How combining description and experience influences the decision-making process

Leonardo Weiss Cohen

A dissertation submitted for the degree of

Doctor of Philosophy (PhD)

of

University College London.

Department of Experimental Psychology

September 2017

I, Leonardo Weiss Cohen, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

Decisions are often made using a combination of descriptive and experiential information. However most of the research in decision-making has either focused on the two paradigms separately, or compared them against each other, rarely combining the two sources of information within the same task. In my research, I will explore how descriptions and experience are integrated into the decision-making process when the two are available concurrently, and how each one influences decisions. I start by showing that descriptions are heavily discounted, with preference given to experiential information, which is easier and more natural to process cognitively. I then explore three moderators of the impact that descriptions have on decisions-from-experience. First, when descriptions are considered implausible, their influence on decisions is reduced. Second, when descriptions are too complex, they become too difficult to decipher, thus reducing their influence. And third, when individuals have more prior experience with a situation, the impact of descriptions is also reduced. Empirical results are supported by cognitive models of how individuals integrate their experience with descriptions, with different weights given to each source of information. Experience was the dominant source of information, but descriptions were taken into consideration, albeit at a discounted level, even after many trials. Models that included representations of the descriptive information fitted the human data more accurately than models that did not. This research has implications on the creation of effective warnings, which can be considered as descriptions which are added to our decision-making processes. More effective warnings can be created by making them plausible, of low complexity, and presented early before experience has been accumulated.

Acknowledgements

I would like to thank my supervisors, Nigel Harvey and Maarten Speekenbrink, for the academic support and very helpful direction, guidance, insight, and ideas.

The inspiration for my research materialised from attending the amazing CoDeS MSc programme at UCL, and for that I thank the wonderful academic and support staff at UCL Experimental Psychology, the brilliant people at LJDM, and my fellow student cohorts: “I can’t get no satisfiction,” CoDeS 2012-13.

I could not have done this without my colleague Manos who introduced me to computational modelling in R, writing manuscripts in LaTeX, read all my drafts, and showed me the ins-and-outs of a PhD. Thanks to my friend Phil for his invaluable comments on an earlier draft of my thesis.

I would not have enjoyed my time at UCL as much without L.L.U.M. and their *llum*, gràcies.

I am and will be eternally grateful to my parents for giving me love, wisdom, education, manners, and support (despite the distance) over the last four decades, «friendship forever.»

Dedico esse trabalho ao meu irmão, Eduardo, o mais sábio de todos.

Published articles

Experiments 1, 2, and 3 also appear in Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2016). Incorporating conflicting descriptions into decisions from experience. *Organizational Behavior and Human Decision Processes*, 135, 55-69. doi: 10.1016/j.obhdp.2016.05.005

Experiments 4, 5, and 6 also appear in Weiss-Cohen, L., Konstantinidis, E., Speekenbrink, M., & Harvey, N. (2017). Task complexity mediates the influence of descriptions in decisions from experience. *Cognition*, 170, 209-227. doi: 10.1016/j.cognition.2017.10.005

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Decisions in everyday life are often made using a combination of descriptions and experience. For example, doctors frequently rely on readings of published literature and research, which can be considered a form of description, and combine it with their own clinical experience, when prescribing drugs or assessing the risk of a medical procedure. Consumers may base their buying decisions on a combination of descriptive reviews and their personal experiences of similar items bought in the past. Drivers pass road signs warning them of traffic queues on a familiar stretch of road in which they have extensive previous direct experience about traffic intensity. The ongoing proliferation of warning signs and labels can be considered as descriptive information that is added to an individual's own experience, typically reminding them of high-loss small-frequency risks that are rarely experienced in person. Passengers frequently run at stations in order to catch their trains, and the overwhelming majority never directly experiences any accidents. But warnings signs are common, reminding individuals that running can be dangerous and cause harm. Road traffic signs warn us of queues, animals, flood or ice in locations where these are rarely experienced. Despite the ubiquitous presence of both sources of information concurrently in daily life, the vast majority of decision-making research has exposed participants either to "decisions from description" or "decisions from experience" separately, very rarely exploring how the two are combined when available simultaneously in the same task.

1.1 Decisions from description vs. experience

Decisions from description (DfD) are those in which information about the choices available are described to participants before any selections are made. In the case of decisions between risky monetary prospects, which is the focus of this dissertation, descriptions will expose participants to the values and frequencies of potential outcomes from each option available. These descriptions are typically provided in writing, for example, the following risky gamble from Kahneman and Tversky (1979, p. 264):

> Which of the following would you prefer?
>
> A: 50% chance to win 1,000, 50% chance to win nothing;
>
> B: 450 for sure.

Although other graphical formats have been used, such as displaying probabilities using pie charts (for an example, see Figure 1 in Ludvig & Spetch, 2011), textual representations such as the one above are undeniably the most common. In the case of probabilistic outcomes, such as the example gamble presented above for option A, their descriptions can be considered abstract, idealised, symbolic, and absolute representations of the underlying noisy stochastic process that is used to generate the actual outcomes. In the most common DfD tasks, participants must inherently inform themselves about the options available to them based on the descriptions, which is all the information they have available, and must do so before they act. They must read, interpret, decode, assess and compare the descriptions, and use this information to decide which option they believe to be most attractive, and act accordingly. Any information provided after a choice is made, for example, in the form of feedback about the outcome of their selection, is received *post-hoc*, too late to be used to inform the decision in such DfD tasks.

However in everyday life individuals are not generally presented with such explicitly perfect and unambiguous descriptions and instead make decisions based on their own direct personal experiences in noisy, inexact and information-deficient environments. In decisions from experience (DfE), individuals are not commonly provided with any information before choices are made and, instead, are required to learn about the potential outcomes from each option via feedback provided after each selection is made, over time. For example, Bechara, Damasio, Damasio, and Lee (1999) used four decks of customised playing cards, turned face down, and did not provide any information about the cards to participants, in their now famous Iowa Gambling Task. Although participants did know that they could choose between four alternative options, and not two or six, they did not have any knowledge about the composition of cards within each deck. As participants selected cards from each deck, and turned them over to reveal their faces, they would learn about the set of available outcomes contained in that deck through the feedback written on the front of the cards (e.g., "You have won 100 dollars", from Bechara et al., 1999, p. 5474). Figure 1.1 shows an example of a computerised DfE paradigm used in Experiment 1 in the next chapter of this dissertation, in which two alternatives were available from which participants could choose.

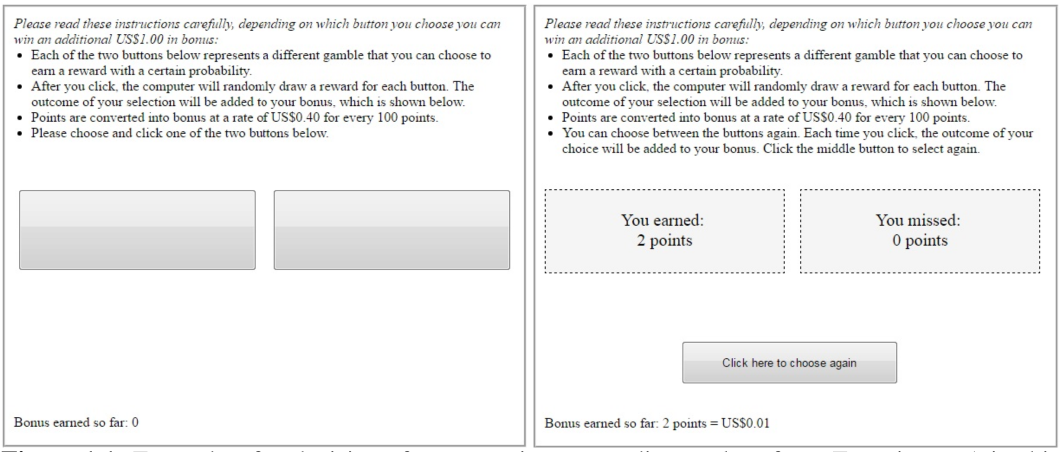

Figure 1.1: Example of a decisions-from-experience paradigm, taken from Experiment 1 in this dissertation. On the left, the screenshot shows the beginning of the task, and also the beginning of every trial, with no information provided to participants about the two alternatives available. On the right, the screenshot shows the feedback provided after the left-hand side button was selected, and participants earned two points. The foregone, or missed outcome, is also shown on the right button.

Because no information is available about the set of potential outcomes and their frequencies at the beginning of DfE tasks, the first choice tends to be an uninformed random selection. Each selection provides additional information in the form of feedback, which can then be used to inform future selections. Each time a participant chooses an option and observes the feedback in a DfE task they gather further information about the processes underlying the generation of outcomes. Individuals therefore learn as they act, sequentially. The amount that is learned after each selection is dependent on certain characteristics of the task.

DfE tasks can be segmented according to the amount of information that is presented after each selection made, either as full-feedback or partial-feedback. Full-feedback tasks provide feedback about all the available alternatives after each selection, the selected option and all the foregone (unselected) options. Partial-feedback tasks provide feedback about the selected option only, and no new information is provided about the foregone options. The experiment with cards by Bechara et al. (1999) mentioned above was a partial-feedback paradigm, as only the selected card was revealed, while the one in figure 1.1 was a full-feedback one, as both outcomes were revealed, after each trial. While it could be said that individuals should learn faster in full-feedback paradigms (Yechiam & Busemeyer, 2006), because more information is provided after each trial, it has been suggested that this additional information could be distracting, perversely slowing learning down (Grosskopf, Erev, & Yechiam, 2006; Yechiam, Druyan, & Ert, 2008; Yechiam & Rakow, 2012).

Apart from the format in which information is presented and learned, DfD and DfE paradigms also typically differ in a key characteristic of the task itself: the number of times individuals are asked to make choices, and the number of outcomes they receive from their choices (Camilleri & Newell, 2013a; Jessup, Bishara, & Busemeyer, 2008). Typically, in DfD tasks, because all information is assessed before any selections are made, participants make a single choice, and receive a single outcome from their choice, for each set of options. In contrast, because information is learned over time in DfE tasks, and the alternatives continuously reassessed, they typically involve multiple choices, and multiple outcomes are experienced over time, for each set of options. The crucial difference between DfD and DfE is therefore the format in which information is presented to participants, and consequently how participants gather information about the choice environment, suddenly via descriptions, or gradually via feedback (Jessup et al., 2008).

For the vast majority of the history of decision-making research, these two paradigms have been explored separately, each in their own individual domain, with dissimilar choice environments adapted to their specific branch of research. Much of the research on DfD has dealt with factors that influence people's risky decision making, by asking individuals to choose between risky gambles based on their descriptions. DfD has been the domain of prospect theory, heuristics, biases, preferences, and axiomatic utility theories of decision making. In fact, the vast majority of human decision-making research to date has been based around descriptive paradigms (Camilleri & Newell, 2009; Erev, Glozman, & Hertwig, 2008; Fantino & Navarro, 2012; Rakow, Demes, & Newell, 2008; Weber, Shafir, & Blais, 2004). Separately, DfE paradigms have been mostly associated with research on learning, exploration, exploitation, and cognitive modelling (e.g., Cohen, McClure, & Yu, 2007; Erev & Haruvy, 2016; Yechiam & Busemeyer, 2005). Recent research comparing indistinguishable sets of choices, although presented either in DfD and DfE, has uncovered unexpected behavioural differences between the two informational presentation approaches.

1.2 The Description-Experience gap

Before Barron and Erev (2003), DfD and DfE had mostly been studied in isolation, with little overlap. Barron and Erev were the first to empirically and systematically compare the DfD and DfE paradigms and study any differences in behaviour, aiming to test if phenomena that had been well-established and extensively studied in DfD tasks would replicate using comparable DfE tasks. Their approach was to present different participants with the exact same set of risky alternatives, either using DfD or DfE paradigms, manipulating whether information was presented via descriptions or via experience. Barron and Erev selected widely-used descriptive paradigms that led to well-established phenomena in the decision-making research: the overweighting of rare events in risky choices; the reflection effect, which is higher risk seeking in the loss than in the gain domain; and the common ratio effect, a preference towards the riskier choice when all outcomes are multiplied by a common ratio (Kahneman & Tversky, 1979). They found that using DfE paradigms, instead of the original DfD approach, these phenomena were reversed: participants underweighted rare events, were more risk seeking in the gain domain, and preferred the safer option after outcomes were multiplied by a common ratio.

One of the limitations of the research by Barron and Erev (2003) was that they used a traditional DfE paradigm, which involves multiple choices and multiple outcomes, over 200 trials, to attempt to replicate phenomena that had been previously established using single-choice, single-outcome DfD paradigms. To address this difference, Hertwig, Barron, Weber, and Erev (2004) and Weber et al. (2004) introduced the sampling paradigm for DfE. In sampling paradigms, individuals can sample freely, without financial consequence, between the available alternatives, learning via the feedback received. These are not choices, but instead only observations. Once participants are satisfied that they have sampled and learned enough about the choices, they can make one single consequential decision, therefore equating the single-choice paradigm more commonly associated with DfD tasks. The earlier findings from Barron and Erev (2003), in particular the underweighting of risky choices in DfE, remained with the sampling paradigm used in Hertwig et al. (2004). However, given that individuals rarely make such inconsequential sampling decisions in real life, and the differences between financially consequent and inconsequent behaviour (Camerer & Hogarth, 1999), differences in number of choices and outcomes remain an ongoing issue when comparing DfD with DfE (Camilleri & Newell, 2011b, 2013a, see also Section 1.2.2 below).

This difference in behaviour between decisions from descriptions and decisions from experience was named the "description-experience gap" by Hertwig and Erev (2009). A substantial body of research has since been dedicated to studying the gap, by presenting different participants with the same choice scenarios, with either descriptions alone or experience alone, and comparing the behavioural results (see Camilleri & Newell, 2013b; Rakow & Newell, 2010, for recent reviews). Despite strong support for the gap, some studies still failed to find any behavioural differences between decisions from description and decisions from experience (e.g., Camilleri & Newell, 2013a; Fox & Hadar, 2006), raising new issues to be explored regarding the mechanisms that contribute towards the appearance of gaps. Initially, it was thought that choice differences were due to informational discrepancies resulting from sampling biases.

1.2.1 Controlling for sampling biases

Sampling biases are the divergence of the outcomes of randomly generated observations during a DfE task, in comparison to their idealised DfD description equivalent. Randomly generated observations are likely to drift away from their associated descriptions unless the number of outcomes observed is substantially large. According to an analysis by Hertwig et al. (2004), 78% of their participants in a DfE task experienced the rare event less frequently than the ideal prescribed by descriptions. Therefore the outcomes that are actually experienced by participants tend to diverge from their perfectly stated verbal descriptions. Sampling biases can strongly influence the statistical information that participants perceive from their observations in DfE, given that observed samples can be biased relative to the population described in DfD. This led many researchers to argue that any differences in behaviour observed between DfD and DfE paradigms were driven by different information being transmitted to participants. If information was different, then different behaviour would be expected.

To exclude the influence of differences in information transmitted in DfD and DfE, researchers have manipulated both the descriptions and the experience, to more closely match the two sources. Rakow et al. (2008) and Hau, Pleskac, and Hertwig (2010) changed the descriptions in DfD to more accurately reflect the average frequency of outcomes actually observed earlier by participants in DfE. Another approach is pseudo-randomising the outcomes observed in DfE, to ensure that participants experience a representative set of outcomes as described in DfD (Ungemach, Chater, & Stewart, 2009). For example, in multi-choice DfE paradigms with 100 trials, and a rare event that is supposed to occur 10% of the time, researchers can ensure that 10 rare events will always be in the stream of outcomes generated. Pseudo-randomisation by blocks can also be employed, splitting the outcomes into blocks, for example 10 blocks of 10 trials each, and ensuring the presence of 1 rare event within each such block. The outcomes within each block are then shuffled, so the actual order in which the rare event appears within each 10 trials is random, but matching the desired distribution (Camilleri & Newell, 2011a). This approach spreads out the appearance of rare events throughout the task.

To ensure that descriptions and experience are as similar as possible, I will employ the pseudo-randomisation approach for the experienced feedback in the experiments throughout this dissertation.[1] While in some studies the gap was eliminated after controlling for sampling differences (Camilleri & Newell, 2011a; Rakow et al., 2008), in other studies behavioural gaps remained even under carefully controlled sampling processes that eliminate informational differences (Camilleri & Newell, 2011b; Hau et al., 2010; Ungemach et al., 2009), which I expect to replicate here. Given that the gap remains, there should be other underlying mechanisms to explain it beyond sampling biases. One of the remaining open discrepancies revolves around differences in the nature of the paradigms.

1.2.2 Single-choice and multi-choice paradigms

Camilleri and Newell (2011b, 2013a) suggested that the gap could also be a result of differences in paradigms between the single-choice and multi-choice environments of DfD and DfE, respectively. They claim that individuals being asked to make a single choice in DfD would choose differently if they were allowed to make multiple choices, as is typical in DfE. This is because in multiple-choice situations results can be more compensatory: losses from one trial can be compensated by gains in another trial (Yechiam, Barron, & Erev, 2005). In single-choice tasks this is not possible, which might explain why individuals overweight rare events in DfD. This concept is also intrinsically connected to sampling biases: the more choices an individual has, the less volatile the outcomes they observe will be, over time. Such differences in approach and consideration between single-choice and multi-

[1] One could argue that pseudo-randomised sequences are unrepresentative of true randomisation, but research has shown that individuals tend to believe that the former appears more random than the latter (Ayton & Fischer, 2004; Bar-Hillel & Wagenaar, 1991; Falk & Konold, 1997; Lopes & Oden, 1987; Peterson & Beach, 1967).

choice decisions have been shown empirically in DfE by Wulff, Hills, and Hertwig (2015a) and in DfD by Joag, Mowen, and Gentry (1990), with single-choice decisions linked to short-run individual non-integrative decisions, and multi-choice as long-run compensatory decisions. The overweighting and underweighting of rare events may not be necessarily linked to the informational source, descriptions or experiences, but instead to the different nature of the paradigms commonly used in each line of research.

DfD tasks are typically single-choice, single-outcome, while DfE tasks are typically multi-choice, multi-outcome, with traditional trial runs of 100 choices. These are the most commonly used approaches, but there have been other, albeit limited, experimental attempts to equate the two paradigms in terms of choice repetition. DfD experiments can be made multi-choice by asking participants to select multiple times form each set of options, without feedback, with a single combined outcome at the end (Jessup et al., 2008). As the authors observed, this can create some very boring tasks, and it is not surprising that it has not been more extensively explored. DfD experiments can also be single-choice, multi-outcome, by telling participants that their single choice will be repeated equally a fixed number of times, providing participants with multiple outcomes (e.g., Camilleri & Newell, 2013a, "Choose 1 option to play from 100 times"). Alternatively, participants can be asked to allocate a certain percentage of their desired selections into each option available, in effect pre-determining the distribution of their selections (e.g., Barron, Leider, & Stack, 2008, Web Appendix A, Experiment 5: "You must allocate 100 choices between the two lotteries below."). In none of these DfD tasks can feedback be used for learning, as it is only provided after all the choices have been made, maintaining the difference that only descriptions, not instances of feedback, are used for learning in DfD.

The sampling paradigm introduced by Hertwig et al. (2004) and Weber et al. (2004) attempts to remove some of the experimental differences between DfE and DfD by making the former into single-choice, single-outcome decisions, closer to the latter (see also Section 1.2 para. 2). However this approach still involves multiple selections by participants from the choices during sampling, which might bias the results of the final choice. While these sampling selections are not financially consequential, and should not be considered actual choices, but instead only observations, the actual cognitive mechanism behind them is still open for discussion. Learning in sampling DfE tasks occurs differently from that in multiple-choice DfE tasks. Not only is the amount of feedback smaller in sampling DfE tasks, with participants relying on relatively small sample sizes (Hertwig & Pleskac, 2010), but the search patterns under sampling are also different from those in multiple-choice tasks (Rakow et al., 2008), which could influence decisions (Hills & Hertwig, 2010). One of the reasons for the difference in observations between sampling DfE and repeated DfE is that in the former case, each choice is free to participants, while in the latter case, each choice is financially consequential.

The description-experience gap still persisted in research attempting to reduce differences in paradigms. In manipulations of the DfD task, for example, the gap remained when comparing traditional multi-choice DfE tasks against multi-choice descriptions without feedback (Jessup et al., 2008), against single-choice descriptions with multiple repeated outcomes (e.g., Barron & Erev, 2003; Lejarraga & Gonzalez, 2011), or when using multiple allocated outcomes (e.g., Barron et al., 2008; Yechiam, Barron, & Erev, 2005). In manipulations of the DfE task, the gap was also still present when comparing traditional single-choice DfD tasks against single-choice DfE using sampling (e.g., Hertwig et al., 2004; Weber et al., 2004), even after forcing participants to observe larger sample sizes (e.g., Hau, Pleskac, Kiefer, & Hertwig, 2008; Hau et al., 2010; Ungemach et al., 2009). A direct and systematic comparison of the paradigms by Camilleri and Newell (2011b, 2013a), confronting single-choice against multi-choice and single-outcome against multi-outcome, has shown that choice patterns are influenced by differences in numbers of choices and outcomes, but behavioural gaps between DfD and DfE remain even after controlling, as much as possible, for such differences.

To avoid any potential confounding effects arising from differences in number of choices and outcomes, the research contained in this dissertation will focus on multi-choice, multi-outcome paradigms. I will accomplish this by employing repeated-choice DfE tasks, which always provides participants with experience via pseudo-randomised feedback (see Section 1.2.1). In addition to the repeated experiential information, my experiments will revolve around controlled introductions of descriptive information within the same paradigm, thus combining descriptions and experience in the same task. These experiments will provide a solid base from which to explore any persistent influences of descriptive and experiential information on the cognitive processes behind decision-making, after controlling for sampling biases and paradigmatic differences.

1.2.3 Differences in cognitive processing

Beyond sampling biases and differences in the paradigm, the description-experience gap might also be a result of differences in how information is cognitively processed, analysed and compared, with different cognitive algorithms applied according to the format in which the information is presented (Camilleri & Newell, 2013b; Glöckner, Fiedler, Hochman, Ayal, & Hilbig, 2012; Hertwig & Erev, 2009). Different ways of presenting descriptive information within variations of traditional DfD tasks had already been shown to influence behaviour in other domains (Lerner, Small, & Loewenstein, 2004; Tversky & Kahneman, 1981; Weber & Johnson, 2009). For example, individuals make different decisions when they are presented with probabilities as natural frequencies (16 out of 20) as opposed to percentages (80%) or standardised frequencies (80 out of 100), as shown in research by Gigerenzer and Hoffrage (1995). Natural frequencies tend to promote more accurate estimates of likelihood and more normative behaviour, leading to fewer biases. Differences have also been found within variations of typical DfE tasks, with observable changes to behaviour according to the way in which the same experiential information is presented, such as feedback delay (Diehl & Sterman, 1995), feedback frequency (Lam, DeRue, Karam, & Hollenbeck, 2011; Lurie & Swaminathan, 2009), the amount of information provided after each trial (Camilleri & Newell, 2011b; Yechiam & Busemeyer, 2006), the format of the feedback (Atkins, Wood, & Rutgers, 2002), and the nature of the feedback (see Gonzalez, 2005, for a review).

According to Hills and Hertwig (2010), information search patterns can also influence decisions. This is likely due to the order in which information is obtained and integrated (Lejarraga, Hertwig, & Gonzalez, 2012), providing stronger evidence to the importance of directly experiencing information sequentially in DfE tasks.

Cognitive theory and experimental evidence suggest that, further to differences within the same informational formats, descriptive and experiential information are also processed differently from one another (Glöckner et al., 2012). For example, the two paradigms of DfD and DfE differ in terms of their underlying sources of uncertainty. According to Knight (1921), decisions under uncertainty can be split into two categories. If the probabilities are known *a priori*, the choice environment is *risky*; and if the probabilities are not known, and must be empirically observed and estimated *statistically*, the choice environment is *ambiguous*. In our daily life, very few situations can be classified as risky in Knight's framework. Examples include flipping a coin or playing a bandit machine in a casino, where the system is pre-programmed to behave in a certain way. In most of our decisions, the probabilities of obtaining each outcome are not previously known, and choices are commonly made in an ambiguous world. Using Knight's terminology, DfD tasks can be considered as decisions under risk, because the probabilities and their outcomes are provided to participants, and known with relative certainty before a decision is made. Participants therefore make a decision under risk. In comparison, DfE tasks can be considered as decisions under ambiguity. Typically in the beginning of the task participants know very little about the choice environment, making it ambiguous. Individuals tend to dislike ambiguous situations (Camerer & Weber, 1992), and they will tend to seek to reduce this ambiguity by exploring and learning about the environment (Güney & Newell, 2015). Neuroimaging research has shown that decisions under risk and decisions under ambiguity are processed by two separable cognitive mechanisms in the brain (Brand, Labudda, & Markowitsch, 2006; Hsu, Bhatt, Adolphs, Tranel, & Camerer, 2005; Huettel, Stowe, Gordon, Warner, & Platt, 2006), although more recent research has also found some overlapping shared systems (Levy, Snell, Nelson, Rustichini, & Glimcher, 2010). Decisions under risk also tend to engage more rational analytical thinking, compared to decisions under ambiguity that elicit more emotional responses (Bechara, Damasio, Tranel, & Damasio, 1997; Bechara & Damasio, 2005; Dunn, Dalgleish, & Lawrence, 2006). This dichotomous relationship is likely to transfer to experiences being more emotionally charged than descriptions (Lejarraga, 2010; Ludvig & Spetch, 2011; Weber, 2006).

Such differences in judgements could translate into differences in the DfD and DfE research, with different types of paradigms leading into different judgements, even when the underlying statistical information is the same. Despite the observed underweighting of rare events in DfE, when prompted for their probability judgements after experiential tasks, participants consistently overestimate the probabilities of rare events (Barron & Yechiam, 2009; Camilleri & Newell, 2011b; Madan, Ludvig, & Spetch, 2014). Therefore in DfE tasks participants appear to overestimate the probabilities of rare events but behave as if they underweight them. This 'overestimation-underweighting paradox' still has not been fully explored (see also Liang, Konstantinidis, Szollosi, Donkin, & Newell, 2017). Hau et al. (2010) suggested that when decisions are made through experience, behaviour will be influenced by different statistical structures within which each individual's mind operates to arrive at their own probability judgements and decisions.

Further theoretical support for the distinction in cognitive processing between DfD and DfE comes from Gigerenzer and Hoffrage (1995), who suggested that sequentially acquired information, such as naturally presented frequencies experienced over time, are easier to track, process and interpret than written percentages and probabilities. Weber et al. (2004) proposed that this is because the former are more closely related to experience-based representations of events, and that brains can more naturally process frequencies than probabilities. Individuals might even be able to learn from experience incidentally, automatically encoding frequencies with minimal effort and attention (Hasher & Zacks, 1984). In comparison, in DfD tasks, decisions can be explained by individuals mentally simulating outcomes from descriptions to arrive at expected values, as shown in a recent computational modelling study (Erev, Ert, Plonsky, Cohen, & Cohen, 2017). Other prominent decision-making theories are built upon the idea that individuals mentally sample information over time until a decision is reached, such as Decision Field Theory (Busemeyer & Townsend, 1993), Decision by Sampling (Stewart, Chater, & Brown, 2006) and the Query Theory of Value Construction (Johnson, Häubl, & Keinan, 2007). These mental simulations push DfD tasks into becoming relatively time consuming and effortful thinking processes. Similarly, it has been suggested that DfD tasks direct individuals to think about the future, as they read about the outcomes and visualise them yet to happen, sampling and mentally simulating outcomes, while in DfE tasks they reflect about the past, thinking about the feedback experienced previously (Ludvig & Spetch, 2011; Wulff, Hills, & Hertwig, 2015b).

Overall, these theories posit that experiences are naturally easier to process than descriptions: descriptions appear to be more cognitively demanding and require more effort to decipher, whereas humans (and all other animals) are more naturally adapted to encode and process experiences. When dealing with descriptions, individuals might engage in explicit and more complex computational processes, mentally simulating outcomes and calculating the expected value of each option; conversely, personal experiences use simpler more instinctive, emotional, implicit, and less demanding integration processes. Research has shown that descriptions can be overwhelmed by experience (Jessup et al., 2008), and decision makers seem to prefer experiential information (Lejarraga, 2010). Concordantly, experience is easier to process cognitively (Glöckner et al., 2012), with personal experiences evoking strong emotional and visceral reactions, vis-à-vis statistical descriptions, which lead to more detached analytical considerations (Weber, 2006).

1.3 Combining description and experience

While there has been extensive research comparing behaviour in DfD and DfE tasks following the initial work by Barron and Erev (2003), there have been very few attempts at combining the two in the same paradigm, and providing participants with both descriptions and experience simultaneously. This is an important area of expansion for the field of decision-making research, given how often we must combine the two sources of information to make decisions in day-to-day life. Instead, the majority of the research in the description-experience gap provided different participants with the same choice but either as DfD or DfE separately, using a between-subjects design.[2] The extant "description-plus-experience" (D+E) research, combining the two sources of information in the same task, so far has provided contradictory results, and a more thorough investigation is warranted. While in traditional DfD paradigms individuals learn about their choices by reading the descriptions alone, and in traditional DfE paradigm they learn about the outcome of their choices by experience alone, in the new combined D+E paradigms, learning can be both via the descriptions, or the feedback after each trial, or a combination of both. The D+E research investigates how the two types of learning interact and influence decisions. Given that D+E tasks combine both description and experience, there are potentially two theoretical directions from which to approach D+E research.

Initially, D+E studies focused on adding experience, in the form of feedback, to DfD tasks, therefore comparing D+E with DfD. The rationale behind this initial stream of research was that experience can be considered normatively irrelevant in D+E tasks since the decision could be made by relying on descriptions alone, when the two are available simultaneously (Newell & Rakow, 2007). There should not be any behavioural differences between D+E and DfD paradigms, especially if taking into account that descriptions are more complete and absolute source of information than noisy and slower experiential feedback: participants wanting to maximise their rewards should base their decisions on descriptions in both D+E and DfD tasks, and ignore experience. This strategy would allow individuals to identify the maximising choice from the first trial, and exploit their preferred alternative repeatedly throughout the task, avoiding the effort of learning via experience, which takes time and detracts from the exploitation of rewards. However, if cognitive effort is in-

[2]Only a handful of studies used within-subjects designs (Camilleri & Newell, 2009; Kudryavtsev & Pavlodsky, 2012; Ludvig & Spetch, 2011; Madan, Ludvig, & Spetch, 2017; Ungemach et al., 2009; Wulff et al., 2015b), where the same participants performed the DfD and DfE tasks sequentially; all of them also identified behavioural gaps.

deed higher for descriptions, and individuals prefer to rely on more natural personal experiences, as presented in the previous section, then the presence of personal experience should influence behaviour, and distract individuals away from analysing descriptions (Yechiam, Barron, & Erev, 2005). In addition, comparing these two types of task creates a confounding effect with regards to the number of choices and outcomes (Camilleri & Newell, 2011b, 2013a), as the early D+E research relied on multiple-choice DfE compared against single-choice DfD.

The first researchers to combine description and experience were Yechiam, Barron, and Erev (2005), who presented different participants with the same two-alternative risky choices either in single-choice DfD or in multiple-choice D+E. The novel D+E task was an extension of a DfE task with repeated-play (multi-choice, multi-outcomes), which instead of having blank buttons as in a traditional DfE paradigm, had the full description of the underlying distribution of outcomes written within each respective button. The same descriptions present within the buttons in DfE were also used in the DfD condition. The authors found a significant difference in behaviour between D+E and DfD choices. Participants took considerably more risk in the D+E condition, which is aligned to underweighting of rare events, while less risk was taken in DfD, equivalent to overweighting of rare events. Another finding by the authors was that the initial behaviour in the D+E condition matched that observed in the DfD condition. This was expected because in the first trial of the D+E task participants had only the descriptive information from which to form their decisions. In comparison, in a traditional DfE task without descriptions, behaviour tends to be random, as participants have no additional information and will most likely choose randomly from the selected options, in order to start gathering information about the environment. There was however a strong sampling bias effect in their study. Only 4 out of 78 participants experienced the rare event, which was programmed to occur with a 0.5% chance. For the majority of participants, the description provided was not a true representation of the experience, as the rare event was not observed 0.5% of times, as described.

In later research by the same authors, Yechiam, Erev, and Barron (2006) proposed that behaviour that was initially informed by descriptions alone, before any direct personal experience, can be influenced and altered once personal experience is accumulated. This is particularly relevant for low-frequency events that rarely occur, and are rarely experienced. In these cases, the description might tell participants of certain rare events, which is taken into account in the decision making process. But once experience is added, and these rare events are not encountered personally, over time, behaviour might change. The same effect was observed in a time-series analysis of the data in the D+E condition of their 2005 research, with risk taking starting much lower (in-line with DfD results) and increasing over time. However the researchers did not run a comparative DfE condition, experience without descriptions, to compare how risk taking would evolve in such tasks. Yechiam and Busemeyer (2006) found similar results for experiences overcoming the initial behaviour as prescribed by descriptions, suggesting that in D+E paradigms, descriptions are used to determine an initial behavioural tendency that is later overcome by the accumulation of experience. In this later research, however, the authors only compared two different D+E paradigms, with and without foregone payoffs, and did not have a comparable DfD or DfE task to assess the differences between the two sources of information.

To overcome the differences in experimental paradigms relating to the number of choices in earlier studies, Jessup et al. (2008) sought to make the two tasks, D+E and DfD, more comparable, by making both repeated-choice, and compared the traditional repeated D+E paradigm with a repeated DfD without feedback. They set out to analyse what happens when feedback is added to DfD, since they compared DfD, which never has feedback, with D+E which includes descriptions and feedback. In their experiments, both the D+E and the DfD conditions included a full descriptions of the choices and their outcomes, and in both of them participants made 100 selections from the two alternatives. The difference in conditions was that in D+E participants had access to both descriptions and experiential feedback after each selection, while in the repeated DfD condition there was no feedback. This way they ensured that learning in the two conditions was controlled, with learning via descriptions and experience in D+E, but learning by descriptions alone in repeated DfD (since there was no feedback). However the two tasks still differed, inasmuch as the D+E condition provided multiple outcomes, while the repeated DfD condition provided a single outcome at the end. While the two tasks were matched in terms of number of choices made by participants, they were not matched in terms of outcomes received. What the authors found was that feedback overwhelmed descriptive information, and behaviour differed significantly between the two conditions. Similar results were found by Newell and Rakow (2007), with a slightly different approach of asking participants for a play strategy, not individual selections, in a probability-matching task.

Subsequent research compared D+E with DfE, hence completely matching the number of choices and outcomes in both tasks, which substantially reduces paradigmatic differences and allows research to more cleanly concentrate on information gathering and learning. This in turn asks the slightly different question of what happens when descriptions are added to DfE. Rakow and Miler (2009) compared two repeated-play tasks, one with experience only (DfE) and one combining descriptions and experience (D+E). However instead of using objective and idealised descriptions of the underlying processes governing the creation of outcomes, the authors presented participants in D+E with a historical summary of all the previously actually observed outcomes, using natural frequencies, moving away from the more widely used probabilities. This information was updated after each trial. As expected, this history differed considerably from the idealised source of outcomes in the beginning of the task, due to sampling biases. Any such deviations should correct themselves after many trials, with the history more closely approximating the real distribution of outcomes over time. However, the authors used non-stationary payoffs that changed at different points of the task, in effect never allowing the historical descriptions to truly inform participants about the future outcomes, and instead, once the payoffs changed, actually providing misleading unhelpful information about the task. What the authors noticed is that participants performed worst in the D+E conditions, since descriptions provided false information. Using a similar approach of providing participants with histories of past outcomes, this time a graphical representation using a series of coloured balls instead of numbers, Barron and Leider (2010) found similar results with a task designed to test the hot-hand effect or gambler's fallacy (Ayton & Fischer, 2004; Bar-Hillel & Wagenaar, 1991).

Barron et al. (2008) were interested in the influence of experiences accumulated before descriptions had been provided, similar to the idea of early or late warning interventions. In their studies, all participants were presented with descriptions and experiences at some point, and made multiple choices with multiple outcomes. However some participants were provided with descriptions from the first trial, in effect a D+E task throughout (early warning), while others were only given descriptions after they had performed half of the task relying on experiences alone, in effect a DfE task initially which later transformed into a D+E task (late warning). The descriptions warned participants that one of the two choices carried a high loss: a low frequency event, with a 0.1% probability. Knowledge of this event should make that alternative less attractive, and therefore chosen less frequently. What the authors showed is that participants took significantly lower risk when exposed to D+E from the beginning. Participants who started with experience only took considerably more risk, with the appearance of descriptions halfway through reducing their risk taking, albeit not completely down to the same levels as the D+E from the beginning. The studies were limited however as the authors did not have a pure DfE condition for comparison. The authors also excluded from the analysis four participants, out of 62, who did observe the high loss rare outcome. Most participants did not experience the high loss effect, and so the description was not a true representation of the experience: it provided conflicting information since it described a rare event that did not actually occur. In summary, the authors showed that providing individuals with conflicting information influenced behaviour.

The studies by Rakow and Miler (2009) and Barron et al. (2008) both provided participants with descriptions that conflicted with experience, even if deception of participants was not the main purpose of those studies, but was instead an unavoidable artefact of the experimental paradigms used. In both studies, the presence of descriptions influenced behaviour. In contrast, the study used by Lejarraga and Gonzalez (2011) provided participants with congruent information via descriptions and experience, which did not influence behaviour. By using a choice paradigm that did not include extremely rare events, or nonstationary payoffs, the two sources of information were more closely matched, avoiding the sampling biases of the experiments above. They compared two experimental conditions: in their DfE condition, participants had to rely on feedback and experience alone without descriptions; while in their D+E condition, participants had descriptions and feedback throughout. The authors did not observe any significant difference in behaviour and stated that descriptions were apparently neglected when experience was also available. These results further supported the previously proposed idea that individuals prefer experience over descriptions, shown empirically by Lejarraga (2010) and Jessup et al. (2008). Lejarraga and Gonzalez proposed that this lack of observable differences in behaviour in their study was due to descriptions being ignored when experience was also available. They also showed how a well-established computational model used in DfE research (Yechiam & Busemeyer, 2005), which relies on integrating experience alone, and does not include any representation for description, can explain behaviour in both DfE and D+E conditions relatively well.

Because Lejarraga and Gonzalez (2011) provided participants with the same information using both descriptions and experience, they could not rely on normative differences in selection behaviour, such as the one used by Barron et al. (2008) who made one option less attractive with descriptions than it was with experiences. Instead, Lejarraga and Gonzalez (2011) had to rely on observing more subtle differences in behaviour based on the existence of a robust description-experience gap and its theoretical predictions of underweighting and overweighting of rare events to test whether description or experience was influencing participants. Behaviour consistent with underweighting of rare events would be expected from participants relying on experiential information, while overweighting would be associated with descriptive information being used. Because the authors observed behaviour consistent with underweighting of the rare event in both D+E and DfE tasks, they suggested this to be evidence that experience was taken into account, but description was neglected. The limitation to their approach is that behaviour consistent with overweighting or underweighting of rare events might actually be the result of differences in paradigms, such as number of choices and outcomes, and not differences in information being provided by description or by experience (Joag et al., 1990; Wulff et al., 2015a). In Lejarraga and Gonzalez (2011), both conditions included multiple choices and outcomes, excluding any potential influence of sampling biases or paradigmatic differences, so perhaps a shift between underweighting and overweighting behaviour was not to be expected.

I believe that the contradictory results provided by Lejarraga and Gonzalez (2011), who did not observe any influence of descriptions when experience was available, and Barron et al. (2008) and Rakow and Miler (2009), who observed behaviour consistent with participants taking into account both descriptions and experience, could be explained by the nature of information being provided by each source: congruent or conflicting. No novel information was provided by descriptions in Lejarraga and Gonzalez (2011), which consequently did not influence the normative maximising choice. In contrast, in the studies with conflicting descriptions by Rakow and Miler (2009) and Barron et al. (2008), one of the choices was made more or less attractive according to whether participants took into account descriptive information or not.

In summary, the evidence available so far in D+E research supports the theory of cognitive differences in processing descriptions and experiences. Participants prefer experiences, in the form of feedback, which are easier and more natural to process, over descriptive information, which is more cognitively demanding. But descriptions are not necessarily fully ignored, and instead influence behaviour, in particular in the first few trials when only descriptions are available and limited experience has been accumulated. Some additional light is provided by Shlomi (2014) in a judgement task, where participants were asked to judge the frequency of certain outcomes. The author observed an effect of discounting of descriptions, when experience was also available, based on computational models applied to the results. If the discounting of descriptions is extreme, it would be indiscernible from neglect, which might be what was observed by Lejarraga and Gonzalez (2011). In this dissertation, I will expand the research on D+E paradigms by further exploring the relationship between descriptions and experience, when the two are available concurrently. The experiments I will employ will be built upon traditional repeated-choice DfE paradigms with multiple outcomes, and participants will be provided either with experience alone, or with a combination of descriptions and experience (D+E), which they can use to inform their decision. The use of multiple-choice approaches across both comparable paradigms is to avoid any confounding influences on behaviour from comparing single-choice against multiple-choice decisions. Pseudo-randomisation of the experienced feedback will also be employed to reduce informational differences between descriptions and experience. I will base my experiments on the tenet that, due to underlying cognitive differences, individuals prefer experiences over descriptions, which likely leads to descriptions being discounted in favour of experiences.

1.4 Computational modelling

Computational modelling is an important tool in the armamentarium of experimental psychology researchers. In addition to relying on behavioural observations alone, computational models can be used to unpack the underlying cognitive processes behind observed human behaviours, acting as a bridge that connects empirical evidence to formal cognitive theories (Shiffrin, Lee, Kim, & Wagenmakers, 2008). The basic proposition of computational modelling is to generate a set of algorithms that can take input from the environment, and return outputs as behaviour based on those inputs, by approaching human cognition as formalised quantitative processes, which are employed by the brain as a computing machine. A good computational model should be able to agree with the observed data (while being anchored in solid cognitive theory), provide valuable insight and understanding on the processes used by the brain, and facilitate *a priori* predictions and generalisations into new environments by steering research into new avenues of exploration

(Busemeyer & Wang, 2000; Erev et al., 2010, 2017; Shiffrin et al., 2008).

