater confidence for the high-v condition over low-c one, even when matched in accuracy (panel a, n = 20). Panel b, trial-to-trial confidence remains unaffected by the orientation variability (v), but is affected in the expected direction by: the average orientation being further away from horizontal (|μθ|), higher stimulus contrast (c), choices being correct (cor), and slower response times (RTs).

Moreover, targeted analyses revealed lower confidence in the low-c than high-v conditions ($t(19)= 3.98$, $p < 0.001$) despite equivalent discrimination accuracy. In other words, participants were more overconfident in the high-v than low-c condition (Figure 4.9), as demonstrated by comparing the extent to which their confidence exceeded their accuracy ($t(19) = 2.66$, $p < 0.05$).

There was a small reduction in confidence in the high-v condition compared to baseline (Figure 4.8, panel a), but this can be explained by participants using their response times as a proxy for their correctness (Zakay & Tuvia, 1998): a regression analysis revealed that trial-to-trial confidence reports decreased with longer response times (RTs) but were unaffected by variability once this additional factor was accounted for (v: $t(19) = 0.38$, $p > 0.7$; all other t-values > 4, all $p < 0.001$; see Figure 4.8, panel b).

Because human confidence reports can sometimes be idiosyncratic, in Exp3 we obtained an alternative, implicit measure of confidence by offering a third alternative that provided reward with 75% probability (Hampton, 2001; Kepecs & Mainen, 2012; Kiani & Shadlen, 2009). This allowed participants to opt-out of trials if they felt uncertain, and us to measure the proportion of opt-out trials in performance-matched low-c and high-v trials. Despite the equivalent accuracy for low-c and high-v conditions ($t(17) = 0.24$, $p > 0.8$), participants preferred to opt-in on high-v than low-c trials ($t(17) = 2.32$, $p < 0.05$; Figure 4.9).

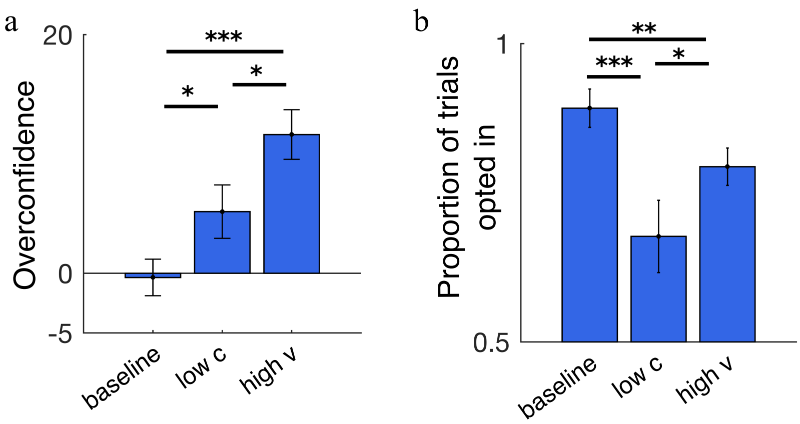

Figure 4.9. Overconfidence and proportion of trials opted-in.

Overconfidence, the difference between average confidence and percentage correct, is highest in the high-v condition (panel a). A significant increase in the probability of engaging with (opting-in) trials in the high-v condition compared to the low-c one and despite matched accuracy.

The full factorial analysis confirmed more frequent opt-in responses with higher contrast ($F(1,17) = 21.2$, $p < 0.001$) but not with lower variability ($F(1.4,23.9) = 3.6$, $p > 0.05$). A logistic regression confirmed that propensity to opt-in decreased for lower contrast ($t(17) = 6.93$, $p < 0.001$) but not for variability ($t(17) = 1.6$, $p > 0.1$). Remarkably, and despite fully informative trial-wise feedback, participants did not incorporate the accuracy cost of integrating more variable information into their sense of certainty about making correct decisions.

4.3.4 Noise blindness for an ideal observer model

To account for these behavioural observations, and to instantiate our explanation for decision suboptimality within a quantitative framework, we built a family of computational models based on the ideal observer framework, and compared them to the human data. We first determined, from participant's "neutral" trials, the level of noise governing choices in each of the experimental conditions. This allowed us to compute likelihood functions that related sensory information to choices in each condition, and to use these functions to predict independent discrimination judgments on "biased" trials (in Exp1-2). All model variants compute the probability that the average orientation is CW given the noisy sensory evidence and the prior cue, using a pair of category-conditioned density functions for each of the six experimental conditions (see methods). Because confidence was reported on a 50-100% scale, we could also predict confidence judgments under the assumption that model choice probabilities translated directly to reported percentage values (in Exp2-3).

We first simulated an optimal model of the task, for which the density of the likelihood functions depends on the combined noise arising during encoding and integration. Because the dispersion of the likelihood functions grows as contrast is reduced and variability increased, this model predicts that on biased trials, participants should adjust their decision criterion on both low-c and high-v conditions. The prediction of this optimal model is incorrect because humans only do so on low-c conditions.

Next, we simulated a "noise-blind" model, in which, the dispersion of the likelihood functions reflected only estimated noise in sensory encoding (not integration noise). Like humans, this model underestimates uncertainty occurring under higher orientation variability and is consequently overconfident about the sensory evidence, and thus systematically ignores the cue on high variability trials.