For example, computational modelling efforts of behaviour in DfD and DfE tasks have strengthened the theoretical framework distinction for two separate models of processing for decision-making using descriptive and experiential information. Computational modelling competitions have shown that the winning models, or those that best fit and predict the data, were significantly structurally different between DfD and DfE tasks, with little to no overlap between them (Erev et al., 2010, 2017). Model competitions are possible because of their quantifiable and therefore comparable nature, allowing for the best underlying cognitive theory to surface above the others in an objective way (Ahn, Busemeyer, Wagenmakers, & Stout, 2008). Typically, Prospect Theory types of models tend to explain behaviour in DfD relatively well (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992), while reinforcement-learning types of expectancy-valence updating models explain DfE results better (Ahn et al., 2008; Busemeyer & Stout, 2002). When comparing the same underlying model structures across DfD and DfE tasks, Kudryavtsev and Pavlodsky (2012) have also shown that the best fitting parameters were considerably different between DfD and DfE tasks, with very limited commonalities.

Within DfD research, Prospect Theory (PT: Kahneman & Tversky, 1979) and its later extension, Cumulative Prospect Theory (CPT: Tversky & Kahneman, 1992) have been shown to very effectively explain and predict human decisions, as well as support many frequently observed behavioural biases. These are extensions of earlier models which used simpler Expected Utility Theory (EUT) to try to model human decisions-making using mathematical axioms from economic theory: individuals would choose the option with the highest expected value (EV), which in turn maximises their rewards (e.g., Friedman & Savage, 1948). However, EUT failed to explain many observed biases and anomalies in human decision-making, and led to several axiomatic violations, with new alternatives to EUT being proposed (e.g., Camerer & Ho, 1994; Camerer & Weber, 1992). PT and CPT successfully overcame many of the limitations of other alternatives, leading to widespread applications, by introducing non-linear value- and probability-weighting functions

(Fennema & Wakker, 1997). These improvements allow CPT to be efficiently used to explain behavioural inconsistencies such as overweighting of rare events, loss aversion, diminishing sensitivities, the endowment effect, and the certainty effect, amongst others (for reviews, see Barberis, 2013; Edwards, 1996).

Within DfE research, one of the most successfully employed families of computational approaches to human behaviour is built upon the theoretical cognitive research on reinforcement learning (RL), transformed into mathematical models. According to RL theories, individuals learn behaviour through trial-and-error interactions with the environment (Kaelbling, Littman, & Moore, 1996). This is commonly represented as a perception-action-perception cycle: the current state of the environment is considered; an action is chosen, which will cause the state to change; feedback from the state change is assessed and used to re-consider the environment; the process is repeated (Sutton & Barto, 1998). This explorative characteristic of RL, which ensures that individuals learn from their own experiences, allows RL models to successfully tackle more complex tasks in unchartered territories. A commonly used RL computational model in DfE tasks, in particular extensively applied to the IGT, is the expectancy-valence learning (EVL) model, which specifies that individuals make their selections based on the expected valence of the outcomes, which is learned and updated after each action is taken (Dai, Kerestes, Upton, Busemeyer, & Stout, 2015; Busemeyer & Stout, 2002; Fridberg et al., 2010; Worthy, Pang, & Byrne, 2013; Yechiam, Barron, & Erev, 2005; Yechiam, Busemeyer, Stout, & Bechara, 2005).

The EVL process behind RL theory can be efficiently summarised using relatively simple mathematical models, which are computationally formalised using three separate components (Ahn et al., 2008; Worthy et al., 2013; Yechiam & Busemeyer, 2005, 2008). Firstly, there is the value function, used to evaluate the environment, which transforms the absolute value of any observed feedback into a subjective utility or valence. For example, the value function can simply take the value of the feedback directly, an EUT approach, or take into account different weights given to gains and losses, or reduce the distances between extreme amounts, similar to variations in EUT, such as using PT or CPT. Secondly, there is a learning or update rule, which takes into account the utility of new post-action observations to update the expectancies attached to each alternative. This component usually involves the integration of prediction errors, which update the pre-action expectancies of the environment with the observed post-action feedback, a cognitive theory borrowed from the neurological research on rewards and dopaminergic cells, which locate RL in the brain (Dayan & Niv, 2008; Niv, 2009; Schultz, 2006). And finally, there is a choice rule, which compares the expectancies of the different alternatives to determine a preferred choice, if deterministic, or more typically, the probability the model would give to choosing each alternative, using probabilistic choice rules.

These EVL models are in effect extensions of EUT approaches to cognition, but applied to repeated DfE tasks. In the case of traditional EUT, typically applied to DfD situations, the expected values of the alternatives are calculated by individuals by reading the description of the options and their outcomes, before any decisions are made, and without relying on any feedback. In the case of EVL models, the expected values are calculated by integrating the feedback from the options when they are received sequentially, typically provided after each selection by the individual. One of the crucial difference between EUT and EVL models is that the former relies on probabilities or percentages assessed holistically, which has been shown to lead to many behavioural biases (Gigerenzer & Hoffrage, 1995), while the latter relies on naturalistic observations of sequential outcomes, which can be more easily processed cognitively (Hasher & Zacks, 1984). In particular when dealing with RL types of models, it allows researchers to explain how information is assessed and integrated by individuals, and used in order to inform their actions, over time.

By focusing on how different components of computational models change according to experimental manipulations, researchers can identify the hidden workings of cognitive processes employed during decision making. For example, computational modelling efforts applied to the IGT have allowed researchers to determine how brain damage affects the cognitive decision processing in such a complex task. Analysing variations in the best fit parameters of the computational models against healthy and clinical populations can expose how different parts of the cognitive process are being affected by damage to the brain (Busemeyer & Stout, 2002; Yechiam, Busemeyer, et al., 2005) or addiction (Stout, Rock, Campbell, Busemeyer, & Finn, 2005). For example, changes in the value function would show if certain clinical populations perceive gains or losses differently, changes in the choice rule parameters would determine which group behaves more randomly or more deterministically, while changes in the updating rule parameters can determine if learning is affected (for a review, see Yechiam, Busemeyer, et al., 2005).

Within the limited research corpus on D+E tasks, there has been only one attempt to computationally model the experimental results, by Lejarraga and Gonzalez (2011). The authors showed that a traditional RL model, the EVL from Yechiam and Busemeyer (2005), explains the experimental behaviour they observed relatively well, even when both descriptions and experience were available. However, the RL model used was a traditional experience-only model that did not include any representations of descriptive information, and led the authors to claim that descriptive information is therefore not used by individuals to arrive at decisions in D+E situations. This warrants further exploration, as their approach was limited. Even though the experience-only model they used, without descriptions, fitted behaviour relatively well, the authors did not attempt to introduce descriptions into the model, in order to determine whether a competing model that included representations of both sources of information would outperform the traditional experience-only model used. I believe that their findings arose only because descriptions are heavily discounted when the task is simple and the information is congruent between descriptions and experience.

So far, computational modelling efforts have been separated between models that rely on information from descriptions and models that rely on information from experience, and applied to DfD and DfE tasks, respectively. The new area of D+E research warrants a new family of models, relying on both descriptive and experiential information. By understanding how these two integrate and interact, further insights can be gathered on the psychological processes underlying the influence of these two sources of information in the decision-making process.

In this dissertation I will introduce a new integrative D+E computational model that includes representations of both descriptions and experience concurrently, allocating different weights to each source. The allocation of different weights to description and experience has been suggested before, and is supported by cognitive theory highlighting differences in processing loads between the two sources (see Section 1.2.3). This new D+E model will also allow for the identification of factors that moderate the allocation of the weight between the two different sources, such as the concepts of plausibility and complexity, studied herein. In the same way that previous computational modelling research looked at the changes in parameters such as attention to losses and learning rates according to experimental conditions, the new model I am proposing will allow me to confirm that descriptions are discounted in D+E tasks, and reveal how this discounting is influenced, by observing changes to the new weighting parameter that combines descriptions and experiences. I will base my new D+E model proposal on the EVL family of RL models, which has been successfully employed before in similar tasks. This way I can confirm if the new D+E model outperforms traditional experience-only EVL models in D+E tasks. I expect that traditional models will be able to explain only behaviour observed in DfE tasks and not the D+E tasks, and that alterations to the model will be needed to allow for the representation of descriptive information to be included. By quantifying the relationship between descriptions and experience with a computational model, I will be able to identify in which situations descriptions are more or less discounted in comparison to experience. In these situations, descriptions can provide useful information for participants to perform better in their tasks.

1.5 Exploration and exploitation

According to RL theory, agents want to repeatedly exploit effective actions which they have previously found to return desired outcomes, in order to maximise their rewards; however in order to be able to find such attractive actions, they must first explore the environment by employing behaviour that they have not tried before, thus fostering learning (Sutton & Barto, 1998). March (1991) has provided separate distinct definitions for classifying behaviour into exploration or exploitation, according to their aims: Exploration involves acting with the specific intention of learning new information about the environment, in particular the outcomes of available actions; Exploitation, on the other hand, involves the use of existing knowledge to obtain an immediate desired outcome such as securing positive rewards. While March postulated these dichotomous designations, RL theory is based on the concept that learning is concurrent with actions: individuals learn as they act, and behaviour is typically both explorative and exploitative, with different strengths. Perhaps a more flexible definition would refer to behaviour that mainly aims to gather new information as exploration, and behaviour that mainly seeks maximum rewards based on the knowledge learned so far as exploitation, viewing them as varying along a continuum (Gupta, Smith, & Shalley, 2006). Optimal behaviour should seek to be both informative and rewarding, optimising the outcomes from decisions while simultaneously improving knowledge (Dam & Körding, 2009). This trade-off between exploration and exploitation is an important challenge in RL, and achieving the perfect balance between the two is often a very difficult and complex dilemma to solve (Cohen et al., 2007). Despite the apparent complexity of this dilemma, individuals effortlessly balance exploration and exploitation in everyday life (Daw, O'Doherty, Dayan, Seymour, & Dolan, 2006; Kaelbling et al., 1996).

The exploration-exploitation trade-off exists in DfE tasks by definition because in those tasks informational learning and actions occur in parallel, via feedback, which is a determining characteristic of such tasks. What is observed experimentally in DfE tasks, and also predicted by RL theory and confirmed with cognitive modelling, is that exploration reduces with time, and gives way to exploitation (Barron & Erev, 2003; Biele, Erev, & Ert, 2009; Erev, Ert, & Yechiam, 2008; Mehlhorn et al., 2015; Speekenbrink & Konstantinidis, 2014). When facing a novel situation, agents typically have little or no information about the environment and will prioritise explorative behaviour to learn as much as possible about the available options and their potential outcomes, with their initial behaviour being mostly uninformed, and therefore random. Over time, as they gather information and learn more, reducing their perceived uncertainty of the environment, agents will then shift behaviour towards more exploitative actions to reap the rewards from the information they have captured so far. In static environments, the theoretical ideal strategy, in order to maximise benefits and reduce costs, is to explore until satisfied with the amount of knowledge they have gathered, thus reducing their uncertainty about the environment to acceptable levels, and upon reaching this stopping point, they will uniquely exploit their preferred (typically the most rewarding) alternative, foregoing costly exploration altogether. However in real life payoffs are typically non-stationary, i.e. they change over time. In these scenarios, purely explorative and purely exploitative behaviours are thus better avoided (Sutton & Barto, 1998). Both have their own associated opportunity costs: Exploration diverts behaviour away from the preferred choice towards sub-optimal alternatives; exploitation provides no additional new information about the options not being exploited, which leads to an increase in uncertainty (Lea, McLaren, Dow, & Graft, 2012; Knox, Otto, Stone, & Love, 2012). While the initial explorative behaviour and consequent shift from exploration to exploitation is expected, agents ideally should never settle on pure exploitation, but instead should periodically explore to check if other states have changed. Pure exploitation without exploration in dynamic settings would impede the agent's ability to realise any changes to the payoffs of the unselected actions.

Since DfD tasks do not provide feedback by design, and all information required to reach a decision is provided before choices are made, there is no opportunity for an exploration-exploitation trade-off in descriptive paradigms. When descriptions are available, individuals should optimally choose their preferred action using the descriptive information, and consistently implement this action thereafter, with no additional learning envisaged or even allowed in typical DfD tasks. Therefore, in theory, all behaviour in DfD tasks can be seen as exploitative. However, this does not appear to happen in practice. Even though individuals cannot learn dynamically in DfD tasks, they seem to behave as if they are exploring the environment, by switching their choices away from their optimal ideal actions, despite there being no opportunity for any additional learning, because there is no feedback (Camilleri & Newell, 2011b, 2013a; Jessup et al., 2008; Newell, Koehler, James, Rakow, & van Ravenzwaaij, 2013; Otto, Taylor, & Markman, 2011). This might be a result of the artificiality of DfD tasks, perturbing the human cognitive processes better adapted to handle the more naturalistic DfE type of task. In fact, because in real life few situations are static, perhaps this is why pure exploitation is seldom observed experimentally even in DfE tasks with feedback (Ashby, Konstantinidis, & Yechiam, 2017; Shanks, Tunney, & McCarthy, 2002): individuals might be better cognitively tuned to handle dynamic situations, where explorative behaviour is advantageous (Stanovich, 2003).

Research using D+E tasks should help shed additional light on exploration-exploitation trade-off behaviour. An ideal agent who wanted to maximise their financial gains in a D+E task could analyse the descriptive information when it becomes available, decide which is their preferred alternative, and exploit that option across all trials, thus negating the need for any further exploration between options, avoiding costly deviations from pure exploitation when descriptions are available. Assuming that descriptions are a true representation of the underlying process from which experiences are generated, feedback should not provide any additional insight about the environment. In fact, relying on experiential information could be detrimental by leading individuals into behavioural biases such as gambler's fallacy or hot-hand effect (Ayton & Fischer, 2004; Bar-Hillel & Wagenaar, 1991; Barron & Leider, 2010; Newell & Rakow, 2007). Naturally, the additional information provided by descriptions should help reduce uncertainty about the environment, and it has been shown that lower uncertainty leads to lower exploration (Dayan & Niv, 2008; Otto, Gershman, Markman, & Daw, 2013), thus reducing exploration in D+E tasks in comparison to DfE tasks without descriptions. However, given that purely exploitative behaviour is seldom observed empirically even when it is the most advantageous strategy for participants to employ, I predict that exploration will be lower when descriptions are present, but do not expect exploration to completely disappear in D+E tasks.

1.6 Overview of the studies

The main theme of the current dissertation is to examine how information acquired via description and via experience interact in decisions from description-plus-experience, in which both sources are available concurrently. The experiments were designed to investigate the hypothesis that descriptions are not necessarily neglected by individuals when both description and experience are available, but that instead descriptions are taken into consideration, albeit at discounted levels in relation to experience, depending on the situation. This research is directly linked to research on warning labels and warning signs, which can be considered as descriptions superimposed over personal experiences, in an attempt to influence behaviour.

I will seek to confirm this discounting effect on descriptions via three potential moderators of its amplitude: conflicting information, task complexity, and prior experience. In Chapter 2, I will explore the influence of descriptions on behaviour when they provide congruent or conflicting information in relation to experience. If descriptions agree with experience and do not add any new information to the participant, they are likely to have limited impact on behaviour. Only when adding novel information should descriptions influence behaviour. Informational conflict is also likely to play a role via the plausibility of descriptive information, with high levels of conflict leading to less plausible, and thus less influential, descriptions, and vice-versa. In Chapter 3, I will focus on how task complexity can moderate the influence of descriptions in D+E tasks. In complex situations that are more difficult to learn experientially, descriptions should play a bigger part in the formation of decisions. However, if descriptions are overly complicated and difficult to decipher, they will be less useful, which should happen for very complex tasks that cannot be described concisely. In Chapter 4, I will analyse how prior experience impacts the relationship between descriptions and experience. Descriptions can be expected to have a lower impact on behaviour in situations where individuals have more prior personal prior experience, and thus have established habits which can interfere with subsequent integration of new information. Descriptions which are made available earlier, before experience, are expected to have the strongest impact on behaviour.

Chapter 2

Conflicting descriptions and warning labels

Situations in which descriptions and experience are available simultaneously, while contradicting each other, are remarkably common in day-to-day decision-making. Differences in information between description and experience might be a result of differences in the samples underlying each source. Descriptions are more likely to be based on a larger set of data, used to build the information presented, such as side effects of medication based on large clinical population, or on-line reviews of products based on a large number of user ratings. In comparison, an individual is unlikely to take certain medications often enough to observe the associated side effects with the same frequencies as the wider population, or experience new products often enough to suffer through some of their faults, leading to considerably smaller sample sizes through personal experience. With small samples and rare events, sample variability can be very large. Even with large samples, the representative set behind a description can differ from an individual's particular experience, depending on the source of the description. Glasgow et al. (2006) and Kamal and Peppercorn (2013) discuss the external validity of medical research findings, which are typically used as reference points for decision-making, but are not always applicable to a doctor's more localised clinical experience. This is especially true for doctors who have to deal with patient populations that are not representative of the reference population in the standard description. Rakow, Vincent, Bull, and Harvey (2005) showed how mortality risk assessments based on reference research conducted in the Unites States differed from personal experience of doctors at a selected hospital in the United Kingdom. The increase in off-label prescriptions of medications (Alexander, Gallagher, Mascola, Moloney, & Stafford, 2011; Radley, Finkelstein, & Stafford, 2006), defined as the atypical usage of drugs for treatment of illnesses beyond those officially approved by the regulators, is a likely source of conflicting descriptions and experiences. Other examples can result from the overzealous usage of warning signs which misrepresent risks, for example by describing a risk as likely when in reality it is rarely experienced. Carson and Mannering (2001) mention the overuse of road traffic ice warning signs in locations where ice is rarely observed.

Differences might also be the result of the more static nature of descriptions, while experiences are more dynamic. Published medical research takes a considerable amount of time to be released to doctors, and once published, tends to be extremely static and resistant to changes and updates. Meanwhile, doctors will continue to innovate and their own personal experiences are constantly being updated. The same applies, for example, to consumer research and reviews, where written information is slower to adapt to changes in the environment. In a dynamic world, adaptive short-term experiences, which are continuously being updated, will be kept up to date; while more static long-term descriptions, which take considerably longer to be updated, will quickly become out of date. This difference between experiences and descriptions would naturally lead to the two diverging over time. Experiments have shown that in naturalistic dynamic environments it is detrimental to rely on static far-sighted descriptions, and individuals in fact are better-off with small myopic experiences (Hertwig & Pleskac, 2010; Rakow & Miler, 2009). And in the new latest world of extremely dynamic on-line information sharing, such divergences can also occur in the opposite direction. For instance, when considering customer reviews on web pages, reviews constantly accumulate, affecting the overall mean rating of a product, leading to more dynamic descriptive information. Conversely, experiences might remain static if a person is simply no longer exposed to similar situations in the future.

If such mismatches between description and experience are encountered frequently, understanding how individuals deal with these situations is crucial for ecologically valid research with real life practical implications. Studies using D+E paradigms are closely related to research on the impact and usage of warning labels and messages. Such warnings are ubiquitous in modern daily life. They might be considered as descriptive information, the warning itself, which is added to an individual's own personal past experience of a situation. The presence of warning labels requires individuals to combine information from description (the warning label) and from personal experience in order to arrive at their decisions. A meta-analysis on warning labels' effectiveness covering 48 studies "suggests that the warning labels' impact on behavioural compliance is not as clear as expected" (Argo & Main, 2004, p. 193). One of the proposed mechanisms for the influence of warning labels is that "[i]f the information in a warning contradicts one's existing beliefs, the warning information might be discounted" (Rogers, Lamson, & Rousseau, 2000, p.130). As such, these warnings are often used to inform us about rare events which are very infrequently experienced but that are typically associated with very high losses, sometimes catastrophic. Frequently they highlight events that never happened to the individual personally, and might never have been observed either. Alternatively, they remind individuals of rarely experienced situations, which might have been forgotten and lost their influence on behaviour (Laughery, 2006). By presenting rare events that are not observed directly and frequently enough by the majority of individuals, warning labels typically carry descriptive information that conflicts with direct personal experience.

Research by Barron et al. (2008) revealed the observable behavioural differences when descriptions did not match experience. While the researchers did not specifically create the conflict between descriptions and experience, it resulted from an experimental artefact. Because of the extremely low probability of the rare event occurring (0.1%), only a few participants actually experienced the rare event, with the majority of participants in their research not experiencing the rare event prescribed by descriptions. For those participants that did not experience the rare event, they were in fact presented with conflicting descriptions in the form of warning labels: “Each time you hit the (Left/Right) button there is a 1 in 1000 chance (.001 probability) that you will lose $15” (p. 128). The described outcomes were considerably less attractive than the observed outcomes, for those who did not experience the rare event. Participants therefore selected the option with the rare event less frequently when the warning was present, which made it less attractive, in line with the normative prediction. Similarly, Rakow and Miler (2009) also showed that conflicting descriptions influence behaviour by using a dynamic task. In their research, participants were provided with an updating history of past observable outcomes, which became conflicting descriptions when there was a regime change to the underlying processes generating future outcomes, making the history no longer applicable. These conflicting descriptions also led to deterioration in participants' choices. Despite their novel findings, their research was limited as neither Barron et al. (2008) nor Rakow and Miler (2009) provided congruent descriptions for comparison.

I believe that in simple tasks combining descriptions and experience, which are almost exclusively the only ones used in the description-experience gap research, differences in behaviour should not be expected if the two sources provide the same information. This is because if both sources are the same, they should lead to the same final behaviour, with no observable differences. I believe that certain characteristics of the description-experience gap, such as the shift between underweighting and overweighting of rare events, are effected by the paradigm, not the type and source of information. By adding descriptions to a simple DfE task, the nature of the paradigm remained unchanged, and Lejarraga and Gonzalez (2011) continued to observe an underweighting of the rare event because their task was still a repeated-choice, repeated-outcome typical DfE task. They proposed that if descriptions were taken into consideration by participants, then a shift towards overweighting should have been observed. Alternatively, I propose that if the paradigm remains the same (i.e., repeated-choice, repeated-outcome), then this should not have been expected.

Instead, in order to check if descriptions are considered and integrated into the decision making process, conflicting information should be used. If the two sources of information are in conflict, and each points towards a different normative behaviour, then by observing how participants behave I should be able to confirm which source of information is being used. A shift in normative behaviour can be expected if one of the options is made substantially less attractive by the information provided in the descriptions in comparison to the information provided by experiences, as observed by Barron et al. (2008), although in their experiment the conflict was not an experimental manipulation, but instead a result of sampling biases. By observing the shift between the choices when such experimentally controlled conflicting descriptions are introduced, the strength of each source could be quantified, in comparison to the addition of congruent descriptions. While tasks with conflicting and congruent information have been separately examined in the past, these have never been combined before in the same paradigm.

The first set of experiments in this dissertation expands the research on D+E paradigms by systematically and explicitly controlling for informational conflict between descriptions and experience. Careful experimental control of the information provided by the two sources of information, conflicting or congruent, should allow us to more closely define which of description and experience is used for decision-making when the two are available simultaneously within the same underlying task, combining the seemingly conflicting findings from Barron et al. (2008), who observed an influence of descriptions added to DfE, and Lejarraga and Gonzalez (2011), who reported that descriptions are ignored in DfE tasks. Therefore in order to tease out any differences in behaviour, I will provide participants with explicitly conflicting or congruent information via descriptions. The conflicting descriptions were created to shift the normative preferred choice according to their expected-values. For example, if congruent descriptions point to the risky choice being more attractive than the safe choice, the conflicting description would reverse that preference. I predict that the addition of congruent descriptive information to simple DfE paradigms will not influence behaviour, replicating the findings observed by Lejarraga and Gonzalez (2011); in contrast, the addition of conflicting descriptive information will shift behaviour towards the normative choice predicted by the new information, as observed in Barron et al. (2008).

Furthermore, I propose that the relationship between descriptions and experience will be moderated by the plausibility of the conflicting information, with less plausible descriptive information exerting a weaker influence on behaviour. This influence of plausibility on conflicting descriptions should be monotonically increasing at the centre, but not at extreme levels of informational conflict. In extreme cases when information is highly implausible, it should be more readily discarded from the decision-making process. This could lead to a reduced marginal influence or even a contrasting effect at the extremes, similar to what has been found in research on anchoring (e.g., Chapman & Johnson, 1994), advice seeking (e.g., Yaniv, 2004), goal setting (e.g., Locke, 1982), and psychophysics (e.g., Brown, 1953). To confirm my hypothesis, in addition to dichotomously providing conflicting or congruent information, different levels of conflict will also be experimentally manipulated. This can be done by creating more or less plausible descriptions, in comparison to the experienced feedback. With this manipulation I can verify the boundaries of the influence of conflicting descriptive information. I expected participants to disregard the descriptions more easily in the implausible conflict conditions, therefore reducing the effect of their influence on their behaviour.

2.1 Experiment 1

The first experiment uses a paradigm similar to the one employed in Lejarraga and Gonzalez (2011), however adding a new experimental condition with conflicting descriptions. If, according to my hypothesis, no behavioural differences were found by the addition of congruent descriptions in Lejarraga and Gonzalez (2011) because the two sources of information contained the same underlying information, not because descriptions are ignored, then by adding conflicting descriptions I will be able to confirm that descriptions actually influence behaviour when providing novel information, as in Barron et al. (2008). Furthermore I predict that the shift in behaviour will be in the direction normatively predicted by the conflicting descriptive information.

2.1.1 Method

Design

The first experiment had a 4 $\times$ 2 between-subjects design: four types of information presentation and two levels of probabilities for the risky option. Information was presented in one of the following: the description-only (D) condition; the experience-only (E) condition; the description-experience-same (DES) condition; and the description-experience-conflict (DEC) condition. The two levels of probabilities referred to the risky option: the 80% probability condition provided participants with a reward 80% of the time (and zero otherwise), and the 20% probability condition had a 20% chance of providing the reward. Each participant was presented with only one type of information presentation and only one level of probability.

Participants

172 participants (67 females; age: *M* = 32 years, *SD* = 10 years) were recruited on-line using Amazon's Mechanical Turk service. Participation was restricted to individuals whose location was defined as in the United States. There were 28 participants in the $D_{20\%}E_{80\%}C$ condition, 24 in the $D_{80\%}$ condition, and 20 each in the remaining six conditions ($D_{80\%}E_{20\%}C$, $D_{80\%}E_{80\%}S$, $D_{20\%}E_{20\%}S$, $E_{80\%}$, $E_{20\%}$, and $D_{20\%}$; subscripts indicate the probability levels used in the description *D* and experience *E*, whilst *C* stands for *conflicting* and *S* stands for *same*). No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment (Bonus: *M* = US$ 0.71, *SD* = US$ 0.12).[1]

[1] Bonus amount in Experiment 1 was not influenced by the probability level condition (80%: 0.71; 20%: 0.70; $F(1,164) = 0.587, p = .44$), but it was influenced by the information presentation condition, with a significantly lower overall bonus in the DEC condition (D: 0.69; E: 0.73; DES: 0.74; DEC: 0.67; $F(3,164) = 4.111, p = .008$).

Task

Participants were initially presented with an instructions screen with information about the task. They were told that the task involved choosing between two on-screen buttons, with each button associated with a gamble paying rewards with a certain chance. The idea of conflicting descriptions was introduced as follows: "Because of the way that computers generate random numbers, sometimes the actual frequency that you will experience of winning rewards might not be the same as the one indicated. Ideally, it should be the same, but sometimes it can fluctuate both up and down. It is up to you to assess how attractive each button is based on the actual rewards you get from clicking it. Choosing wisely between the two gambles, in order to maximise your points, will help you increase your bonus."

After reading the instructions, participants were then presented with two buttons side by side on screen: one button provided the participant with the sure outcome of two points 100% of the time, and the other button was a risky gamble which gave participants five points either 20% or 80% of the time, depending on the experimental condition, and zero points otherwise (see Figure 1.1 on page 13). Safe and risky button locations were counterbalanced between participants. Choices were made using the mouse. All of the participants' choices between the two options were financially consequent and accumulated towards their final pay. Points were converted to money at a rate of US$ 0.20/100 points in the 80% condition and US$ 0.40/100 points in the 20% condition.[2] Accumulated amounts in points and US dollars were shown on-screen and updated after each choice was made. Participants completed the task in an average of 7.0 minutes (SD = 3.5 minutes).

In the description-only (D) conditions, each button had a label that provided participants with a description of the underlying distribution of outcomes, as detailed in Table 2.1. Participants were told to choose one button once, and that their selection would be repeated by the computer 100 times, each time drawing from the underlying distribution of the option chosen, to calculate their total bonus (Camilleri

[2]Different exchange rates were used according to condition to ensure that all subjects could earn a similar amount of money for their participation in the task, by keeping the financial expected value of the highest earning gamble equal to US$0.80 over 100 trials, regardless of condition, across all experiments in this chapter.

& Newell, 2013a).

Table 2.1: Experiment 1: Button labels according to experimental condition.

Condition	Safe choice	Risky choice
$E_{80\%}$, $E_{20\%}$	(blank)	(blank)
$D_{80\%}$, $D_{80\%}E_{80\%}S$, and $D_{80\%}E_{20\%}C$	2 points with 100% probability.	5 points with 80% probability; Zero otherwise.
$D_{20\%}$, $D_{20\%}E_{20\%}S$, and $D_{20\%}E_{80\%}C$	2 points with 100% probability.	5 points with 20% probability; Zero otherwise.

The experiential conditions (E, DES and DEC) involved 100 repeated individual choices. After each trial, participants were given full feedback, with both the earned and foregone outcomes displayed in the relevant buttons, and asked to choose again (see Figure 1.1 on page 13). In the E conditions, the two buttons were always blank. In the DEC and DES conditions, the buttons contained descriptive labels, as detailed in Table 2.1. In the DES conditions, the description matched the experience: the outcomes of each choice were drawn from the same distribution as that described in the button. In the DEC conditions, the description for the risky choice showed a probability level opposite to the one used to draw the experiential outcomes after each choice. For example, participants in the $D_{80\%}E_{20\%}C$ condition were shown a risky button with a description (D) that indicated an 80% probability of winning five points. However the actual gains experienced (E) by participants for the risky choice were drawn with a 20% probability distribution. In this condition, the conflicting description made the choice more attractive than it was in reality. The situation was reversed in the $D_{20\%}E_{80\%}C$ condition, with the conflicting description making the risky choice appear less attractive (D=20%, E=80%).

In order to avoid sampling biases, samples were pseudo-randomised in groups of 10 outcomes each, for each participant. Within each 10 outcomes, the samples were yoked to perfectly represent the exact appropriate level of reward events expected in the underlying distribution, either eight or two observations (80% and 20% conditions respectively), in a randomised order (see Section 1.2.1).

2.1.2 Results

The main dependent variable was the proportion of maximisation choices (Max-rate). The Max-rate was calculated as the average proportion of times that participants selected the option with the highest actual expected value (EV) according to experienced feedback, but ignoring descriptions, for each block of 20 trials. In the 80% probability condition that was the risky choice, and in the 20% probability condition, that was the safe choice.[3]

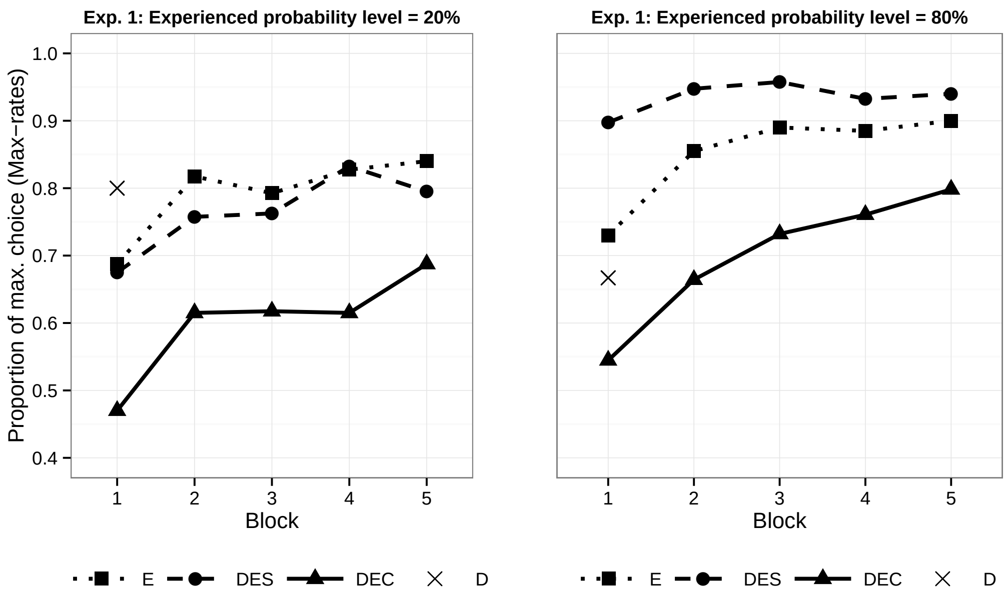

Figure 2.1: Evolution of the maximisation choice rate (Max-rate) for each block of 20 trials for Experiment 1. The left panel shows the results when the risky choice paid a reward of five points 80% of the time (Max choice = risky choice), while in the right panel the same reward was paid 20% of the time (Max choice = safe choice). The lines refer to the different descriptions presented to participants (E: experience only, no description; DES: description and experience same; DEC: description and experience conflicting). The X refers to the description-only condition, which involved a single choice.

A single analysis combined the two experienced probability conditions (80% and 20%), by allowing the maximisation choice in each to change according to the condition (Figure 2.1). The Max-rates in each block were analysed with a linear mixed-effects model using the lme4 package (Bates, Maechler, Bolker, & Walker, 2014), and post-hoc analyses using the lsmeans package (Lenth, 2016), with Tukey

[3]Using Prospect Theory (Tversky & Kahneman, 1992) instead of simple EV did not change the maximisation choices in any of the experiments in this chapter.

adjustments, in R (R Core Team, 2014). The between-subjects factors were the probability level (80% or 20%) and types of information presentation (E, DES or DEC). The D condition was excluded from the quantitative analysis because of the different nature of the paradigm, since it involved only one decision without feedback, which is not directly comparable to the repeated decisions of the other conditions (Mean Max-rates: $D_{80\%}$: 67%; $D_{20\%}$: 80%). The within-subjects conditions were the blocks of 20 choices each. The model also contained a random intercept and a random slope across blocks, for each participant. This approach was used to capture the nested structure of the data. The random intercepts and slopes account for differences in individual levels of risky choices, and their changes over time.

The main effect of information presentation condition on Max-rates was significant. Max-rates were lowest in the DEC condition (DEC: 65.0%; DES: 85.0%; E: 82.2%; $\chi^2(2) = 29.51$, $p < .001$). Descriptions influenced behaviour, with the conflicting DEC information impairing performance and reducing the Max-rates. The main effect of probability level was significant, with lower Max-rates in the 20% condition (20% condition: 72.0%; 80% condition: 82.9%; $\chi^2(1) = 9.72$, $p = .002$). The main effect of block was significant, with a positive slope and Max-rates increasing over time (Block1: 66.7%; Block5: 82.7%; $b = 0.037, \chi^2(1) = 58.66$, $p < .001$). The interaction between block and information condition was significant ($\chi^2(2) = 8.68$, $p = .01$), with different slopes as a result of steeper increase in Max-rates over time in the DEC condition, and almost no change over time in the DES condition (slopes b for each condition: DES=0.019, E=0.034, DEC=0.052). No other interactions were significant. The patterns of behaviour observed in the DES condition were also similar to those in the work by Jessup et al. (2008).

A post-hoc analysis for the last block of 20 trials was also conducted, in order to compare the effects of the three different types of information presentation conditions (DEC, DES and E), with Tukey corrections. It is of particular interest to look at the Max-rates in the last block since by then participants can be expected to have stabilised in their preferred choice (Bechara, Damasio, Damasio, & Anderson, 1994; Ert & Erev, 2007). There was no difference between the DES and E conditions (E: 87.0%; DES: 86.8%; $t(204.67) = 0.05$, $p = .999$). The presence of descriptive information congruent to experience influenced behaviour weakly or not at all, consistent with the results observed by Lejarraga and Gonzalez (2011).

However, in the DEC conditions, which combined conflicting information from description and from experience, participants' behaviour was shifted towards the choice predicted from the descriptive button labels, as observed in the relevant E and DES conditions, and away from the choice observed in the other conditions. For example, if the $D_{20\%}$ description in the $D_{20\%}E_{80\%}C$ condition was influencing behaviour, I would expect a shift away from the behaviour observed in the $D_{80\%}E_{80\%}S$ and the $E_{80\%}$ conditions and towards what was observed in the $D_{20\%}E_{20\%}C$ and $E_{20\%}$ conditions, and vice-versa for the $D_{80\%}E_{20\%}C$ condition. This effect was observed. In the 80% condition, the DEC condition made the risky choice less attractive by describing a lower probability of rewards than experienced, and in the 20% it made it more attractive. Therefore I would expect a decrease in Max-rates in both conditions. And this was observed: Max-rates in the DEC condition (DEC: 74.3%) were significantly lower than in the other two conditions (against DES: $t(204.67) = 2.56$, $p = .03$; against E: $t(204.67) = 2.62$, $p = .03$).

2.1.3 Discussion

The conflicting descriptive information influenced behaviour significantly and in the direction predicted by the misleading information provided, reducing maximisation rates. This behaviour could be explained by participants taking into account the descriptive information and integrating it with the experiential information into their decision-making process. If participants were disregarding the descriptive information completely, the Max-rates should not have differed between the comparable DEC, DES and E conditions for each probability level, since the experienced feedback in these three conditions was the same.

The behaviour observed in the DEC conditions was also shifted towards what could be interpreted as a more random pattern, with Max-rates closer to 50% than in the respective DES and E conditions. I propose that this shift is towards more predicted behaviour as inferred by the conflicting descriptions, with participants being influenced by the content of the information available in the descriptions. However, other reasons might cause a similar shift towards random behaviour, as a result of the increase in uncertainty and confusion in the DEC conditions. The conflicting information introduced further uncertainty into the task, and uncertainty can make participants believe less in their own experience and also explore more often, which could lead to more random-like behaviour (Erev & Barron, 2005; Knox et al., 2012; Speekenbrink & Konstantinidis, 2015). In order to test if participants were being influenced by the content of the conflicting description, or simply behaving more randomly, I devised Experiment 2 in which the conflicting information should influence participants away from randomness.

2.2 Experiment 2

While in Experiment 1 the conflicting information influenced participants towards what could potentially be interpreted as more random behaviour, Experiment 2 was designed so that the conflicting information should influence participants behaviour away from randomness. In both conditions in Experiment 1, a decrease in the Max-rates was expected, and observed. The new design for Experiment 2 uses a paradigm in which an increase in Max-rates is expected by making the conflicting descriptions point to a more attractive maximisation choice. In this way, conflicting descriptions should move the observed behaviour away from randomness.

2.2.1 Method

Design

These experiments used a 3 $\times$ 2 between-subjects designs with three types of information presentation and two levels of probabilities in the risky option. Information was presented in one of the following: the experience-only (E) condition; the description-experience-same (DES) condition; or the description-experience-conflict (DEC) condition. The two levels of probabilities of the risky option were changed from Experiment 1: in Experiment 2a the probabilities used were 80% and 40%, and in Experiment 2b they were 40% and 20%. Each participant was presented with one type of information presentation and one level of probability.

Participants

Participants were recruited on-line using Amazon's Mechanical Turk service. Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment.

Experiment 2a 120 individuals participated (63 females; age: *M* = 34 years, *SD* = 11 years), 20 in each experimental condition. Average bonus paid was US$ 0.90 (*SD* = US$ 0.06).[4]

Experiment 2b 120 individuals participated (43 females; age: *M* = 31 years, *SD* = 9 years), 20 in each experimental condition. Average bonus paid was US$ 0.55 (*SD* = US$ 0.18).[5]

Task

The experimental paradigm was similar to that of Experiment 1, the only differences being the new values and probabilities for the risky option. The safe button still paid a sure outcome of two points 100% of the time in both experiments. In Experiment 2a, the risky choice paid rewards of six points either 80% or 40% of the time, according to the probability level condition. In Experiment 2b, the risky choice paid four points either 40% or 20% of the time. The new outcomes and button labels can be seen in Table 2.2. Points were converted to money at a rate of US$ 0.20/100 points in Experiment 2a and US$0.40/100 points in Experiment 2b. Participants completed the task in an average of 6.1 minutes (SD = 2.6 minutes).

In Experiment 1 the conflicting condition shifted the maximisation option, for example, if the maximisation option was the risky choice according to experience, the conflicting description would inform participants that the safe choice was the maximisation one, according to expected value (EV). In Experiment 2, this shift did

[4]Bonus amount in Experiment 2a was not influenced by information presentation ($F(2,114) = 0.08, p = .92$), but it was influenced by probability level, with a significantly lower overall bonus in the 40% condition (80%: 0.92; 40%: 0.89; $F(1,114) = 7.54, p = .007$).

[5]Bonus amount in Experiment 2b was not influenced by information presentation ($F(2,114) = 0.73, p = .48$), but it was influenced by probability level, with a significantly lower overall bonus in the 40% condition (40%: 0.37; 20%: 0.72; $F(1,114) = 1163.9, p < .001$).

Table 2.2: Experiment 2: Button labels according to experimental condition.

Condition	Safe choice	Risky choice
Experiment 2a		
$E_{80\%}$, $E_{40\%}$	(blank)	(blank)
$D_{80\%}E_{80\%}S$, and $D_{80\%}E_{40\%}C$	2 points with 100% probability.	6 points with 80% probability; Zero otherwise.
$D_{40\%}E_{40\%}S$, and $D_{40\%}E_{80\%}C$	2 points with 100% probability.	6 points with 40% probability; Zero otherwise.
Experiment 2b		
$E_{40\%}$, $E_{20\%}$	(blank)	(blank)
$D_{40\%}E_{40\%}S$, and $D_{40\%}E_{20\%}C$	2 points with 100% probability.	4 points with 40% probability; Zero otherwise.
$D_{20\%}E_{20\%}S$, and $D_{20\%}E_{40\%}C$	2 points with 100% probability.	4 points with 20% probability; Zero otherwise.

not occur. The maximisation option was always the same regardless of descriptions. In Experiment 2a, the maximisation option was always the risky choice, while in Experiment 2b, the maximisation option was always the safe choice. However, the conflicting description did make the maximisation option more or less attractive. For example, in Experiment 2a, in the 40% condition, the EV of the risky choice according to experience was 2.4 points, while conflicting description would say that the EV of the risky choice was actually 4.8 point, in contrast with the EV of the safe choice which was fixed at 2 points. In Experiment 2b, the safe choice, always the maximisation option, was made more or less attractive by the conflicting label (EV of the risky choice was 0.8 or 1.6 points according to the labels). Using Prospect Theory (Tversky & Kahneman, 1992) instead of EV leads to the same maximisation options.

Therefore, in the 40% condition of Experiment 2a and the 40% condition of Experiment 2b, the conflicting descriptions made the maximisation options (risky choice in the case of 2a and safe choice in the case of 2b), more attractive. While in the 80% condition of Experiment 2a and 20% condition of Experiment 2b, the conflicting descriptions made the maximisation options seem less attractive. I am going to group these conditions across experiments for the analysis below.

2.2.2 Results

As in Experiment 1, the main dependent variable was the average proportion of times individuals selected the maximisation choice (Max-rate), in blocks of 20. The same analysis used in Experiment 1 was conducted for Experiment 2, with the same fixed and random components. I combined conditions 40% of Experiment 2a and 40% of Experiment 2b and analysed them together as the "up" conditions (Figure 2.2), and combined conditions 80% of Experiment 2a and 20% of Experiment 2b together as the "down" conditions (Figure 2.3), according to the predicted direction of change in maximisation rates.

"Up" conditions

In the 40% probability level condition for Experiment 2a and 40% probability level condition for Experiment 2b, an increase in the maximisation rates was predicted for the conflicting descriptions (Figure 2.2).

The main effect of information presentation condition was significant. Max-

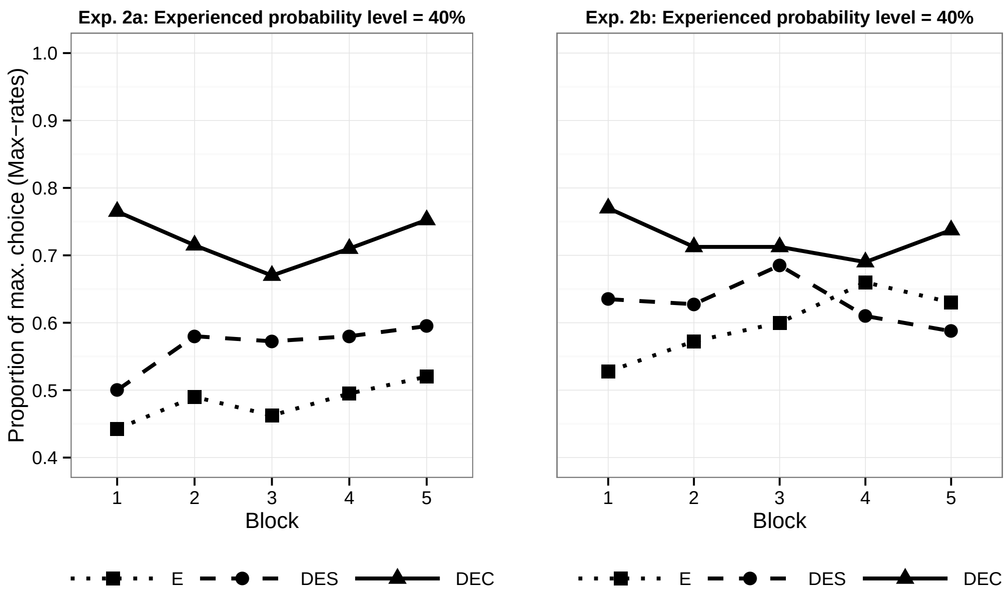

Figure 2.2: Evolution of the maximisation choice rate (Max-rate) for each block of 20 trials for Experiment 2 "up" conditions. The left panel shows the results when the risky choice paid a reward of six points 80% of the time, while in the right panel the same reward was paid 40% of the time. The lines refer to the different descriptions presented to participants (E: experience only, no description; DES: description and experience same; DEC: description and experience conflicting).

rates were higher in the DEC condition (DEC: 72.4%; DES: 59.7%; E: 54.0%; $\chi^2(2) = 18.11$, $p < .001$), and away from randomness. The main effect of block was not significant, with Max-rates relatively constant over time (Block1: 60.7%; Block5: 63.7%; $\chi^2(1) = 1.40$, $p = .24$). There was no significant difference between the two experiments (Exp2a: 59.0%; Exp2b: 65.1%; $\chi^2(1) = 2.74$, $p = .10$). None of the interactions were significant.

As in Experiment 1, a post-hoc analysis was conducted at Block 5 to look at more stable behaviour. The same pattern was observed, with no significant difference in Max-rates between E and DES, as expected (E: 57.5%; DES: 59.1%; $t(202.88) = 0.271$, $p = .96$). The Max-rates for DEC (74.5%) were significantly higher than both other conditions (against E: $t(202.88) = 2.84$, $p = .01$; against DES: $t(202.88) = 2.57$, $p = .03$). Conflicting descriptions influenced behaviour significantly, in the direction predicted by the descriptive information provided. The descriptive information presented a probability of rewards for the risky option higher than experienced, so an increase in Max-rates was expected and observed, replicating the results from Experiment 1, but in the opposite direction.

"Down" conditions

In the 80% probability level condition for Experiment 2a and 20% probability level condition for Experiment 2b, a decrease in the maximisation rates was predicted for the conflicting descriptions (Figure 2.3).

The main effect of information presentation condition was not significant. Max-rates were not different across the conditions (DEC: 87.5%; DES: 87.9%; E: 88.5%; $\chi^2(2) = 0.78$, $p = .68$). They were also considerably high and close to the ideal maximum of 100%. The main effect of block was significant, with Max-rates increasing over time, with a positive slope (Block1: 80.3%; Block5: 92.3%; $b = 0.022, \chi^2(1) = 79.01$, $p < .001$). There was a significant difference between the two experiments (Exp2a: 92.7%; Exp2b: 83.1%; $\chi^2(1) = 15.86$, $p < .001$). There was a significant interaction between Block and information presentation ($\chi^2(2) = 13.61$, $p = .001$), due to the steep increase from the first to second blocks in the DEC condition (because of the conflicting descriptive information) in com-

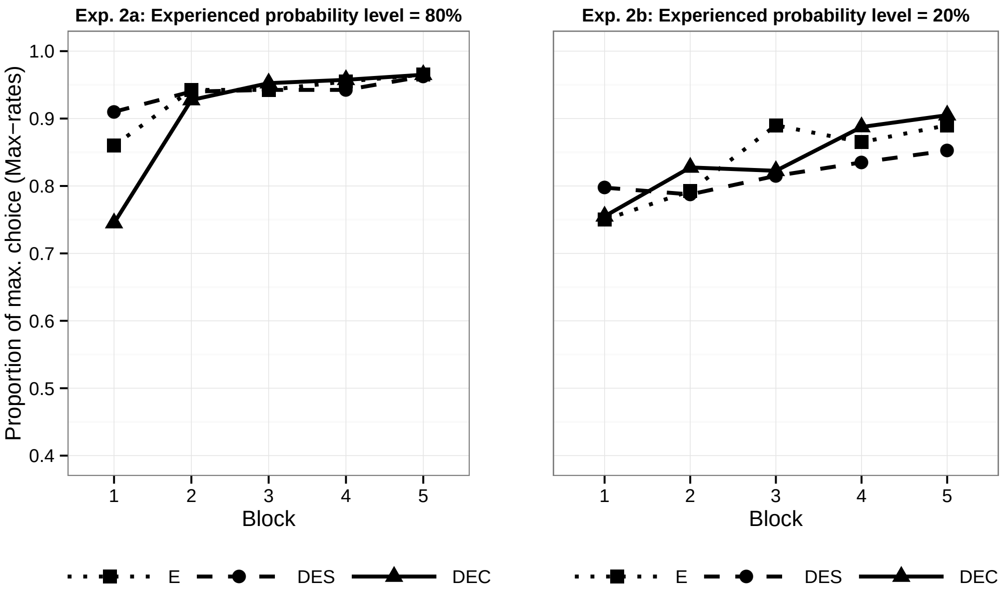

Figure 2.3: Evolution of the maximisation choice rate (Max-rate) for each block of 20 trials for Experiment 2 "down" conditions. The left panel shows the results when the risky choice paid a reward of four points 40% of the time, while in the right panel the same reward was paid 20% of the time. The lines refer to the different descriptions presented to participants (E: experience only, no description; DES: description and experience same; DEC: description and experience conflicting).

parison to the flatter more stable behaviour across the other conditions. None of the other interactions were significant. A post-hoc analysis at Block 5 showed no significant differences (ps $> .70$).

2.2.3 Discussion

The influence of conflicting description on behaviour was observed in the 40% condition of Experiment 2a and the 40% condition of Experiment 2b. In these "up" conditions, the conflicting description made the maximising choice more attractive by increasing the difference in expected value between the risky and safe choices, and led to significant increases in the Max-rates, away from random behaviour, in the direction expected from the conflicting descriptions. While in the previous experiment the conflicting descriptions led to a reduction in performance, and more random selections, in the current experiment the conflicting descriptions led to an increase in performance, making participants select the maximising choice more frequently. Thus confirming that the result observed in Experiment 1 was not simply due to confusion and participants behaving more randomly.

However no significant influence of conflicting description (relative to the E and DES conditions) was observed in the 80% condition of Experiment 2a and the 20% condition of Experiment 2b. Instead I observed ceiling effects, with Max-rates close to 100% regardless of experimental condition. In comparison to Experiment 1, which provided the same probabilities of reward (five points with 80% and 20%), in Experiment 2a participants could earn more points (six points) and in Experiment 2b fewer points (four points). These changes led to increases of Max-rates in comparison to Experiment 1, resulting in the ceiling effects.

The opportunity cost of complying with warning labels seems to moderate behaviour (Argo & Main, 2004) and also needs to be taken into account: In the 80% condition of Experiment 2a and the 20% condition of Experiment 2b, the expected value of one option was considerably higher than the other, and hence complying with the misleading description (which would lead to deviation from optimal behaviour) was more costly. In addition, the influence of descriptive information might have been given less importance than that of experiential information, as previously suggested (Jessup et al., 2008; Lejarraga, 2010; Rogers et al., 2000; Shlomi, 2014). Descriptions, if discounted, would have less influence on behaviour when the difference in expected values is higher, and this would explain the behaviour observed. The discounting of descriptions will be further investigated in Experiment 3 and in the cognitive modelling section below.

2.3 Experiment 3

In Experiment 3, I aimed to verify the boundaries of the influence of conflicting descriptive information. I propose that the influence of conflicting descriptions would be monotonically increasing at the centre, but not at extreme levels of informational conflict. In extreme cases when information is highly implausible, it should be more readily discarded from the decision-making process. This could lead to a reduced marginal influence or even a contrasting effect at the extremes, similar to what has been found in research on anchoring (e.g., Chapman & Johnson, 1994), advice seeking (e.g., Yaniv, 2004), goal setting (e.g., Locke, 1982), and psychophysics (e.g., Brown, 1953). In order to check how a description's plausibility influences behaviour, I created experimental conditions with varying levels of conflict: no conflict, plausible conflict and implausible conflict.

Plausibility was manipulated via the difference between the actual experienced against the described frequencies of rewards. Across all conditions, the risky option returned rewards 50% of the time experientially. In the two plausible conflict conditions, descriptions informed participants that rewards were paid either with a 25% or 75% probability, relatively close to the true experienced frequency of 50%, making the descriptions plausible explanations for the experience. In the two implausible conflict conditions, descriptions informed participants that rewards were paid either with a 1% or 99% probability, which were highly implausible given the experience. I expected participants to disregard the descriptions more easily in the implausible conflict conditions, therefore reducing their effect on behaviour.

2.3.1 Method

Design

This experiment followed a between-subjects design with six experimental conditions, with manipulations of the descriptions that were provided to participants. Each participant was presented with only one type of description, assigned randomly from the following options: the experience-only (E) condition; the description-experience-same (DES) condition, in which the description matched the experience at 50% probability of receiving a reward; two *plausible* conflict (DEC_p) conditions, with descriptive probabilities of 25% and 75%; and two *implausible* conflict (DEC_i) conditions, with descriptive probabilities of 1% and 99%.

Participants

240 participants (110 females; age: M = 33 years, SD = 11 years) were recruited on-line using Amazon's Mechanical Turk service. Participation was restricted to individuals whose location was defined as in the United States. There were 40 participants in each experimental condition. No participants were excluded from

the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment. Average bonus paid was US$ 0.80 (*SD* = US$ 0.03).[6]

Task

Table 2.3: Experiment 3: Button labels according to experimental condition.

Description condition of risky option	Safe choice button label	Risky choice button label
Experience-only ($E_{50\%}$)	(blank)	(blank)
No conflict ($D_{50\%}E_{50\%}S$)	2 points with 100% probability.	4 points with 50% probability; Zero otherwise.
Plausible conflict 25% ($D_{25\%}E_{50\%}C_p$)	2 points with 100% probability.	4 points with 25% probability; Zero otherwise.
Plausible conflict 75% ($D_{75\%}E_{50\%}C_p$)	2 points with 100% probability.	4 points with 75% probability; Zero otherwise.
Implausible conflict 1% ($D_{1\%}E_{50\%}C_i$)	2 points with 100% probability.	4 points with 1% probability; Zero otherwise.
Implausible conflict 99% ($D_{99\%}E_{50\%}C_i$)	2 points with 100% probability.	4 points with 99% probability; Zero otherwise.

The experimental paradigm was similar to that of Experiments 1 and 2, the only differences being the new values and probabilities for the risky option. In this experiment, the experience drew from the same underlying distributions across all conditions, with the risky option always returning 4 points with 50% probability, and the safe option always returning 2 points with 100% probability. The expected values of the risky and safe options were the same. The only between-subjects manipulation was the descriptive information. The new outcomes and button labels can be seen in Table 2.3. Points were converted to money at a rate of US$ 0.40/100 points. In addition, after participants had finished selecting between the two choices, the descriptions (if previously present) were hidden and a blank text box appeared under each button. Participants were asked to input their judgements for the actual experienced frequencies of rewards for each button, in the range of 0-100%, using the numbers in their keyboards. The removal of the descriptions from the screen was done to reduce any potential anchoring effect, and avoid participants

[6]Because the expected value of the safe and risky options were matched, there was no significant difference in the bonus paid according to experimental condition ($F(1,238) = 0.259, p = .611$).

from simply copying the descriptions as their answers. Participants completed the task in an average of 8.0 minutes (SD = 5.6 minutes).

2.3.2 Results and Discussion

In Experiment 3, the expected values (EV) of the two options were always the same, according to experience, with no obvious maximisation choice according to experienced EV alone. The main dependent variable was the average proportion of times individuals selected the risky choice (R-rate) in each block of 20 trials (Figure 2.4A shows the last block). The same analysis that was used in Experiments 1 and 2 was conducted for Experiment 3, with the same components. The average judgements of the frequency of reward appearances were also analysed using a one-way ANOVA by experimental condition (Figure 2.4B).

The main effect of information presentation condition was significant. R-rates increased with higher description levels ($\chi^2(5) = 52.35$, $p < .001$). A linear con-

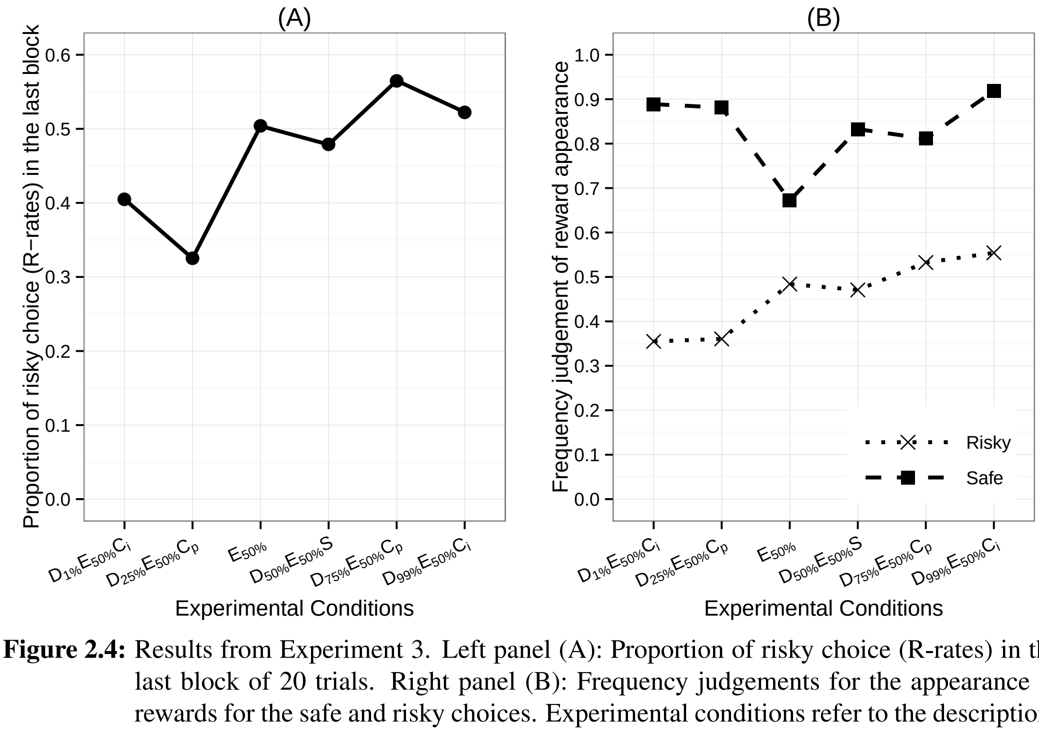

Figure 2.4: Results from Experiment 3. Left panel (A): Proportion of risky choice (R-rates) in the last block of 20 trials. Right panel (B): Frequency judgements for the appearance of rewards for the safe and risky choices. Experimental conditions refer to the descriptions presented to participants for the risky option ($E_{50\%}$: experience-only, no description, $D_{50\%}E_{50\%}S$: same description (50%), $D_{25\%}E_{50\%}C_p$: plausible conflicting description (25%), $D_{75\%}E_{50\%}C_p$: plausible conflicting description (75%), $D_{1\%}E_{50\%}C_i$: implausible conflicting description (1%), $D_{99\%}E_{50\%}C_i$: implausible conflicting description (99%)).

trast ordered according to the level of descriptions presented was positive and significant ($b = 1.09, t(234) = 3.18, p = .002$). Among the higher-level contrasts, the quadratic and cubic ones were also significant (quadratic: $b = 1.15, t(234) = 3.07, p = .002$; cubic: $b = -2.17, t(234) = 3.97, p < .001$), defining the S-shape seen in Figure 2.4A. The main effect of block was not significant ($\chi^2(2) = 0.02$, $p = .90$), however the interaction of information and block was significant ($\chi^2(5) = 11.34$, $p = .045$), with an increase in Max-rates over time in the $D_{1\%}E_{50\%}C_i$ condition ($b = 0.028$), a decrease in the $D_{99\%}E_{50\%}C_i$ condition ($b = -0.019$), and relatively flat in the other conditions, with slopes close to zero. I also looked at plausibility as an independent variable, using 3 factors: plausible, implausible, and base, the latter applying to the DES and E conditions. The effect of plausibility on R-rates was significant ($\chi^2(2) = 11.32$, $p = .003$). A post-hoc analysis at the last block shows that, as in the previous experiments, presenting participants with congruent descriptions did not influence behaviour in relation to no descriptions (DES: 47.9%; E: 50.4%; $t(369.21) = 0.38$, $p = .99$). In the plausible conflict conditions DEC_p of 25% and 75%, I observed significantly different R-rates from the base conditions ($t(371.84) = 2.79, p = .02$). In the implausible conflict conditions DEC_i of 1% and 99%, I observed a reversal in the influence of the conflicting descriptions, with R-rates no longer significantly different from the base conditions ($t(371.84) = 1.47, p = .31$). This reversal can be seen in Figure 2.4A.

The frequency judgement of reward appearances was also analysed, using a one-way ANOVA by experimental condition (Figure 2.4B). All participants experienced frequencies of rewards of 100% for the safe choice and 50% for the risky choice, which would have been their unbiased correct answers. For the safe choice, participants' judgements were not different across the individual description conditions ($M = 87\%; F(4, 195) = 1.587, p = .179$); in the E condition however, their judgements were significantly lower (E: 67%; $F(5, 234) = 5.851, p < .001$). There was also a significant difference in the judgements for the risky choice ($F(5, 234) = 8.952, p < .001$). This effect was analysed with five polynomial contrasts: The linear, quadratic and cubic contrasts were all significant (ps<.001); while the two remaining higher order contrasts were not significant (*p*s>.52). This would indicate a sigmoid-shaped monotonically increasing judgement in relation to the description: participants presented with descriptions of higher probabilities of rewards responded with higher frequency judgements of the observed rewards, with diminishing sensitivities at the extremes (see Figure 2.4(b)). The individual frequency judgements for the risky choice were also significantly correlated to the individual R-rates ($r = .50, n = 238, p < .001$).

The lack of influence of congruent descriptions and the significant influence of conflicting descriptions on behaviour were again observed in Experiment 3, replicating the results found in Experiments 1 and 2. In addition, plausible conflicting descriptions influenced R-rates in a monotonic way, with high described probabilities of rewards increasing R-rates, and vice-versa. However, in the case of implausible conflicts, a contrasting effect was observed. A more extreme and more implausible described probability had a weaker effect on behaviour than a less extreme but plausible one. If the description is highly unlikely to be a true representation of the experience, participants give it lower weight in their decision-making process. These differences in decision weights will be specified with a cognitive model in the next section.

2.4 Cognitive modelling

To further test if participants integrate the descriptive information into their decision-making processes, a set of cognitive computational models was fitted to the experimental data. If the descriptive information influenced human behaviour, then a model that includes representations of both description and experience should fit better than a model that relies on experience alone. I therefore compared *experience-only* against *description-experience* models. Within the description-experience models, I tested two different approaches: a *fixed-weight* approach, in which the weights given to description are fixed over time and over conditions, and a *Bayesian-updating* approach, in which the weights given to description change over time according to the plausibility of the evidence observed in contrast with the description.

2.4.1 The Models

The aim of fitting a cognitive model to the data was to assess and formalise how the two sources of information, descriptive and experiential, are combined. I did not aim to have an extensive comparison between different decision-making models in decision making paradigms (such as the one in Yechiam & Busemeyer, 2005).

I fitted three models to the behavioural data. They all share the same basic structure, which is defined by the final expected value $FEV_j(t)$ of each choice j available to participants at time t:

$$FEV_j(t) = \xi_j(t) \cdot D_j + [1 - \xi_j(t)] \cdot E_j(t).$$

I propose that the two sources of information are combined via $\xi_j(t)$, a parameter which determines the weight given to description at each point in time for each option. A representation of the descriptive information is included via D_j, which is the expected value calculated from the descriptive information available to participants, using cumulative prospect theory (CPT: Tversky & Kahneman, 1992). The experience is represented by $E_j(t)$, which is the expected value calculated from the experiential information received by participants in the form of feedback up to trial t, based on a delta-rule reinforcement learning model.

2.4.2 Description (D_j)

The subjective expected value of the descriptive information for choice j, D_j, was fixed over time and calculated as the CPT value based on the descriptions provided to participants in the button labels. According to Tversky and Kahneman (1992), the CPT value is calculated using the curvature parameter for values and weighting parameter for probabilities, ν and ω respectively,

$$D_j = \sum_m W(p_{jm}) V_{jm}^{\nu},$$

where p_{jm} are the probabilities and V_{jm} are the potential values for each outcome m of option j; ν is the free parameter that determines the curvature of the value function ($0 \leq \nu \leq 1$), with lower values reducing the distance between extreme values of rewards; and $W(\cdot)$ is the probability weighting function. $W(\cdot)$ is defined as:

$$W(p) = \frac{p^{\omega}}{(p^{\omega} + (1-p)^{\omega})^{\frac{1}{\omega}}},$$

where ω is the free parameter ($0 \leq \omega \leq 5$) that determines the curvature of the probability weighting function. Values of ω below 1 lead to overweighting of rare events, while values above 1 lead to underweighting of rare events.

2.4.3 Experience ($E_j(t)$)

A simple reinforcement learning model was used to fit the experimental choice data. This model has extensively been used before in research on repeated decisions from experience (Erev & Barron, 2005; Lejarraga & Gonzalez, 2011; Yechiam & Busemeyer, 2005; Yechiam & Rakow, 2012).

Firstly, observed outcomes are evaluated by a prospect-theory type of utility function (Kahneman & Tversky, 1979). The utility function $v_j(t)$ of option j is defined as:

$$v_j(t) = [payoff_j(t)]^{\nu},$$

where $payoff_j(t)$ is equal to the payoff, in points, at each trial t for each option j, and ν is the same parameter that determines the curvature of the utility function for the description.

Secondly, expectancies for the value of rewards for each option are formed via a learning rule, which integrates the experienced feedback after each trial. The learning rule used was a delta rule, which uses a learning rate that determines how much the new information gathered via feedback, in the form of prediction error, influences the updating of the expectancies at each trial (Sutton & Barto, 1998; Yechiam & Busemeyer, 2005). Feedback observed is integrated after each trial, to arrive at the experienced expectancy $E_j(t)$ for option j at time t:

$$E_j(t) = E_j(t-1) + \phi \cdot \left\{\delta_j(t) + \gamma \cdot [1-\delta_j(t)]\right\} \cdot [v_j(t) - E_j(t-1)],$$

where ϕ is the free learning rate parameter $(0 \leq \phi \leq 1)$, which is a weight given to new information observed, with lower values resulting in slower learning. The free parameter γ, $(0 \leq \gamma \leq 1)$, denotes the weight associated with the feedback of the foregone option, such that when γ=1 the foregone and observed payoffs are weighted the same, and when γ=0 foregone payoffs are disregarded. The variable $\delta_j(t)$ is a dummy variable, which is equal to one if option j was chosen on trial t, and zero otherwise.

2.4.4 The weight given to description ($\xi_j(t)$)

Three different approaches were used to combine description and experience by manipulating the parameter $\xi_j(t)$, which is the weight given to description: experience-only, fixed-weight and Bayesian-updating. Each of these approaches was fitted to the behavioural data individually, and the fit results were then compared.

Experience-only model

In the experience-only model, no representation of description was included, with $\xi_j(t)$ fixed to zero across all trials and options. Therefore, the final expectancy $FEV_j(t)$ was defined to be equal to the expectancy derived from experience alone, via the reinforcement learning model, $E_j(t)$. This model assumes that descriptions do not influence the decision making process.

Fixed-weight model

In the fixed-weight model, the weight given to description, $\xi_j(t)$, was set as a free parameter, ξ, constant across all trials, options and conditions $(0 \leq \xi \leq 1)$. This model assumes that description influenced the individual choices, but with a fixed weight that did not depend on the experimental condition and did not change over time.

Bayesian-updating model

Based on my behavioural results, I observed that descriptions influenced the decision-making processes in different ways in each experimental condition. More plausible descriptions seemed to have a stronger influence on decisions than less plausible ones. I propose a Bayesian-updating model in which the weight given to description, $\xi_j(t)$, equals the subjective probability that the description is true on that trial, given the evidence observed thus far. In this model, the weights given to description will differ for each option and change over time, according to the experimental conditions.

The Bayesian model assumes that, at each trial, either the description is true, denoted as $\mathscr{D}_j(t)$, or the task is in a different state, denoted as $\mathscr{E}_j(t)$, where the probabilities of rewards are not as described but instead are as experienced. From trial to trial, the state of the task can change, such that if the description was true on the previous trial, it no longer is true on the next. The task of a participant is twofold: to determine whether the task is in state $\mathscr{D}_j(t)$ or $\mathscr{E}_j(t)$, and to estimate the relevant probabilities of the outcomes when the task is in state $\mathscr{E}_j(t)$. The probabilities in state $\mathscr{D}_j(t)$ do not need to be estimated as they come from the description itself. Initially it makes sense to rely on the description, as there is no information to estimate the probabilities of winning in the other state. Over time however, it is possible to learn that the true probabilities of the outcomes are different than those described, in which case the weight given to description should diminish. With this model, less plausible conditions lead to lower weights given to description than more plausible conditions, thus making this approach more adaptive to the non-monotonicity observed in the behavioural results than using a fixed weight (Figure 2.5).

I set the weight given to the description as the Bayesian predicted probability that the description is true on that trial, $\xi_j(t) = p(\mathscr{D}_j(t)|k_j(1:t-1))$. This probability is based on all the observations made up to the previous trial $t-1$, denoted here as $k_j(1:t-1)$, which are all the observed outcomes of option j, whether rewards were obtained or not, from trial 1 to trial $t-1$. I will simplify this notation by

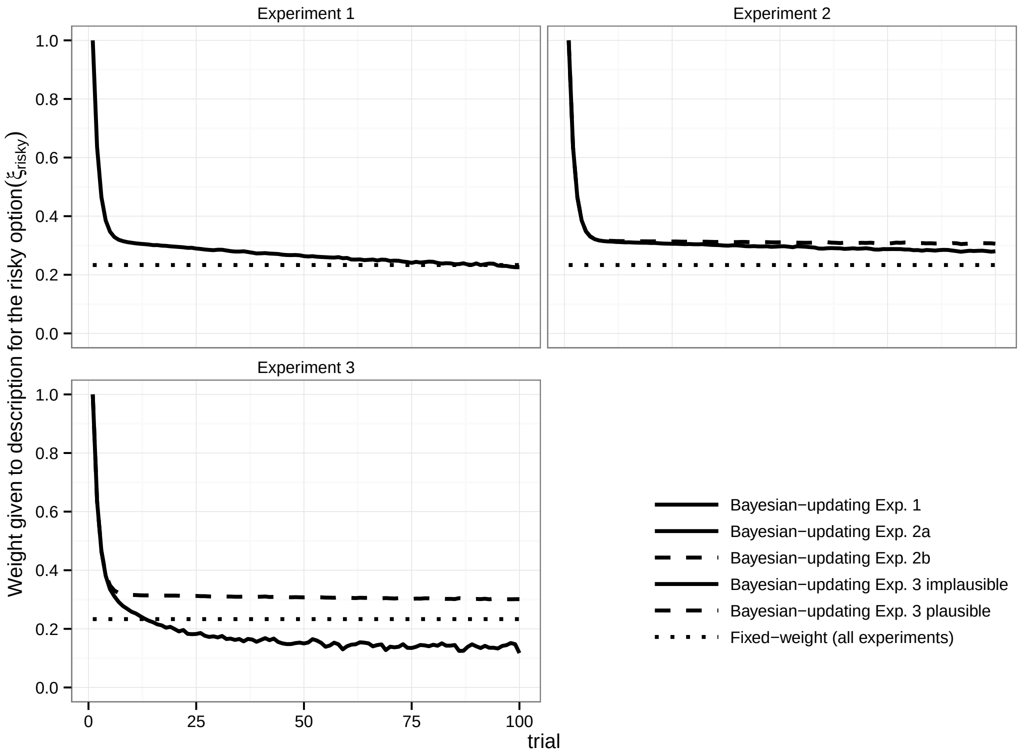

Figure 2.5: Average weights given to description for the risky option (ξ_{risky}) in the two different description-experience models, based on the best fit parameters, for the conflicting (DEC) conditions only. The Bayesian-updating approach shows how the weight changes over time according to the accumulation of evidence, and the plausibility of the description, as measured by the difference between descriptive and experiential information.

using a subscript to denote the information used to calculate the probability, with $p_{t-1}(\mathscr{D}_j(t)) = p(\mathscr{D}_j(t)|k_j(1:t-1))$, and analogously using p_t to include all the information from trial 1 to trial t, $k_j(1:t)$. I also assume that the state of the task can change over trials according to the transition probabilities $\kappa_{\mathscr{D}}$ and $\kappa_{\mathscr{E}}$ which are free parameters (range: 0-1). $\kappa_{\mathscr{E}}$ is the probability that state $\mathscr{D}_j$ is true in trial t if state $\mathscr{E}_j$ was true in trial $t-1$:

$$p(\mathscr{D}_j(t)|\mathscr{E}_j(t-1)) = \kappa_{\mathscr{E}}$$

and conversely for $\kappa_{\mathscr{D}}$:

$$p(\mathscr{E}_j(t)|\mathscr{D}_j(t-1)) = \kappa_{\mathscr{D}}.$$

The effect of the transition probabilities is to change the prior distributions at each trial. The prior probability that the description is true on trial $t+1$ for option

j is then:

$$p_t(\mathscr{D}_j(t+1)) = p_t(\mathscr{D}_j(t)) \cdot (1-\kappa_{\mathscr{D}}) + p_t(\mathscr{E}_j(t)) \cdot \kappa_{\mathscr{E}}.$$

Since one of the two states has to be true at any point, the two probabilities $p_{t-1}(\mathscr{D}_j(t))$ and $p_{t-1}(\mathscr{E}_j(t))$ are complimentary and the prior probability that state $\mathscr{E}_j$ is true is simply $p_{t-1}(\mathscr{E}_j(t)) = 1 - p_{t-1}(\mathscr{D}_j(t))$. The prior $p_{t-1}(\mathscr{D}_j(t))$ and its compliment $p_{t-1}(\mathscr{E}_j(t))$ can be used to calculate the Bayesian posterior:

$$p_t(\mathscr{D}_j(t)) = \frac{p(k_j(t)|\mathscr{D}_j(t))p_{t-1}(\mathscr{D}_j(t))}{p(k_j(t)|\mathscr{D}_j(t))p_{t-1}(\mathscr{D}_j(t)) + p(k_j(t)|\mathscr{E}_j(t))p_{t-1}(\mathscr{E}_j(t))}.$$

I set the initial prior, $p_0(\mathscr{D}_j(1))$ at the first trial, to be equal to one, since at that point there was no information experienced so far, and participants had to rely solely on the descriptive information provided to base their decisions.

According to the descriptive information $\mathscr{D}_j(t)$, the probability of a win $k_j(t)$ observed in t if the description holds on trial t is

$$p(k_j(t)|\mathscr{D}_j(t)) = pd_j^{k_j(t)}(1-pd_j)^{(1-k_j(t))},$$

where pd_j is the probability of obtaining a reward for option j as provided by the description. For example, in Experiment 1, pd_j could be either 0.2 or 0.8 for the risky options, depending on the experimental conditions, and 1.0 for the safe option.

The relevant probabilities for state $\mathscr{E}_j(t)$ have to be learned from experience. Assuming that people start with a Beta prior over these probabilities, the posterior distributions over these probabilities are also Beta distributions, and the likelihood of outcome $k_j(t)$ if the task is in state $\mathscr{E}_j(t)$ is

$$p(k_j(t)|\mathscr{E}_j(t)) = \frac{\mathrm{B}(\alpha_j(t-1)+k_j(t), \beta_j(t-1)+1-k_j(t))}{\mathrm{B}(\alpha_j(t-1), \beta_j(t-1))},$$

where B is the Beta function, and $\alpha_j(t)$ and $\beta_j(t)$ are its parameters, which are updated as follows, with new experiential evidence $k_j(t)$ weighed by the probabilities given to $\mathscr{E}_j(t)$, for $t \geq 1$:

$$\alpha_j(t+1) = \alpha_j(t) + p_t(\mathscr{E}_j(t)) \cdot k_j(t);$$

$$\beta_j(t+1) = \beta_j(t) + p_t(\mathscr{E}_j(t)) \cdot (1 - k_j(t)).$$

For $t = 1$ the values of $\alpha_j(1)$ and $\beta_j(1)$ are defined by the initial expected value of the Beta distribution which is set to the probability provided in the description, $\alpha_j(1)/(\alpha_j(1) + \beta_j(1)) = pd_j$. I constrained the total weight of the initial Beta prior as a free parameter $\mathscr{S} = \alpha_j(1) + \beta_j(1)$ for all options j ($1 \leq \mathscr{S} \leq 500$).[7] Higher values of $\mathscr{S}$ led to a slower accumulation of new evidence towards $p(\mathscr{E}_j(t))$.

2.4.5 Choice rule

After description and experience were integrated into the final expectancy calculation, $FEV_j(t)$, the choice rule used was a time-independent soft-max rule (Yechiam & Busemeyer, 2005) that combined the FEV_j across all options at each trial to determine the probability of choosing option j among all options J:

$$Pr[Choice(t+1) = j] = \frac{e^{\theta \cdot FEV_j(t)}}{\sum_J e^{\theta \cdot FEV_J(t)}},$$

where θ is the choice sensitivity free parameter, $(0 \leq \theta \leq 1)$. If $\theta = 0$, the model randomly guesses between the expectancies regardless of their values, while higher values of θ will lead to more deterministic maximisation behaviour.

2.4.6 Model fitting

Data sets containing 100 simulated participants were generated for each of the 4 experiments $\times$ 6 conditions,[8] with the same methodology used to generate actual data sets for the experiments. A total of 2,400 modelled simulated participants were confronted with 608 observed human participants. All simulated participants across all experiments and underlying experimental conditions shared the same set of free

[7] I also fitted a model in which the experience was initially set to follow a Beta(1,1) distribution, $\alpha_j(1) = \beta_j(1) = 1$, with an expected value of 0.5. This model also outperformed the fixed-weight model but by a smaller margin.

[8] The description-only (D) condition from Experiment 1 was excluded from the cognitive modelling, as it included a single-shot decision without experience.

parameters. The best fit parameters were found by minimising the log-likelihood between the average observed proportions of risky choice and the average model-predicted risky choice for each of the individual conditions separately, with each condition receiving the same weight (Erev & Barron, 2005). Because of the different number of parameters between the models, the Bayesian Information Criterion (BIC), which penalises for additional parameters, was calculated to compare the models, $BIC = LL + f \cdot ln(N)$, where f is the number of free parameters and N is the number of fitted observations for each evaluation. Lower BIC values represent better fitting models.

2.4.7 Model Evaluation and Results

Three models were evaluated: the experience-only model, which did not account for the influence of descriptive information, with four free parameters; the fixed-weight model, which assumed a single fixed weight across all trials and all conditions, with six free parameters; and the Bayesian-updating model, with its reducing weight given to description over time according to the feedback received, with eight free parameters.

The best fit parameters were relatively consistent across the different models (Table 2.4). In the best-performing Bayesian-updating model, the modelled influence of description started at 1 and converged towards a stable level ranging be-

Table 2.4: Best fit parameters of the three cognitive models. n.a. = not applicable. Note: the weight given to description (ξ) in the experience-only model was fixed to zero.

Parameter	Bayesian updating	Fixed weight	Experience only
ν (curvature of value function)	0.91	0.93	0.92
ω (curvature of probability function)	1.02	1.35	n.a.
ϕ (learning rate)	0.08	0.08	0.09
γ (foregone's weight)	0.99	0.96	1.00
θ (choice sensitivity)	0.83	0.80	0.74
ξ (description's weight)	n.a.	0.23	*zero (fixed)*
$\kappa_{\mathscr{D}}$ (switch rate for description)	0.17	n.a.	n.a.
$\kappa_{\mathscr{E}}$ (switch rate for experience)	0.36	n.a.	n.a.
$\mathscr{S}$ (Initial Beta prior sum for experience)	454	n.a.	n.a.
Total number of free parameters	8	6	4

tween 0.15 and 0.30 according to the plausibility of the descriptions (Figure 2.5 on page 70). It seems that even after many trials and despite the large amount of evidence gathered via feedback, participants were still taking the descriptive information into consideration, albeit at a discounted level; experience never gained participants' full attention. This could be explained by the constant presence of the descriptive information on the buttons, which might have continuously reinforced its influence. Barron et al. (2008) found a similar lingering influence of descriptive information even when descriptions were only presented briefly. In comparison, the fixed-weight model predicted a constant weight given to description of $\xi = 0.23$ throughout all trials and conditions.

Table 2.5: Mean BIC values for the experience-only and description-experience cognitive models fitted on the three groups of experimental conditions: description-experience-conflict (DEC), description-experience-same (DES) and experience-only (E). Values in brackets are the differences in relation to the base model at the top of each column. Lower BIC values represent better fitting models.

Model	Overall	DEC	DES	E
Experience-only (Base)	503	627	417	412
Fixed-weight	425 (-16%)	454 (-28%)	400 (-4%)	410 (-1%)
Bayesian updating	417 (-17%)	437 (-30%)	394 (-6%)	410 (-1%)

The results of the model fitting analysis were in line with the behavioural results (Table 2.5). The mean Bayesian Information Criteria (BIC) values for the two description-experience models were substantially smaller than those for the experience-only model (range: 16-17% smaller overall; lower BIC values represent better fitting models). Therefore, models that included the descriptive information provided a better fit for the observed behaviour than a model that did not include the influence of description (Figure 2.6). The Bayesian-updating model, which allows for the plausibility of the description as a source of information to influence the weights given to description and experience, was the best fitting model overall, with a 17% lower mean BIC value than the experience-only model, and 2% lower than the fixed-weight model. Among individual conditions, the highest reduction in mean BIC values was obtained in the DEC conditions, with the Bayesian-updating model 30% lower than experience-only and 4% lower than fixed-weight models, showing the influence of the conflicting information on participants' choices.

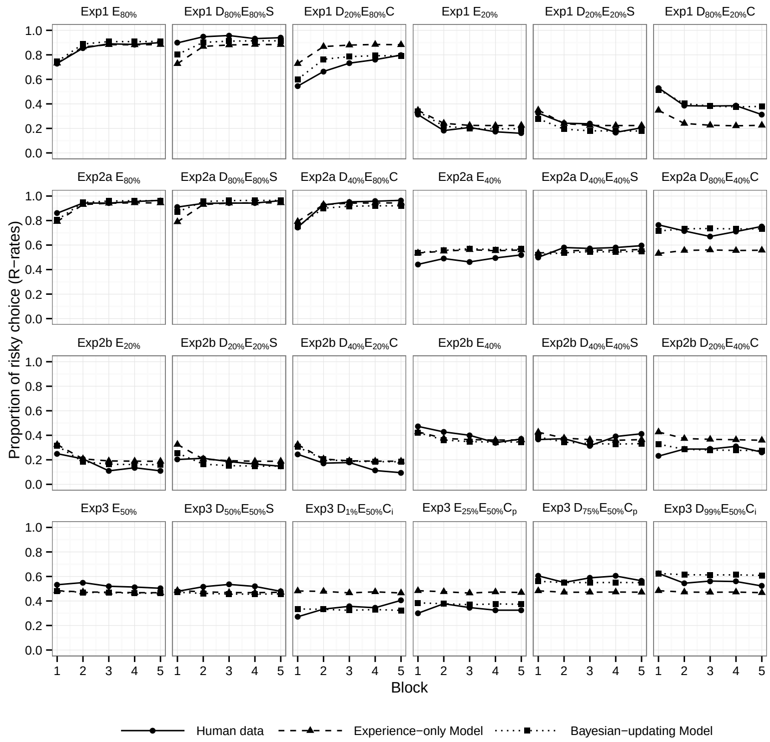

Figure 2.6: Comparison of observed human data, and two cognitive models, for each of the 24 experimental conditions. The description-experience model shown is the Bayesian-updating model, which is closer to human behaviour than the experience-only model, in particular in the conflicting conditions. Each row shows data for a separate experiment: from top to bottom, Experiments 1, 2a, 2b and 3. The subscripts after each D and E indicate the probability of rewards for the risky choice in that condition for the description and experience respectively, while S stands for same description, C for conflicting description, C_p for plausible conflict and C_i for implausible conflict.