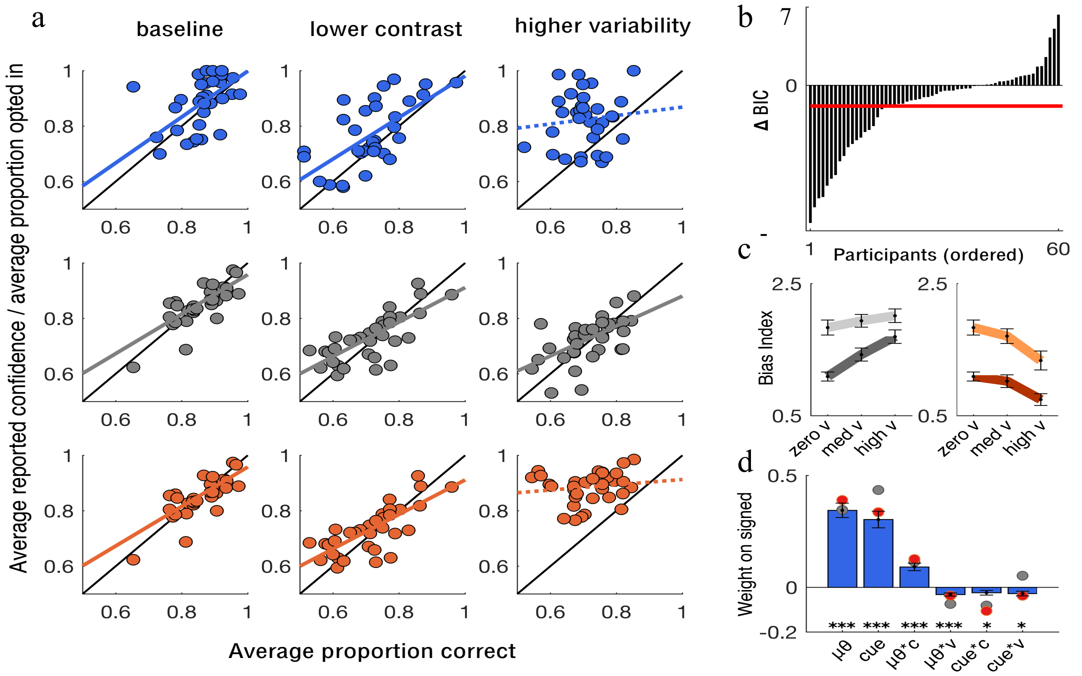

Figure 4.10. Comparison of model and human behaviour.

Panel a, human participants (blue), optimal observers (grey) and integration-noise blind observers (orange), with higher accuracy in the baseline and low-c conditions, report higher certainty in their choices (i.e. confidence reports and probability of opting in). Only optimal observers show this pattern in the high-v condition (best fitting slopes; solid for *p* < 0.001, dotted for *p* > 0.4). Seven participants excluded from this analysis (excessive opt-out), but results when including them were almost identical (N=33). Panel b, average ΔBIC across participants (-18.5) suggests very strong evidence for the noise-blind model. Panel c, lines show opposite influence of prior cue on choices from optimal (grey; positive slopes) and noise-blind models (orange, negative slopes), for increased variability but same pattern for reduced contrast (lighter lines above darker ones). Panel d, impact on signed confidence for humans (blue), optimal model (grey) and noise-blind model (orange). Humans and the noise-blind model predict lower influence of the prior cue for trials with higher variability (cue*v), a pattern which is in opposition with what is observed from the optimal model.

The noise-blind model (but not the optimal model) successfully captures three key features of the human data: (i) overconfidence on high-variability trials, (ii) a breakdown in the relationship between mean accuracy and the metacognitive measures of certainty (Figure 4.10, panel a) and (iii) the diminished influence of prior information in higher variability trials, as seen by the reversal in the bias index (Figure 4.10, panel c), and in the regression predicting signed confidence (Figure 4.10, panel d). Finally, quantitative comparison yielded very strong evidence supporting the noise-blind over the optimal model (Figure 4.10, panel b).

4.3.5 Is noise blindness the simplest explanation to the observed suboptimal behaviour?

One recent study has reported that humans and monkeys discriminating the direction of dot-motion stimuli are both faster to respond and more confident when the strength (coherence) of the frame-to-frame motion energy signals is more volatile (Zylberberg, Fetsch, & Shadlen, 2016). As the authors note, this can be explained under an (otherwise optimal) sequential sampling account if participants are unaware of the volatility in the signal, i.e. they adopt a common decision threshold for low- and high-volatility trials. This prompted us to examine an alternative explanation for our findings: participants are not suboptimal, but are simply unable to detect if the grating orientations were more (or less) variable on each trial (i.e. whether they are "variability-blind" rather than "noise-blind").