The reduction in mean BIC values across all the DES conditions was lower at only 6% when comparing the Bayesian-updating with the experience-only model. This finding is similar to that found in Lejarraga and Gonzalez (2011), where a model that did not include descriptions fit the behavioural data relatively well, likely due to the equivalent information provided by both sources, and where the addition of description did not improve the model fit substantially. While I did observe a small improvement with the description-experience model, most of the reduction in BIC values came from the first few trials. Comparing the mean BIC values of the DES conditions between the Bayesian-updating and the experience-only models, there was a 39% reduction in the first five trials, a 9% reduction in the next

15 trials, and only a small 0.1% reduction in the last 80 trials. While the experiential information had to be learned over many trials, the descriptive information was available from the beginning of the task. Thus, if participants were relying only on experience, they should have chosen randomly in the first few trials until enough information was learned to steer their decisions, while if participants used descriptions, they could rely on that information to direct their earlier choices. The reduction in BIC values observed in the DES condition comes from this earlier availability of descriptive information which helps explain participants' choices in the initial trials. According to the behavioural and modelling results, participants chose in accordance with the descriptions available, showing the influence of the congruent information as well on their choices.

The largest improvement in model fit among the two description-experience models was in the DEC conditions of Experiment 3: The mean BIC values of the Bayesian-updating model were 8% lower than those in the fixed-weight model. In this experiment, I manipulated the plausibility of the descriptive information. While a fixed-weight model had a single ξ parameter for all conditions, the Bayesian-updating model was able to adapt to the different levels of plausibility of the information given (Figure 2.5 on page 70), and therefore did a better job to explain behaviour than the fixed-weight model. This is further evidence that, when the descriptions are implausible, participants give them lower weights in their decision processes, and vice versa, as predicted by the best fit Bayesian-updating model.

2.5 Discussion on conflicting descriptions

The aim of the experiments in this chapter was to identify one of the potential sources of the apparent contradictory results provided in earlier D+E research. By adding descriptions to traditional DfE paradigms, Lejarraga and Gonzalez (2011) and Lejarraga (2010) showed that these descriptions did not influence behaviour, and claimed that descriptions appear to be neglected, in comparison to Barron et al. (2008) and Rakow and Miler (2009) who showed that descriptions did significantly alter behaviour. The current experiments have shown that the choice behaviour of individuals exposed to a combination of description and experience is indeed influenced by both sources of information. In support of this, a cognitive model that included both the descriptive and experiential information fitted the observed behaviour better than a model that relied on experience alone. However, both observed behaviour and fitted cognitive models were significantly influenced by descriptions only when they provided novel information, as in Barron et al. (2008)'s earlier research, and not when the description and experience transmitted the same information, as in Lejarraga and Gonzalez (2011), thus reconciling their findings in a single study.

Previous research has shown that information provided by experience can overwhelm descriptions, making experience the preferred source of information (Jessup et al., 2008). When Lejarraga and Gonzalez (2011) exposed participants to a combination of descriptive and experiential information, they observed that choice behaviour could be explained by experience alone, as if the descriptive information was neglected. Based on the findings from the current experiment, an alternative explanation can be considered. Because in Lejarraga and Gonzalez's experiment both the description and the experience carried the same information, the description might not have been actively disregarded. Instead, it is possible that because it did not add any relevant information that could not be inferred from feedback, it did not lead to any observable differences in behaviour. In contrast, Barron et al. (2008) used a paradigm with partial descriptions, which alerted participants to the presence of a rare negative event. Since most of the participants in their study did not experience this rare event, the description provided novel information, which in turn influenced behaviour. The addition of descriptions that carried the same information as experience to the paradigm was therefore not enough to shift behaviour, especially under the assumption that it is ultimately the choice mechanism paradigm (e.g., one-choice/single-outcome versus repeated-choice/multiple-outcomes), not the type of information presentation, that generates the behavioural differences in the description-experience gap (Camilleri & Newell, 2013a; Jessup et al., 2008).

Instead of individuals neglecting specific sources of information, as proposed by Lejarraga and Gonzalez (2011), participants might integrate experience with prior beliefs about the outcomes of their choices (Rakow & Newell, 2010), such that different weights are given to each source of information, depending on their relevance. These experimental and computational results indicate that participants were influenced by descriptive information, albeit at a discounted rate in comparison to experience. Even though experiential information was dominant, the discounted influence of descriptive information still remained after many trials. The influence of experience appeared to grow steeply in the first few trials, but quickly reached an asymptotic level where it remained for the remainder of the task. Even after many trials, participants still behaved overall as if description received around a quarter of their decision weight, although this proportion was influenced by the plausibility of the information, with implausible descriptions receiving lower weights. Experiment 3 showed that the plausibility of the description has an effect on behaviour: only plausible descriptions influenced behaviour monotonically, with a reversal of the effect in the case of highly implausible descriptions.

Overall, the experiments in this chapter have shown that descriptions are taken into account and do influence behaviour when they provide novel information. Albeit this influence is heavily discounted and preference seems to be given to experience over descriptions, and the weight given to each source can itself be moderated by other factors such as the plausibility of the descriptions. One of the limitations of this first set of experiments was that they all relied on very simple tasks, with two alternatives providing only three outcomes in total. While these are commonly used in decision-making research, they are not very cognitively demanding and are relatively simple and quick to learn, either via experience or via descriptions. In the next chapter, I will explore the influence of descriptions using more complex tasks.

Chapter 3

Task complexity

Lejarraga (2010) and Lejarraga and Gonzalez (2011) explored situations where descriptions were made less attractive to participants by increasing their perceived complexity, therefore making them even harder to process cognitively while keeping the underlying experiential task unchanged. The authors showed that, by increasing the cognitive cost of processing descriptions, there was an increase in preference for experiences. Their findings indicate that even though cognitive theory may point towards individuals preferring experiences over descriptions, the strength of this preference may not necessarily be static, and could change according to the situation, such as task complexity manipulations. In Lejarraga and Gonzalez (2011), the complexity of the description was increased by changing from a relatively simple text (e.g., "win 4 cents with an 80% chance or win 0 cents otherwise.", p. 288) into a more complex one involving conditional probabilities but with the same underlying distribution of outcomes (e.g., "with an 80% chance you play lottery 1 and with a 20% chance you play lottery 2. Lottery 1 pays 4 cents with a 90% chance or 0 cents otherwise. Lottery 2 pays 4 cents with a 40% chance or 0 cents otherwise.", p. 288). A similar approach was taken by Lejarraga (2010, see Figure 1). One limitation of these previous studies, however, was that the researchers did not change the complexity of the task itself, only the complexity of the descriptions used to label the same underlying processes by using simpler or more complex notation.

While task complexity can be a subjective construct, and significantly dependent on individual psychological differences, such as interest, stimulation and cognitive capabilities, it is also related to certain underlying structural task characteristics that can be defined objectively (Campbell, 1988; Wood, 1986). One such structural dimension of complexity relates to the number of different alternatives from which participants can select, and the number of possible outcomes available from each alternative (Payne, 1976, 1982). The influence of task complexity on decision behaviour, using manipulations on the number of alternatives and potential outcomes, has been explored in decision-making research, albeit separately using DfD and DfE tasks. Payne (1976) observed that participants will engage in more simplified decision strategies for tasks of higher complexity, by manipulating the number of alternatives (2, 4, 8, or 12) and the number of attributes per alternative (4, 8, or 12). Kahneman and Tversky (1979) proposed that an increase in the number of alternatives from which to choose, due to the large amount of information available, leads to the engagement of simplified choice heuristics that ignore some of the information available, through an editing process, thus reducing cognitive demands. Thorngate (1980) demonstrated how the simulated performance of different "unsophisticated" decision heuristics reduces when applied to tasks with an increasing number of alternatives (2, 4, or 8), or an increasing number of unique potential outcomes from each alternative (also 2, 4, or 8), with a reduction in the selection of the best alternative according to their expected values. Similar results were found by Johnson and Payne (1985), who calculated that more complex tasks (again 2, 4, or 8 alternatives $\times$ outcomes) reduce the proportion of accurate choices made by simulated heuristics. According to Ert and Erev (2007), increasing the number of alternatives (either 2, 6, or 50) creates confusion, reduces performance, and leads to higher risk-seeking behaviour. Noguchi and Hills (2016) also observed higher risk-taking due to an increase in complexity of DfE tasks, measured by the number of alternatives available (either 2 or 32). Increasing the complexity of the task makes learning slower, more costly, and more cognitively demanding (Ashby et al., 2017; Fasolo, Hertwig, Huber, & Ludwig, 2009; Frey, Mata, & Hertwig, 2015).

Increasing the number of alternatives and outcomes increases the entropy of the task, which can be associated with higher task complexity (Fasolo et al., 2009). Entropy is an objective measure that has been used to quantify task complexity, based on information theory, with higher entropy associated with higher complexity and increased information load (Swait & Adamowicz, 2001). Research on information load has uncovered a non-monotonic inverted U-shape relationship between amount of information available and decision accuracy (Eppler & Mengis, 2004; Hwang & Lin, 1999). It has been shown empirically that an increase in information aids the decision-making process initially, but only up to a certain point, after which any additional information is actually detrimental and reduces the quality of the decisions (Jacoby, Speller, & Berning, 1974; Jacoby, Speller, & Kohn, 1974). This is commonly called "information overload", and describes the negative effects of receiving too much information. Potential causes for information overload include complexity of information provided, number of items of information and number of alternatives, among others (see Eppler & Mengis, 2004, p. 332). In simpler tasks, individuals tend to use full processing strategies, but in more complex task, some information is discarded and heuristics might be employed to reduce cognitive effort (Paquette & Kida, 1988; Streufert & Driver, 1965). This decrease in the use of available information leads to decisions of lower quality (Chewning & Harrell, 1990; Payne & Braunstein, 1978). Similar U-shaped patterns peaking for medium complexity tasks have been found across other dimensions, such as choice satisfaction (Reutskaja & Hogarth, 2009), purchasing intentions (Shah & Wolford, 2007), the ability to accurately assess values (Keller & Staelin, 1987), the extent of information processing (Paul & Nazareth, 2010), and overall effort allocation (Swait & Adamowicz, 2001).

Despite some research on the influences of task complexity on behaviour, such as the ones mentioned above, the overwhelmingly majority of the research in decision-making uses relatively simple tasks (Hertwig & Erev, 2009; Rakow & Newell, 2010). The most commonly used decision-making environments involve two alternatives from which participants can choose, with each alternative providing two potential outcomes, for a total of four outcomes. Even simpler tasks have been employed, when one such alternative only provides one potential outcome, a sure option that always returns the same amount. In such tasks, there are only three potential outcomes (e.g., the experiments in the previous chapter in this dissertation). Building upon these simple tasks commonly used to study general decision-making, the extant D+E research has also exclusively focused on similarly simple tasks, for example by employing paradigms with two alternatives with a total of three (Barron et al., 2008; Jessup et al., 2008; Lejarraga & Gonzalez, 2011) or four (Rakow & Miler, 2009; Yechiam & Busemeyer, 2006) potential outcomes.

This canonical preference for simple tasks, with their associated low costs of learning from experience, might be the driver behind the limited influence of descriptions on D+E research so far, and might explain why participants have shown preference for experiences over descriptions. This unexplored space of task complexity is likely to be one of the moderators of the influence that descriptions and experiences have on D+E tasks. In simple tasks, participants quickly learn to identify the structure of the environment experientially, differences in cognitive processing efforts are likely to be negligible, and no additional information is needed or desired. Descriptions, if available, are not very useful and do not help participants. One of the reasons for the lack of observable behavioural differences in the research by Lejarraga and Gonzalez (2011) is likely due to the simplicity of the task used. Complex tasks, however, are more difficult and take longer to learn experientially. This should make descriptions relatively more attractive than before, as relying on descriptions should provide an advantage to participants by giving them additional information that reduces learning time by lowering the need to explore the environment experientially. Higher task complexity should lead to situations in which engaging the extra processing effort associated with descriptions becomes cost-efficient. Therefore, an increase in task complexity should lead to an increase in the influence of descriptions on behaviour.

Furthermore, this relationship between task complexity and influence of descriptions need not be monotonic. Task complexity is also closely linked to information processing, with more complex tasks being defined as those in which there is increased information load, diversity and rate of change (Campbell, 1988). More alternatives (load) and more outcomes (diversity) require more attention, as they need to be evaluated, increasing information load, which in turn increases task complexity. Consequently, the influence of descriptions is likely to reduce in very complex tasks, after peaking in medium complexity tasks. If the task becomes too complex, then the descriptions required to summarise the task also become overly complex. The excess of information available, both in experience and description, should lead to information overload. Descriptions might also become too unwieldy to process cognitively, reducing their attractiveness as a source of information. Fantino and Navarro (2012) found that descriptions can lead to suboptimal choices if they are too complex to be understood and employed correctly. The maximum influence of descriptions could be expected in tasks of medium complexity, at the point where performance starts to suffer but descriptions are still not too complex. It is at this point that descriptions should be able to provide the most assistance. Complex tasks, taking longer to learn experientially, might also lead to an artificially perceived conflict between the two sources of information. For example, in a task with several alternatives from which an individual can choose, it would take substantially longer to eliminate noise, reduce variance and establish a good unbiased overview of the choice environment, in comparison to a simple task with only two choices. In such situations, descriptions might provide novel information.

The second set of experiments in this dissertation will explore how task complexity moderates the integration of descriptions and experience as part of the decision-making process. I predict that the addition of congruent descriptions will influence behaviour in more complex D+E tasks, even when both in theory provide the same underlying information, unlike what was observed before using very simple tasks by Lejarraga and Gonzalez (2011) and in Chapter 2 of this dissertation. For this purpose, I will initially introduce congruent descriptions, concurrently to experiences, to a relatively complex task that has been widely explored in DfE research before, but never before with descriptions, the Iowa Gambling Task (IGT: Bechara et al., 1994). The IGT was originally designed to mimic typical real-life DfE situations. It attempts to do so by using a relatively complex choice environment without descriptions. The IGT's complexity derives from the larger number of alternatives and outcomes, when compared to more traditional decision-making tasks. While most research tends to use sets of two choices with two outcomes each, the IGT provides participants with four choices, in the forms of decks of cards, and each deck of cards has between two and six potential outcomes. The IGT increases the traditional total of four possible outcomes in a typical DfE task to fourteen possible outcomes, making it much more difficult to learn and to find the most attractive alternatives. In the traditional IGT, all learning is via experience alone, in the form of feedback after each selection, with no additional information about the values and frequencies of the outcomes for each deck being provided. I believe that the IGT, which is based around a peculiarly complex paradigm, can be better exploited by individuals with the benefit provided by the presence of descriptions. In this case, I expect congruent descriptions to influence behaviour in this more complex D+E task, speeding up learning and helping participants to perform better on the task, with higher financial rewards. I also expect participants to gather less information experientially, by exploring less when descriptions are available, as the added descriptions will provide additional information to participants, reducing uncertainty.

In addition to exploring the influence of descriptions in one such complex experiential task, the IGT, I will also directly manipulate task complexity, from simple to complex, in a later experiment. Task complexity will be experimentally controlled across two separate dimensions, by altering the number of alternatives and the number of outcomes from each alternative. I will fine tune the experiment to create simple tasks very similar to ones used before in D+E research (Lejarraga & Gonzalez, 2011), tasks of complexity similar to the IGT, as well as considerably more complex tasks that have not been explored in this type of research before. By directly controlling for task complexity through these experimental manipulations, I expect to replicate the results from previous research in simple tasks, while observing the non-monotonic inverted U-shaped relationship proposed above in the complex tasks, with the beneficial influence of descriptions peaking in medium-complexity tasks, and deteriorating when tasks become overly complex, thus verifying the boundaries of the effect and the influence of complexity on descriptions.

3.1 Experiment 4

The first experiment in this chapter introduced congruent descriptions to a relatively complex DfE task, the Iowa Gambling Task (Bechara et al., 1994). Given the high complexity of this task, which is difficult and slow to learn through experience alone, I expected the addition of descriptions to help participants perform better, finding the more attractive alternatives earlier, and earning higher financial rewards, in comparison to the traditional task without descriptions.

3.1.1 Method

Design

Experiment 4 was based on the original Iowa Gambling Task, with the addition of descriptive information for half of the participants in a new experimental manipulation, creating a described IGT. The experiment followed a two-way between-subjects design controlling for the presence or absence of descriptions: in the experience-only (E) condition, participants relied on experience alone (in the form of feedback after each trial) to learn about the options available to them, without any descriptions; and in the description-plus-experience (DE) condition, participants were shown a full description of the distribution of outcomes available for each option, in addition to the experiential feedback after each trial.

Participants

100 participants were recruited on-line using Amazon's Mechanical Turk service (47 females; age: M=31.0 years, SD=10.5 years), half in each experimental condition. Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment, with an average bonus of US$ 0.55 (SD=US$ 0.28).[1]

[1] Participants in the DE condition received a significantly higher bonus than those in the E condition, according to a Wilcoxon rank sum test (DE=US$0.68, E=US$0.41, W=1919, z=4.47, $p < .001$).

Task

The task closely followed the original IGT (Bechara et al., 1994). The instructions closely matched the original wording apart from changes needed for the computerised on-line delivery of this version of the task (Chiu & Lin, 2007). Participants were presented with four decks of cards (decks A, B, C, and D), side by side on the screen, with the backs of the cards displayed and their faces hidden (Figure 3.1). The faces of the cards provided the feedback after each selection, with the number of points earned or lost associated with each individual card. The naming of the decks given here was used for analysis only and not shown to participants. The order of the decks from left to right was randomised for each participant, as well as the pattern on the back of each deck. Choices were made using the mouse. To avoid rapid sequential clicking of the same choice repeatedly, participants were required to move the mouse cursor to a button at the bottom of the screen between selections.

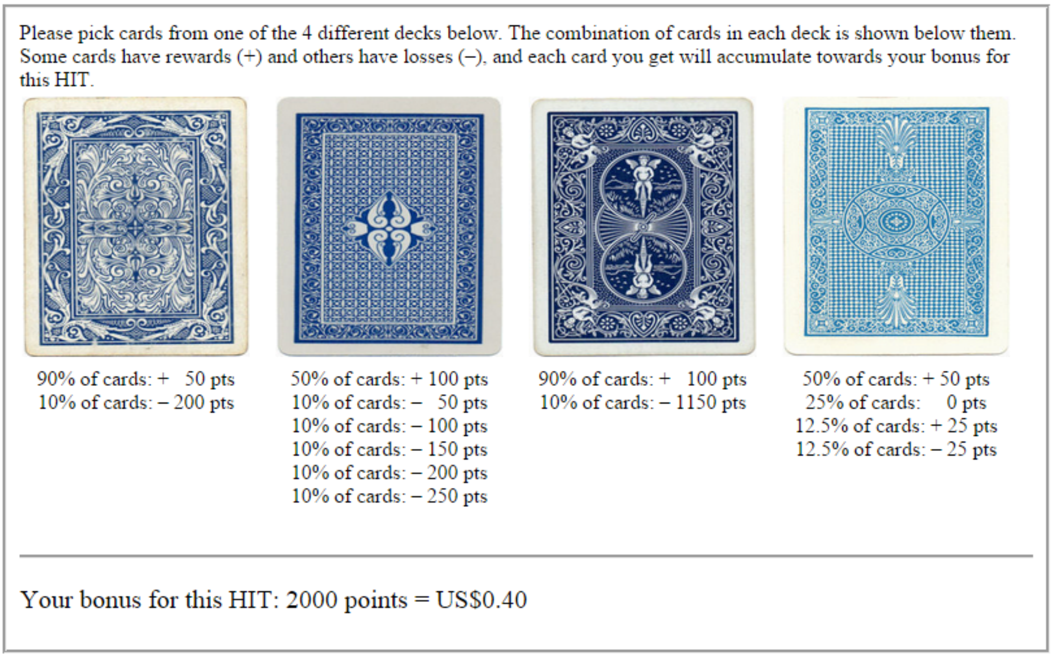

Figure 3.1: Screenshot of the first trial of Experiment 4 in the description-plus-experience (DE) condition, showing the described IGT, with a random presentation of the decks in the order D, A, B, C from left to right. The patterns on the backs of each deck were also randomised. In the experience-only (E) condition, the space underneath the cards was left blank, with no descriptions shown, and the second sentence in the title regarding the combination of cards in each deck was replaced with "Each deck contains a different combination of cards". HIT stands for Human Intelligence Task, a terminology employed by Amazon's Mechanical Turk service to designate individual tasks.

Participants' choices were financially consequent and accumulated towards their final pay. Participants started the task with 2000 points and points earned or lost after each selection were added to or deducted from their total. Points were converted to money at a rate of US$ 0.20/1000 points. Accumulated amounts in points and U.S. dollars were shown on-screen and updated after each choice was made.

The schedule of outcomes from each deck was the same as in the original IGT: The order of the cards within each deck was not random but instead followed the fixed order given in the original task, with a repeating pre-defined sequential pattern of 40 cards for each deck (Bechara et al., 1994, Figure 1). In contrast to the original IGT, which showed rewards and losses separately for each card (e.g. "You have won 100 points, but you also have lost 150 points"), I opted to summarise outcomes as single net values (e.g. "–50 points"). This made the task simpler to describe, and circumvented the predictability of rewards associated with the original study (Steingroever, Wetzels, Horstmann, Neumann, & Wagenmakers, 2013). Decks A and B have a negative expected value of –25 points for each card, while decks C and D have a positive expected value of +25 points for each card (Table 3.1). Hence decks A and B are referred to as the *disadvantageous* decks, and decks C and D are the *advantageous* decks. In order to maximise their bonus, participants have to select more often from the advantageous decks and avoid the disadvantageous ones.

In the experience-only (E) condition, participants did not receive any further

Table 3.1: Actual card composition and wording of descriptions shown underneath each deck in Experiment 4. The expected value for each individual card in decks A and B was –25 points and in decks C and D was +25 points.

Experience-only condition (E), *N*=50			
Deck A	Deck B	Deck C	Deck D
(blank)	(blank)	(blank)	(blank)
Description-plus-Experience condition (DE), *N*=50			
Deck A	Deck B	Deck C	Deck D
50% of cards: +100 pts 10% of cards: -50 pts 10% of cards: -100 pts 10% of cards: -150 pts 10% of cards: -200 pts 10% of cards: -250 pts	90% of cards: +100 pts 10% of cards:-1150 pts	50% of cards: +50 pts 25% of cards: 0 pts 12.5% of cards:+25 pts 12.5% of cards: -25 pts	90% of cards: +50 pts 10% of cards: -200 pts

information about the decks, and had to learn which decks were the advantageous ones via the feedback provided after each selection, in what was a close replication of the traditional IGT paradigm. In the description-plus-experience (DE) condition, a description of the cards contained in each deck (Table 3.1) was permanently displayed underneath the relevant deck, across all trials (Figure 3.1), in addition to the feedback provided, resulting in a described IGT. After each selection, participants were shown only the outcome in points of the card from the deck they selected (i.e., partial feedback), and the card selected was replaced at the end of that deck, with no changes to the card order in the non-selected decks. Therefore participants learned only about the deck they selected on each trial, with no new feedback information for unselected decks. Participants were not told beforehand how many cards they would get to choose, and instead were instructed to choose cards from decks repeatedly until told to stop, which was after 100 choices. The task was self-paced and was completed on average in 8.30 minutes (*SD*=4.35).

3.1.2 Results

Selections from advantageous decks

The main dependent variable was the frequency of cards selected from the advantageous decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 3.2A). They were analysed with a linear mixed-effects model using the lme4 package (Bates et al., 2014) and post-hoc analyses using the lsmeans package (Lenth, 2016), with Tukey adjustments, in R (R Core Team, 2014). The between-subjects conditions were the presence or absence of descriptions (E vs. DE). The within-subjects conditions were the blocks of 20 choices each. The model also contained a random intercept for each participant, and random slopes across blocks for each participant.

The main effect of the presence of descriptions was significant, with participants selecting from advantageous decks significantly more frequently in the DE condition than in the E condition, across all blocks of 20 trials on average, with a large effect size (DE=74.14%, E=55.90%, $\chi^2(1) = 27.78, p < .001, d = 0.89$). The presence of descriptions helped participants identify the advantageous decks and

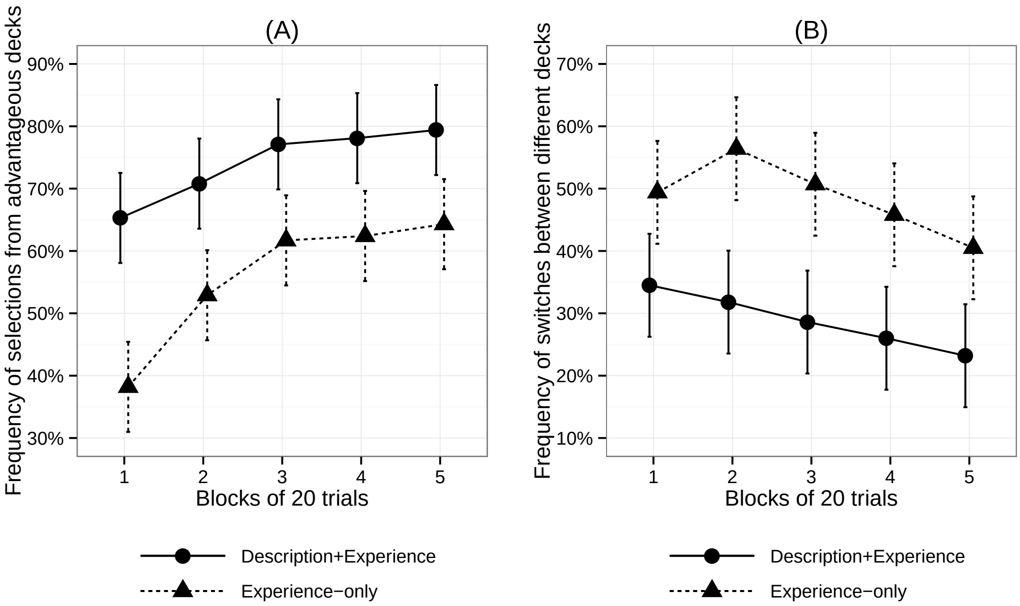

Figure 3.2: Experiment 4. (A) Evolution of the average frequency of selection of advantageous decks (Decks C + D), as a percentage of total for each block. (B) Evolution of the average frequency of switches between different decks, as a percentage of total for each block. Each block contains 20 trials. Error bars represent the 95% confidence interval around the mean. Experience-only results are for the traditional IGT, while the description+experience results are for the described IGT.

select from them more often.

The main effect of block was also significant, with a positive slope ($b = 0.035, \chi^2(1) = 38.10, p < .001$), indicating a higher selection of advantageous decks over time. Post-hoc analyses with Tukey adjustments showed a significantly higher selection of advantageous decks in the last block compared to the first block (Block1=51.75%, Block5=71.85%, $t(392) = 8.01, p < .001, d = 1.13$). This increase in selections from the advantageous decks over time is a result of participants gradually learning the task, and being able to identify which decks are the advantageous ones, as well as learning to avoid the disadvantageous ones, in order to extract higher rewards from the task. Most of the learning however appears in the first two blocks (Block1 vs. Block2: $t(392) = 4.02, p < .001$; Block2 vs. Block3: $t(392) = 3.01, p = .02$), with no significant differences when applying sequential pairwise comparisons between the last three blocks (ps>.97). This early stabilisation of choice preferences is consistent with previous research (Bechara et al., 1997; Ert & Erev, 2007).

The interaction between presence of description and block was not significant ($\chi^2(1) = 2.77, p = .10$), suggesting that the selection rate of advantageous decks across time is similar between the two conditions. In order to exclude this effect of learning, and to focus on stable behaviour, a post-hoc analysis with Tukey adjustments was performed comparing the two conditions at the last block: in the last 20 trials, the presence of descriptions led to a 23% significant increase in the selection of advantageous decks, with a medium effect size (Block5 only: DE=79.40%, E=64.30%, $t(230.1) = 2.91, p = .004, d = 0.58$). Even after participants had had a chance to learn about the task experientially, the presence of descriptions still significantly helped them select from the advantageous decks more often.

Switching rates

In addition to the frequency of deck selection, the switching rates between decks were also analysed (Figure 3.2B). A selection was classified as a switch every time a card was picked from a different deck to that from which the previous card had been selected. The same model structure was used as in the previous analysis.

The main effect of description was significant, with switching rates being 40% lower in the DE condition compared to the E condition, in each block of 20 trials (DE=28.82%, E=48.56%, $\chi^2(1) = 14.48, p < .001, d = 0.76$). Overall participants seemed more uncertain in the E condition and explored more among the different decks, while in the DE condition they exploited more their preferred options, switching less often.

The effect of block was also significant, with a negative slope ($b = -0.028, \chi^2(1) = 19.15, p < .001$), indicating a reduction in switching rates over time. Post-hoc analyses with Tukey adjustments showed that switching rates were lower in the last block compared to the first block (Block1=41.95%, Block5=31.85%, $t(392) = 4.51, p < .001, d = 0.64$). As participants gathered more information from the task, they explored less and exploited their preferred choices more.

The interaction between presence of description and block was not significant ($p > .20$), suggesting that the reduction in switching rates was the same across the conditions and were not influenced by the descriptions. A post-hoc analysis with Tukey adjustments between the two description conditions at the last block showed that participants switched 57% less often in the DE condition, a significant difference (Block5 only: DE=23.20%, E=40.50%, $t(161.4) = 2.93, p = .004, d = 0.59$).

3.1.3 Discussion

When presented with descriptions in the DE condition, participants selected from the advantageous decks more often than when they had to rely on experience without descriptions in the E condition. Therefore, the presence of descriptions influenced behaviour and helped participants to find the rewarding cards and avoid the loss-generating ones, leading to 66% higher financial bonuses. This difference in behaviour is indicative of participants integrating the descriptive information into their decision making processes, as descriptions informed participants about the potential outcomes of their choices, and could be used to identify the advantageous decks. Furthermore, the behaviour observed in the E condition (without descriptions, therefore a replication of the traditional IGT paradigm) was similar to that found in previous studies using this task. Frequency of advantageous (good) deck selection across all trials in this experiment, *M*=56%, was similar to a weighted mean from a meta-analysis of 39 studies covering 1,427 healthy participants, *M*=57% (Steingroever et al., 2013, Table 5).

Better performance across all trials can be partially explained by the availability of additional descriptive information from the beginning of the task in the DE condition, which provided an advantage to participants, and could be used to make an initial informed choice among the available options. Those in the E condition lacked any information about the composition of cards within each deck, so their first selection was necessarily a random uninformed choice between the four decks available. This advantage can be removed by comparing the behaviour after it has stabilised. Even after many trials, in the last block of 20 trials, deck selection still differed significantly between the E and DE conditions, with participants in the DE condition choosing 23% more often from the advantageous decks than in the E condition. At this point, choice behaviour had mostly stabilised in their preferred choices.

I also expected exploration to reduce when descriptions were available, and this was observed with lower switching rates in the DE experimental condition. Switching rates can be seen as proxies for exploration (Ert, Erev, & Roth, 2011), as individuals who exclusively exploit their preferred option would not need to switch between the options. In decisions from experience, without descriptions, participants must learn about the decision environment through exploration and feedback. If descriptions are present, by providing additional information about the available options, they offer an alternative avenue for comparing them and finding the most attractive one, reducing the need for exploration. Exploration still remains however, as uncertainty is not fully eliminated, and participants still need to confirm that descriptions are true throughout the task. In addition, participants might be exploring to avoid the boredom of selecting the same alternative repeatedly, or to select a mixed strategy across their preferred alternatives.

Overall, the presence of descriptions influenced behaviour in a complex task such as the described IGT. However, I was concerned that the usage of a pre-determined fixed schedule of outcomes, as in the original IGT, was not being truly represented by the descriptions. While the descriptions were a true representation of the frequency of the cards within each deck, there was no mention of the actual sequence in which the cards appeared, which might have led participants to believe that cards were shuffled and their order was random. While the original pre-determined sequence is one of the many potential sequences in which the cards would appear if the outcomes were truly randomised, the descriptions could have also provided participants with the actual sequence of cards, since they were previously known. In order to make descriptions a truer representation of the experience, in the next experiment I will replace the fixed schedule with a pseudo-randomised ordering of cards within each deck.

3.2 Experiment 5

3.2.1 Method

Design

The aim of Experiment 5 was to replicate the results found in Experiment 4, and confirm that the presence of descriptions influence behaviour in complex tasks with congruent descriptions. As before, there were two conditions: experience-only (E) and description-plus-experience (DE). The only alteration to the paradigm was in the ordering of cards within each deck. Instead of using the original pre-determined sequence of cards from Bechara et al. (1994), I employed a pseudo-randomised approach within blocks of 40. This approach should make the experience a truer representation of the descriptions (see also footnote 1 in Chapter 1). Since the actual order of cards was not known until the computer randomised it, information about the sequence could not have been provided in the descriptions to participants.

Participants

100 participants were recruited on-line using Amazon's Mechanical Turk service (42 females; age: *M*=36.7 years, *SD*=11.8 years), 49 in the experience-only (E) condition and 51 in description-plus-experience (DE). Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment, with an average bonus of US$ 0.52 (*SD*=US$ 0.33).[2]

Task

The task was a replication of Experiment 4, and closely followed the original IGT (Bechara et al., 1994). Participants in the experience-only (E) condition did not receive any additional information, while participants in the description-plus-

[2]Participants in the DE condition received a significantly 78% higher bonus than participants in the E condition, according to an asymptotic Wilcoxon rank sum test (DE=US$0.66, E=US$0.37, W=1922, z=4.49, $p < .001$). In addition, there was no significant difference in bonus between Experiments 1 and 2 (W=5122.5, z=0.72, p=.76)).

experience (DE) condition were shown a description of cards underneath each deck. The only alteration from Experiment 4 was in the ordering of cards within each deck. In Experiment 4, the order was always fixed and known beforehand, following the set sequence from the original IGT study, with each participant observing the same ordering of cards. In Experiment 5, this pre-determined fixed ordering was abandoned and replaced with a pseudo-randomised approach (see Section 1.2.1). The frequency of cards within each deck was the same as in Experiment 4 (Table 3.1), but their order was shuffled. Each participant observed a newly randomised ordering of cards. Pseudo-randomisation was used to ensure that within each set of 40 cards, participants experienced the same frequency of cards as that in the descriptions, but in a random order. For example, for Deck D, there would always be 36 cards of +50 and 4 cards of –200 points within each set of consecutive 40 cards, with a newly randomised order in each set. The actual sequence was not known until the computer randomised it. The reason for choosing 40 cards is to replicate the original IGT which was also based on sets of 40 cards. This approach is similar to using a deck of 40 cards, which is initially shuffled and revealed in order without replacement. Once all 40 cards of a deck have been shown, the computer would re-shuffle and start again. The task was self-paced and was completed on average in 8.82 minutes (*SD*=4.12).

3.2.2 Results

As before, the main dependent variable was the frequency of cards selected from the advantageous decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 3.3). They were analysed with a linear mixed-effects model as in Experiment 4, with the same fixed and random components.

The main effect of the presence of descriptions was significant, with participants selecting from advantageous decks significantly more frequently in the DE condition than in the E condition, across all block of 20 trials on average, with a large effect size (DE=74.67%, E=54.02%, $\chi^2(1) = 18.49, p < .001, d = 0.82$). The presence of descriptions helped participants identify the advantageous decks and

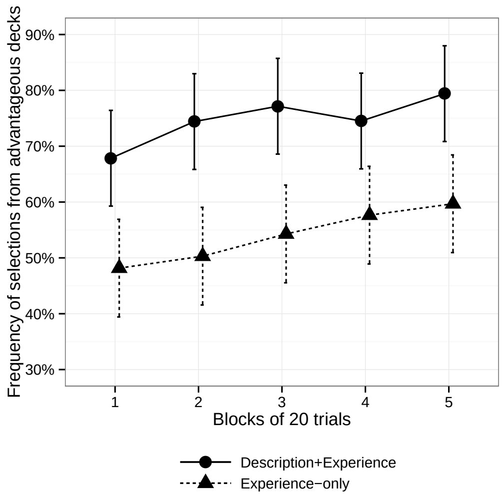

Figure 3.3: Experiment 5. Evolution of the average frequency of selection of advantageous decks (Decks C + D), as a percentage of total for each block. Each block contains 20 trials. Error bars represent the 95% confidence interval around the mean. In Experiment 5 the order of the cards within each deck was pseudo-randomised, the only change in the task in comparison to Experiment 4.

select from them more often.

The main effect of block was also significant, with a positive slope ($b = 0.023, \chi^2(1) = 9.69, p = .002$), with an increase in selection of advantageous decks over time (Block1=58.00%, Block5=69.55%, $t(392) = 4.03, p < .01, d = 0.57$), albeit with a smaller effect size when compared to Experiment 4. The interaction between presence of description and block was not significant ($\chi^2(1) = 0.17, p = .68$).

A post-hoc analysis with Tukey adjustments was also performed comparing the two conditions at the last block: in the last 20 trials, the presence of descriptions led to a 33% significantly higher selection of advantageous decks, with a medium effect size (Block5 only: DE=79.41%, E=59.69%, $t(211.3) = 3.71, p = .004, d = 0.63$). Even after participants had had a chance to learn about the task experientially, the presence of descriptions still significantly helped them select from the advantageous decks more often.

3.2.3 Discussion

The findings from Experiment 4 were replicated in Experiment 5, with very similar effect sizes for the behavioural impact of the presence of descriptions. Participants presented with descriptions selected from advantageous decks more often, and obtained higher financial bonuses.[3] In this experiment the descriptions were a truer representation of the experience, since the actual order of cards was not previously known, until the computer shuffled and randomised them. Only the frequency of the cards within each deck was known, but not their ordering. Because of pseudo-randomisation, the frequency described was an exact representation of the actual experience, within each set of 40 cards revealed by participants. While in Experiment 4 a true description could have included the actual sequence of cards, this was not possible in Experiment 5.

Across Experiments 4 and 5, which were complex decisions-from-experience tasks based around the IGT, the presence of congruent descriptions influenced behaviour and helped participants, regardless of whether the sequence of card was pseudo-randomised or followed the original fixed schedule. These findings initially appear to go against previous studies using simpler tasks that have shown no influence of congruent descriptions on behaviour, such as the experiments in the previous chapter in this dissertation, as well as Lejarraga and Gonzalez (2011). I propose that it was the increased complexity of the task, with its four options and multiple outcomes, that led to descriptions being taken into account by participants in these two experiments, while in previous studies the tasks were simpler, using two options with fewer outcomes. In the next experiment, I sought to analyse how task complexity influences the behavioural impact of descriptions in a more controlled experimental set-up, by creating a task that reconciles my results in Experiments 4 and 5 with those in previous research on description-plus-experience.

[3]The data from Experiments 4 and 5 were combined into one single analysis to evaluate the main effect of Experiment as a proxy for the ordering of cards. There was no significant effect of Experiment ($\chi^2(1) = 0.09, p = .76$). There was no significant interaction between Experiment and presence of descriptions ($\chi^2(1) = 0.02, p = .88$). Therefore, the use of a pre-determined schedule or a pseudo-randomised order of cards had no influence on the selection of advantageous decks overall (Exp4=65%, Exp5=64%), or on the impact of descriptions (E: Exp4=56%, Exp5=54%, $t(196) = 0.41, p = .68$; DE: Exp4=74%, Exp5=75%, $t(196) = 0.12, p = .91$). The combined impact of descriptions across the two experiments was significant, with a large effect size (E=55%, DE=74%, $t(196) = 5.98, p < .001, d = 0.85$).

3.3 Experiment 6

3.3.1 Method

Design

In the first two experiments in this chapter I observed the influence of descriptions in a relatively complex task, and considered the contrast between my results and those obtained in previous research using simpler tasks. In the current experiment I manipulated complexity directly. The aim was to start with simple tasks, similar to those used in earlier description-plus-experience research, and then to increase the complexity within the same experimental framework, therefore directly observing how task complexity moderates the influence of descriptions on behaviour. To achieve this, Experiment 5 was modified by manipulating the complexity of the task while maintaining the same basic set-up of selecting cards from different decks with and without descriptions throughout. The task followed a 3 $\times$ 3 $\times$ 2 between-subjects experimental design. I controlled task complexity across two different dimensions: the number of decks of cards available for participants to choose, which was 2, 4, or 6; and the number of potential outcomes within each choice (i.e., the number of different types of card that composed each deck), which was also 2, 4, or 6. This created a matrix of 3 $\times$ 3 tasks (see Figure 3.5 on page 101). Within each cell of this matrix, participants were given either an experience-only task (E), with no descriptions, or a description-plus-experience task (DE), with descriptions.

When comparing the complexity of this experiment with the previous ones, Experiments 4 and 5 closely matches the central cell of the new experimental matrix of Experiment 6. The previous experiments, based on the IGT, had a total of 14 potential outcomes across its 4 different choices, each choice having an average of 3.5 outcomes. The central cell in the new experiment has a total of 16 outcomes split across 4 different choices with 4 outcomes each (see Table 3.2). Therefore the new experiment creates both a simpler task (2 choices $\times$ 2 outcomes) and a more complex task (6 choices $\times$ 6 outcomes) in comparison to Experiments 4 and 5 (4 choices $\times$ 3.5 outcomes on average). The simplest task of the new experiment, with 2 choices and 2 outcomes is similar to previous research in the field of description-plus-experience, such as the ones in Chapter 2. The most complex task, with 6 choices and 6 outcomes is considerably more complex than what has been researched before in this field.

The reason for expanding the experiment into highly complex tasks is because I believe that the relationship between task complexity and the influence of descriptions is non-monotonic. As observed in earlier research, in simple tasks, descriptions have no perceptible influence on behaviour (see Chapter 2). In Experiments 4 and 5, I noticed that by increasing task complexity, descriptions provide useful information to participants and assist behaviour, because the task is now more complex and learning experientially is no longer trivial. However I also believe that when the task becomes overly complex, the descriptive information becomes extensively verbose in order to explain it, and therefore also difficult to decipher. In this case, I expected overall performance to suffer in the experience-only condition, and also did not expect much improvement due to the addition of description.

Participants

540 participants were recruited on-line using Amazon's Mechanical Turk service (239 females; age: *M*=33.2 years, *SD*=10.1 years), 30 in each experimental condition. Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment, with an average bonus of US$ 0.44 (*SD*=US$ 0.32).[4]

Task

The task closely followed Experiment 5, with randomised outcomes, apart from changing the number of options available and number of potential outcomes from each option. Each participant was allocated to a single experimental condition

[4]Participants in the DE conditions received a significantly higher bonus than participants in the E condition, according to an asymptotic Wilcoxon rank sum test (DE=US$0.50, E=US$0.39, W=44825, z=4.47, $p < .001$).

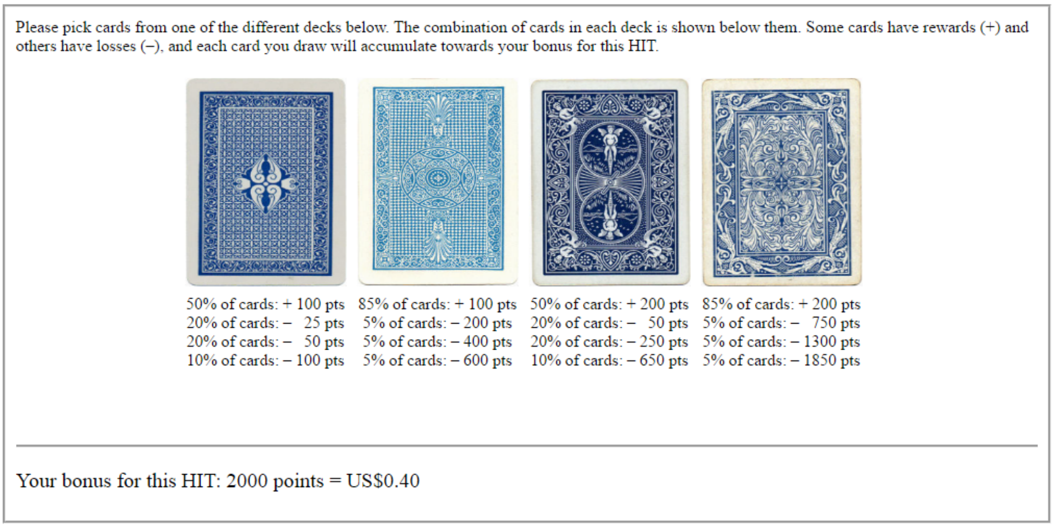

Figure 3.4: Screenshot of the first trial of Experiment 6 in the 4-deck × 4-outcome condition with description-plus-experience (DE). In the 2-deck condition only the two middle card positions were used, while in the 6-deck condition an additional two cards were shown in the leftmost and rightmost empty spaces. The order of the decks and the patterns on the back of each deck were randomised. In this example, from left to right, the decks are D, F, A, and C. Descriptions were not shown in the E condition, and the second sentence in the title was also changed to "Each deck contains a different combination of cards".

across number of decks, number of outcomes, and presence of descriptions. Participants were presented with either 2, 4, or 6 decks of cards. A total of 6 decks of cards were created, named A to F for analysis (Table 3.2). Participants in the 2-deck condition were presented with decks A and D; participants in the 4-deck condition were presented with decks A, B, D, and E; and all decks were presented to participants in the 6-deck condition. The order of presentation of the decks was randomised, as well as the patterns on the back of the decks. Decks of cards were shown side by side, with the 2- and 4-deck conditions using only the central 2 and 4 positions, respectively (Figure 3.4). To ensure that participants could see all the decks at the same time, the size of the window used was recorded, and no participant had a window size smaller than the minimum required.

Each deck had either 2, 4, or 6 potential outcomes according to the experimental condition. In contrast to Experiment 5 in which each deck had a different number of outcomes, resembling the payoff schedule of the original IGT, all decks in Experiment 6 had the same number of outcomes within each condition (either 2, 4, or 6 outcomes). The outcomes within each deck were adapted from Chiu and

Table 3.2: Schedule of outcomes used in Experiment 6, written as pairs of "probability: points". In the 2-choice conditions, decks A and D were used. In the 4-choice conditions, decks A, C, D and F were used. In the 6-choice conditions, all decks were used. The actual description text presented to participants followed that of Experiments 4 and 5, in the form of "–% of cards: – pts". The expected value for each individual card in decks A, B, and C was –25 points, and in decks D, E, and F it was +25 points.

2 outcomes					
Deck A	Deck B	Deck C	Deck D	Deck E	Deck F
50%: +200	70%: +200	85%: +200	50%: +100	70%: +100	85%: +100
50%: –250	30%: –550	15%: –1300	50%: –50	30%: –150	15%: –400
4 outcomes					
Deck A	Deck B	Deck C	Deck D	Deck E	Deck F
50%: +200	70%: +200	85%: +200	50%: +100	70%: +100	85%: +100
20%: –50	10%: –150	5%: –750	20%: –25	10%: –50	5%: –200
20%: –250	10%: –550	5%: –1300	20%: –50	10%: –150	5%: –400
10%: –650	10%: –950	5%: –1850	10%: –100	10%: –250	5%: –600
6 outcomes					
Deck A	Deck B	Deck C	Deck D	Deck E	Deck F
20%: +255	30%: +250	40%: +280	20%: +135	30%: +125	40%: +140
20%: +200	30%: +200	25%: +200	20%: +100	30%: +100	25%: +100
10%: +90	10%: +50	20%: +40	10%: +30	10%: +25	20%: +20
20%: –50	10%: –150	5%: –750	20%: –25	10%: –50	5%: –200
20%: –250	10%: –550	5%: –1300	20%: –50	10%: –150	5%: –400
10%: –650	10%: –950	5%: –1850	10%: –100	10%: –250	5%: –600

Lin (2007) and are shown in Table 3.2. Decks A, B and C have a negative expected value of -25 points for each card, while decks D, E and F have a positive expected value of +25 points for each card. Hence decks A, B and C were considered the disadvantageous decks, and decks D, E and F were considered the advantageous decks. The schedule of outcomes was pseudo-randomised in sets of 20: within each 20 cards, there was a full representation of all cards for that deck, in the correct proportions, but in randomised order. For example, for Deck A with 2 outcomes, there would always be 10 cards of +200 and 10 cards of -250 within each set of consecutive 20 cards, with a newly randomised order in each set. Participants were either given description-plus-experience or experience-only. Each participant made 100 selections and the task was completed on average in 8.4 minutes (*SD*=4.1).

3.3.2 Results

The main dependent variable was the frequency of cards selected from the advantageous decks, calculated as the total number of cards selected from decks D, E and F for each sequential block of 20 choices (Figure 3.5). It was analysed with a linear mixed-effects model as in Experiments 4 and 5. The between-subjects conditions were the number of decks (2, 4, or 6), the number of outcomes in each deck (2, 4, or 6), and the presence or absence of descriptions (E or DE). The within-subjects conditions were the blocks of 20 choices each, as categorical variables. The model also contained a random intercept and a random block slope for each participant. Post-hoc analyses were Tukey adjusted.

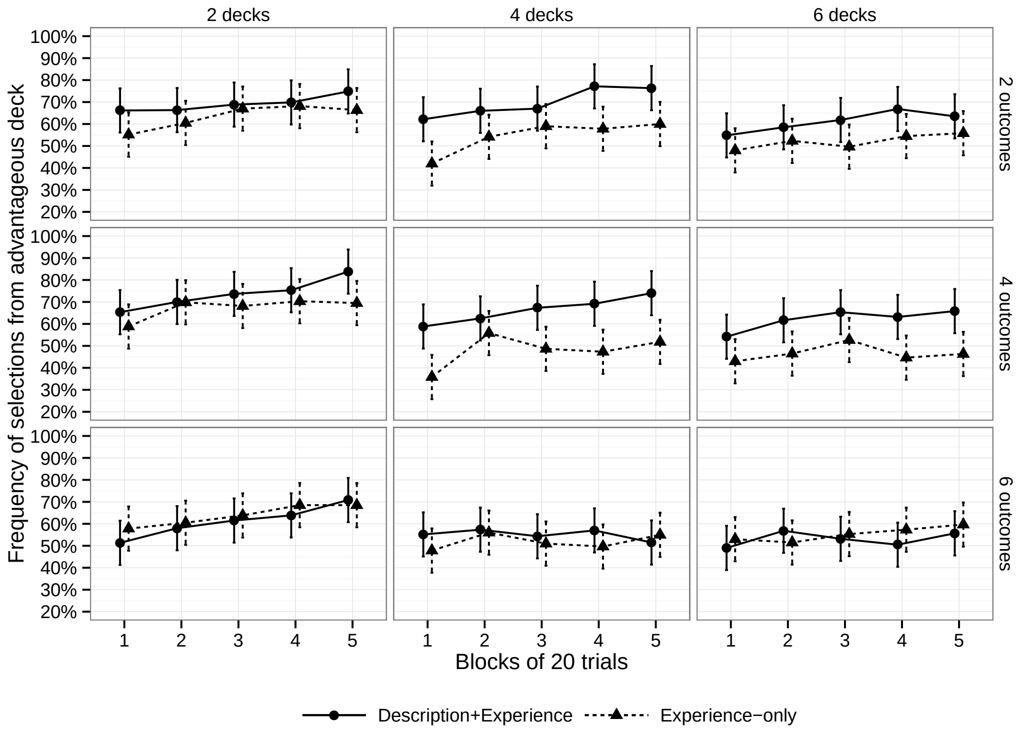

Figure 3.5: Experiment 6. Evolution of the average frequency of selection of advantageous decks (Decks D + E + F), as a percentage of total for each block. Each block contains 20 trials. Each data point represents an average of 30 participants, for a total of 60 participants in each plot. Error bars represent the 95% confidence interval around the mean.

The main effect of the presence of descriptions was significant, with participants selecting from advantageous decks significantly more frequently in the DE condition than in the E condition, on average across each block of 20 trials

(DE=63.47%, E=55.90%, $\chi^2(1) = 14.72, p < .001, d = 0.33$). The presence of descriptions helped participants identify the advantageous decks and select from them more often. There was also a significant main effect of number of decks ($\chi^2(2) = 22.18, p < .001$), indicating that an increase in task complexity, as measured by number of decks, led to a decrease in selection from advantageous decks, as the task became harder and it was more difficult to identify the advantageous decks (2-decks: 66.40%, 4-decks: 57.60%, 6-decks: 55.05%). There was no significant main effect of number of outcomes ($\chi^2(2) = 3.53, p = .17$), with no influence of number of outcomes on selections from advantageous decks (2-outcomes: 61.70%, 4-outcomes: 60.65%, 6-outcomes: 56.70%). There was a learning effect, leading to performance improvement over time, as seen by the main effect of block, with a positive slope ($b = 0.021, \chi^2(1) = 62.75, p < .001$), confirmed by a post-hoc analysis that showed higher selection of advantageous decks in the last block when compared to the first block (Block1=53.25%, Block5=63.85%, $t(2,234.5) = 9.64, p < .001, d = 0.59$).

To further elucidate what drives the influence of descriptions, I analysed the 2-way interactions between descriptions and the two complexity manipulations separately. The interaction between number of outcomes and presence of description was significant ($\chi^2(2) = 8.27, p = .016$). In a post-hoc analysis for number of outcomes, the largest difference in selection of advantageous decks between E and DE were observed in the middle 4-outcome condition, a smaller but still significant difference in the 2-outcome condition, and no significant difference in the 6-outcome condition (DE–E difference in each condition: 2-outcomes=9.95%, $t(558.6) = 2.92, p = .004, d = 0.44$; 4-outcomes=13.40%, $t(558.6) = 3.92, p < .001, d = 0.58$; 6-outcomes=–0.65%, $t(558.6) = 0.18, p = .85, d = 0.06$). There was no significant interaction between number of decks and presence of description ($\chi^2(2) = 4.51, p = .10$). In a post-hoc analysis for number of decks, the same pattern of a larger difference in selection of advantageous decks between E and DE was observed in the middle 4-deck condition (DE–E difference in each condition: 2-decks=3.11%, $t(558.6) = 0.91, p = .36, d = 0.14$; 4-decks=12.26%, $t(558.6) =$

$3.56, p < .001, d = 0.53$; 6-decks=7.37%, $t(558.6) = 2.16, p = .03, d = 0.32$).

None of the other 2-, 3-, and 4-way interactions were significant (all *p*s $\geq .07$). In particular, as in Experiment 4, an interaction between presence of description and block was not present ($\chi^2(1) = 0.55, p = .46$), which would indicate no difference in learning due to the presence or absence of descriptions. In other words, changes in selections of the advantageous options progressed similarly across blocks regardless of the presence or absence of descriptions.

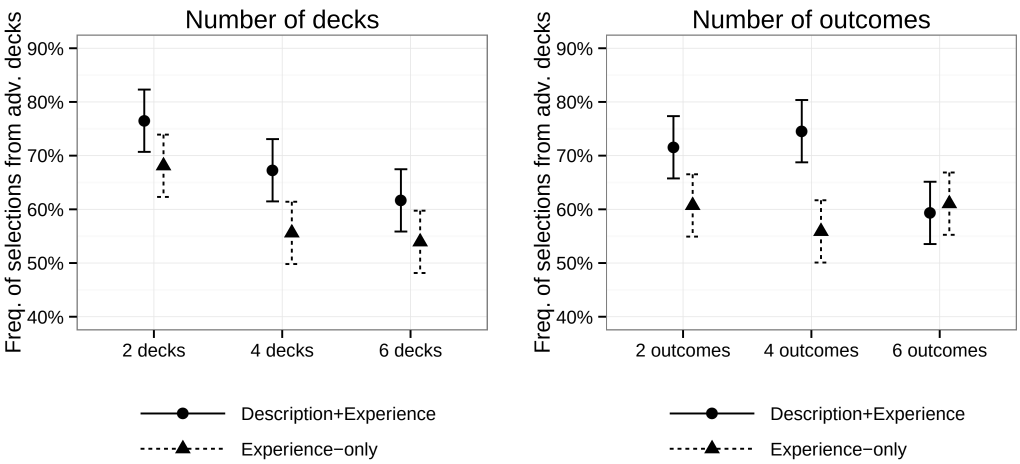

Figure 3.6: Experiment 6. Average frequency of selection of advantageous decks (Decks D + E + F) in the last 20 trials (Block 5). Error bars represent the 95% confidence interval around the mean.

To exclude the effect of learning, and to focus on more stable behaviour, a post-hoc analysis was performed at the last block only (Figure 3.6), with Tukey adjustments. In the last 20 trials, the presence of descriptions led to an overall 16% significant increase in the selection of advantageous decks, with a medium effect size (Block5 only: DE=68.50%, E=59.20%, *t*(1,178.6)=3.84, *p*=.005, *d*=0.33). The small difference and low effect size hide the underlying interactions between the complexity manipulations and the presence of description. The influence of description was highest in the middle condition for number of outcomes (Block5 only, DE–E difference in each condition: 2-outcomes=10.85%, *t*(1,178.6)=2.59, *p*=.01, *d*=0.39; 4-outcomes=18.65%, *t*(1,178.6)=4.46, $p < .001$, *d*=0.67; 6-outcomes=-1.70%, *t*(1,178.6)=0.41, *p*=.68, *d*=0.06). The influence of description was also highest in the middle condition for number of decks (Block5 only, DE–E difference in each condition: 2-decks=8.40%, $t(1,178.6)=2.01$, $p=.05$, $d=0.30$; 4-decks=11.65%, $t(1,178.6)=2.79$, $p=.005$, $d=0.42$; 6-decks=7.70%, $t(1,178.6)=1.85$, $p=.07$, $d=0.28$). A non-monotonic inverted U-shaped pattern of influence of description across both complexity manipulations was observed (Figure 3.7).

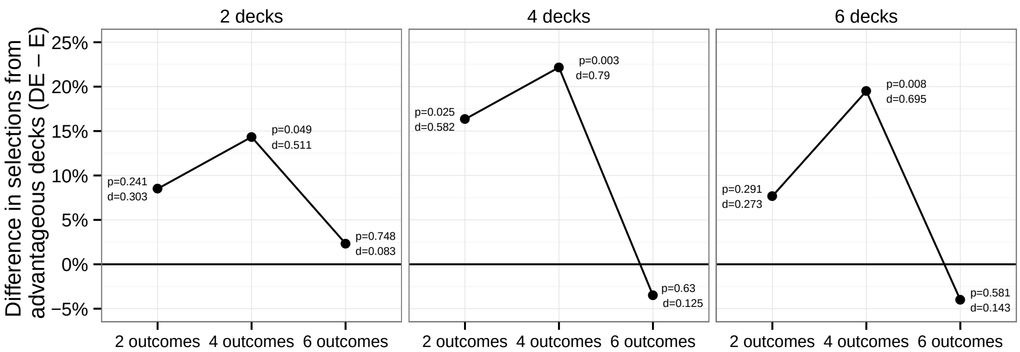

Figure 3.7: Experiment 6. Difference in frequency of selection of advantageous decks (Decks D + E + F) in the last 20 trials (Block 5) due to the presence of description (DE–E). The values for the significance level (p-values) and effect size (Cohen's d) for each difference are shown next to each data point. The highest influence of description can be observed in the middle condition: 4-decks x 4-outcomes.

3.3.3 Discussion

The same overall influence of the presence of descriptions observed in Experiments 4 and 5 was replicated in Experiment 6. As before, participants who were presented with descriptions selected from the advantageous decks 14% more often than participants who did not receive descriptions. These figures are lower than the comparable results in Experiments 4 and 5 because they hide the intricate underlying relationship between task complexity and the influence of descriptions, which followed a non-monotonic inverted U-shaped pattern (Figure 3.7), as predicted.

In simple tasks, in which the payoffs were relatively easy to learn experientially, participants did not benefit from the presence of descriptions. They performed well with experience alone, finding the advantageous decks, and performance did not improve significantly by adding descriptions. This lack of influence from descriptions in simple tasks replicated findings in previous similar research using two alternatives (Lejarraga & Gonzalez, 2011). In very complex tasks, the payoff structure was not only considerably more difficult to learn experientially, but it also required very long and verbose descriptions. Reading, analysing, and deciphering the written descriptions very likely presented a substantial cognitive challenge. In these situations, participants were not as proficient in finding the advantageous decks as in the simpler tasks, leading to lower performance via experience alone, and complicated descriptions did not provide any significant assistance.

It was in the medium complexity tasks that descriptions most helped participants. In the middle cell of Experiment 6 (4 choices $\times$ 4 outcomes), descriptions led to a 39% increase in overall selections of advantageous decks (43% in the last block), and an 82% increase in the final financial bonus, the highest across all conditions and similar to the results from Experiments 4 and 5, which closely matches this task in terms of overall complexity. These medium complexity tasks were too complex to be learned via experience alone, but not overly complex to be described succinctly and for these descriptions to be analysed by and provide useful information to participants. Descriptions are most influential in decision making when they provide additional information to participants that cannot be as easily and efficiently gathered in a timely fashion via experience alone.

3.4 Cognitive modelling

In order to theoretically model the role of descriptions in decisions from experience, and their interaction with task complexity, a set of cognitive computational models was fitted to the experimental data. My aim was to evaluate how descriptions influence traditional reinforcement learning models, and how the descriptive information is represented and integrated into the decision-making process.

I started by fitting a reinforcement learning (RL) model to the observed human behaviour in the experience-only conditions of Experiment 6. RL models that rely solely on experience via feedback have been extensively and successfully used to explain behaviour in decisions from experience in the past (Erev & Barron, 2005; Yechiam & Busemeyer, 2008; Yechiam & Ert, 2007), and the IGT is a paradigm commonly explored with these models (Ahn et al., 2008; Busemeyer & Stout, 2002; Dai et al., 2015; Worthy et al., 2013; Yechiam & Busemeyer, 2005).

Based on their past performance in similar tasks, I would expect RL models to fit my behavioural data well in the experience-only conditions of the current experiments. With regards to tasks combining descriptions and experience, Lejarraga and Gonzalez (2011) have shown that a simple RL model can also explain behaviour well in simple tasks with descriptions, and I expect to replicate their findings here in my simple tasks. However, given the observed difference in behaviour when description was added to the these paradigms in more complex tasks, and my previous modelling efforts (see Chapter 2), I predict that traditional experience-only RL models will perform poorly in the description-plus-experience conditions. As shown here by the observed empirical results, descriptions can sometimes influence behaviour and can provide additional useful information for participants to perform better in their tasks. These situations should be conducive for a model that combines descriptions and experience.

Below, I present a description-plus-experience model that combines both descriptive and experiential information, which should help explain the observed differences in behaviour in the description-plus-experience conditions. My previous attempt at a description-plus-experience model combined the two sources of information with different weights, and the weights determined the importance given to each source depending on the experimental condition (see Chapter 2). Since the influence of descriptions in Experiment 6 appeared to have been moderated by the complexity of the task, I will vary these weights according to a complexity measure, based on entropy. In comparison with traditional experience-only RL models, I expect the description-plus-experience model to provide a better fit for the observed human behaviour in description-plus-experience tasks, in particular in more complex tasks, with little or no difference in the simpler tasks.

3.4.1 The Models

The aim of fitting a cognitive model to the data was to assess and formalise how the two sources of information, descriptive and experiential, are combined. Two models were fitted to the behavioural data: an *experience-only* prospect-valence learning (PVL) model and a *description-plus-experience* adaptation of that model

(D-PVL). The PVL model is a reinforcement-learning model that relies on experiential information alone, using the feedback provided after each trial (Ahn et al., 2008; Fridberg et al., 2010). The D-PVL model built upon that, combining the experience-only RL component from the PVL model with a representation of the descriptive information. The descriptive component was calculated as the expected value of the information presented to participants underneath each choice. The two sources of information were combined using a weight, which was determined via entropy, a proxy for task complexity. Crucially, the experiential part of the D-PVL model was based around the same RL model as the PVL model. Therefore the D-PVL model added descriptions to a traditional experience-based RL model, and I was particularly interested in how this integration was performed. I start by describing the PVL model, which formed the basis of both models.

3.4.2 Experience-only model (PVL)

The models were built upon one successful RL model from the literature, a prospect-valence learning (PVL) model using a prospect-theory utility function and a delta-learning rule. This model has been extensively and efficiently used in the decisions-from-experience literature, in particular using the IGT, and shown to perform better than competing models when fitting experimental data to simulated participants, which is the approach used here (Ahn et al., 2008; Ahn, Krawitz, Kim, Busemeyer, & Brown, 2011; Dai et al., 2015; Fridberg et al., 2010; Worthy et al., 2013; Yechiam & Busemeyer, 2005, 2006).

Firstly, observed payoffs are evaluated by a prospect-theory type of utility function (Kahneman & Tversky, 1979), $U(\cdot)$, defined as:

$$U(r_j(t)) = \begin{cases} (r_j(t)/100)^{\alpha}, & \text{if } r_j(t) > 0. \\ -\lambda(-r_j(t)/100)^{\beta}, & \text{if } r_j(t) < 0. \end{cases}$$

where $r_j(t)$ is the payoff received from option j at time t. Payoff values were divided by 100 to reduce the magnitude of the observed feedback and realign them closer to their monetary payoffs. The free parameters α and β, both ranging between 0 and

2, determine the curvature of the value function for positive and negative payoffs, respectively. Lower values of α and β reduce the distances between extreme values of payoffs, while higher values magnify the distances. The loss aversion parameter, λ, is the free parameter $(0 \leq \lambda \leq 10)$ that determines higher sensitivity to losses in comparison to gains. The higher the value of λ, the higher the importance given to losses over gains.

Secondly, expectancies for the value of rewards for each option are formed via a learning rule, which integrates the experienced feedback after each trial. The learning rule used was a delta rule, which uses a learning rate that determines how much the new information gathered via feedback, in the form of prediction error, influences the updating of the expectancies at each trial (Speekenbrink & Konstantinidis, 2015; Konstantinidis, Ashby, & Gonzalez, 2015; Sutton & Barto, 1998; Yechiam & Busemeyer, 2005, 2006). Feedback observed is integrated after each trial, to arrive at the experienced expectancy $E_j(t)$ for option j at time t:

$$E_j(t) = E_j(t-1) + \phi \cdot \delta_j(t) \cdot [U(r_j(t)) - E_j(t-1)],$$

where ϕ is the free learning rate parameter $(0 \leq \phi \leq 1)$, which is a weight given to new information observed, with lower values resulting in slower learning. The variable $\delta_j(t)$ is a dummy variable, which is equal to one if option j was chosen on trial t, and zero otherwise. The model only updates the value $E_j(t)$ of an option when that option has been selected and its feedback has been observed. When the option has not been selected, the $E_j(t)$ remains unchanged. The initial value for E_j was set to zero.[5]

Finally, the model-predicted probability of selecting a given option j at time t is determined by a time-dependent Softmax rule (Sutton & Barto, 1998) that combined the expected values E_j across all options J:

[5]Attempts to change this to the value of descriptions, D_j, in the description-plus-experience model, led to worse fitting models, as it resulted in more constant behaviour over time with a flatter curve. This is likely due to the participants exploring their options in the beginning of the task even when descriptions were present, a behaviour that would have been suppressed by a model with a non-zero starting E_j.

$$\hat{P}_{jt} = \frac{e^{\Theta \cdot E_j(t)}}{\sum_J e^{\Theta \cdot E_J(t)}},$$

where Θ is the choice sensitivity. If $\Theta = 0$, the model randomly guesses between the options regardless of their expectancies, while higher values of Θ will lead to more deterministic maximisation behaviour. Θ itself is time-dependent and varies according to t, and is determined by the free parameter θ, $(0 \leq \theta \leq 2)$:

$$\Theta = (t/10)^{\theta},$$

this allows choice sensitivity to increase over time, making selections more random in the beginning and more deterministic as time progresses, to reflect the natural tendency of individuals to explore more in the beginning of tasks and less as the task progresses and they have gathered more information from the environment. Values of θ below 1 make the shape of the choice sensitivity over time concave, while values above 1 make it convex, and linear when $\theta = 1$.

3.4.3 Description-plus-experience model (D-PVL)

In the description-plus-experience (D-PVL) model, a representation of descriptions for each choice j, D_j, was combined with the experience, E_j, at each trial t, as follows:

$$ED_j(t) = \omega_c \cdot D_j + (1 - \omega_c) \cdot E_j(t).$$

The experience component, $E_j(t)$, was calculated using the same PVL approach as in the experience-only model, although new parameters were fitted. A representation of the descriptive information is included via D_j as the subjective expected value of the descriptive information for choice j, calculated using cumulative prospect theory (CPT), based on the descriptions provided to participants underneath each alternative. According to Tversky and Kahneman (1992), the CPT value is calculated using a value and a probability-weighting function, $W(\cdot)$ and $U(\cdot)$ respectively,

$$D_j = \sum_m W(p_{jm})U(v_{jm}),$$

where p_{jm} are the probabilities and v_{jm} are the potential values for each outcome m of option j. $U(\cdot)$ is the same function as defined above for the PVL model, using the same parameters. $W(\cdot)$ is the probability weighting function, defined as:

$$W(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{\frac{1}{\gamma}}},$$

where γ is the free parameter $(0 \leq \gamma \leq 2)$ that determines the sensitivity to probabilities via the curvature of the probability weighting function. Values of γ below 1 lead to overweighting of rare events, while values above 1 lead to underweighting of rare events.

Experience, $E_j(t)$, and description, D_j, are combined using ω_c which determines the weight given to description, and its compliment given to experience (see Chapter 2). The ω_c weight changes according to experimental condition c, and is calculated as follows:

$$\omega_c = 1 - e^{(-\xi/S_c)},$$

where ξ is a free parameter which determines the strength of the weight given to descriptions ($0 \leq \xi \leq 3$), divided by the entropy S_c, for each condition c, which was calculated according to the choices available to participants. Entropy has been used before to quantify task complexity, with higher entropy associated with higher complexity (Swait & Adamowicz, 2001). According to the weighting formula used, the weight given to descriptions decreases when entropy increases, therefore ω_c is higher in simpler experimental conditions such as 2×2 and lower in more complex ones such as 6×6. An exponential relationship was used to ensure that the weight ω_c remained bounded between 0 and 1, regardless of the values of S_c.

Entropy for condition c, denoted by S_c, was calculated in two different ways, and the two approaches will be compared in the results section. Entropy was initially defined according to the probabilities displayed to participants in the descriptions:

$$S_c = -\sum_{jm}[p_{jm} \cdot log_2(p_{jm})],$$

for all probabilities p_{jm} for every outcome m of option j for that condition.

However while this approach provided a relatively good fit, it did not provide the best fit to the observed behavioural results. I believe this is because the basic RL model with partial feedback already captures some of the idiosyncrasies of having different numbers of alternatives from which to choose. The more alternatives available, the less an individual learns about the environment after each selection, since only information about one choice is revealed at each trial. These findings will be discussed in more detail in the results section below.

I therefore propose an alternative approach to calculating entropy. I divided the total entropy for each condition by the number of alternatives, or number of decks of cards, in that condition, denoted as A_c, which resulted in an average entropy per alternative:

$$S'_c = -\frac{1}{A_c}\sum_{jm}[p_{jm} \cdot log_2(p_{jm})].$$

Alternatively, S'_c can be considered as the entropy of one of the alternatives, chosen randomly between the ones that were available. The model using this averaging entropy approach will be denoted as D-PVL$'$. As an example, the value of S'_c for the simplest condition in Experiment 6, condition 2×2, which had two options and each option had two outcomes with 50% probability each, was $S'_c = -1/2 \cdot (0.5 \cdot log_2(0.5) + 0.5 \cdot log_2(0.5) + 0.5 \cdot log_2(0.5) + 0.5 \cdot log_2(0.5)) = 1.0$. In comparison, entropy for the middle condition, 4×4, was $S'_c = 1.3$, and for the most complex condition, 6×6, it was $S'_c = 2.3$. S'_c was mostly influenced by the number of alternatives. In comparison, the values for the total S_c were higher, and increased much faster as task complexity increased.

The same Softmax choice rule from the PVL model is used for the D-PVL models, replacing E_j with ED_j, although as before, new parameters are fitted.

3.4.4 Model fitting

Data sets containing 100 simulated participants were generated for each of the 9 experimental conditions in Experiment 6 (number of decks: 2, 4, or 6 $\times$ number of outcomes: 2, 4, or 6), with the same pseudo-randomised methodology used to generate actual data sets for the experiments in blocks of 20.

I started by fitting the experience-only PVL model to the observed human behaviour in the 9 different experience-only (E) experimental conditions of Experiment 6. This model was not allowed to take into consideration the descriptive information, as the participants also did not have access to any descriptions. A total of 900 modelled simulated participants were confronted with 270 observed human participants. All simulated participants across all underlying experimental conditions shared the same set of free parameters. The best fit parameters were found by minimising the multinomial negative log-likelihood (*LL*) between the average observed proportions of choice from each deck and the average model-predicted proportions for each of the individual conditions separately, with each condition receiving the same weight (Erev & Barron, 2005):

$$LL_c = -2\sum_{jt} ln\left(\frac{N_c!}{\prod n_{jt}!}\prod(\hat{P}_{jt})^{n_{jt}}\right)$$

where N_c is the total number of participants in each condition, n_{jt} is the number of participants who chose option j at trial t, and $\hat{P}_{jt}$ is the model-predicted probability of choosing option j at trial t.

To allow for behavioural differences between the E and DE experimental conditions, I also fitted the PVL model to the DE conditions in Experiment 6. Any changes in the parameters could be explained by a different approach that individuals might have taken towards the task when description was available. However this model still does not allow for the descriptive information itself to be integrated into the decision-making process. I check if the descriptive information was used by participants by fitting the two alternative D-PVL models (with total entropy and with average entropy) against the description-plus-experience (DE) experimental conditions of Experiment 6. The same 100 simulated participants were used as above, but now the model was also allowed to take into account the descriptive information, since this was available to participants. The D-PVL models were fitted in the same way as the PVL model above against the observed human behaviour, minimising the LL.

Because of the different number of parameters between the models, the Bayesian Information Criterion (BIC), which penalises for additional parameters, was calculated to compare the models, $BIC_c = LL_c + f \cdot ln(N)$, where f is the number of free parameters and N is the number of fitted observations for each evaluation (100 trials). Lower BIC values represent better fitting models. The mean BIC_c is reported, which is the mean across all 9 conditions.

3.4.5 Model Evaluation and Results

Three models were evaluated, with four sets of parameters fitted in total: the experience-only PVL model, which did not account for the influence of descriptive information, with five free parameters, was fitted twice, against human behaviour in the E (called PVL_e) and the DE conditions (PVL_{de}), separately, allowing for two sets of different parameters; and the two description-plus-experience (D-PVL) models, one with total entropy and one with the alternative average entropy approach (D-PVL$'$), which combined both descriptive and experiential information, both with seven free parameters, were fitted against human behaviour in the DE conditions only.

The results of the model fitting analysis were in line with the behavioural results (Figure 3.8). The PVL_e model fitted against the human behaviour in the E experimental conditions proved a relatively good fit (mean BIC by condition $M_{BIC} = 1,283$). This model was considerably better than a base model, which randomly selects decks of cards at each trial among the available options, returning an $M_{BIC} = 1,484$ in the E conditions. As expected, when comparing the PVL_e model to the human behaviour in the DE conditions, the fit was substantially worse overall ($M_{BIC} = 1,403$). This is because of the behavioural differences observed in the experiment, likely a result of the introduction of descriptive information, while the model was not allowed to integrate that new information. It was still better than

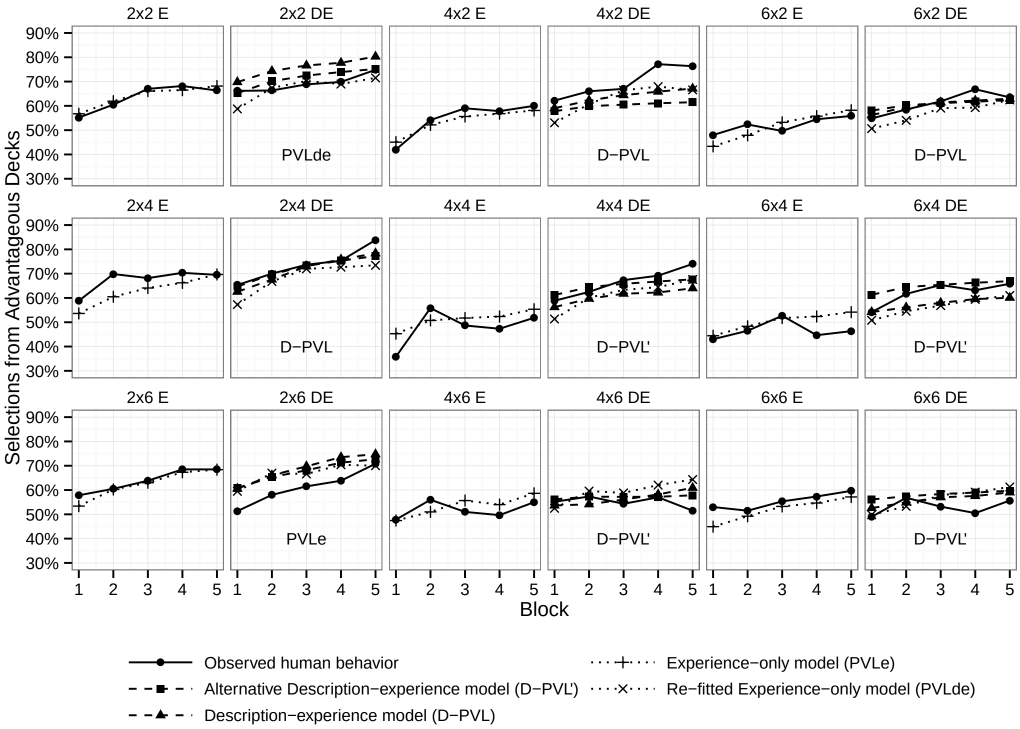

Figure 3.8: Comparison of observed human data, and the three cognitive models fitted against them, for each of the 18 experimental conditions. The experimental conditions above each cell can be identified as number of decks × number of outcomes followed by E for experience-only and DE for description-plus-experience. The text within the cells identify which was the best fitting model for that particular condition (only one PVL$_e$ model was fitted against the E conditions). The best-fitting model overall was the alternative D-PVL$'$ model using the average entropy approach, in particular in the higher complexity conditions. In medium-complexity conditions, the total entropy model D-PVL provided a better fit. The PVL models fitted better in one condition each, both simpler conditions.

the random behaviour base model in the DE conditions ($M_{BIC} = 1,624$). The higher random BIC for the base model in the DE conditions indicates that participants were behaving less randomly when description was present, and therefore a model that predicts random behaviour is a poorer predictor of human behaviour in the DE condition, but a better predictor in the E conditions, when participants were behaving closer to random.

The fit results were substantially improved by refitting the experience-only model to the behaviour in the DE conditions (PVL$_{de}$), with new parameters ($M_{BIC} = 1,316$), as show in Table 3.3. The new PVL$_{de}$ model still did not include any descriptive information, and relied on experience alone. While the original parameters fitted against the E conditions would provide a poor prediction for behaviour in the

DE conditions, the newly fitted parameters accommodate for some of the differences in observed behaviour. This could be a result of over-fitting, and will be checked in a cross-validation generalisation analysis against Experiments 4 and 5 in the next section.

Table 3.3: Best fit parameters of the three cognitive models, PVL (fitted twice, against E and DE observed human data in Experiment 6), D-PVL and D-PVL$'$ (fitted against DE data only), and mean BICs. Lower BICs represent better fits. n.a.=not applicable.

Free parameter	Exp. only E data (PVL$_e$)	Exp. only Refit DE data (PVL$_{de}$)	Description+ Experience (D-PVL)	Alternative Descr.+Exp. (D-PVL$'$)
α (curvature of pos. values)	1.23	0.51	1.60	1.26
β (curvature of neg. values)	0.44	0.47	0.53	0.60
λ (weight of neg. values)	1.83	1.82	9.58	2.73
γ (curvature of probabilities)	n.a.	n.a.	0.92	0.84
ϕ (learning rate)	0.31	0.27	0.02	0.06
θ (choice sensitivity)	0.14	0.06	0.22	0.20
ξ (description's weight)	n.a.	n.a.	0.59	2.12
No. of free parameters	5	5	7	7
Mean BIC	1,403	1,316	1,294	1,271
No. of conditions best fit	1	1	3	4

Finally, by fitting the D-PVL model that was allowed to take descriptions into account, and using total entropy to moderate the weights given to description, there was an additional improvement in the fit against the DE conditions, with a 2% reduction in BIC ($M_{BIC} = 1,294$), compared to the PVL$_{de}$ model. The alternative model D-PVL$'$, using the average entropy approach, resulted in an even better fit ($M_{BIC} = 1,271$), with an additional 2% reduction in BIC. While the parameters did not change across experimental conditions, the results can be split according to them, and verify in which conditions each model performed better. The D-PVL models were the best performing models with lower BICs in 7 out of the 9 experimental conditions, in particular the conditions with higher task complexity. The two conditions in which they were outperformed by the PVL models were both 2-deck conditions (with 2 and 6 outcomes). In these relatively simple 2×2 and 2×6 conditions, the experience-only PVL model proved a better fit for the observed behaviour, replicating the finding from Lejarraga and Gonzalez (2011), who also showed that a simple experience-only RL model without descriptions provided a good fit for human behaviour in simple tasks. It is only in more complex tasks, where I predicted that descriptions would be more useful for participants, that the D-PVL models outperformed the PVL models.

The alternative model using average entropy, D-PVL$'$, returned considerably better fits in comparison to the total entropy D-PVL model. Average entropy can be interpreted as the entropy of a single alternative, or deck of cards, selected at random from the ones available. Because all the alternatives in my paradigms contained the same number of potential outcomes, their average entropies did not differ considerably. Comparing the models using Schwarz weights (Wagenmakers & Farrell, 2004) showed a strong preference for D-PVL$'$, with $w(BIC_{DPVL'}) > .9999$, which can be interpreted as the probability that this is the best model among the models presented here (Lewandowsky & Farrell, 2011). I believe that employing average entropy as a moderator of weights given to descriptions yielded better fitting models because the experience component of my models (E_j) already incorporated the deleterious influence of additional alternatives, but it is not influenced by the number of outcomes. The performance of traditional experience-only RL models already deteriorates considerably when dealing with a larger number of alternatives (Konstantinidis et al., 2015). This is specially the case when only partial feedback is available, as the RL model can only update the expectancy of the most recently selected alternative, for which feedback was presented. This results in a smaller reduction of uncertainty about the environment after each trial when more alternatives are available, slowing down the differentiation between the alternatives and, consequently, the ability to identify the better ones. Similar patterns were shown by Ashby et al. (2017). When there are only two options, with each selection, half of the alternatives are updated. With six options, only one sixth is updated. Therefore it takes longer to reduce uncertainty with a traditional RL model when there are more alternatives. There is no similar mechanism for number of outcomes from each alternative, with the RL model incorporating new information in the same manner regardless of the number of outcomes. Traditional RL models such as the ones used here are therefore sensitive to number of alternatives but insensitive to number of outcomes. Therefore I believe that by adding an entropy measure that is mostly related to the number of outcomes, not alternatives, I have breached this remaining gap in the D-PVL models, without taking into consideration the effect of number of outcomes twice, which is the case when total entropy is used.

Since I observed a non-monotonic U-shaped curve in the relationship between presence of description and task complexity in the human behaviour from Experiment 6, a cognitive model that captures human behaviour appropriately should also replicate that finding. I compared the modelled predictions for the last block of trials for the PVL$_e$ and the D-PVL$'$ models (Figure 3.9). The inverse U-shaped pattern that was observed in the human behaviour was also replicated with cognitive models. Increasing the number of outcomes increased the entropy monotonically, but influence of description, moderated by a weight determined by the inverse of entropy, was non-monotonic.