Intuitively, the latter explanation seems unlikely in Exp1-3, where our analyses compare "high" and "zero" variability trials, which should be clearly distinguishable by human participants. Nevertheless, to test this directly, we ran a further experiment (Exp4) that was very similar to Exp1, with the exception that after the orientation judgment participants were additionally asked to report, on each trial, whether the contrast was "high" or "low" (rms = 0.6 vs. 0.15) or whether the orientation variability was "high" or "low" (angle s.d. = 0 or 10), see Figure 4.2, panel c and methods. Performance in low-c and high-v conditions in the neutral context were again matched ($t(23) = 1.16$, $p > 0.25$), allowing us to compare bias indices on those biased trials where the variability discrimination judgments were correct and incorrect (Figure 4.11). We observed that participants still failed to adjust their criterion appropriately even on those trials where they correctly discriminated the variability, with more bias on low-c than high-v trials for both correct variability discriminations ($t(23) = 3.21$, $p < 0.01$) and incorrect variability discriminations ($t(23) = 4.05$, $p < 0.001$; see Figure 4.11). Levels of bias did not differ between correct and incorrect variability discrimination trials ($F(1,23) = 2.0$, $p > 0.1$).

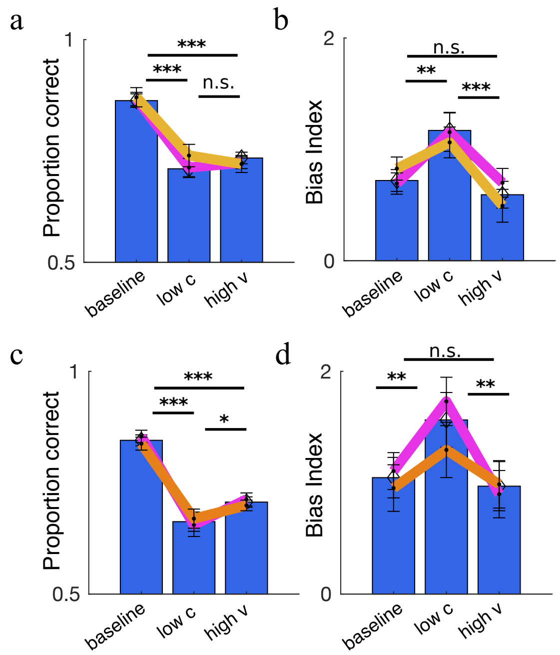

Figure 4.11. Behavioural results of experiments 4 and 5.

Behavioural effects of stimulus contrast and variability on choices for experiment 4 (n = 24) and experiment 5 (n = 24), for all trials (blue bars). For experiment 4, coloured lines indicate trials where the subsequent discrimination of contrast was correct (pink lines) or where the subsequent discrimination of variability was correct (ochre lines). For experiment 5, lines indicate trials where the contrast level was blocked (pink lines) or when the variability level was blocked (ochre lines). Panel a, accuracy for the baseline, reduced contrast (low-c) and increased variability (high-v) conditions across participants for experiment 4. Panel b, for experiment 4, the influence of prior information is highest when encoding noise is high (low-c) and lowest when integration noise is high (high-v). Panel c, same as panel A but for experiment 5. Panel d, same as panel B but for experiment 5.

Further, in experiment 5, we decided to block the manipulation of contrast and variability, thus freeing participants from the demand of estimating which condition they were in. Even when blocking the variability level participants were not more biased with higher variability compared to baseline ($t(23) = 0.32$, $p > 0.7$), whereas when blocking contrast, participants were more biased in the low-contrast condition compared to baseline ($t(23) = 3.31$, $p < 0.01$). Together, these results strongly suggest that "variability-blindness" is an unlikely explanation for participants' failure to adjust their decision criterion.

4.3.6 Neither a drift diffusion model or a variability blind model capture human behaviour

Heightened volatility of motion energy signals accelerates response times (Zylberberg et al., 2016), and so we additionally asked whether participants were faster when variability was higher. In our task, however, they were not. Consistent with other previous reports (De Gardelle & Summerfield, 2011), we observed significantly longer response times for increased variability conditions ($F(1.53,29.08) = 28.55$, $p < 0.001$; Figure 4.12). As such, we fitted a diffusion model to capture participant's accuracy, under the assumption that contrast reduces the drift rate and variability increases the noise in evidence accumulation, but the predictions of response times for the model were the opposite of what is observed from participants (Figure 4.12).

Finally, we fit another version of the ideal observer model to our data. This new model encapsulates the idea that participants could be blind to stimulus variability and not to integration noise. For this model, we assume a single category-conditioned density function marginalized over the different variability conditions (equivalent to the optimal model collapsing its density functions across the different variability conditions). Quantitative comparison across participants revealed strong support for the noise-blind model compared to the variability-blind model (ΔBIC= -9.2). Simply put, this result shows that people are not simply unable to assign the noise to its most likely cause, instead, people behave as if this extra source of uncertainty did not exist.

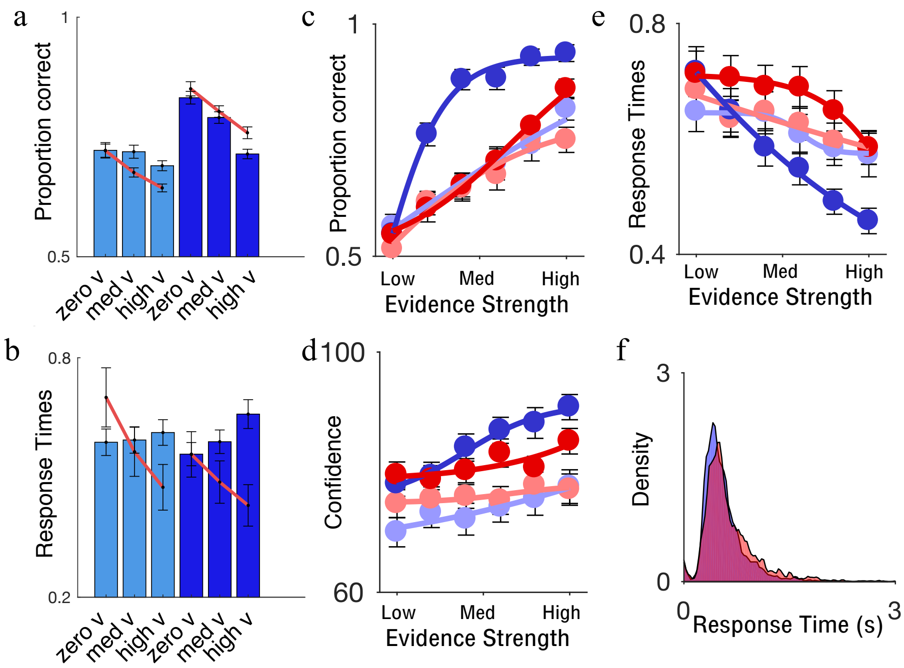

Figure 4.12. Drift Diffusion Model cannot account for data.

A simple drift diffusion model cannot account for slower responses when variability is high (experiments 1 and 2, n = 40). Error bars represent standard error of the mean across participants. Panel a, accuracies for human (blue bars) and drift diffusion model (red lines) are lower when contrast is low (pale blue bars shorter than dark blue bars), and when variability is high (negative slopes). Panel b, response times for humans and drift diffusion model show opposing effects of variability (positive slopes for humans but negative ones for the model). Panel c, choices are more often correct with increasing evidence (positive slopes n = 40), particularly in the baseline condition (zero variability and high contrast; dark blue curve); note that the conditions of increased variability (high variability high contrast; dark red curve) and reduced contrast (low contrast, zero variability; faint blue curve) have equated slopes. Panel d, stronger evidence leads to higher confidence reports (positive slopes, n = 20). The condition of increased variability (high variability high contrast; dark red curve) has higher confidence reports than the reduced contrast (low contrast, zero variability; faint blue curve) despite being equated in panel c. Panel e, evidence strength also increases the speed of choices (negative slopes n = 40). Response times are delayed by stimulus variability (red curves above blue ones). Panel f, collapsing the response times across all participants (n = 40) for the conditions of high contrast and zero variability (blue shade) and high contrast and high variability (pink shade) shows that increased variability shifts the response times to be slower.

4.3.7 Could people simply be sampling a subset of the stimuli?

Sampling a subset of the stimuli, that is, subsampling the stimuli, can in principle provide an alternative explanation as to: (i) why people have lower accuracy with higher variability, and (ii) why they do not adjust their confidence for higher variability stimuli. Thus, under a subsampling account, accuracy in high variability trials does not drop because of difficulty in integrating variable pieces of evidence but simply because of the increased mismatch between the real mean orientation of the stimulus and the mean orientation of the subset. However, we believe subsampling is not what drives our effects for the following reasons.

First and foremost, while subsampling could be an alternative explanation why people are blind to the variability of the trial, it does not explain why they are blind to the noise arising from subsampling. Second, while accuracy is expected to go down with high variability stimuli, the subjective difficulty for categorizing the subset of items should increase: the mean orientation of a subsample will tend to lie further away from the category boundary than the mean of all items. This decreased subjective difficulty would reflect in faster response times, which as shown before runs against our behavioural results.

Third, while the first two points suffice to show that subsampling cannot explain noise-blindness, we ran a further analysis and experiment to show that subsampling is not even likely to happen in our task. Thus, we simulated an agent that would have the same amount of encoding noise (as estimated from our participants' choices) as when discriminating high-contrast stimuli with zero variability (but we eliminated integration noise). We then simulated the accuracy of this agent would have in the condition of high-contrast and high-variability, when making decisions based on subset sizes of the stimulus (from one sample to all samples). As shown in Figure 4.13, panel a, the actual performance of our participants would require perfect integration of between three and four items. Finally, in experiment 6, we manipulated the size of the stimulus sample (four or eight items, with the actual standard deviation equated across set-sizes). If participants were subsampling, we would have seen a strong impact on high-variability trials for the stimuli with 8 items: where the maximum mismatch between the "actual mean orientation of the sample" and the "mean orientation of the subset of sampled items" is expected. However, the saw no significant difference from the set-size manipulation ($F(1,20) = 0.006$, $p > 0.93$), but strong differences driven by contrast ($F(1,20) = 40.9$, $p < 0.0001$) and variability ($F(1,20) = 30.50$, $p < 0.0001$).