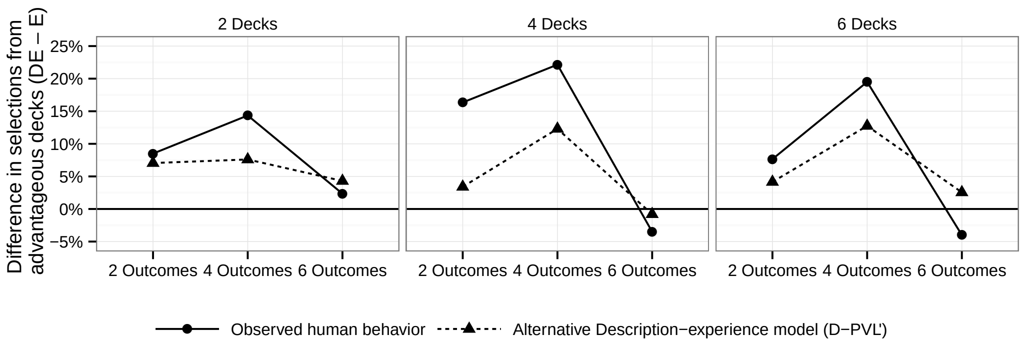

Figure 3.9: Difference between selection of advantageous decks as predicted by the cognitive models and human behaviour, for each experimental condition, in the last block of 20 trials. The modelled difference is the prediction of the experience-only model subtracted from that of the description-plus-experience model.

The best fit parameters were relatively consistent across the different models (Table 3.3). A few parameters changed in reaction to the presence of descriptions. In particular, the learning rate ϕ was lower in the D-PVL$'$ model ($\phi_{DPVL'} = 0.06$) compared to the PVL$_e$ model ($\phi_{PVLe} = 0.27$). As participants had access to descriptive information, they did not need to learn as much from feedback after every trial, and updated their expectancies more slowly in the reinforcement learning component of the model. I also observed that the weight given to negative rewards λ, was higher in the D-PVL′ model ($\lambda_{DPVL'} = 2.73$) than in the PVL$_e$ model ($\lambda_{PVLe} = 1.83$). I believe that this is due to the increased relevance of losses when constantly presented in textual descriptions, as they appear more salient and ever-present than the occasionally observed feedback, similar to the "mere presentation" effect (Erev, Glozman, & Hertwig, 2008). The choice sensitivity parameter θ was higher in the D-PVL′ model ($\theta_{DPVL'} = 0.20$) than in the PVL$_e$ model ($\theta_{PVLe} = 0.14$). Higher choice sensitivity translates into more deterministic behaviour, and it is likely that the presence of descriptions reduced uncertainty, and allowed participants to be more secure in their decisions, behaving less randomly in their choices. Finally, the weights given to description for the D-PVL′ model, based around the best fitting ξ parameter and calculated according to the ω_c formula above with average entropy, varied between 0.57 for the most complex conditions and 0.92 for the simplest conditions.

3.4.6 Cross-validation with the IGT

I also ran a cross-validation analysis of the models fitted against Experiment 6 on the observed human data from Experiments 4 and 5, using the original parameters without refitting, to confirm if my findings could be generalised. I simulated the model predicted outcomes with the parameters from Table 3.3 using the IGT paradigms from Experiments 4 and 5, and compared it against the observed human behaviour (Figure 3.10). One hundred simulated participants were created using the outcomes from the IGT, both using a fixed schedule and a random schedule. Because no new parameters were fitted, the BIC values shown here have a penalty factor $f = 0$ for all models. The PVL$_e$ model was a relatively good predictor for the E conditions in Experiments 4 and 5 ($M_{BIC} = 1,807$). This is not surprising as this model has been extensively used before to predict behaviour in the traditional IGT. As in Experiment 6, the PVL$_e$ model was not a good predictor for the DE conditions, with a considerably worse BIC ($M_{BIC} = 1,906$). However, in comparison with Experiment 6, the re-fitted PVL$_{de}$ model did not show the same level of improvement against the

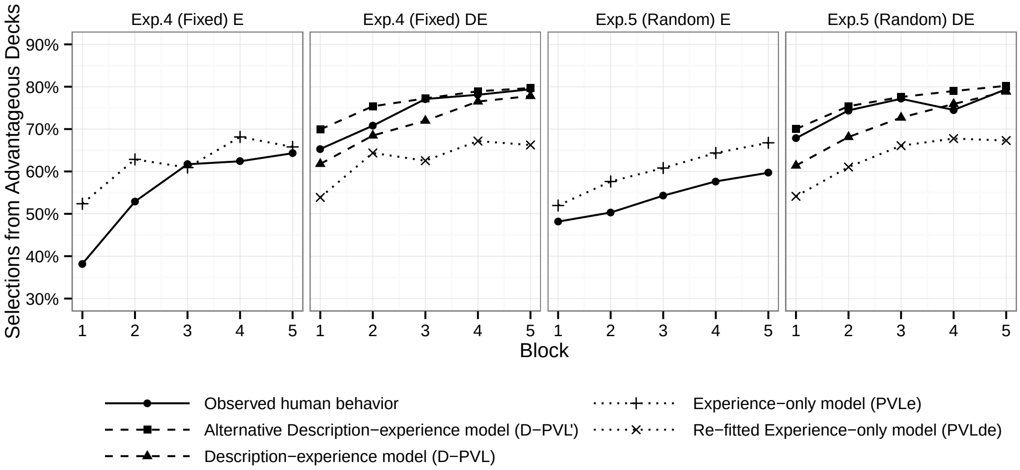

Figure 3.10: Model predictions applied to the IGT simulated data from Experiments 4 and 5, both with fixed and random schedule of outcomes, using the parameters fitted against human data from Experiment 6. Against the IGT, the PVL$_{de}$ did not show the same improvement as in Experiment 6. The D-PVL$'$ model returned much better predictions, ahead of the D-PVL model.

human behaviour in the DE experimental conditions $(M_{BIC} = 1,854)$. This could be interpreted as a result of over-fitting of the model against behaviour in Experiment 6. The D-PVL models were considerably better predictors of the behaviour in the DE conditions of Experiments 4 and 5 overall. As before, the alternative model D-PVL$'$ with $(M_{BIC} = 1,542)$ provided a better fit than the total model D-PVL with $(M_{BIC} = 1,633)$. Overall, the D-PVL$'$ model was a much better predictor of human behaviour in the IGT of Experiments 4 and 5 than the PVL model, which was expected given the complexity of this task and the empirical differences in human behaviour observed between the E and DE conditions.

3.4.7 Modelling discussion

Overall, the description-plus-experience D-PVL$'$ model proved a better fit of the observed human behaviour, in particular in the higher complexity experimental conditions. While previous attempts to model description-plus-experience tasks with congruent information had shown that a traditional experience-only model PVL could be used to explain behaviour relatively well (Lejarraga & Gonzalez, 2011), I believe that this was only the case because the tasks used were simple. Simple tasks should be easy to learn experientially, and the addition of descriptive information did not influence behaviour, or modelling predictions significantly. This finding has been replicated here in my simplest experimental condition 2×2 (and also 2×6), with the experience-only PVL models providing good fits for the observed human behaviour and little difference between predicted model results between PVL and D-PVL (see Figure 3.9). However when complexity is increased, then descriptions can provide useful information for participants, helping them make better decisions. This was shown empirically in the behaviour data, and confirmed with cognitive modelling. If descriptions had not been taken into account by participants, then the best fitting parameters for the PVL_e model from the E conditions should predict behaviour in the DE conditions well, which was not the case. Even the re-fitted experience only models PVL_{de} did not provide a much improved fit: this might have been a result of over-fitting, as shown by the generalisation analysis of the models against the IGT tasks in Experiments 4 and 5.

3.5 Discussion on task complexity

Previous research had shown that introducing congruent descriptions, those that provide the same underlying information as the experienced feedback, to decisions from experience, did not influence behaviour (Lejarraga & Gonzalez, 2011), a result that was replicated in the previous chapter in this dissertation. However I believe that this was because the tasks used previously were relatively simple tasks, and that descriptions might only be taken into consideration by individuals when it is advantageous to do so, given the higher cognitive cost associated with processing them, compared to the easier processing of experiential information (Glöckner et al., 2012; Lejarraga, 2010). The aim of the experiments in this chapter was to show that congruent descriptions can influence behaviour in more complex tasks, where the addition of descriptions is advantageous given the higher cognitive effort required to decipher the task. The presence of congruent descriptions helped participants perform better in the described IGT in Experiments 4 and 5, as well as in medium-complexity tasks in Experiment 6, enabling them to choose the advantageous decks more often and sooner. In Experiment 4, given the higher complexity of the task, it was cognitively advantageous for participants to use descriptions, which in turn influenced behaviour. This did not occur in previous research, where the task was relatively simple and relying on experience alone was sufficient. This effect was replicated in Experiment 5. The higher selection of advantageous decks in the description conditions of Experiments 4 and 5 showed that the availability of descriptions in complex choice environments such as the IGT helps individuals identify the decks with the higher long-term earning potential. Not only was this identification made earlier, but even after many trials and extensive experiential learning, participants did not reach the same level of performance without descriptions.

Experiment 6 showed that the influence of descriptions on behaviour was moderated by task complexity. Descriptions helped participants improve their performance the most in tasks of medium complexity. When the task was very simple, participants were able to learn about the task experientially, which requires lower cognitive effort than analysing the descriptions. Participants mostly seemed to have neglected descriptions, replicating results observed in previous research conducted with similarly simple tasks, such as Lejarraga and Gonzalez (2011) and those in the previous chapter. Increasing the task complexity too much, however, led to a situation in which both experience and descriptions were overly complex. Learning via experience in complex tasks is more difficult, but processing the complicated written information required to describe such a complex task is also overly taxing and demanding. In these very complex tasks, the addition of descriptions did not help participants' performance. Descriptions were most useful in medium complexity tasks, where the experience is relatively too complex to be learned easily and efficiently, but descriptions are still relatively simple and can be processed without too much additional effort. This created a non-monotonic inverted U-shaped pattern for the relationship between task complexity and influence of descriptions: highest for tasks of middle complexity, and lower in both extremes of low and high complexity (see Figure 3.7 on page 104). Similarly shaped relationships between task complexity and decision performance had been observed before in other domains (Eppler & Mengis, 2004; Hwang & Lin, 1999; Streufert & Driver, 1965). They suggest that too much information can lead to cognitive overload, with deleterious influences on performance, similar to what was observed here.

A cognitive model that combined representations of both descriptive and experiential information was also fitted to the behavioural data in Experiment 6. The combined description-plus-experience model provided better fitting results than a more traditional experience-only model that relied on experience alone, and did not consider any additional descriptive information, in particular in the more complex experimental conditions. In the simpler conditions of Experiment 6, the combined model was no better than the traditional model, a result that was previously shown both in Lejarraga and Gonzalez (2011) and Chapter 2, where an experience-only model fit the behavioural data relatively well for simple tasks. It was observed that, as task complexity increased, the addition of descriptive information into the model led to better fitting results. The combined description+experience model also returned the same inverted U-shaped pattern for the relationship between task complexity and performance improvement due to the addition of descriptions, as observed behaviourally. The description-plus-experience combined model fitted against Experiment 6 also generalised well, providing good predictions for the behaviour observed in the DE conditions of Experiments 4 and 5. An analysis of the best fitting parameters showed that individuals pay more attention to losses, learn more slowly and choose less randomly when descriptions are available to them. The proxy for complexity used in the model was entropy, and while an overall task entropy measure was envisaged initially, the best fitting model resulted from the use of an average entropy for each alternative, which varied mostly due to the change in the number of outcomes. I believe this to be the case because traditional reinforcement learning models using partial-feedback already indirectly take into account the number of alternatives available, since only one can be updated at a time.

The description-plus-experience tasks presented in this chapter have shown, both behaviourally and computationally, that the influence of descriptions on decisions from experience is moderated by task complexity. While in simple tasks, explored in previous research, descriptions seemed to be completely neglected, a result also replicated in this chapter, I have also shown that in medium-complexity tasks descriptions are taken into account, improve performance, and reduce exploration. However, when the task is overly complex, descriptions are not so useful. I propose that the lack of usefulness of descriptions in overly complicated tasks might be due to the complexity of the descriptions themselves which are used to describe these tasks. Perhaps if simpler descriptions could have been provided, then these would influence behaviour more, even in complex tasks.

Chapter 4

Prior experience

Most of the description-plus-experience (D+E) research so far, including the studies reported in the previous chapters in this dissertation, has looked at the interaction between descriptions and experience when the two are available concurrently from the beginning of the task. When descriptions are always available, the first selection made by participants is already influenced by the descriptive information, prior to any experiential feedback. This was observed empirically in Chapters 2 and 3 of this dissertation. The extant D+E research has thus mostly investigated the behavioural influence of experiences posterior to the influence of descriptions. This might be a common situation, as individuals might read instructions manuals before using a new device, read reviews before going to a new restaurant, or read the patient information leaflet before taking a new medication.

The reversed situation should not be overlooked, when prior experience is available before descriptive information. Notable examples include driving before the introduction of seatbelt legislation, drinking alcohol before the affixing of beverage warning labels on bottles, and smoking before plain packaging legislation. The literature on safety warnings includes familiarity as a potential moderating factor for their effectiveness, referring to prior personal experience and expectations with a product or the environment (Laughery, 2006; Rogers et al., 2000). The question remains as to how prior experiences, before the presentation of descriptions, would affect the dynamic balance of influence between the two sources of information. This can be investigated empirically by allowing participants to accumulate experiences in a DfE-type task without descriptions initially, and introduce the descriptive information only after a certain amount of experience has accumulated. Moreover, the amount of prior experience can also be experimentally controlled.

Barron et al. (2008) have conducted an initial exploration of how prior experience can influence the behavioural impact of descriptive information, provided in the form of warnings. In their study, the authors demonstrated how presenting descriptions before any prior experience, at the inception of the task, had a stronger impact on behaviour compared to participants who received the descriptive information only after they had the chance to accumulate experience, by allowing them to perform the first half of the task without descriptions. Participants who were only shown the descriptions later, and were allowed prior experience with the task before the appearance of descriptions, displayed behaviour befitting a more subdued influence of the warning description, which can be explained cognitively via stronger discounting of descriptions due to prior experience.

A similar field-study by Miron-Shatz, Barron, Hanoch, Gummerum, and Doniger (2010) showed the importance of being exposed to a warning label before any personal experience, in order to increase compliance. They showed that even a small amount of previous direct personal experience is sufficient to trigger behavioural inertia and leads to lower adherence to warning labels. In their study, experienced parents who had previously used a certain medication with their older children, also administered it with their younger children, even after a new warning label about dangerous side effects of giving this medication to young children had been introduced in the interim. Inexperienced parents who had not used the drug prior to the warning label were more reluctant to give it to their young children.

An interesting stream of research remains relatively unexplored with regards to controlling the amount of prior experience, and its impact on the influence of descriptions. I believe that there are two interconnected cognitive mechanisms behind the influence of prior experiences on the strength of the impact of descriptions on DfE tasks: behavioural inertia and learning interference. Both can lead to sub-optimal decision-making, by hindering the exploration of new alternatives and consideration of new information that could lead to better choices (Dutt & Gonzalez, 2012). The two phenomena need not be mutually exclusive, and I expect a combination of both to operate behind the moderating influence of prior experiences on the impact of descriptions.

Behavioural inertia

Behavioural inertia can be defined as the tendency to repeat previous choices, even when those choices might lead to lesser rewards than could have been obtained otherwise (Biele et al., 2009; Erev & Haruvy, 2005). One of the drivers behind inertia is the inherent preference that individuals show towards maintaining their current status, even when faced with options that would allow them to improve on it, also called the "status quo bias" in decision-making (Samuelson & Zeckhauser, 1988). One of the reasons for this behaviour, according to the authors, is a psychological commitment for consistency in people's choices that might make it more difficult to accept a shift in behaviour, promoting the repetition of previous actions. While the status quo bias can reflect a preference that individuals often show to avoid making decisions altogether if possible, it can also refer to continuously repeating one's previous decisions or strategy, even if that is no longer ideal. In cases where repeated decisions are to be made, such as in DfE tasks, this decision avoidance can promote behavioural inertia, because once a certain number of initial decisions have already been made earlier, repetition of previous choices can be seen as avoiding further decision making, which is cognitively less stressful (Anderson, 2003). Behavioural inertia can be frequently observed empirically, and its influence can be strong, insofar that a reliable predictor of future choice can usually be made based on simple repetition of the preceding choice (Avrahami & Kareev, 2011; Dutt & Gonzalez, 2012; Ert et al., 2011).

Using computational models, Ert et al. (2011) demonstrated that inertia is inversely proportional to the level of surprise caused by the feedback received: more surprise triggers behavioural change, while lack of surprise leads to more inertia (see also Nevo & Erev, 2012). The authors defined surprise as the gap between new information being learned and old information that has already been learned, a concept that is also related to reward prediction errors (Dayan & Niv, 2008). These are important components of the learning mechanism, as learning is highest when reward prediction errors, and surprises, are higher (Niv, 2009). Surprise can be seen as a proxy for potential gains to be learned from the environment, with situations with low potential for surprises being those in which most has been learned already and with little uncertainty remaining. When individuals must learn about a new environment, and little is previously known about it, most information received will be relatively new and highly surprising, and prediction errors are high. Over time, and in particular in stable scenarios, information that is received after a gathering considerable data is less likely to be surprising, and prediction errors go down. According to the surprise-triggers-change theory, at the beginning of new tasks surprise is highest and inertia is lowest, and as time progresses and information is captured, surprise reduces and inertia increases. Therefore inertial behaviour is better aligned with situations in which the desire to learn is low, and surprise is also low, and repeating previous actions is not detrimental. Avrahami and Kareev (2011) observed how inertia increases over time, with the accumulation of experience. Erev and Haruvy (2005) have found that individuals display lower inertia in new task situations, where more learning is desirable. In fact, contrary to general reinforcement learning theory that supports the idea that information from different options is compared and a new decision is made at each trial, Erev and Haruvy (2005) suggest that individuals might not be making new decisions, and instead are only following their previously decided strategy, repeatedly. Inertia might be an adaptive artefact derived from the experience of auto-correlation in everyday life, reflecting the fact that the past is highly correlated with the future. The weather today is a relatively good predictor of the weather tomorrow, so a new decision about what to wear is typically not required every day (Erev & Haruvy, 2005).

Therefore prior experience might lead to behavioural inertia, and the larger the amount of prior experience, the higher the expected inertia, by reducing surprise. However, if the descriptive information that is made available posterior to experience contains new conflicting information that triggers substantial surprise, for example, by warning individuals of a rare event that has not yet occurred experientially, this increase in surprise should in theory lead to a reduction in inertia, and behavioural change, when novel descriptions are introduced later.

Learning interference

Even when the descriptive information is remarkably surprising, prior experiences can also shape future behaviour by interfering with the learning process, altering if and how new information is taken into consideration. "The knowledge the agent brings to the task at the start – either from previous experience with related tasks or built into it by design or evolution – influences what is useful or easy to learn . . ." (Sutton & Barto, 1998, p. 7).

Prior experiences can thus influence decisions causing new information, which would be used for learning, to be ignored. If new information is ignored, then old behaviour is more likely to persist due to the lack of new information being accumulated that could shift behaviour. One of the reasons for ignoring information as a result of prior experience supports the idea that individuals seek to avoid cognitive dissonance (Anderson, 2003). Gathering new information that might lead to changing behaviour could require admitting that their previous choices were wrong. "Thus, an individual tends to discard or mentally suppress information that indicates a past decision was in error" (Samuelson & Zeckhauser, 1988). Confirmation bias is a related phenomenon that also leads to informational neglect: individuals will attend mostly to new information that confirms their previously established beliefs, and ignore those that go against it (Klayman & Ha, 1987; Krizan & Windschitl, 2007).

Even when new information is not completely ignored, it might be discounted or distorted. For example, research has shown how the amount of prior experience influences learning, via the reduction of uncertainty (Daw et al., 2006; Speekenbrink & Konstantinidis, 2014). New information might be misinterpreted or misconstrued, for example to reduce cognitive dissonance. Surprising information that is received after inertia has set in might be discounted as a one-off or outlier that is unlikely to repeat. Thus, the interference provided by prior experience can lead to a dampening of the perception of surprises. New descriptive information that is provided after prior experience, such as late warnings, must compete against this stronger tendency to behave more inertially, and has to be considered in the context of all the prior information that has already been observed.

Experiments

The third set of experiments in this dissertation will focus on how the amount of prior experience moderates the influence of descriptions on behaviour. In this chapter, I propose to extend the earlier research by Barron et al. (2008) and Miron-Shatz et al. (2010) by controlling for the amount of prior experience before descriptions are revealed to participants: they will be given either at first trial, never, or at different points during the task. By deferring the introduction of descriptions with different delays, the amount of prior experience that participants will accumulate before encountering the descriptions will be manipulated, and the relationship between prior experience and impact of descriptions can then be analysed.

I predict that the effect of descriptions will be highest when there is no prior experience and they are presented from the beginning of the task, replicating the results in Barron et al. (2008). In addition, I expect the amount of prior experience to moderate the influence that descriptions have on behaviour. An increase in prior experience before a warning appears should lead to behavioural inertia and interfere with the effect of descriptive information on the learning process, thus reducing its subsequent impact on behaviour. Descriptions that appear later, when participants have a large amount of prior experience with the task, should not influence behaviour as strongly as descriptions that appear early. If descriptions provide useful information that can help participants perform better in the task, such as in the form of warnings about risks, then early presentation of descriptions should help performance more than later ones. Furthermore, the results from Miron-Shatz et al. (2010) might sustain the idea that even the smallest amount of prior experience is necessary to settle into inertial behaviour. In this case, in addition to the predicted relationship that increases in prior experience will reduce the impact of descriptions, there might also be a fundamental step-change in behaviour between no prior experience at all and any amount of prior experiences, small or large.

4.1 Experiment 7

This first experiment in this chapter introduces delayed descriptions to the Iowa Gambling Task (IGT: Bechara et al., 1994), and is based on Experiment 5 in Chapter 3. In contrast to that earlier experiment, in which descriptions were either displayed from the first trial or not available at all (both experimental conditions which are also duplicated here), in the current experiment I will introduce delayed descriptions, which are only shown later at different points during the task, in a new experience-before-description (ED) experimental manipulation. In the two experimental conditions that are duplicated from Experiment 5, I expected to replicate the earlier findings from that experiment, where participants who received written descriptions about the decks of cards performed better than participants who did not. With the new ED experimental conditions, created by allowing participants to accumulate different levels of prior experience before the descriptions are presented at different points during the task, I expected to replicate the findings of Barron et al. (2008), with participants who receive the descriptions earlier in the task taking less risk and performing better than participants who receive descriptions later. In addition, I expected the amount of prior experience to moderate the impact that descriptions had on the performance of participants in the task. I expected descriptions to have a stronger impact on the behaviour of participants who observe the descriptions earlier, with a lower amount of prior experience, than participants who only receive the descriptions later in the task, after accumulating a larger amount of personal experience.

4.1.1 Method

Design

Experiment 7 was a between-subjects design with six experimental conditions manipulating when descriptions were presented to participants: one experience-only condition, in which descriptions were not presented (E); one description-before-experience condition, with descriptions available from the first trial (DE); and four experience-before-description conditions, with descriptions appearing at different points during the task (ED20, ED40, ED60, and ED80). The E and DE conditions are close replications of those in Experiment 5 (Section 3.2). While in the E condition participants learned the outcomes of their choices through experience via feedback alone, with blank buttons providing no descriptions, in the DE and ED conditions participants were also presented with a written description of the potential outcomes for each option, in addition to feedback. In the DE condition, descriptions were displayed from the first trial, therefore participants had this information available before any feedback was shown and any experience accumulated. In the ED conditions, participants first went through a certain number of trials without any descriptions. The numbers after the ED in each condition represent the trial in which descriptions first appeared: In condition ED20 descriptions were shown from trial 20 onward, in ED40 from trial 40 onward, and so on for ED60, and ED80. After appearing, descriptions remained present until the end of the task, which lasted 100 trials. Each participant was allocated to only one experimental condition.

Participants

195 participants (108 females; age: M = 35.5 years, SD = 11.8 years) were recruited on-line using Amazon's Mechanical Turk service, with an average of 32.5 participants per experimental condition (N for each condition: DE=33, ED20=34, ED40=33, ED60=32, ED80=30, E=33). Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.20 for participating and an additional bonus according to the outcomes of the choices they made during the experiment (Bonus: M = US$ 0.24, SD = US$ 0.14).

Task

The task was a close reproduction of Experiment 5 in Chapter 3, which was accomplished by adding descriptions to the Iowa Gambling Task (IGT). The original IGT schedule of outcomes was used (see Table 3.1 on page 87), although not using a pre-determined order as in the original task (Bechara et al., 1994, Figure 1), but using the same pseudo-randomisation approach as in Experiment 5 (Section 3.2).

Participants could select from four decks of cards: two decks with short-term gains but overall negative long-term expected values, decks A and B, commonly named as the disadvantageous decks; and two decks with low short-term gains but overall positive long-term expected values, decks C and D, called the advantageous decks. In the current task, the experimental conditions also manipulated the presence of the descriptions displayed underneath the decks of cards on screen, as in the experiments in Chapter 3. Two of the experimental conditions of Experiment 5 were replicated here: the experience-only condition (E) in which participants had to rely on experience alone, via feedback, to inform their decisions; and the description-experience condition (DE), in which participants had access to both experience and descriptions, and descriptions were available from the beginning of the task (Figure 3.1 on page 86), with no prior task experience before descriptions appeared. Four additional experimental conditions were introduced in the current task, grouped as the experience-before-descriptions (ED) conditions, in which the descriptions were only displayed later during the task, after different levels of prior experience was gathered during the task. In these conditions the task began initially with no descriptions underneath the decks, and participants accumulated experience via feedback alone. At different points during the task, after trial 20, 40, 60, or 80, depending on the experimental condition, the descriptions were revealed. After appearing, descriptions remained on screen until the end of the task. As in the original IGT, participants started with 2000 points (worth US$ 0.20) and a running total of both points and money was always displayed on screen, with points converted into bonus at a rate of US$ 0.10/1000 points. The task was self-paced over 100 trials and was completed on average in 8.78 minutes (*SD*=4.17).

4.1.2 Results

Selections from advantageous decks

The main dependent variable was the frequency of cards selected from the advantageous decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 4.1). If descriptions were influencing behaviour and helping participants' performance, as predicted, then the selection of

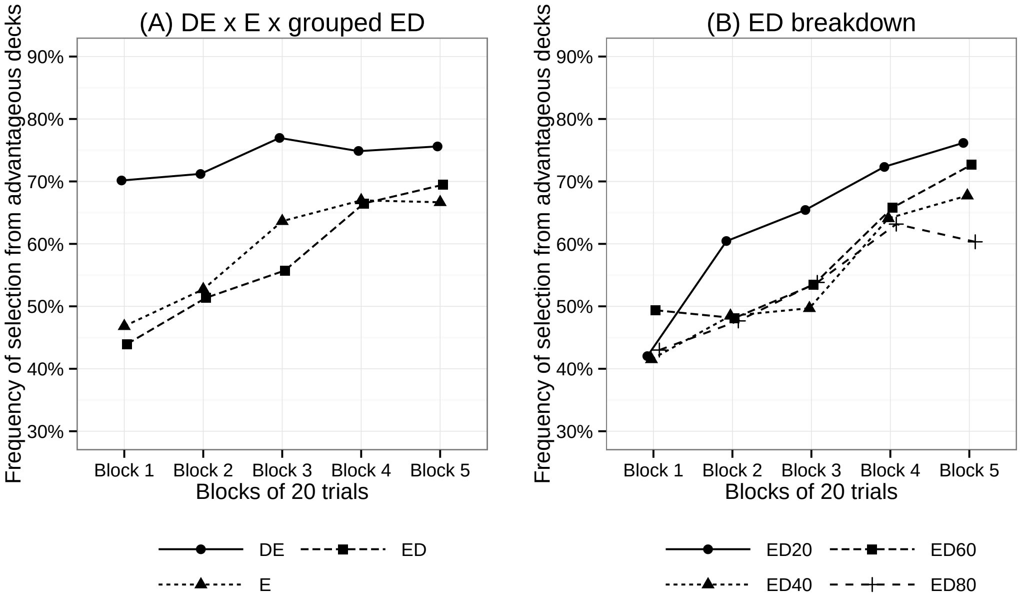

Figure 4.1: Evolution of the selection of advantageous decks for each block of 20 trials for Experiment 7. (A) High level comparison between E, DE and the average collapsed ED conditions, showing that most of the differences are a result of the presence of descriptions from the first trial (DE condition). (B) Detailed breakdown of the underlying ED conditions, showing that the behaviour progressed very similarly in the different ED conditions. E refers to the experience-only condition (no descriptions), DE refers to description-before-experience and ED refers to experience-before-description. The number after the ED is the trial in which descriptions were first revealed to participants.

advantageous decks should be higher when descriptions were present. They were analysed with a linear mixed-effects model using the lme4 package (Bates et al., 2014) and post-hoc analyses using the lsmeans package (Lenth, 2016), with Tukey adjustments, in R (R Core Team, 2014). The between-subjects conditions were the timing of descriptions (DE, ED20, ED40, ED60, ED80, E). For simplicity, in some analyses, the ED conditions were collapsed and averaged together, making the comparisons easier to interpret (DE vs. ED vs. E), as in Figure 4.1A, with further details amongst the underlying ED conditions in a separate analysis, as in Figure 4.1B. The within-subjects conditions were the blocks of 20 choices each. The model also contained a random intercept, and a random slope for the blocks over time, for each participant.

The manipulation for the timing of descriptions overall, across all six experimental conditions, was significant ($\chi^2(5) = 30.117, p < .001$). A post-hoc pairwise comparison with Tukey adjustments showed that the difference was mostly driven by a difference between the DE and E conditions (DE=73.75%, E=59.35%, $t(192) = 2.761, p = .017$), and between the DE and average of the collapsed ED conditions (ED=57.40%, $t(192) = 3.961, p < .001$). There was no significant difference between the averaged ED and E conditions ($p = .88$). A comparison between the four different ED conditions alone also showed no significant difference between them (D20=63.33%, D40=54.33%, D60=57.85%, D80=53.66%, $\chi^2(3) = 4.00, p = .26$). Participants performed better in the condition in which they had descriptions throughout (DE) in comparison to the condition in which they did not have access to descriptions and had to rely on experience alone (E), a replication of the findings in Chapter 3. However the lack of significant difference overall between condition ED and E indicates that later appearances of descriptions led to same behaviour over the entire task as no descriptions at all (Figure 4.1A).

The effect of block was significant with a positive slope ($b = 1.10, \chi^2(1) = 70.63, p < .001$), which shows that participants learned how to identify the better decks over time and improved their performance by selecting more from advantageous decks as the task progressed (Block1=48.80%, Block5=69.85%, $t(228.82) = 7.60, p < .001$). The interaction between experimental condition and block was not significant ($\chi^2(5) = 10.00, p = .075$), although it appears that in the DE condition participants kept the same level of performance throughout, and most of the performance improvement and learning was in the other experimental conditions, driven by learning via experiential feedback (Figure 4.1A).

This was confirmed with a separate analysis including the presence of descriptions (instead of experimental condition) and block. This analysis showed a significantly higher selection from advantageous decks in blocks in which descriptions were available, compared with blocks in which descriptions were not present (Descr.=65.25%, No-Descr.=57.85%, $t(958.53) = 3.44, p < .001$). The interaction between presence of description and block was also significant ($\chi^2(1) = 7.54, p = .006$), with the slope for block being higher when descriptions were not present ($b = 1.20$) compared to when descriptions were present ($b = 0.53$), a result indicating increased learning over time via experience alone without descriptions, and more stable behaviour when descriptions were also available, in particular in experimental condition DE which can be seen in Figure 4.1A.

Switching rates between decks

An analysis on the switching rates between the different decks was also conducted to verify the exploration and exploitation behaviour of participants in response to descriptions (Figure 4.2). As in Experiment 4 in Chapter 3, switching rates were used as proxies for exploration (Ert et al., 2011). A selection was classified as a switch every time a card was picked from a different deck to that from which the previous card had been selected. The same model structure was used as in the previous analysis.

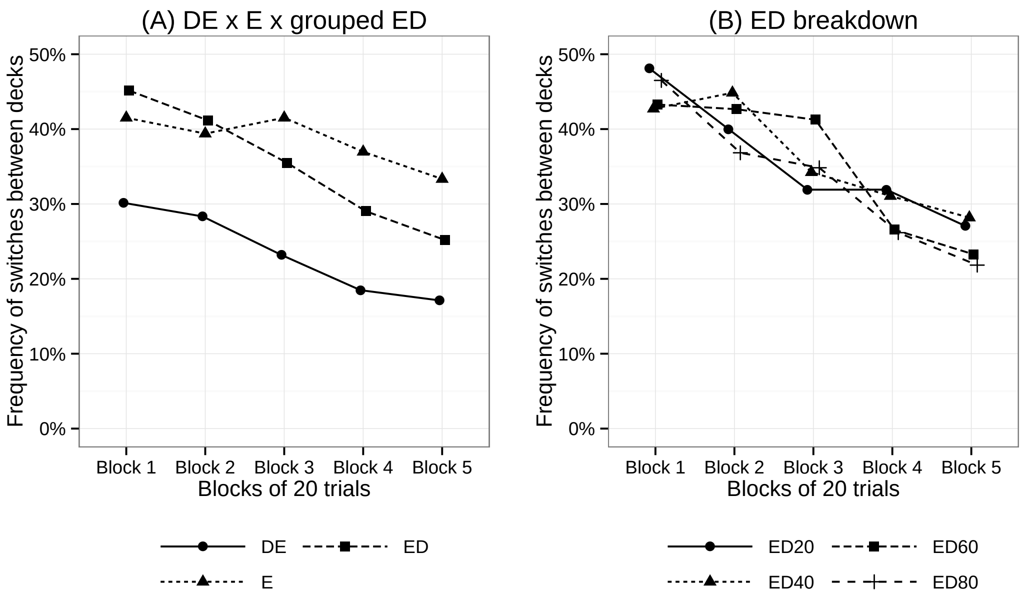

Figure 4.2: Evolution of the switching rates between decks in each block of 20 trials for Experiment 7. (A) High level comparison between E, DE and the average combined ED conditions, showing that most of the differences are a result of the presence of descriptions from the first trial (DE condition). (B) Detailed breakdown of the underlying ED conditions, showing that switching rates progressed very similarly in the different ED conditions.

The single analysis including all experimental conditions returned no significant effect for the presence of descriptions ($\chi^2(5) = 9.87, p = .08$). As before, this was a combination of a significant difference between the grouped high-level experimental conditions (DE=23.45%, ED=35.20%, E=38.55%; $\chi^2(2) = 9.71, p = .008$), but no significant difference between the underlying individual ED conditions

($\chi^2(3) = 0.32, p = .96$). A post-hoc analysis, and Figure 4.2B, show that the effect is a result of the significantly lower switching rates in the DE condition in comparison to the other conditions (against ED: $t(192) = 2.78, p = .02$, against E: $t(192) = 2.83, p = .01$), with no significant difference between the E and ED conditions ($t(192) = 0.79, p = .71$). Descriptions present from the beginning of the task had the most effect on switching rates, with no significant differences between no descriptions and delayed descriptions on the explorative behaviour.

There was a significant effect of block overall, with a negative slope ($b = -0.72, \chi^2(1) = 89.36, p < .001$), with a lower switching rate over time, consistent with Experiment 4 in Chapter 3, showing that participants reduce switching rate over time as they learn about the task and reduce uncertainty and change from exploration into exploitation (Block1=42.05%, Block5=25.10%, $t(293.66) = 8.60, p < .001$). The interaction was not significant ($\chi^2(5) = 8.84, p = .12$).

Impact of descriptions on behaviour in the ED conditions

To verify how prior experience influenced behaviour at the point of appearance of descriptions, a metric for the impact of descriptions on behaviour was generated by calculating the difference in selection of advantageous decks in the blocks immediately before and after the trial in which descriptions appeared for each condition. For example, in condition ED60, in which descriptions were first displayed after trial 60, the behavioural impact of descriptions was calculated by taking the frequency of selection of advantageous decks in block 3 (trials 61–80), after descriptions appeared, minus that in block 2 (trials 41–60), before descriptions appeared (Figure 4.3). Positive figures indicate that participants shifted their selections towards more advantageous decks, or a positive impact from observing the descriptions. Only the ED conditions were analysed in this way, since for the DE and E conditions there were no blocks of trials before and after descriptions, respectively, for comparison.

The impact of descriptions was analysed with a one-way ANOVA, with the four ED conditions as between-subjects factors, using a polynomial contrast. Overall, there was a significant difference in the impact of descriptions between the

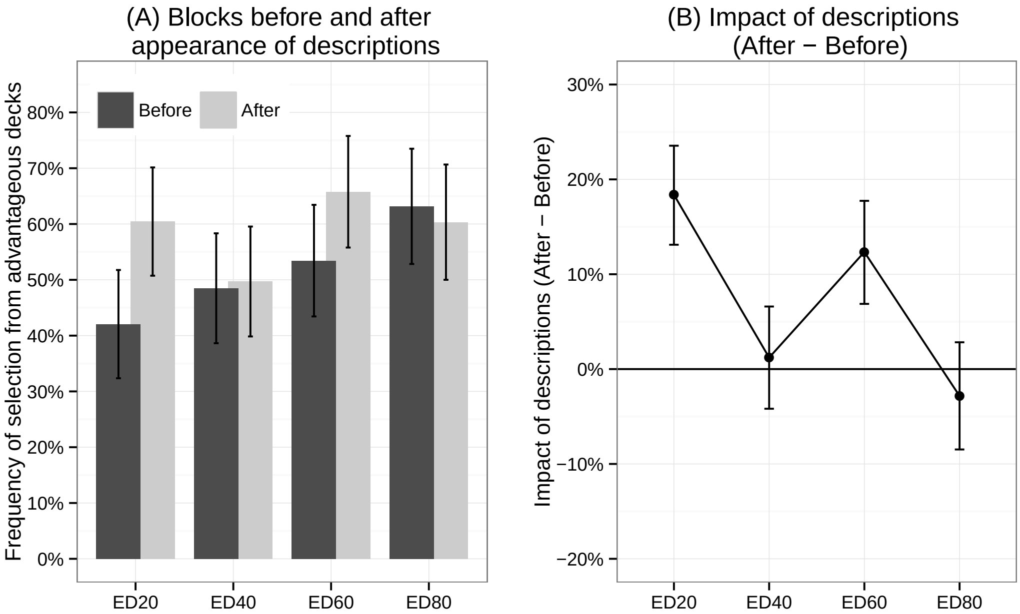

Figure 4.3: Impact of descriptions on selection of advantageous decks in Experiment 7. (A) Comparison of the blocks immediately before and immediately after the appearance of descriptions in each experimental condition. (B) Impact of descriptions (block after minus block before descriptions) in each experimental condition, showing the diminishing impact of descriptions on behaviour with an increase in prior experience.

different ED conditions ($F(3,125) = 2.87, p = .039$). The linear contrast for the influence of descriptions against the experimental condition was negative and significant ($b = -2.35, t(125) = 2.02, p = .046$). This indicates that with higher levels of prior experience, descriptions had lower impact, as predicted (Figure 4.3).[1]

4.1.3 Discussion

In the current experiment, participants in the DE experimental condition who had access to descriptions performed better in the task than participants who in the E experimental condition had to rely on experience alone, replicating the results observed in Chapter 3. Participants also switched less often between options in the DE condition in comparison to the E condition. In a complex task such as the IGT, descriptions are taken into consideration by participants and integrated into the decision-making process, helping them identify the advantageous decks (see

[1] In the next experiments I will show how the impact of descriptions was also influenced by the accumulation of points throughout the task. However in this experiment the accumulated points were not influential across the ED conditions because they did not significantly differ between the experimental conditions ($p = .20$) or over time between blocks ($p = .74$).

also Experiments 4 and 5). The findings from Barron et al. (2008), who showed that participants who had access to descriptions prior to any experience, from the first trial, performed better than participants who received descriptions halfway through the task, were also replicated here with the significant difference between the DE and ED experimental conditions.

There was no observable difference in average behaviour overall, both in selection from advantageous decks and switching rates, across all blocks, between participants in the E experimental condition who did not receive any descriptions and participants in the ED conditions who were presented with delayed descriptions. This highlights the importance of early presentation of descriptions to ensure their maximum influence, before any prior experience, and is likely due to strong primacy and inertia effects. It appears that presentation of descriptions before any prior experience (condition DE) is the main driver of the influence of descriptions in selections and switching rates observed in this task.

When comparing the different ED experimental conditions, and the behaviour observed before and after the appearance of descriptions for each participant, the hypothesis that more experience leads to a lower impact of descriptions was confirmed. Overall, across all conditions, the appearance of descriptions led to an increase in selection from advantageous decks, but the impact of descriptions was lower for participants who were presented the descriptions later in the task, and therefore had more time to accumulate larger amounts of prior experience. Thus confirming the importance of early descriptions and warnings, before experience, as proposed by Barron et al. (2008). The positive influence of descriptions on behaviour observed earlier in this dissertation with between-subjects conditions was also confirmed here in a within-subjects analysis in the ED conditions, by showing that the same participants selected more often from the advantageous decks after the descriptions have been made available to them.

The findings in the current task might be a result of learning effects. It has been shown previously that behaviour in DfE tasks stabilises relatively early, at around trial 50, as seen in the sequential block analysis in Experiment 4 (see also

Ert & Erev, 2007). This has been confirmed with cognitive models which show that most of the learning happens when uncertainty is highest, such as the beginning of tasks or when encountering new scenarios (Dayan & Niv, 2008; Speekenbrink & Konstantinidis, 2015). Similar findings of early learning are reported in the IGT literature, with most participants reporting being generally aware of the contents of the decks and confident of their choices by trial 50 (Bechara et al., 1994, 1997). In the current task, because descriptions provided the same information as that which could have been learned experientially, then perhaps the lower impact of descriptions associated with an increase in prior experience in the current task was a result of participants having learned the task well at that point, and descriptions are not providing any novel information, reducing their potential helpfulness and consequent impact on behaviour over time. To control for the potential confounding effect of learning, in the next experiments in this chapter, I will modify the paradigm to ensure that participants cannot learn the task via experience alone, and that descriptions always provide novel information at the point they first appear.