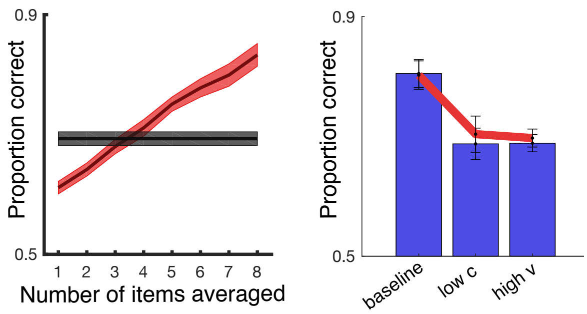

Figure 4.13. Participants are unlikely to engage in subsampling.

Panel a, average accuracy for human participants (grey horizontal shade) and for a perfect averaging model (red shade), i.e. one with no integration noise but the same amount of encoding noise as recovered from participants' responses (in the lowest noise condition). An observer would need to perfectly average the orientation of 3 to 4 items before reaching the same level of performance as our participants. Panel b, proportion correct for Exp6. Participants (n = 21) had the same performance when required to integrate 4 items (blue bars) compared to 8 items (red line). Both reducing contrast (high encoding noise), and increasing orientation variability (high integration noise) achieved the same level of performance cost compared to the baseline condition regardless of the number of items that the stimulus array was composed of (main effect of contrast, $F(1,20) = 40.9$, $p < 0.0001$; main effect of variability, $F(1,20) = 30.50$, $p < 0.0001$). If participants were not integrating most items in some conditions then performance in those conditions would be lower when the stimulus was composed of 8 items, however, performance was unaffected by the set-size manipulation ($F(1,20) = 0.006$, $p > 0.93$).

4.4 Discussion

Understanding why humans make good decisions in one setting and poor decisions in others is a longstanding problem in the behavioural sciences. Here, we propose a new explanation for the previously reported "gap" in optimality between perceptual and cognitive decisions. We found, akin to what has been reported in previous perceptual tasks (Ernst & Banks, 2002; Körding & Wolpert, 2004), that humans are sensitive to noise arising during sensory encoding, and adjust their decision policies accordingly, furnishing optimal decisions about stimuli that are low in contrast or otherwise difficult to perceive. Critically, however, we observed that humans are “blind” to integration noise, arising when multiple potentially conflicting pieces of information are integrated to make a choice. In our experiment, this noise blindness engenders two prominent biases that are characteristic of human judgments in real-world settings: overconfidence and base rate neglect. Firstly, even expert humans are overconfident about their performance when stimuli provide conflicting information about how to respond, for example when making medical (Saposnik, Redelmeier, Ruff, & Tobler, 2016) or financial (Frydman & Camerer, 2016) decisions. Secondly, when stimuli are characterised by multiple attributes, humans fail to adequately incorporate additional contextual information into their choices, such as information about stimulus base rates (Bar-Hillel, 1980). Indeed, humans seem to be particularly prone to economic biases when multiple attributes of a stimulus provide discordant or conflicting information, for example when comparing a high-value, high-risk prospect to a safer option with lower return (Pratt, 1964). We thus argue that the failure of metacognitive sensitivity to integration noise may be a key determinant of the suboptimal biases that pervade cognitive and economic judgments in naturalistic settings.

The effects of conflicting information on choice have mainly been investigated in a cognitive setting (MacLeod, 1991), although they have recently received attention in perceptual tasks too (De Gardelle & Summerfield, 2011; Zylberberg et al., 2016; Zylberberg, Roelfsema, & Sigman, 2014). A recent study using discrimination of dot motion signals has similarly reported that humans and monkeys are more confident when signals are more volatile (variable in time), even though this manipulation leads to a small but significant decrement in accuracy (Zylberberg et al., 2016). These authors can elegantly capture this finding within the framework provided by the drift-diffusion model (which offers an approximation to sequential Bayesian integration), by assuming that humans employ a common threshold-setting policy under high and low volatility. In their experiment, it may simply be that observers (who had been extensively trained on a low-volatility condition before exposure to volatile energy signals) may have failed to detect the heightened variability in motion energy.

By contrast, in the current report, humans were overconfident even when the differences in stimulus variability were perceptually salient, as well as when prompted to attend to the variability as a secondary, decision-relevant factor. Importantly, their biases were not influenced by whether variability discrimination was accurate or not. Control analyses allowed us to establish that participants were integrating most items in the stimulus array consistently across experimental conditions, and that participants are not simply blind to the variability of orientations in the stimulus array. Instead, we argue that participants are metacognitively blind to the cost to performance provoked by more variable stimuli. This is particularly surprising in our experiment, given that participants received fully informative feedback on every trial, which should have made it relatively easy to learn a likelihood function mapping of stimulus features onto the probability of a correct response.

We do not know why humans are “blind” to integration noise. One possibility is that resolving conflict among multiple attributes of a stimulus is highly computationally costly, and basing decision policies on the true noise estimate would lead to unduly prolonged deliberation that slows responses and thus reduces reward-harvesting rates. However, this remains to be tested. Nevertheless, the current work offers a new approach to understanding cognitive biases that corrupt decision-making. Other phenomena may also be within the theoretical reach of our account. For example, humans enjoy a rich and vivid subjective impression of cluttered natural scenes, even when a failure to detect salient change suggests that processing is highly limited (Simons & Levin, 1997). Human information processing is sharply limited by capacity – but as agents we may not be fully aware of the extent of these limitations.