4.2 Experiment 8

To exclude the confounding effect that learning might have had on the impact of descriptions in Experiment 7, by reducing their potential usefulness as a source of information, the next experiment was designed to be a task which could not be learned via experience alone, without the help of descriptions. This was accomplished by introducing high-loss rare events, and controlling the outcomes so that these are only experienced later in the task, and never before the appearance of the descriptions. Since these descriptions provided information about all the potential outcomes for each option, including the rare events which had not been personally experienced by participants before being revealed, they acted as warnings against them, and always provided novel information. Thus any observed differences on the impact of descriptions due to the experimental manipulations can no longer be attributed to participants learning the task or to the reduction of the usefulness of descriptions. The outcomes of the new task were concocted to create the same conflict as in the IGT: Via experience alone, and in the short term, decks A and B were more attractive; however if participants took into account the descriptions, which warned them against large rare losses which occurred later in the task, then decks C and D were the long-term advantageous choices (cf. the IGT in Bechara et al., 1994, and in Experiment 4 in Chapter 3).

4.2.1 Method

Design

Experiment 8 was a between-subjects design with five experimental conditions, manipulating when descriptions were presented to participants as in the prior experiment: one description-before-experience condition, with descriptions available from the first trial (DE); and four experience-before-description conditions, with descriptions appearing at different points during the task (ED20, ED40, ED60, and ED80). The experimental conditions were named as in Experiment 7, with the numbers after the ED representing the trial at which descriptions were first shown. Each participant was allocated to only one experimental condition. The task lasted for 120 trials. Crucially, the descriptions warned participants of high-loss low-frequency events that only occurred in the last 20 trials of the task, and therefore always provided novel information, since these events never occurred before the descriptions appeared.

Participants

150 participants (92 females; age: M = 36.5 years, SD = 11.3 years) were recruited on-line using Amazon's Mechanical Turk service, with an average of 30 participants per experimental condition (N for each condition: DE=31, ED20=29, ED40=30, ED60=31, ED80=29). Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.20 for participating and an additional bonus according to the outcomes of the choices they made during the experiment (Bonus: M = US$ 0.22, SD = US$ 0.43).

Task

The overall format of the task was based on the IGT and closely followed the previous experiments in this dissertation, with four decks of card on screen, with their horizontal on-screen positioning counterbalanced. Cards were shown face-down, and after each trial a card from the selected deck only was revealed, using a partial-feedback approach. Decks A and B were the High-Risk options, because their selection throughout the task would lead to short-term gains but overall losses in the long-term, with large losses in the last 20 trials, and decks C and D were the Low-Risk options, because they led to lower short-term gains, but overall gains in the long-term.

Table 4.1: Actual card composition and wording of descriptions shown underneath each deck in Experiment 8. The descriptions provided a true representation of the outcomes across all trials, however the rare events were concentrated in the last 20 trials. The expected value for each individual card for the first 100 trials, excluding the rare events, was 62.5, 50, 42.5 and 35 points for decks A, B, C, and D, respectively. Across all trials, including the rare events, the EV in decks A and B was –46.5 and –48 points respectively, and in decks C and D was +25 points for both.

Actual experienced outcomes			
First 100 trials:			
Deck A	Deck B	Deck C	Deck D
90% of cards: +75 pts	80% of cards: +75 pts	90% of cards: +50 pts	80% of cards: +50 pts
10% of cards: -50 pts	20% of cards: -50 pts	10% of cards: -25 pts	20% of cards: -25 pts
Last 20 trials:			
Deck A	Deck B	Deck C	Deck D
90% of cards: +75 pts	80% of cards: +75 pts	90% of cards: +50 pts	80% of cards: +50 pts
10% of cards: -5500pts	20% of cards: -2500pts	10% of cards: -900 pts	20% of cards: -275 pts
Description labels shown underneath each deck			
From the first trial in condition DE and after trial *n* in condition ED*n*			
Deck A	Deck B	Deck C	Deck D
90% of cards: +75 pts	80% of cards: +75 pts	90% of cards: +50 pts	80% of cards: +50 pts
8% of cards: -50 pts	16% of cards: -50 pts	8% of cards: -25 pts	16% of cards: -25 pts
2% of cards: -5500 pts	4% of cards: -2500 pts	2% of cards: -900 pts	4% of cards: -275 pts

The task was structured so that in the first 100 trials the High-Risk choices had higher EV than the Low-Risk choices. In order to ensure that the descriptions provided novel information to participants, so that they could not simply learn experientially as in Experiment 7, all of the high-loss low-frequency events occurred in the last 20 trials of the task. In the first 100 trials of the task, decks A and B had a higher Expected Value (EV) of +62.5 and +50 points respectively, while decks C and D had lower EV of +42.5 and +35 points respectively. Therefore based on experience alone, decks A and B could be considered the most attractive decks. The purpose of this set-up was for participants to learn experientially that decks A and B are more attractive, until the appearance of the descriptions. Based on the information provided by descriptions, and if all 120 trials are considered, decks A and B had negative EV, of –46.5 and –48 points respectively, in comparison to the positive EV of decks C and D, both +25 points. These EV were the true EV of each deck based on all trials, as the last 20 trials included all the high-loss low-frequency events to compensate for the fact that they did not occur earlier in the task. For example, according to the descriptions in Deck A, across 120 trials, participants should have observed 2% of cards, or 2.4 cards on average, with –5500 points. These cards never appeared in the first 100 trials, but instead all 2.4 cards appeared in the last 20 trials, thus completing the overall distribution of cards. This was done by replacing the –50 points cards in the last block with –5500 points cards in Deck A, and so on accordingly for the other decks. Participants started the task with 2000 points, worth US$ 0.20. The task was self-paced over 120 trials and was completed on average in 9.29 minutes (*SD*=3.31).

Before the appearance of descriptions, the feedback experienced by participants should result in a higher selection from the High-Risk options, as these provided higher EV. After the appearance of descriptions, if these influenced behaviour and participants took into account the novel information displayed, then behaviour should shift towards the Low-Risk options. The amount of this shift in behaviour can be analysed to determine the impact of descriptions across the different experimental conditions. The reason for constraining all the high-loss low-frequency events to the last 20 trials was to ensure that across all conditions, a clean comparison could be made in the 20 trials before and after descriptions appeared, without being contaminated by the high-loss events. In comparison, Barron et al. (2008) allowed the rare event to occur randomly throughout the task but later eliminated those participants who had observed the rare event, due to its unclear influence on behaviour, in an approach they called "potentially objectionable" (p. 129) and also claimed to have "remained somewhat concerned that the results could be biased by the elimination of the few participants who observed the rare outcome" (p. 131). With the current approach used here, no participant was excluded from the analysis (for additional thoughts on the behavioural effects of experiencing rare negative events, see also Weinstein, 1989; Yechiam, Barron, & Erev, 2005; Yechiam et al., 2006).

4.2.2 Results

Selections from Low-Risk decks

The main dependent variable was the frequency of cards selected from the Low-Risk decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 4.4). If descriptions influenced be-

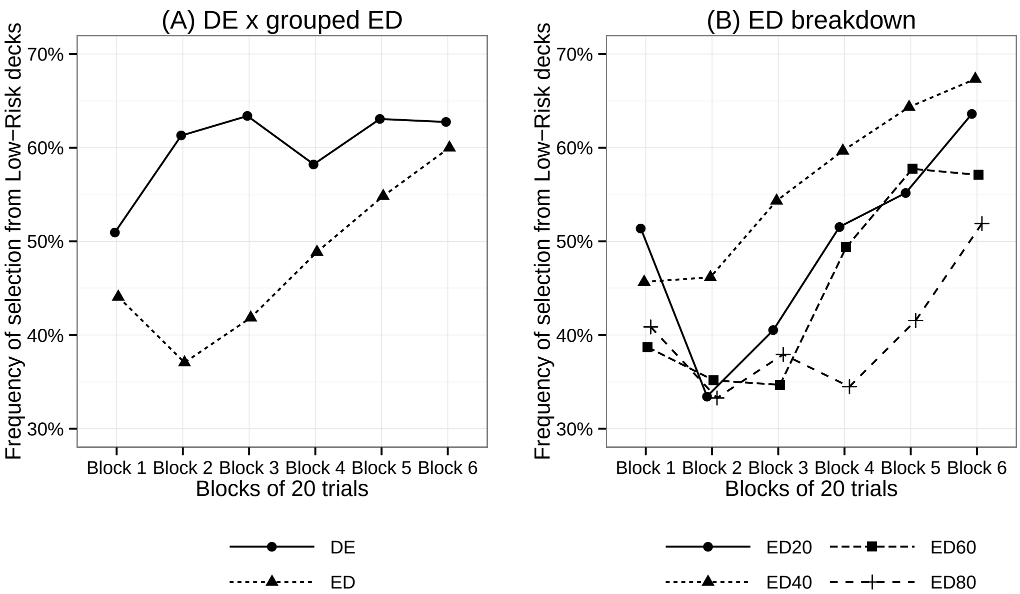

Figure 4.4: Evolution of the selection of Low-Risk decks for each block of 20 trials for Experiment 8. (A) High level comparison between the DE and the average combined ED conditions, showing that most of the differences are a result of the presence of descriptions from the first trial (DE condition). (B) Detailed breakdown of the underlying ED conditions, showing that the behaviour progressed very similarly in the different ED conditions, with a negative slope when descriptions were not present, changing to a positive slope after they were displayed.

haviour, as observed in the previous chapters in this dissertation, then participants should select more from Low-Risk choices when descriptions are present, since their informational contents warn the subjects about the high-loss low-frequency events which are more severe in High-Risk decks, which should consequently normatively move behaviour away from those options. A similar linear mixed-effects model approach as in the previous experiment was employed here, with the same fixed and random components, with one fewer experimental condition (the E condition was not included in the current task).

The manipulation for the timing of descriptions overall was significant, across all five conditions ($\chi^2(4) = 16.93, p = .002$). A post-hoc pairwise comparison with Tukey adjustments showed that the difference was mostly driven by a difference between the DE and average across all ED conditions (DE=59.95%, ED=47.80%, $t(148) = 2.62, p = .01$). A comparison between the four different ED conditions was not significant (D20=49.30%, D40=56.25%, D60=45.45%, D80=40.00%, $\chi^2(3) = 7.65, p = .054$). As before, participants performed significantly better in the task, selecting more from Low-Risk decks, when they had descriptions available from the first trial, before any prior experience, but the differences in overall performance according to the amount of prior experience was not observed overall across all trials.

The effect of block was significant with a positive slope ($b = 0.70, \chi^2(1) = 21.99, p < .001$), as before showing that participants learned the composition of the decks over time and improved their performance by selecting more from Low-Risk decks as the task progressed (Block1=45.50%, Block6=60.54%, $t(177.93) = 3.80, p = .003$). The interaction between experimental condition and block was not significant ($\chi^2(4) = 3.27, p = .51$). A separate analysis including the presence of descriptions (instead of experimental condition) and block showed a significantly higher selection from Low-Risk decks in blocks in which descriptions were available, compared with blocks in which descriptions were not present (Descr.=50.54%, No-Descr.=40.42%, $t(738.26) = 3.57, p < .001$). As in the previous experiment, there was a significant interaction between the presence of description (collapsed across all experimental conditions) and block ($\chi^2(1) = 19.62, p < .001$). As expected, due to the way that experience and descriptions were structured, in the current experiment the slope was negative when descriptions were not present ($b = -0.43$), as participants were learning via experience alone and selecting more often from the High-Risk decks before the appearance of the high-loss rare events; and when descriptions were present, the slope was positive ($b = 0.82$), as participants then shifted their behaviour towards more Low-Risk choices over time (Figure 4.4B).

Impact of descriptions on behaviour in the ED conditions

As in the previous experiment, a metric for impact of descriptions on behaviour was calculated for each participant. This metric evaluated the difference in selections from Low-Risk decks, by subtracting the selection rate in the block immediately preceding the appearance of descriptions from that of the block immediately after. Positive figures indicate that participants shifted their selections towards Low-Risk

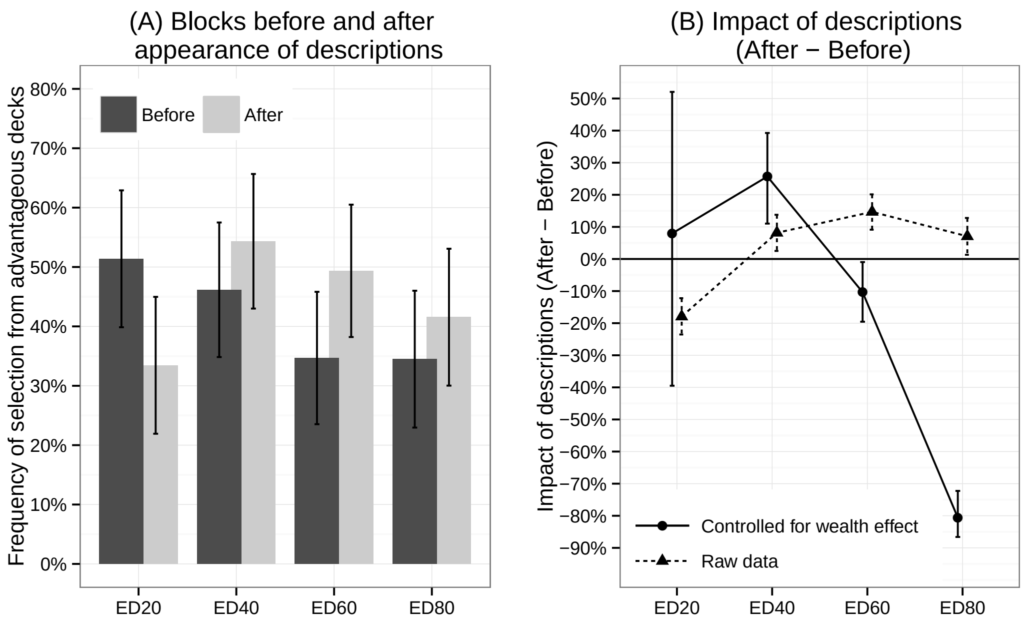

Figure 4.5: Impact of descriptions in selection of Low-Risk decks in Experiment 8. (A) Comparison of the blocks immediately before and immediately after the appearance of descriptions in each experimental condition. (B) Impact of descriptions (block after minus block before descriptions) in each experimental condition, showing the reduced impact of descriptions on behaviour in relationship to prior experience after controlling for accumulated wealth, but a positive relationship in the raw data.

decks in response to the descriptions, which warned them that the Low-Risk decks were more attractive overall. Only the ED conditions were analysed in this way, since the DE condition did not have any trials before descriptions, for comparisons (Figure 4.5).

The impact of descriptions was analysed with a one-way ANCOVA, with the four ED conditions as between-subjects factors, using a polynomial contrast. A covariate variable for the total amount of points that participants had accumulated at the trial in which descriptions first appeared was also included in the model. This was done to exclude the influence that accumulated wealth had on risk taking. I expected participants who had accumulated more points at the trial in which the descriptions first appeared to be more observant of the large losses indicated in the text attached to each deck, and shift more towards Low-Risk decks to preserve their winnings, similar to an endowment effect (Kahneman, Knetsch, & Thaler, 1991). In comparison, participants who had not had much success and had not accumulated many points should be more likely to engage in additional risk to enlarge their winnings, when presented with descriptions.

Overall, there was a significant difference in the impact of descriptions between the different ED conditions ($F(3,111) = 10.92, p < .001$). The linear contrast for the impact against the experimental condition was negative and significant ($b = -50.26, t(111) = 2.67, p = .009$). This indicates that with higher levels of prior experience, descriptions had a lower impact, for the same amount of accumulated points (Figure 4.5B). As predicted, the relationship between accumulated points and impact of descriptions was also significant and with a positive coefficient ($b = 0.008, F(1,111) = 37.66, p < .001$). Across all conditions, participants with higher accumulated points shifted more towards the Low-Risk decks, when descriptions appeared, in order to preserve their winnings reacting to the high-losses warnings in the text. The interaction between timing of descriptions and accumulated points was not significant ($p = .45$).

It is possible to also evaluate the actual raw impact of descriptions before controlling for accumulated wealth (as shown using the dotted line in Figure 4.5B). In this case, the contrast is still significant but positive ($b = 3.65, t(115) = 2.39, p = .02$). This indicates a significant positive correlation between accumulated points and the trial in which descriptions first appeared in the actual observed behaviour without controlling for accumulated wealth. This is because participants in the later experimental conditions had accumulated more points by the time descriptions first appeared, and therefore their reaction to the warnings was to shift more towards the Low-Risk decks to protect their winnings. Participants in the ED20 condition had not accumulated many points yet, and therefore had little to lose. This might also explain why in that experimental condition the appearance of descriptions, warning them against the high-losses associated with the High-Risk decks, actually led to a significant reduction in selections from the Low-Risk decks (Before=51.35%, After=33.45%, $t(115) = 2.34, p = .02$), equivalent to an increase in selection from High-Risk decks. This might be seen as a perverse effect of descriptions, which led to the opposite desired behaviour expected from such a warning.

4.2.3 Discussion

Similarly to the findings from the previous experiments in this dissertation, the presence of descriptions in the current task did influence behaviour in the predicted normative direction. Descriptions helped participants identify and select the Low-Risk decks more often, with their higher long-term rewards, therefore avoiding the large losses associated with the High-Risk decks, according to the warnings contained in the text of the descriptions. Over the course of the task, as the trials progressed, participants selected increasingly more High-Risk decks when descriptions were not present, and shifted towards increasingly more Low-Risk decks when descriptions were displayed on screen, aligned with the differences in information between experience and descriptions.

The endowment effect refers to the finding that individuals allocate higher values to what they already own (Kahneman et al., 1991; Rick, 2011), and thus are more reluctant to part with their belongings, via associations with the self (Morewedge, Shu, Gilbert, & Wilson, 2009). This well-explored phenomenon implies that the impact of descriptions should be strongly moderated by the amount of accumulated wealth for each participant, in the form of points which translated into financial rewards for their participation in the task. Participants who had accumulated more points at the time that descriptions first appeared reacted more strongly to descriptions and shifted towards the Low-Risk decks more, in order to safeguard their financial earnings, as they had more to lose: Descriptions warned them of potential high losses that they wanted to avoid. The perverse effect of this was that many participants decreased their selections from Low-Risk decks when descriptions first appeared, taking more risk, the opposite of the desired result from the descriptions, in particular those who had not accumulated many points before descriptions first appeared (Figure 4.6). I believe this might have been a result of some participants myopically focusing on the positive rewards which were also indicated in the descriptions, and those were higher in the High-Risk than in the Low-Risk choices. Perhaps those participants were content with taking the additional risk to gather some of those higher rewards, despite the presence of the descriptions and their warnings, as a short-term strategy. While this was an initial reaction immedi-

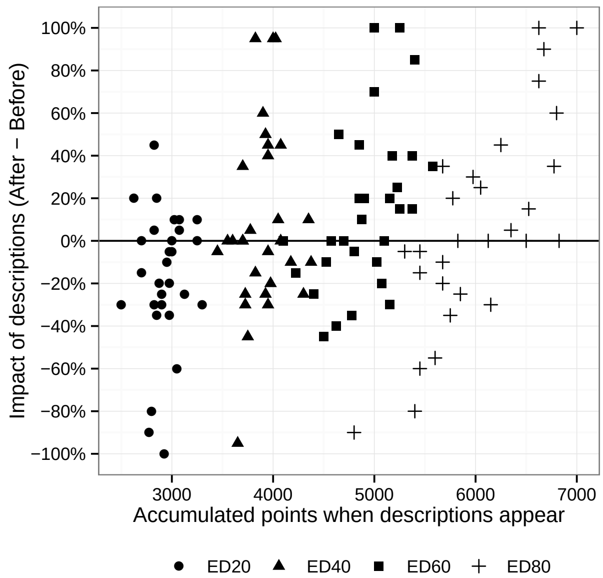

Figure 4.6: Impact of descriptions in selection of Low-Risk decks in Experiment 8 in relationship to the amount of accumulated points at the point descriptions first appeared. The scatterplot shows that participants who had accumulated more points, shifted more towards the Low-Risk decks. In addition, because points were accumulated over time, participants in later experimental conditions had accumulated more points, as can be seen by the clear separation between experimental conditions across the x axis.

ately following the appearance of descriptions, with participants later reducing their risk-taking, in accordance with the descriptive warnings, as observed in condition ED20 for example (see Figure 4.4B), this finding has important implications for the creation of effective warnings that reduce risk taking. Experiment 9 in the next section will test this hypothesis by removing the positive aspects of the descriptions, making them more akin to warning labels by referring to the large losses only.

After controlling for the effect of accumulated wealth, the hypothesis for the negative relationship between the impact of descriptions on behaviour and prior experience predicted and observed in Experiment 7 was replicated here in Experiment 8 (Figure 4.5B). In the previous experiment, employing the IGT, this relationship could have been a result of incremental learning via experience over time, with descriptions not providing any novel information to participants who had already the chance and time to learn the task experientially. This potentially confounding effect was eliminated in the current experiment, in which the descriptions always provided novel information that could not have been learned experientially, and this additional information acted as warnings against low-frequency high-loss events that never occurred experientially before the descriptions appeared, with the same behavioural results. This is analogous to many real-life warning situations, where individuals might be presented with a warning sign or label informing them of a risk which they never experienced before, but that it would be beneficial for them to take into account when considering the risk-reward balance of the situation.

The raw behavioural data, before controlling for wealth, showed an increase in the impact of descriptions with more prior experience. This is likely a result of the progressive accumulation of wealth over time, as in the task all the options had positive EV until trial 100, which means that participants on average accumulated points as the task evolved, and later appearance of descriptions was associated both with an increase in prior experience and also an increase in accumulated wealth. In the current experiment, increases in prior experience were confounded with increases in accumulated wealth. In this case it is expected that any descriptions that warn individuals of high losses must be seen in the context of the wealth they have, and how much they could lose. What the analysis shows is that the increase in the impact of descriptions over experimental conditions was not enough to have been predicted by the associated increase in wealth over time for those conditions, and therefore the controlled analysis confirmed the opposite, negative, relationship. In Experiment 10, I will eliminate this confounding effect of wealth accumulation by creating a task that has zero EV over time, therefore not allowing participants to accumulate any points as the task progresses, and ensuring that all participants are at the same level of wealth when descriptions first appear across the different experimental conditions, dissociating wealth accumulation from prior experience.

4.3 Experiment 9

Experiment 9 was designed to verify the hypothesis that the behaviour observed in Experiment 8 was partly a result of participants also taking into account the descriptive information about the positive rewards and the high-frequency events, which might have reduced the impact of the descriptions' warnings against the high-loss low-frequency events. This was accomplished with a simple change to the paradigm in comparison with the previous experiment, which was to include only information about the high-loss low-frequency events in the descriptions, making it closer to a warning, as opposed to a full description of outcomes. Such warnings might be more representative of warnings used in real life scenarios, which tend to highlight the potential for high losses, and should be more efficient in influencing behaviour and reducing risk taking.

4.3.1 Method

Design

The same design from Experiment 8 was used, albeit with only four between-subjects experimental conditions: the experience-before-description (ED) conditions, with descriptions appearing at different points during the task (ED20, ED40, ED60, and ED80), allowing for different levels of accumulation of prior experience. All participants started the task without descriptions, with experience only. The numbers after the ED represent the trial at which descriptions were first shown to participants. Each participant was allocated to only one experimental condition.

Participants

122 participants (73 females; age: *M* = 35.8 years, *SD* = 11.4 years) were recruited on-line using Amazon's Mechanical Turk service, with an average of 30.5 participants per experimental condition (*N* for each condition: D20=31, D40=30, D60=31, D80=30). Participation was restricted to individuals whose location was defined as in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.20 for participating and an additional bonus according to the outcomes of the choices they made during the experiment (Bonus: *M* = US$ 0.22, *SD* = US$ 0.44).

Task

The only changes in the task in comparison to Experiment 8 were the alteration to the labels used in the descriptions, and the exclusion of the DE condition. Instead of full descriptions of the outcomes within each deck, in the current experiment only the high-loss low-frequency event for each deck was shown, making the descriptions more closely resemble warnings (Table 4.2). The composition of cards and set-up was the same as in Experiment 8. As in the previous experiment, the rare outcomes indicated in the descriptions only appeared in the last 20 trials, therefore when descriptions were first displayed, the information conveyed was always new to participants. The task was self-paced over 120 trials and was completed on average in 9.32 minutes (*SD*=4.10).

Table 4.2: New wording of descriptions shown underneath each deck in Experiment 9.

Description labels			
Deck A	Deck B	Deck C	Deck D
Large losses in this deck (out of 100 cards) 2 cards: -5500 pts	Large losses in this deck (out of 100 cards) 4 cards: -2500 pts	Large losses in this deck (out of 100 cards) 2 cards: -900 pts	Large losses in this deck (out of 100 cards) 4 cards: -275 pts

4.3.2 Results

Selections from Low-Risk decks

As in the previous experiment, the main dependent variable was the frequency of cards selected from the Low-Risk decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 4.7). A similar linear mixed-effects model approach as in the previous Experiment was employed here, with one fewer experimental condition (the DE condition was not included in the current task).

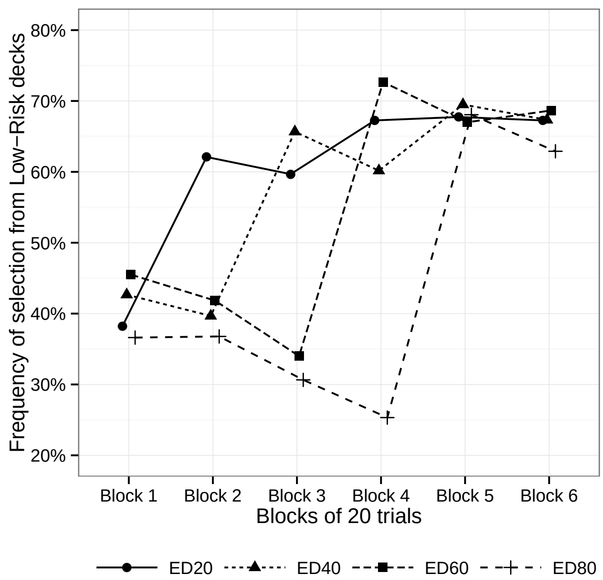

Figure 4.7: Evolution of the selection of Low-Risk decks for each block of 20 trials for Experiment 9, across the different ED experimental conditions, showing a clearer effect of the influence of the warning labels on behaviour when they were first displayed, by shifting selections from High-Risk towards Low-Risk options. The number after the ED identifies the trial in which descriptions were first revealed to participants.

The manipulation for the delayed presentation of descriptions overall was significant ($\chi^2(3) = 16.05, p = .001$). A linear polynomial contrast was significant with a negative slope ($b = -10.70, t(118) = 3.34, p = .001$), but not the quadratic or cubic ones ($ps > .23$), showing that selection from low-risk decks reduced monotonically with increases in prior experience (D20=60.35%, D40=57.50%, D60=54.95%, D80=43.40%). However this was most likely a result of participants with less prior experience having fewer trials with descriptions. The effect of block was significant and positive ($b = 1.18, \chi^2(1) = 54.40, p < .001$), which shows that, as in the previous experiments, participants were learning the composition of the decks over time and improving their performance by selecting more from Low-Risk decks as the task progressed (Block1=40.75%, Block6=66.55%, $t(147.65) = 6.09, p < .001$). The interaction between experimental condition and block was not significant ($\chi^2(3) = 0.67, p = .88$).

As before, the analysis on the presence of descriptions and block showed a significantly higher selection from advantageous decks in blocks in which descriptions were available, compared with blocks in which descriptions were not present (Descr.=50.54%, No-Descr.=65.10%, $t(583.25) = 12.21, p < .001$). There was again a significant interaction between the presence of descriptions (collapsed across all experimental conditions) and block ($\chi^2(1) = 15.02, p < .001$). As in Experiment 8, the slope was negative when descriptions were not present ($b = -0.90$) and positive when descriptions were present ($b = 0.20$). This interaction can be seen in Figure 4.7.

Impact of descriptions on behaviour

A metric for the impact of descriptions on behaviour was again calculated, comparing the selection of Low-Risk decks in the block immediately before and after the first appearance of the labels. In the current experiment the impact of the appearance of descriptions shifting behaviour towards higher selection from Low-Risk decks can be clearly seen in Figure 4.7. Positive figures indicate that participants shifted their selections towards Low-Risk decks in response to the descriptions, which warned them that the those decks were more attractive overall (Figure 4.8). The impact of descriptions was analysed with the same model from Experiment 8, using a one-way ANCOVA, with the four ED conditions as between-subjects factors, and a covariate variable for the total amount of points accumulated when descriptions first appeared.

Overall, there was a significant difference in the impact of descriptions between the different ED conditions ($F(3,114) = 5.28, p = .002$). The linear contrast for the impact of descriptions against the experimental condition was significant

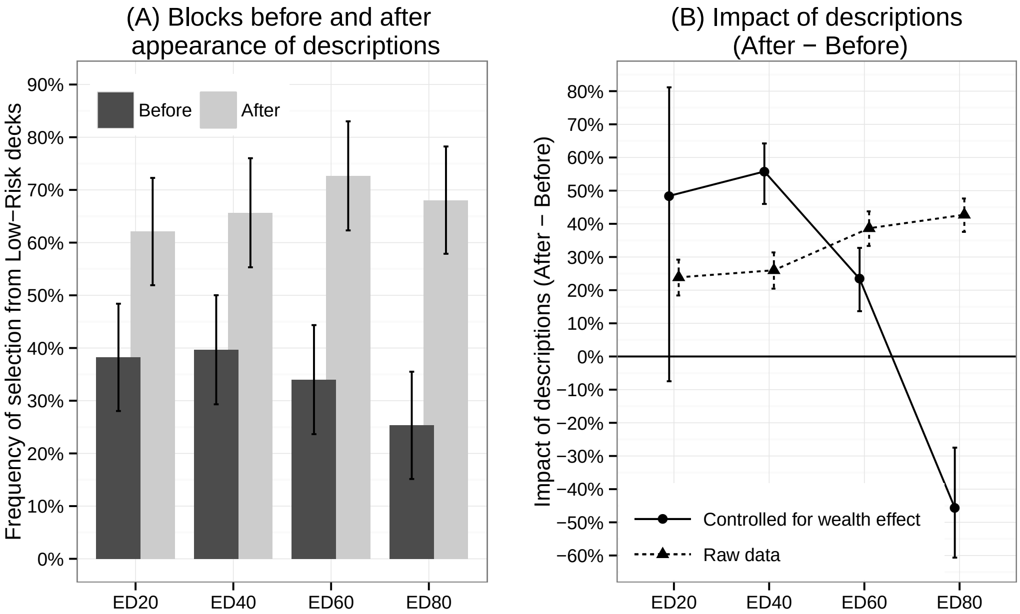

Figure 4.8: Impact of descriptions in selection of Low-Risk decks in Experiment 9. (A) Comparison of the blocks immediately before and immediately after the appearance of descriptions in each experimental condition. (B) Impact of descriptions (block after minus block before descriptions) in each experimental condition, showing the negative relationship between the impact of descriptions on behaviour and prior experience after controlling for accumulated wealth, but a positive slope in the raw data.

with a negative slope ($b = -3.6, t(114) = 3.50, p < .001$), replicating the results from the previous experiment. This indicates that with higher levels of prior experience, descriptions had a lower impact on behaviour (Figure 4.8). As predicted, the relationship between accumulated points and the difference was also significant and with a positive coefficient ($b = 0.008, t(114) = 2.94, p = .004$). Participants with higher accumulated points shifted more towards the Low-Risk decks. The interaction between timing of descriptions and accumulated points was not significant ($p = .69$). As before, the raw data before controlling for accumulated wealth returned a significant but positively sloped linear contrast for the ED conditions ($b = 3.10, t(118) = 2.26, p = .026$), with no increases in risk taking this time.

Comparing Experiment 9 with Experiment 8

In order to analyse how the different wording of the labels influenced behaviour, the impact of descriptions in Experiment 9 was compared with those from the equivalent ED experimental conditions of Experiment 8, by combining the two experiments into a single ANCOVA for the impact of descriptions.

The effect of the wording of the labels was significant, with a larger impact of descriptions in Experiment 9 than in Experiment 8, and a large effect size (Exp9=32.81%, Exp8=2.99%, $F(1,233) = 34.01, p < .001, d = 0.76$). The descriptions in Experiment 9 providing warnings only, and no information about the gains, were more efficient in directing a desired behavioural shift towards the safer Low-Risk choices. The combined analysis also returned a significant positive relationship between accumulated wealth and impact of descriptions ($b = 0.008, t(225) = 3.28, p = .001$), a significant linear and negative contrast for the experimental conditions after controlling for accumulated wealth ($b = -50.26, t(225) = 2.70, p = .007$) and positive linear contrast before controlling for it ($b = 3.65, t(233) = 2.48, p = 0.014$), replicating the results observed in the underlying individual experiments.

4.3.3 Discussion

As predicted, presenting descriptions in the form of warnings, which refer only to the large rare losses, increased the magnitude of the predicted behavioural shift towards Low-Risk decks, which traditionally is the desired reaction when deploying such warnings. By not mentioning the positive outcomes associated with the choices, the perverse effect of the descriptive labels observed in the previous experiment, where participants who had not accumulated many points, in particular those in the early-presentation ED20 condition, was mostly eliminated, with few participants now taking more risk when descriptions were first shown (Figure 4.9).

The descriptions used in this experiment were more efficient than in the previous experiment. The impact of descriptions were significantly higher in Experiment 9 than the ones in the equivalent experimental conditions of Experiment 8, with a large effect size. Perhaps the reduction in information, by making participants focus only on the losses, made them more salient and efficient. Alternatively, the addition of positive information provides a perverse effect as it might have made the risky choices more interesting, at least in the short term, when comparing between the available alternatives, especially given that the High-Risk options had larger indi-

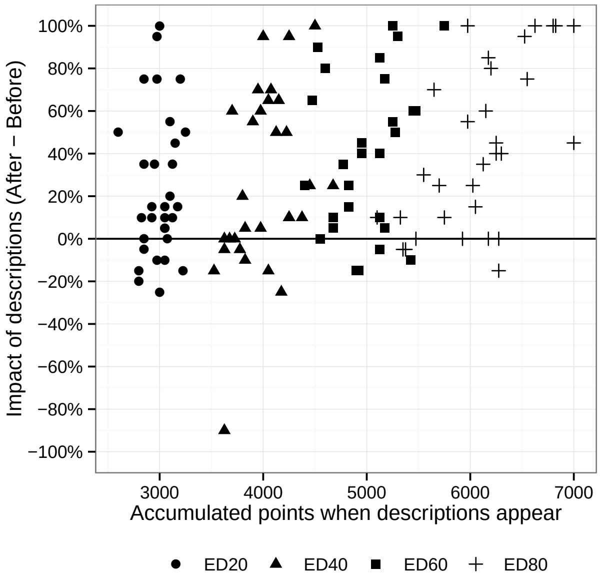

Figure 4.9: Impact of descriptions in selection of Low-Risk decks in Experiment 9 in relationship to the amount of accumulated points at the point descriptions first appeared. The scatter-plot shows that few participants shifted away from the Low-Risk decks in Experiment 9, in comparison with Experiment 8, as a result of the new descriptions used in the form of warnings.

vidual frequent positive outcomes (+75 against +50). If individuals were comparing these frequent positive outcomes only, the High-Risk could appear more attractive. By eliminating this information in Experiment 9, participants could no longer make this inference using descriptions. Further research in these dynamics is required.

In both Experiments 8 and 9, there was a significant correlation between block and accumulated wealth, which can be seen by the clear separation between the experimental conditions and the amount of accumulated points in Figures 4.6 and 4.9. Participants who were shown the descriptions later, and had accumulated more prior experience, had also accumulated more wealth. These two were therefore confounding effects. The next experiment will remove this relationship between experience and wealth by ensuring that participants do not accumulate points over time, only experience.

4.4 Experiment 10

Experiment 10 was designed to remove the effect of the accumulated wealth observed in the previous experiments, by creating a task with an initial expected value

(EV) of zero for each of its options. This will ensure that, on average, participants do not accumulate points over the course of the task, and instead stay at the same level of wealth throughout. During the first 100 trials, participants thus accumulated only experience and information, in the form of feedback after each trial, but not wealth. As in the previous experiments, the extreme rare outcomes were constrained to the last block of 20 trials, differentiating between the attractive and unattractive options, and allowing participants to accumulate (or lose) financial rewards, depending on their selections after the appearance of descriptions. By excluding the confounding effect linking prior experience with accumulated wealth, I expected to show the predicted negative relationship between the impact of descriptions and prior experience in the raw behavioural results, without the need for controlling for accumulated wealth differences, which should not be present in this experiment.

4.4.1 Method

Design

Experiment 10 was a between-subjects design with four experimental conditions, manipulating when descriptions were presented to participants: four experience-before-description conditions, with descriptions appearing at different points during the task (ED20, ED40, ED60, and ED80). The experimental conditions were named as in the previous experiments in this chapter, with the numbers after the ED representing the trial at which descriptions were first shown. Each participant was allocated to only one experimental condition. As before, the task lasted for 120 trials, with all high-loss low-frequency events mentioned in the descriptions restricted to the last 20 trials of the task, ensuring that the labels always provided novel information, since these events never occurred before the descriptions appeared.

Participants

135 participants (75 females; age: M = 35.4 years, SD = 12.1 years) were recruited on-line using Amazon's Mechanical Turk service, with an average of 33.75 participants per experimental condition (N for each condition: D20=35, D40=34, D60=33, D80=33). Participation was restricted to individuals whose location was defined as

in the United States. No participants were excluded from the analysis. Participants were paid a fixed amount of US$ 0.25 for participating and an additional bonus according to the outcomes of the choices they made during the experiment (Bonus: *M* = US$ 0.32, *SD* = US$ 0.57).

Task

The overall format of the task closely followed that of the previous experiments in this chapter, with four decks of cards on screen. Cards were shown face-down, and after each trial a card from the selected deck only was revealed, using a partial-feedback approach. As before, the low-frequency events were constrained to the last block of 20 trials. In the first 100 trials, all decks had an EV of zero, with the introduction of the larger low-frequency events in the last block of 20 trials differentiating between the decks. Two of the decks (A and B) were High-Risk decks with negative low-frequency events and overall negative EV of –60 points, and the other two (C and D) were Low-Risk decks with positive low-frequency events and overall positive EV of +25 points (Table 4.3).

Table 4.3: Actual card composition and wording of descriptions shown underneath each deck in Experiment 10. The descriptions provided a true representation of the outcomes across all trials, however the rare events were concentrated in the last 20 trials. The expected value for each individual card for the first 100 trials, excluding the rare events, was zero for all decks. Across all trials, including the rare events, the EV in decks A and B were –60 points, and in decks C and D they were +25 points.

Actual experienced outcomes			
First 100 trials:			
Deck A	Deck B	Deck C	Deck D
90% of cards: +250 pts	80% of cards: +250 pts	90% of cards: +100 pts	80% of cards: +100 pts
10% of cards: -2250 pts	20% of cards: -1000 pts	10% of cards: -900 pts	20% of cards: -400 pts
Last 20 trials:			
Deck A	Deck B	Deck C	Deck D
90% of cards: +250 pts	80% of cards: +250 pts	90% of cards: +100 pts	80% of cards: +100 pts
10% of cards: -5250pts	20% of cards: -2500pts	10% of cards: +350 pts	20% of cards: +225 pts
Description labels shown underneath each deck after trial *n* in condition ED*n*			
Deck A	Deck B	Deck C	Deck D
90% of cards: +250 pts	80% of cards: +250 pts	90% of cards: +100 pts	80% of cards: +100 pts
8% of cards: -2250 pts	16% of cards: -1000 pts	8% of cards: -900 pts	16% of cards: -400 pts
2% of cards: -5250pts	4% of cards: -2500pts	2% of cards: +350 pts	4% of cards: +225 pts

When all trials are taken into account, the descriptions are true representations of the cards contained within each deck. However prior to trial 100, the descriptions contradicted the experience. As in Experiments 8 and 9, the descriptions always provided novel information to participants, since they included the rare events, which never appeared before the descriptions were shown. The purpose of keeping the EV equal to zero up to trial 100 was to ensure that participants did not accumulate points over time, so that the behavioural response to the descriptions could be dissociated with any wealth accumulation effect, allowing for a cleaner influence of prior experience. Participants started the task with 5000 points, and points were converted into money at a rate of US$0.10 per 1000 points. The task was self-paced, lasted 120 trials and was completed on average in 9.17 minutes (*SD*=3.58).

4.4.2 Results

Selections from Low-Risk decks

The main dependent variable was the frequency of cards selected from the Low-Risk decks, calculated as the total number of cards selected from decks C and D for each sequential block of 20 choices (Figure 4.10). As before, if descriptions influenced behaviour normatively, then participants should select more from Low-Risk choices when descriptions are present. A similar linear mixed-effects model approach as in the previous experiment was employed here, with the same fixed and random components.

The experimental manipulation for the timing of presentation of descriptions overall was significant ($\chi^2(3) = 12.41, p = .006$). A linear polynomial contrast was significant with a negative slope ($b = -7.95, t(131) = 2.83, p = .005$), but not the quadratic or cubic ones (ps$>$ .18), showing that selection from low-risk decks reduced monotonically with increases in prior experience (D20=57.85%, D40=46.91%, D60=48.69%, D80=44.02%). This is most likely a result of participants with less prior experience having fewer trials with descriptions. The effect of block was significant and positive ($b = 0.56, \chi^2(1) = 19.13, p < .001$), indicating an improvement in performance by selecting more from Low-Risk decks as the task progressed, although this is most likely due to the appearance of descriptions

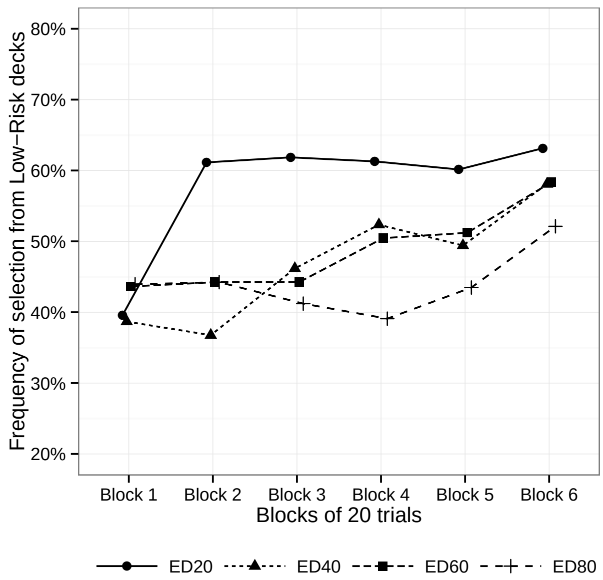

Figure 4.10: Evolution of the selection of Low-Risk decks for each block of 20 trials for Experiment 10, across the different ED experimental conditions, showing the impact on behaviour when descriptions are first shown. The number after the ED identifies the trial in which descriptions were first revealed to participants.

during the task. The interaction between experimental condition and block was not significant ($\chi^2(3) = 2.87, p = .41$).