References

Ackermann, J. F., & Landy, M. S. (2015). Suboptimal decision criteria are predicted by subjectively weighted probabilities and rewards. *Attention, Perception, & Psychophysics*, *77*(2), 638–658.

Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, *44*(3), 211–233.

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron*, *74*(1), 30–39.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, *108*(3), 624.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, *340*(6128), 95–98.

Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., De Gardelle, V., Castañón, S. H., & Summerfield, C. (2014). Adaptive gain control during human perceptual choice. *Neuron*, *81*(6), 1429–1441.

De Gardelle, V., & Summerfield, C. (2011). Robust averaging during perceptual judgment. *Proceedings of the National Academy of Sciences*, *108*(32), 13341–

13346.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.

Frydman, C., & Camerer, C. F. (2016). The Psychology and Neuroscience of Financial Decision Making. *Trends in Cognitive Sciences*, *20*(9), 661–675.

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, *349*(6245), 273–278.

Gurland, J., & Tripathi, R. C. (1971). A simple approximation for unbiased estimation of the standard deviation. *The American Statistician*, *25*(4), 30–32.

Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, *98*(9), 5359–5362.

Hertwig, R., & Erev, I. (2009). The description–experience gap in risky choice. *Trends in Cognitive Sciences*, *13*(12), 517–523.

Hunt, L. T. (2014). What are the neural origins of choice variability? *Trends in Cognitive Sciences*, *18*(5), 222–224.

Jarvstad, A., Hahn, U., Rushton, S. K., & Warren, P. A. (2013). Perceptuo-motor, cognitive, and description-based decision-making seem equally good.

Proceedings of the National Academy of Sciences, *110*(40), 16271–16276.

Juslin, P., & Olsson, H. (1997). Thurstonian and Brunswikian origins of uncertainty in judgment: a sampling model of confidence in sensory discrimination. *Psychological Review*, *104*(2), 344.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*(4), 237.

Kahneman, D., & Tversky, A. (1982). Variants of uncertainty. *Cognition*, *11*(2), 143–157.

Kepecs, A., & Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *367*(1594), 1322–1337.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annu. Rev. Psychol.*, *55*, 271–304.

Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, *324*(5928), 759–764.

Knill, D. C., Kersten, D., & Yuille, A. (1996). Introduction: A Bayesian formulation of visual perception. *Perception as Bayesian Inference*, 1–21.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, *10*(7), 319–326.

Levitt, S. D., & List, J. A. (2007). What Do Laboratory Experiments Measuring Social Preferences Reveal About the Real World? *The Journal of Economic Perspectives*, *21*(2), 153–174.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432–1438.

Ma, W. J., & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience*, *37*, 205–220.

MacLeod, C. M. (1991). Half a century of research on the Stroop effect: an integrative review. *Psychological Bulletin*, *109*(2), 163.

Macmillan, N. A., & Creelman, C. D. (2004). *Detection theory: A user's guide*. Psychology press.

Mamassian, P., Landy, M. S., & Maloney, L. T. (2002). Bayesian modelling of visual perception. *Probabilistic Models of the Brain: Perception and Neural Function*, 13–36.

Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society*, 122–136.

Renart, A., & Machens, C. K. (2014). Variability in neural activity and behavior. *Current Opinion in Neurobiology*, *25*, 211–220.

Saposnik, G., Redelmeier, D., Ruff, C. C., & Tobler, P. N. (2016). Cognitive biases associated with medical decisions: a systematic review. *BMC Medical Informatics and Decision Making*, *16*(1), 138.

Scott, B. B., Constantinople, C. M., Erlich, J. C., Tank, D. W., & Brody, C. D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *Elife*, *4*, e11308.

Simons, D. J., & Levin, D. T. (1997). Change blindness. *Trends in Cognitive Sciences*, *1*(7), 261–267.

Sonnemans, J., Theo Offerman, T., & others. (2001). Is the quadratic scoring rule behaviorally incentive compatible?

Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, *31*(1), 137–149.

Summerfield, C., & Tsetsos, K. (2015). Do humans make good decisions? *Trends in Cognitive Sciences*, *19*(1), 27–34.

Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience*, *1*(3), 183–184.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, *12*(8), 291–297.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, *113*(11), 3102–3107.

Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124–1131.

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

Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron*, *36*(5), 955–968.

Wyart, V., & Koechlin, E. (2016). Choice variability and suboptimality in uncertain environments. *Current Opinion in Behavioral Sciences*, *11*, 109–115. https://doi.org/10.1016/j.cobeha.2016.07.003

Wyart, V., Nobre, A. C., & Summerfield, C. (2012). Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proceedings of the*

National Academy of Sciences, *109*(9), 3593–3598.

Zakay, D., & Tuvia, R. (1998). Choice latency times as determinants of post-decisional confidence. *Acta Psychologica*, *98*(1), 103–115.

Zylberberg, A., Fetsch, C. R., & Shadlen, M. N. (2016). The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *ELife*, *5*, e17688.

Zylberberg, A., Roelfsema, P. R., & Sigman, M. (2014). Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition*, *27*, 246–253.

5 DISCUSSION

The ultimate function of the brain is to allow us to interact with the world. However, inferring the best course of action from the noisy and uncertain signals acquired from the environment is a challenging computational task. To facilitate the decision for the best course of action, the brain makes use of contextual information. For instance, a rattling sound can simply inform us that our cat is playing with its favourite toy, when we hear it at home, but it can alert us about the presence of a snake in our path if we hear it as we walk in the mountains. Interestingly, in the scientific literature, context dependent decisions can be seen under two very different perspectives.

On the one hand, the optimality of our perceptual inferences hinges on the ability of our brains to make use of context dependent information (Green & Swets, 1974; Sun & Perona, 1998). For instance, by using prior information about the probability of encountering a rattlesnake at home in the example above, or by means of delaying a decision until more information is acquired, as when trying to recognize someone's face in the darkness.

On the other hand, in the economic literature, the lack of consistency in choices, provoked by contextual decision-making, is seen as evidence of irrational behaviour (Luce, 1959; Von Neumann & Morgenstern, 1945). For instance, why would you prefer apples over pears when given only those two options, but pears over apples when also given the option of choosing oranges?

In this thesis, I present three data chapters that explore the influence of context on decision-making in different scenarios all of which require the integration of multiple pieces of information. In the first two data chapters I explore the influence on the choices that are made between two (high valued) options, by having a third option available for choice. In the third data chapter I explore the influence of prior information on the ability to categorize the average value of a stimulus array. A central topic throughout the thesis is the distinction between uncertainty arising at different stages of processing of information. With a particular emphasis given to the consequences to (i) the computational nature of the inference problem, and (ii) our interpretation of the consequences, that uncertainty arising at different stages carries with it.

In the **first data chapter**, we focus on the distractor effect: a change in the sensitivity for choosing the best of two high valued options that is provoked by the value of a third option. Two recent studies have focused on variants of the distractor effect as a signature for irrational choices, and propose two separate neural mechanisms that account for decreases (Louie, Khaw, & Glimcher, 2013) and increases (Chau, Kolling, Hunt, Walton, & Rushworth, 2014) in the sensitivity of the choices. Using computer simulations, we provide a systematic exploration of the cases in which a distractor effect can be recovered under a variety of simple models. One crucial result of this exploration is the fact that a distractor effect can be recovered even under very simple assumptions about the sources of variability in choices, such as Gaussian noise in the representation of the value of the options. Another important result is that, under a variety of models, the direction of the distractor effect (whether it is an increase or a decrease in sensitivity) across an experiment, is dependent on the design of the choice set used by the experimenters. We thus call for a reconsideration of the validity of using distractor effects as a signature for irrational behaviour or for arbitrating between competing mechanisms for decision-making.

In the **second data chapter**, we continue the investigation of decision-making in the context of three options, but from an experimental point of view. Most accounts of multi-alternative decision-making rely on mechanisms that affect how the value of the options is represented in the brain but pay little attention to how that value is constructed. We devised a three-option sequential integration paradigm, whereby three simultaneous streams of information provide the decision-maker with the evidence about the value of each option (one per stream), by sequentially presenting the value on different attributes for all options. Based on previous work, we hypothesized two mechanisms that would dictate the construction of value by means of prioritizing information processing in two parallel frames of reference: (i) at a local level, by means of selectively promoting the processing of pieces of evidence that are most informative for the choice out of all simultaneously available ones (Tsetsos et al., 2016; Tsetsos, Chater, & Usher, 2012); and (ii) at a global level, by means of promoting the processing of ranges of values that are most informative for the choice (Li, Herce Castañón, Solomon, Vandormael, & Summerfield, 2017).

We found behavioural, computational and physiological evidence that showing the influence that these two mechanisms have during the construction of value. And how a distractor option can affect decisions by means of truly distracting the process of information prioritization. Finally, we were able to show via computer simulations, how these two mechanisms have protective effects against noise arising during the process of integrating information or after. Reinforcing the view that seemingly irrational mechanisms can actually optimize accuracy by exploiting statistical regularities in the environment in order to boost the signals prior to the occurrence of noise (Li et al., 2017; Tsetsos et al., 2016).

In the **third data chapter**, we explore the effect of prior information on choices under two distinct sources of uncertainty. The aim of this chapter was to provide a novel hypothesis about why people's beliefs comply with optimality so well in some scenarios but depart so drastically from optimality in others. For instance, in perceptual and sensorimotor tasks people make optimal use of prior probabilities (Ernst & Banks, 2002; Körding & Wolpert, 2006; Saunders & Knill, 2001), but in more cognitive tasks, people tend to systematically neglect the base rate (prior probability) of events (Tversky & Kahneman, 1975). We hypothesized that the difference in optimality for these different domains was driven by the metacognitive sensitivity people have about the sources of noise that typically govern choices in the two domains: in the perceptual domain noise typically arises from poor representation of the sensory signals, whereas in the cognitive domain noise typically arises from information overload or imperfect integration of evidence.

Across six behavioural experiments, we manipulated simultaneously and orthogonally noise arising at these two different stages: (i) lowering the contrast of an array of eight stimuli, to achieve poorer representations of each piece of evidence, and (ii) increasing the variability in the signals provided by the eight pieces of evidence, to increase the difficulty for averaging their signals. We were able to show that an optimal use of prior information in choices dominated by the first source of noise, but a suboptimal neglect of prior information for decisions dominated by the second source of noise. In line with our hypothesis, we were able to show that sub-optimality in participants' choices was tightly linked to their inability to recognize the increased level of noise, an effect that we were able to account with a computational model that neglects integration noise but is otherwise optimal.

Across the thesis, a general view was favoured of the brain as an optimal inference machine. We were able to distinguish the concepts of rationality and optimality, emphasizing that seemingly irrational choices can actually be derived from mechanisms that are in place to optimize inferences in the presence of distinct sources of noise. Interestingly, a paradox arises from the findings of the second and third data chapters. The mechanisms proposed, in the second data chapter, seem to make use of statistical properties of the environment that emerge only when noise occurs during the process of integrating different pieces of evidence or after. However, in the third data chapter, we show that people may actually be blind to these "later" sources of noise. How can people, then, be blind to their noise but still implement mechanisms that allow them to deal with this noise? The short answer is that we do not know yet.

One important difference in the findings that might shed light on this question, is that the types of noise that affect the decisions in the two tasks have quite different sources. On the one hand, when making decisions between three options in the information integration paradigm (data chapter two), one of the sources of noise is almost imposed by nature of the task, there is more information available for the decision-makers that they can actually attend to. This source of information loss is distinct from noise that will arise as the distinct pieces of evidence for each option get integrated in time. The prioritization mechanisms we described allow participants to minimize the loss of information, which conveys protective properties against noise occurring during integration. However, we did not manipulate the noise occurring at these later stages to see if we could provoke a corresponding change in the protective mechanisms. Thus, the protective mechanisms described in the second data chapter, could have developed without concomitant awareness of the costs that they protect choices against.

On the other hand, in the third data chapter we manipulated noise (across different trials) occurring precisely at the stage where the several pieces of evidence get integrated. This manipulation exposed: (i) the lack of a mechanism at place to deal with these variable levels of noise, and (ii) a lack of awareness of the costs themselves.

So, what is special about noise at the integration stage and why are we blind to it? Again, we are able provide nothing but speculations about it. In the studies were humans are shown to be near-optimal (Ernst & Banks, 2002; Körding & Wolpert, 2004; Saunders & Knill, 2001), the crucial result is that *on average* the judgments from participants (about a combination of signals) are almost exactly where the *average* expected judgment under optimality should be if they weighted each piece of evidence by their reliability. However, our experimental paradigm allowed us to ask whether people are aware of (and correct for) the variability around that near-optimal average judgment.

Our brains need remarkable mechanisms at work in order to generate the smooth interaction we have with our world. We are often fooled by this seemingly smooth interaction: for instance, our visual experience of the world seems rich, detailed and colourful, when in fact we have very scarce and corrupted sensory evidence to construct our experience (Cohen, Dennett, & Kanwisher, 2016). So, perhaps, our brains usually carry out such a great job in combining and interpreting sensory evidence, that in most situations we are better-off trusting them blindly. Imagine being aware of all the conflicting information that drives our choices, we would probably vacillate interminably before committing to any action.

References

Chau, B. K., Kolling, N., Hunt, L. T., Walton, M. E., & Rushworth, M. F. (2014). A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nature Neuroscience*, *17*(3), 463–470.

Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the bandwidth of perceptual experience? *Trends in Cognitive Sciences*, *20*(5), 324–335.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–433.

Green, D. M., & Swets, J. A. (1974). Signal detection theory and psychophysics (Rev. ed.). *Huntington, NY: RF Krieger*.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, *427*(6971), 244–247.

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, *10*(7), 319–326.

Li, V. L., Herce Castañón, S., Solomon, J. A., Vandormael, H., & Summerfield, C. (2017). Robust Averaging Protects Decisions from Noise in Neural

Computations. *BioRxiv*, 147744.

Louie, K., Khaw, M. W., & Glimcher, P. W. (2013). Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences*, *110*(15), 6139–6144.

Luce, R. D. (1959). *Individual Choice Behavior a Theoretical Analysis*. John Wiley and sons.

Saunders, J. A., & Knill, D. C. (2001). Perception of 3D surface orientation from skew symmetry. *Vision Research*.

Sun, J., & Perona, P. (1998). Where is the sun? *Nature Neuroscience*, *1*(3), 183–184.

Tsetsos, K., Chater, N., & Usher, M. (2012). Salience driven value integration explains decision biases and preference reversal. *Proceedings of the National Academy of Sciences*, *109*(24), 9659–9664.

Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., & Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. *Proceedings of the National Academy of Sciences*, *113*(11), 3102–3107.

Tversky, A., & Kahneman, D. (1975). Judgment under uncertainty: Heuristics and biases. In *Utility, probability, and human decision making* (pp. 141–162). Springer.

Von Neumann, J., & Morgenstern, O. (1945). *Theory of games and economic behavior*.

Princeton University Press Princeton, NJ.