As before, the analysis on the presence of descriptions and block showed a significantly higher selection from advantageous decks in blocks in which descriptions were available, compared with blocks in which descriptions were not present (Descr.=53.88%, No-Descr.=43.45%, $t(676.05) = 3.77, p < .001$). However there was no significant interaction between presence of description and block, with similar slopes in both cases (Slope for block: Descr.=0.10, No-Descr.=0.14, $t(553.76) = 0.15, p = .88$). This is likely a result of increase in selection from Low-Risk decks being concentrated at the point in which the blocks changed from no-description to description, with apparently limited change in selections in the other blocks, as can be seen in Figure 4.10. Because the options had an EV of zero until the last block, no specific trend over time was expected, as experiential feedback would not lead participants to prefer any specific decks in terms of EV comparisons alone.

Impact of descriptions on behaviour

As in the previous experiment, a metric for the impact of descriptions on behaviour was calculated, comparing the selection of Low-Risk decks in the block immediately before and after the first appearance of the labels. Positive figures indicate that participants shifted their selections towards Low-Risk decks in response to the descriptions, which warned them that the those decks were more attractive overall (Figure 4.11). The impact of descriptions was analysed with a one-way ANCOVA, with the four ED conditions as between-subjects factors, using a polynomial contrast, with a covariate variable for the total amount of points that participants had accumulated at the trial in which descriptions first appeared.

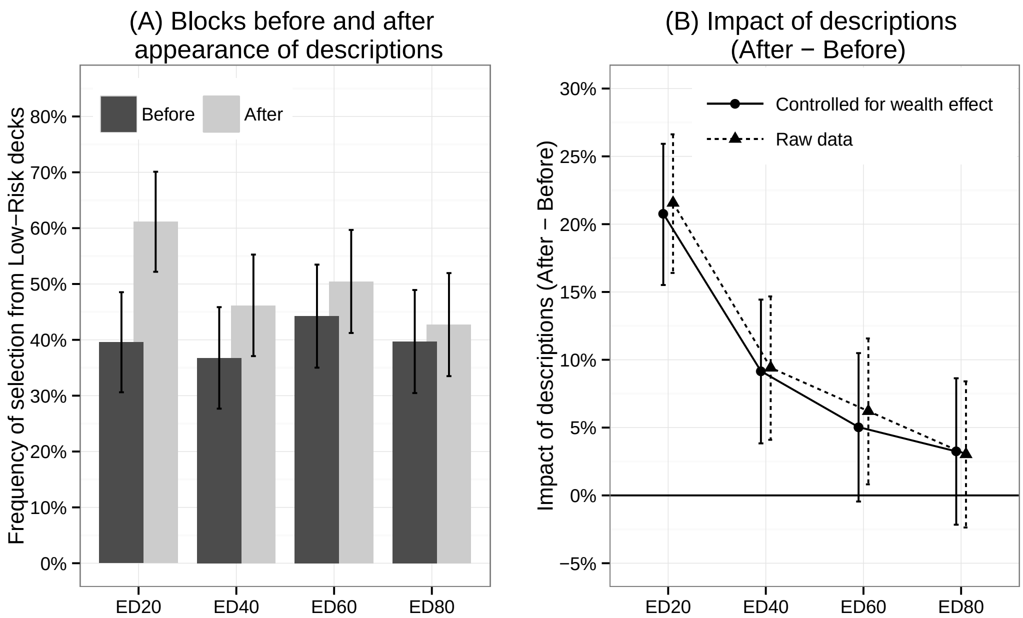

Figure 4.11: Performance improvement in selection of Low-Risk decks in Experiment 10. (A) Comparison of the blocks immediately before and immediately after the appearance of descriptions in each experimental condition. (B) Performance improvement (block after minus block before descriptions) in each experimental condition, showing the negative slope of the relationship between performance improvement and prior experience. There was no significant differences in accumulation of wealth across conditions hence the same results are present when these are controlled for.

Overall, the linear contrast for the effect of prior experience on impact of descriptions was significant with a negative slope ($b = -2.63, t(131) = 2.11, p = .04$), with the quadratic and cubic contrasts not significant (ps$>$ .48), although there was no significant difference between the individual ED conditions ($F(3,131) =$

$1.71, p = .17$). This indicates that with higher levels of prior experience, descriptions led to a reduction on the impact of descriptions (Figure 4.11B). As predicted, the relationship between accumulated points at the point of appearance and the impact of descriptions was not significant ($b = 0.0001, t(127) = 0.52, p = .60$), as accumulated points also did not differ between the experimental conditions (ED20=4777, ED40=4560, ED60=3968, ED80=4253, $F(3,131) = 0.49, p = .69$) and the linear contrast was not significantly different from zero ($b = -484, t(131) = .96, p = .34$). The interaction was not significant ($p = .45$).

4.4.3 Discussion

In this experiment, the predicted negative relationship between the impact of descriptions on behaviour and prior experience was observed, without the need to control for the effect of accumulated wealth. This was achieved by using options that returned an EV of zero throughout the first 100 trials, thus ensuring that the amount of points accumulated did not differ significantly between the experimental conditions. The data showed that when the presentation of descriptions are delayed, allowing for an increase in prior experience, the behavioural impact of the information contained in those descriptions in reduced, thus supporting my initial hypothesis that the amount of prior experience moderates how descriptive information is integrated into the decision making progress.

These findings are an extension of the earlier work by Barron et al. (2008) and the experiments elsewhere in this dissertation. They suggest that descriptions have their strongest impact on behaviour when shown earlier, highlighting the importance of early warnings. As participants accumulate more experience with a task, it is likely that behavioural inertia sets in, thus reducing the potential for new information to influence behaviour. The surprise-triggers-change theory (Ert et al., 2011; Nevo & Erev, 2012) would suggest that these descriptions should indeed influence behaviour as they contained surprising information, so it is possible that individuals are discounting the information contained in the descriptions more strongly when they appear later in the task, thus reducing their surprise factor.

4.5 Discussion on prior experiences

The objective of the experiments in this chapter was to assess how the accumulation of prior experiences moderates the impact of descriptions on behaviour, and the importance of early warnings. Earlier research by Barron et al. (2008) had shown that warnings presented at the beginning of the task, before any experience, had a greater impact on behaviour than experience presented later, after participants had personally experienced half of the task without descriptions. These results were replicated in all the experiments in this section, with an apparent higher discounting of the effect of the descriptions in the ED conditions, with delayed presentation of descriptions, than the DE conditions, which presented descriptions from the beginning. Results from the earlier chapters in this dissertation, which showed that descriptions influence behaviour in the direction normatively predicted by their content when they provide novel information compared to experience (Chapter 2) and in complex DfE tasks (Chapter 3), were also replicated here, with participants exhibiting better performance overall when descriptions were available to them. The experimental manipulations in this chapter also allowed for the overall impact of descriptions on behaviour to be confirmed here using within-subjects data, with the same participants performing part of the task without and part with descriptions, eliminating some of the potential individual differences variability associated with between-subjects designs.

In addition, the new hypothesis that the amount of prior experience moderates the influence of descriptions, with a negative relationship, was also confirmed. Descriptions had a higher impact on behaviour, with participants shifting their choices to those normatively predicted by the descriptive information more, in the experimental conditions that presented descriptions earlier, before the accumulation of large amounts of prior experience. In the conditions with later presentation of descriptions, and larger accumulation of prior experience, descriptions had a lower impact. This reinforces the importance of earlier descriptions to influence behaviour, and that early warnings lead to higher compliance and increased avoidance of risk. This finding supports the hypothesis that the amount of prior experience interferes with learning, when taking into account that switching rates did not differ between the different ED conditions, which is a good proxy for differences in behavioural inertia (see Figure 4.2B on page 135).

In Experiments 8 and 9, it initially appeared that participants with more prior experience responded more strongly to the descriptions' appearances, with a higher impact of descriptions when these were presented later. However this was shown to be an artefact of wealth accumulation during the task through statistical analyses, and later confirmed empirically in Experiment 10. In most experiments in this dissertation, which is also typical in general DfE experimental tasks, participants accumulated more points and financial rewards over time. Therefore participants who were presented the descriptions later in the task, with more prior experience, had also accumulated more wealth, confounding the two.

According to the endowment effect (Kahneman et al., 1991; Morewedge et al., 2009; Rick, 2011), individuals are more reluctant to lose wealth that they own. Therefore those with more accumulated points would be more sensitive to the large losses of which the descriptions warned them. The relationship between wealth and reaction to the losses in the descriptions is aligned with predictions in the well-explored research on the endowment effect, and was replicated here in the experiments. Across all conditions, the impact of descriptions was stronger for participants who had accumulated more points. A similar effect with gambling warning messages had been observed before by Ginley, Whelan, Keating, and Meyers (2016), who noticed that warning messages may encourage winning gamblers to stop playing, but not losing players.

This leads to two interesting phenomena. Firstly, there is a perverse effect of warnings, with an increase in risk-taking in participants who had not accumulated much wealth by the time that descriptions first appeared, the opposite effect desired by such warnings. Some participants chose the options which were associated with disastrous financial losses more frequently (see Figure 4.6 in page 148). This was likely a result of those participants noticing the frequent positive rewards associated with each option, and how, if they ignored the large losses, the High-Risk decks actually might have returned higher rewards if they could avoid those losses. Perhaps in those cases participants myopically took more risk in order to inflate their financial rewards, thinking they could afford more risk-taking in the short-term. The effect was short lived, contained mostly in the few trials following the appearance of the descriptions; later in the task participants reduced their risk-taking. This can be clearly seen in condition ED20 of Experiment 8 (see Figure 4.4B in page 143). Also, the effect was mostly eliminated by descriptions that did not mention the positive rewards but focused instead on the large losses, in Experiment 9 (see Figure 4.4B in page 143). This is a perverse effect of warnings, leading to more risk-seeking behaviour. Understanding how individuals integrate descriptive risk information is crucial to avoid creating such perverse warnings that might make the risky choices more attractive to individuals (see also Ben-Ari, Florian, & Mikulincer, 1999; Ferraro, Shiv, & Bettman, 2005; Hansen, Winzeler, & Topolinski, 2010).

Secondly, there seems to be a closely linked important relationship between experience, wealth and warnings, and it might be impossible to dissociate the accumulation of wealth and experience in real life situations. Frequently, as individuals progress through life, they accumulate experience with a certain situation, as well as wealth of several forms associated with it. This is an important relationship to be taken into account when deploying efficient warnings to reduce risk-taking and increase behavioural compliance. When facing warnings, individuals might always be balancing on one hand their prior experience, which increases the discounting of a warning and reduces its effectiveness, with accumulated wealth, which has the opposite effect. Prior experience is a force towards reducing the efficiency of warnings, while accumulated wealth might make them more efficient. Experience is closely related to wealth. When designing warnings, these two conflicting but closely connected forces must be taken into consideration. Individuals who observe a warning might be torn between two internal forces when considering how to react and comply to them: their prior experience might push them towards lower compliance, while higher accumulated wealth means that they have more to lose, therefore increasing compliance. Prior experience and wealth (not only of the financial sort) are often confounded in real life, and difficult to dissociate.

Finally, while the impact of descriptions, as measured by comparing the behaviour in the block immediately preceding against that immediately following the appearance of descriptions, was moderated by prior experience, there did not appear to be an overall long-term influence in behaviour. While in the previous chapters I have shown that even after many trials, the difference in behaviour between no descriptions and descriptions from the beginning of the task remained, in trials where description presentation was delayed, overall behaviour did not differ significantly between the different delays and between ED and E. Only in the DE conditions in the current chapter, those which presented descriptions before any prior experience, did long-term differences in behaviour remain. It seems that any amount of previous personal experience, no matter how little, has a strong moderating and long-lasting impact on behaviour. As previously proposed by Miron-Shatz et al. (2010), even a few trials of personal experience seem to be enough to reduce the impact of any later warnings. This might be a result of strong inertia effect driven by the first few choices made by an individual when facing a new unknown task. While descriptions were more impactful when presented to individuals with lower amounts of prior experience, there appears to have been a step-change, with the strongest long-term effect in situations where there was no prior experience at all. This reinforces the importance of early warnings. Overall, the experiments presented here support both the presence of behavioural inertia and learning interference as a result of prior experience. Inertia was observed via the step-change between conditions DE and ED, and learning interference was observed via the different magnitude of the impact of descriptions across the different underlying ED conditions. Further research in this area is needed to understand and perhaps quantify the relative impact on behaviour of these two cognitive phenomena in more detail.

Chapter 5

General Discussion

The aim of this dissertation was to expand research on decisions from description-plus-experience (D+E), in order to better understand how descriptive and experiential information are combined when the two are available in the same task concurrently. In particular, the experiments reported here investigated the impact on behaviour resulting from the addition of descriptions to typical DfE tasks, employing paradigms with repeated choices leading to multiple outcomes. These situations are akin to real-life scenarios, in which individuals might receive information in the form of written descriptions, for example from restaurant reviews, patient information leaflets, and health-and-safety warning labels, while simultaneously collecting experiential information via an action-feedback loop from repeatedly making choices in those same environments.

Previous research using D+E tasks had led to contradictory results. Some studies, such as Lejarraga (2010) and Lejarraga and Gonzalez (2011), suggested that descriptions are ignored in tasks that combine both descriptions and experiences, with individuals relying on experience alone, even when descriptions are also available. Other studies, such as Rakow and Miler (2009) and Barron et al. (2008), showed that descriptions do influence behaviour when combined with experiences. However the latter studies which detected behavioural differences provided information in which descriptions and experience were in conflict. In comparison, the former studies that revealed no observable behavioural differences presented the same information in both descriptions and experience. These studies and their discrepancies formed the background to the research presented here.

I proposed a comprehensive theory for the integration of descriptive and experiential information that can support the seemingly contradictory results found previously, and which has been tested and confirmed by the experiments in this dissertation. Instead of descriptions being ignored when experience is also available, I propose that the two sources of information, descriptions and experience, are both taken into account in the decision-making process. The integration of the two sources of information is made using different relative weights, with experience typically exerting higher influence in the decisions, while descriptions are heavily discounted. This proposition is mainly based on two related cognitive theories. First, the premise that the description-experience gap is also driven by differences in paradigms, such as single-choice or repeated-choice, not differences in how the information is presented; this suggests that, in some situations, presenting individuals with the same information via description or via experience would not lead to observable differences in behaviour, if the nature of the paradigm is kept constant (Camilleri & Newell, 2011b, 2013a). Second, the suggestion that descriptions and experiences are processed differently, with experience being processed more naturally and automatically, and descriptions being more cognitively costly (Glöckner et al., 2012; Hasher & Zacks, 1984; Weber et al., 2004), leads to individuals preferring experience over descriptions (Lejarraga, 2010). This, I believe, supports my suggestion that descriptions are discounted by individuals. Furthermore, I set out to test if the relative weights given to each source can vary according to different situations, and explored three potential moderators for the impact of descriptions on behaviour. If descriptions are discounted at different levels according to certain moderating factors, the contradictory results in the earlier D+E literature can be reconciled, with some situations reducing the impact of descriptions to the point of apparent neglect (e.g., Lejarraga & Gonzalez, 2011), whilst not in others (e.g., Barron et al., 2008).

5.1 Summary of findings

The first set of experiments, in Chapter 2, showed that descriptions are not completely ignored when presented in combination with experience (as suggested by Lejarraga & Gonzalez, 2011), but instead that they are taken into consideration by individuals and integrated into the decision-making process, albeit with descriptions being heavily discounted and receiving relatively lower importance in comparison to experiential information. By manipulating the descriptions so that they provided participants with both congruent and conflicting information, Chapter 2 reconciled the previously seemingly contradictory results in the D+E research. When congruent descriptions were added to DfE tasks, there was no observable differences in behaviour. This finding was similar to the one from Lejarraga and Gonzalez (2011), and, in fact, the paradigm used here was inspired by the one they used. Congruent descriptions did not seem to provide any novel information to participants, and there was no expected discernible difference in behaviour if participants had relied on either descriptions or experience. While Lejarraga and Gonzalez (2011) expected that by adding descriptions to experience the behaviour would shift from underweighting to overweighting of rare events, this did not happen, likely because it is not the nature of the source of the information, but perhaps the simple or repetitive characteristics of the task, which drive these behavioural differences. It was only by providing conflicting information that an impact on behaviour was observed due to the presentation of descriptions, in the normative direction as predicted by the novel information they contained, thus confirming that descriptions are integrated into the decision-making process. This had been shown previously in Rakow and Miler (2009) and Barron et al. (2008), and was also replicated in Chapter 2 in the experimental conditions which displayed conflicting descriptions. In addition, an experimental manipulation controlling for the plausibility of the conflict provided by the descriptions showed how less plausible descriptions had a lower impact on behaviour than more plausible ones, confirming that the strength of preference for experiences over descriptions can vary according to the situation.

The second set of experiments, in Chapter 3, moved the D+E research away from simple tasks, which is the focus of the vast majority of research in decision-making, and into the domain of more complex tasks, which remains mostly unexplored, by increasing the number of choices available and number of potential outcomes from each choice. While in Chapter 2, which used the more traditional simple tasks, congruent descriptions did not influence behaviour, only conflicting descriptions did, in Chapter 3, with more complex tasks, congruent descriptions which provided the same underlying information as experiences did influence behaviour. In the case of more complex tasks, which are harder to learn via experience alone, descriptions can add useful novel information, and influence behaviour, even when the descriptions are a true representation of the experience. Specifically, the addition of descriptions to a relatively complex DfE task, the Iowa Gambling Task, showed that descriptions did influence behaviour, helping participants perform better and learning faster in the task. A direct manipulation of task complexity, by changing the number of potential alternatives and number of outcomes from each alternative, showed that the relationship between task complexity and influence of descriptions follows a non-monotonic inverted U-shape. In simple tasks, which are easy to learn and in which descriptions do not provide much useful additional information, their influence is limited, as shown also in Chapter 2. When the task is too complex, then the descriptions needed to explain this task can also become too complex, unwieldy and difficult to decipher, thus also limiting their usefulness and capacity to aid behaviour. It was in medium complexity tasks explored in Chapter 3, such as the IGT, that descriptions provided the most benefit and helped participants achieve higher financial rewards.

Finally, the third set of experiments, in Chapter 4, explored how prior experience moderates the influence of descriptions. In the other experiments in this dissertation, descriptions were either not presented at all, or presented from the beginning of the task, before participants had any chance to accumulate any knowledge about the task. In an earlier experiment, Barron et al. (2008) showed how early presentation of warnings, before any experience, had a stronger impact on behaviour with higher compliance and higher avoidance of risk, when compared to warnings that are presented later, halfway through the task, after some personal experience has been accumulated via the feedback during the task, without descriptions. In Chapter 4 I replicated the results from Barron and colleagues, showing that descriptions presented at the beginning of the task have a stronger impact on behaviour than descriptions presented later. In addition, I expanded their earlier research by manipulating the amount of prior experience, presenting the descriptions at different points during the task according to the experimental condition. The experiments showed that the amount of prior experience does moderate the impact of descriptions, in a negative relationship. Overall, participants who were presented the descriptions later in the task, when they had time to accumulate more prior experience, reacted less strongly to the descriptions. The descriptions had the strongest impact when presented earlier in the task, with little prior experience. It also appears that in addition to the negative relationship, there is a stronger step-up difference between zero experience (descriptions before any trials) and any different amount of non-zero experience. Even little experience seems to undermine the influence of descriptions in comparison to descriptions that appear at the beginning of the task, thus highlighting the importance of presenting descriptive information, such as warnings, early. The experiments in Chapter 4 also highlighted the important, perhaps dissociable, relationship between accumulated wealth and by prior experience. It appears that participants were being influenced by both the endowment effect of their accumulated wealth, which dictates that descriptions will have a stronger influence in behaviour, and prior experience, which has the opposite effect. This led to a perverse influence of descriptions for some participants, who decided to take more risk after descriptions appeared, opposite to the effect that warnings aim to produce.

5.2 Exploration, exploitation, and inertia

As predicted, the additional presence of descriptive information in the D+E tasks in this dissertation led to a reduction in switching between options, which can be seen as a proxy for exploration, in comparison with experience-only tasks. This is likely a result of the additional information provided by descriptions reducing the uncertainty contained in the task, with lower uncertainty leading to reduced exploration (Dayan & Niv, 2008; Otto et al., 2013). Moreover, the presence of descriptions might trigger the idea that the information provided descriptively is more complete, including all the information needed to make decisions, reducing motivation to search beyond that (Schmidt & Spreng, 1996). Learning from experience, on the other hand, is inferential, does not rely on clearly defined information-sets, and naturally induces broader information searches, fostering exploration (Hadar & Fox, 2009; Rakow & Newell, 2010). However, the purely exploitative behaviour that would maximise financial gains for participants in the tasks presented here was still not observed empirically. In static environments such as the ones employed here, it is expected that after individuals are satisfied with exploring the alternatives, their behaviour should ideally converge into pure exploitation (Dam & Körding, 2009). While participants switched less often between decks when description was present, it did not completely eliminate the switching behaviour associated with exploration. This lingering exploratory behaviour, even after a considerable amount of information had been collected about the task, has been observed before both in DfE (Ashby et al., 2017; Barron & Erev, 2003; Biele et al., 2009; Erev, Ert, & Yechiam, 2008; Mehlhorn et al., 2015; Speekenbrink & Konstantinidis, 2014), and repeated DfD tasks (Camilleri & Newell, 2011b, 2013a; Jessup et al., 2008; Newell et al., 2013; Otto et al., 2011), with behaviour rarely converging into pure exploitation even when that is the most beneficial strategy, with many underlying explanations proposed.

Perhaps individuals are ethologically protecting themselves against potential changes in the payoff structure, which would go unnoticed in exploitative only behaviour in dynamic environments, leaving individuals ignorant of changes in the reward structures of non-explored options (Knox et al., 2012; Speekenbrink & Konstantinidis, 2015). In such dynamic situations, static descriptions are actually deleterious to performance (Hertwig & Pleskac, 2010; Rakow & Miler, 2009). While the tasks in this dissertation were all static and not dynamic, perhaps human cognitive mechanisms for gathering information are better adapted to dynamic situations, which are more often encountered in real life, and not to static environments, which normally characterise artificial laboratory experiments. In addition, Shanks et al. (2002) proposed that this sub-optimal type of exploratory behaviour could be due to boredom, as participants did not want to choose the same option repeatedly, despite of the costs associated with diverging from optimal behaviour (see also Wilson et al., 2014). Alternatively, participants might have been selecting a mixed strategy where their preferred selection pattern was to diversify across different options with a certain frequency, rather than having a single favourite option (Ashby et al., 2017; Konstantinidis et al., 2015). In decisions combining description and experience, exploration might be partly a result of participants' need to confirm the veracity of descriptions via direct personal experience. Perhaps this was driven by limited trust in the descriptions, and further research is necessary to establish the relationship between the influence of descriptions and the trustworthiness of their sources.

Because in traditional DfE tasks based on experience alone participants do not have any prior information about the task and must learn experientially over time via feedback, their first choice in such DfE tasks must necessarily be a random uninformed selection. In comparison, if descriptions are added from the beginning of the task, then the first choice can be informed by such descriptions. Such behavioural differences were observed in the experiments presented in this dissertation. This difference in initial behaviour also appeared to lead to inertia effects and be behind the lingering effect of conflicting descriptions, even after many trials, as well as the continuous differences in switching rates over time. This potentially inertial behaviour can explain why there appeared to be a step-change producing a large difference between descriptions that are presented before any experience, and those given after even a small amount of experience, as observed in Chapter 4 and suggested by Miron-Shatz et al. (2010).

Even though experiential information was dominant, the discounted influence of descriptive information still remained after many trials, in particular in the experiments in Chapters 2 and 3 where the differences between experimental conditions persisted until the last trial. Descriptions helped participants identify their most advantageous selections earlier, but even after many trials and extensive learning participants did not seem to reach the same level of performance without descriptions. These findings were replicated in the experiments in Chapter 4, although not in the delayed description conditions, in which behaviour was more closely determined by experience throughout. In Chapter 4, the behavioural results showed that early presentation of descriptions is crucial for influencing behavioural patterns. While in the earlier chapters the differences in behaviour between the DE and E conditions remained after many trials, never converging, in the delayed description ED conditions of Chapter 4, those conditions were similar to the E conditions in the long term.

5.3 Cognitive modelling

In support of the empirical results showing that descriptions influence behaviour in accordance with their informational content, cognitive models that included representations of both the descriptive and experiential information fitted the observed behavioural results in D+E tasks better than models that ignored descriptions and relied on experience alone. These description-plus-experience models significantly outperformed the experience-only models in particular in tasks in which descriptions provided novel information for participants, which could not be easily learned by experience alone. In the case of conflicting information in Chapter 2, where descriptions provided participants with information different to that gathered experientially, this contradictory information interfered with human and model behaviour, leading to lower performance. In Chapter 3, results showed that, despite providing congruent information, in more complex tasks it takes longer and it is more difficult to learn by experience alone, thus making descriptions useful again, and influencing behaviour helping participants and model improve their selection performance.

Descriptions might have been useful in the complex tasks in Chapter 3 because in more complex situations people are slower to learn by experience alone; this might have created informational discrepancies between descriptions and experience, even though the two contained the same information in the long run. Descriptions were not beneficial to the models (and to behaviour) when they provided only redundant information that could also be easily learned by experience alone. In simple tasks in Chapters 2 and 3, in which descriptions did not provide any significant advantage to participants, there were no observable differences in behaviour when descriptions were added to task, nor when they were added to the models, similar to the modelling efforts reported by Lejarraga and Gonzalez (2011), who showed that an experience-only model predicted behaviour well in their task. In the overly complex tasks in Chapter 3, descriptions again became too complex to be easily processed by participants, reducing their usefulness and their impact on behaviour, with participants again preferring to rely more heavily on experience.

Overall, the influence of experience dominated the influence of description, with higher weights given to the former over the latter according to the models. This finding is supported by the cognitive theory concerning differences in processing between the two types of information, with a natural preference towards experiences, which are easier to process, over descriptions, which are more cognitively demanding. The cognitive models proposed in this dissertation integrated descriptions and experience by allowing each to take a relative weight against each other. This computational weight was also influenced by the situations tested here experimentally, namely the plausibility of descriptions, with more plausible descriptions receiving higher weights than less plausible descriptions, and the complexity of the task, with descriptions receiving lower weights when they were too complex to be swiftly decoded and utilised by individuals. The models presented here were thus able to integrate the characteristics observed empirically, such as the influence of plausibility in Chapter 2 and of task complexity in Chapter 3, via the different weights allocated to descriptions. In addition to description weights, an analysis of the best fitting parameters showed that individuals pay more attention to losses, learn more slowly and choose less randomly when descriptions are available to them. Attention to losses is likely a by-product of the "mere presentation" effect (Erev, Glozman, & Hertwig, 2008), with individuals more attentive to negative events when they have them presented to them descriptively. Slower learning and less random behaviour is aligned with the expectation that descriptions provide additional useful information to participants, reducing uncertainty, which helps them learn faster and make more deterministic decisions.

5.4 Limitations and further research

Given the ubiquity of situations in which descriptions and experience are available concurrently, in myriad of different combinations, approaches, and sequences, further exploration of description-plus-experience paradigms, looking at how descriptive and experiential sources of information are integrated, is crucial for understanding decision-making in realistic settings outside of the psychology lab, in particular given the lack of research dedicated to this area so far. While my research has shown that descriptions are taken into account when combined with experience, and that descriptions receive a lower weight in comparison with experience, there are still many moderators of this balance between description and experience to be researched, beyond the three already scrutinised here.

One area of further research directly related to the current dissertation is the trustworthiness of the descriptions, which likely influenced the exploration-exploitation relationship observed here. Continuous exploration even when complete descriptions were available might be a sign of lack of trust in the information provided. Descriptions are inherently communicative, and whether individuals accept them or not depends on how much they trust the source. In contrast, experiences are inherently personal, thus avoiding issues of trust. The trust that individuals allocate to the source of the descriptions is likely to influence the weight allocated to the information conveyed by them (Birnbaum, Wong, & Wong, 1976), with higher weights allocated to one's own opinions (Yaniv, 2004). Confidence could also be measured, by evaluating how much participants trust their ability to learn from the environment experientially. In some situations, individuals might not be confident in their own experiences, and thus prefer to rely more on descriptions. Sources that are more trustworthy might lead to stronger influences in behaviour, while sources that are not trusted might be more easily disregarded. Research measuring or manipulating the trustworthiness of each source could explore this important relationship.

Another related area of further investigation is the representativeness of the descriptions, as individuals might believe that descriptions do not necessarily apply to them (Kahneman & Tversky, 1972). Descriptions are likely seen as information collected by other individuals, perhaps in other situations, further compounding this effect (Pronin, Gilovich, & Ross, 2004). If individuals do not believe that the descriptions are representative of their own personal situation, they might more easily disregard them. The perception of descriptions as either representations of the future or of the past can be another confounding factor. Individuals might believe that descriptions are representations of the past, and not applicable to future events. Conversely, individuals might believe that descriptions are idealised targets of future outcomes, and disconnect them from past outcomes. Both approaches are likely to influence behaviour differently, in comparison to the assumption that descriptions and experience are timelessly connected both to the past and the future.

The relationship between risk appetite, risk behaviour, and differences in DfD and DfE tasks also warrants further exploration. Research has shown that individuals' judgements of probabilities can differ between those solicited using descriptive and experiential information, leading to the "overestimation-underweighting paradox", whereas individuals demonstrate behaviour consistent with both underweighting and overweighting of rare events (Barron & Yechiam, 2009; Liang et al., 2017; Madan et al., 2014). As shown here in Figure 2.4B on page 63, participants' judgements of the frequency of the events were not aligned with their behaviour. Risk preferences have typically been measured using written questionnaires, which is similar to soliciting one's preferences via descriptions (Charness, Gneezy, & Imas, 2013). In practice, risk taking is more often actualized in real-life DfE situations, thus invalidating many of these risk preference metrics. Individuals might also not be able to predict their own risk preferences in real-life situations based on descriptive prompts, and even be confused by such risk assessments (Kahneman & Snell, 1992; Slovic, 1987). Perhaps a D+E approach to measuring risk preferences might return more realistic predictions of risk taking.

In all the experiments presented here, once descriptions were revealed to participants, they remained always available, until the end of the task. This constant availability of the descriptive information might have strengthened their impact on decisions. This might not always necessarily be the case for real-life decisions involving descriptions and experience. For instance, a patient who consults the side effects of medication by initially reading the informational leaflet, and then collects experience by taking the medication, is unlikely to frequently return to the descriptive information. This means that descriptions might also often be available at a single point in time, not continuously present. One research question that remains open for exploration is how the results would change if descriptions were not constantly present, but only briefly available, and then removed. Given the results observed in Chapter 4 on the strong effect of early descriptions before any experience, I would still expect a behavioural impact of such disappearing descriptions, due to decision inertia. However, over time, their impact might reduce or even extinguish. In addition, the constant availability of descriptions might induce information overload, or highlight any discrepancies between the two sources of information (especially in the case of conflicting descriptions), reducing the impact of descriptions. In this case, sporadically presented descriptions might have a higher impact on behaviour than constantly present ones. The interplay between these two potential effects warrants further investigation.

Finally, one crucial methodological limitation in this area of research is the utilisation of between-subjects designs, which I have used here exclusively, and is particularly difficult to circumvent. In fact within-subjects experiments are very rarely employed in the description-experience gap research (see also footnote 2 on page 25). A within-subject design would allow for a cleaner comparison of the effects of descriptions on behaviour, excluding individual differences. The experiments on Chapter 4 were able to compare behaviour with and without descriptions within the same participants, and confirm the impact of descriptions on behaviour using within-subjects analyses; however this was confounded with the amount of prior experience and knowledge accumulated by participants within the same task.

5.5 Conclusion and implications

Throughout this dissertation I have shown that descriptions influence behaviour when available in combination with experience, despite preference being allocated to the latter. Using behavioural and cognitive modelling results, descriptions were shown to influence behaviour in the normative direction as predicted by their informational content, proving that they are integrated into the decision-making process, with different weights according to the situation.

These findings have important implications in the field of decision-making research. Decision-making biases originally found in descriptive paradigms are not always replicated in experiential paradigms, and in some instances reversals are observed (e.g., Barron & Erev, 2003; Fantino & Navarro, 2012; Hau et al., 2010). These decision-making biases have been widely explored in applications such as behavioural interventions, social marketing and governmental policy-making. Perhaps the reason why some of these attempted applications are not successful is because the research behind them is based on descriptive paradigms and applied in experiential settings. More relevant research would be based on real-life experiential situations and would allow for the influence of descriptions and other prior beliefs by giving them appropriate weightings.

The results from the experiments in this dissertation can also help with ongoing open issues around the description-experience gap. In particular, the idea that the gap is driven by differences in the paradigm, with traditional DfD tasks relying on single-choice single-outcome tasks and DfE tasks relying on multi-choice multi-outcome. Throughout this dissertation, all of the tasks were DfE-type tasks with multi-choice and multi-outcomes, with the addition of descriptions as an experimental manipulation. The finding that cognitive models that include descriptions are not better than cognitive models that rely on experience alone when the two provide the same information is additional support for the theory that differences in paradigm are an important part of the description-experience-gap (Camilleri &

Newell, 2011b, 2013a).

There are also important ramifications in the field of safety communication and warning label research. Warnings can be seen as descriptions introduced to well-experienced situations, in order to influence behaviour. But research has shown that warning information is commonly disregarded by individuals, or has limited impact on behaviour, with its overall effectiveness still unproven (Argo & Main, 2004; Laughery, 2006; Rogers et al., 2000; Wagenaar, Hudson, & Reason, 1990). If descriptions are discounted, their reduced salience might not be enough to sway behaviour as intended by warnings. To counteract this apparent discounting effect, designers of warning labels might be tempted to exaggerate risks and appeal to emotions and personal experience in order to increase compliance. However highly exaggerated descriptions might become implausible and could have the opposite effect on behaviour, as observed in Chapter 2. Simpler warnings should also be more efficient, as shown here by the lower influence of complex descriptions on behaviour in Chapter 3. While most real-life tasks are considerably more complex than the experiments presented here, it might be that individuals perceive certain tasks to be simpler than they are, perhaps by habituation, reducing their acceptance of and compliance with additional information in the form of descriptions or warnings. More effective warnings and better safety behaviour compliance is likely to benefit from further understanding of the relationship between descriptions and experience.

References

Ahn, W.-Y., Busemeyer, J. R., Wagenmakers, E.-J., & Stout, J. C. (2008). Comparison of decision learning models using the generalization criterion method. *Cognitive Science*, *32*, 1376–1402. doi: 10.1080/03640210802352992

Ahn, W.-Y., Krawitz, A., Kim, W., Busemeyer, J. R., & Brown, J. W. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics*, *4*, 95–110. doi: 10.1037/a0020684

Alexander, G. C., Gallagher, S. A., Mascola, A., Moloney, R. M., & Stafford, R. S. (2011). Increasing off-label use of antipsychotic medications in the United States, 1995-2008. *Pharmacoepidemiology and Drug Safety*, *20*, 177–184. doi: 10.1002/pds.2082

Anderson, C. J. (2003). The psychology of doing nothing: Forms of decision avoidance result from reason and emotion. *Psychological Bulletin*, *129*, 139–167. doi: 10.1037/0033-2909.129.1.139

Argo, J. J., & Main, K. J. (2004). Meta-analyses of the effectiveness of warning labels. *Journal of Public Policy & Marketing*, *23*, 193–208. doi: 10.1509/jppm.23.2.193.51400

Ashby, N. J., Konstantinidis, E., & Yechiam, E. (2017). Choice in experiential learning: True preferences or experimental artifacts? *Acta Psychologica*, *174*, 59–67. doi: 10.1016/j.actpsy.2017.01.010

Atkins, P. W., Wood, R. E., & Rutgers, P. J. (2002). The effects of feedback format on dynamic decision making. *Organizational Behavior and Human Decision Processes*, *88*, 587–604. doi: 10.1016/S0749-5978(02)00002-X

Avrahami, J., & Kareev, Y. (2011). The role of impulses in shaping decisions. *Journal of Behavioral Decision Making*, *24*, 515–529. doi: 10.1002/bdm.707

Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, *32*, 1369–1378. doi: 10.3758/BF03206327

Barberis, N. C. (2013). Thirty years of prospect theory in economics: A review and assessment. *Journal of Economic Perspectives*, *27*, 173–196. doi: 10.1257/jep.27.1.173

Bar-Hillel, M., & Wagenaar, W. A. (1991). The perception of randomness. *Advances in Applied Mathematics*, *12*, 428–454. doi: 10.1016/0196-8858(91)90029-I

Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, *16*, 215–233. doi: 10.1002/bdm.443

Barron, G., & Leider, S. (2010). The role of experience in the Gambler's Fallacy. *Journal of Behavioral Decision Making*, *23*, 117–129. doi: 10.1002/bdm.676

Barron, G., Leider, S., & Stack, J. (2008). The effect of safe experience on a warnings' impact: Sex, drugs, and rock-n-roll. *Organizational Behavior and Human Decision Processes*, *106*, 125–142. doi: 10.1016/j.obhdp.2007.11.002

Barron, G., & Yechiam, E. (2009). The coexistence of overestimation and underweighting of rare events and the contingent recency effect. *Judgment and Decision Making*, *4*, 447–460.

Bates, D., Maechler, M., Bolker, B., & Walker, S. (2014). *lme4: Linear mixed-effects models using Eigen and S4.* R package version 1.1-6. Retrieved from `http://cran.r-project.org/package=lme4`

Bechara, A., & Damasio, A. R. (2005). The somatic marker hypothesis: A neural theory of economic decision. *Games and Economic Behavior*, *52*, 336–372. doi: 10.1016/j.geb.2004.06.010

Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, *50*, 7–15. doi: 10.1016/0010-0277(94)90018-3

Bechara, A., Damasio, H., Damasio, A. R., & Lee, G. P. (1999). Different contributions of the human amygdala and ventromedial prefrontal cortex to decision-making. *The Journal of Neuroscience*, *19*, 5473–5481.

Bechara, A., Damasio, H., Tranel, D., & Damasio, A. R. (1997). Deciding advantageously before knowing the advantageous strategy. *Science*, *275*, 1293–1295. doi: 10.1126/science.275.5304.1293

Ben-Ari, O. T., Florian, V., & Mikulincer, M. (1999). The impact of mortality salience on reckless driving: A test of terror management mechanisms. *Journal of Personality and Social Psychology*, *76*, 35–45. doi: 10.1037/0022-3514.76.1.35

Biele, G., Erev, I., & Ert, E. (2009). Learning, risk attitude and hot stoves in restless bandit problems. *Journal of Mathematical Psychology*, *53*, 155–167. doi: 10.1016/j.jmp.2008.05.006

Birnbaum, M. H., Wong, R., & Wong, L. K. (1976). Combining information from sources that vary in credibility. *Memory & Cognition*, *4*, 330–336. doi: 10.3758/BF03213185

Brand, M., Labudda, K., & Markowitsch, H. J. (2006). Neuropsychological correlates of decision-making in ambiguous and risky situations. *Neural Networks*, *19*, 1266–1276. doi: 10.1016/j.neunet.2006.03.001

Brown, D. R. (1953). Stimulus-similarity and the anchoring of subjective scales. *The American Journal of Psychology*, *66*, 199–214. doi: 10.2307/1418726

Busemeyer, J. R., & Stout, J. C. (2002). A contribution of cognitive decision models to clinical assessment: Decomposing performance on the Bechara gambling task. *Psychological Assessment*, *14*, 253–262. doi: 10.1037/1040-3590.14.3.253

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

Busemeyer, J. R., & Wang, Y.-M. (2000). Model comparisons and model selections based on generalization criterion methodology. *Journal of Mathematical Psychology*, *44*, 171–189. doi: 10.1006/jmps.1999.1282

Camerer, C. F., & Ho, T.-H. (1994). Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty*, *8*, 167–196. doi: 10.1007/BF01065371

Camerer, C. F., & Hogarth, R. M. (1999). The effects of financial incentives in experiments: A review and capital-labor-production framework. *Journal of Risk and Uncertainty*, *19*, 7–42. doi: 10.1023/A:1007850605129

Camerer, C. F., & Weber, M. (1992). Recent developments in modeling preferences: Uncertainty and ambiguity. *Journal of Risk and Uncertainty*, *5*, 325–370. doi: 10.1007/BF00122575

Camilleri, A. R., & Newell, B. R. (2009). Within-subject preference reversals in description-and experience-based choice. In N. Taatgen, J. van Rijn, J. Nerbonne, & L. Schomaker (Eds.), *Proceedings of the 31st annual conference of the cognitive science society* (pp. 449–454). Austin, TX: Cognitive Science Society.

Camilleri, A. R., & Newell, B. R. (2011a). Description- and experience-based choice: Does equivalent information equal equivalent choice? *Acta Psychologica*, *136*, 276–284. doi: 10.1016/j.actpsy.2010.11.007

Camilleri, A. R., & Newell, B. R. (2011b). When and why rare events are underweighted: A direct comparison of the sampling, partial feedback, full feedback and description choice paradigms. *Psychonomic Bulletin & Review*, *18*, 377–384. doi: 10.3758/s13423-010-0040-2

Camilleri, A. R., & Newell, B. R. (2013a). The long and short of it: Closing the description-experience "gap" by taking the long-run view. *Cognition*, *126*, 54–71. doi: 10.1016/j.cognition.2012.09.001

Camilleri, A. R., & Newell, B. R. (2013b). Mind the gap? Description, experience, and the continuum of uncertainty in risky choice. In N. Srinivasan & P. Chandrasekhar (Eds.), *Progress in brain research* (pp. 55–72). Oxford, UK: Elsevier. doi: 10.1016/B978-0-444-62604-2.00004-6

Campbell, D. J. (1988). Task complexity: A review and analysis. *The Academy of Management Review*, *13*, 40–52. doi: 10.2307/258353

Carson, J., & Mannering, F. (2001). The effect of ice warning signs on ice-accident frequencies and severities. *Accident Analysis & Prevention*, *33*, 99–109. doi: 10.1016/S0001-4575(00)00020-8

Chapman, G. B., & Johnson, E. J. (1994). The limits of anchoring. *Journal of Behavioral Decision Making*, *7*, 223–242. doi: doi: 10.1002/bdm.3960070402

Charness, G., Gneezy, U., & Imas, A. (2013). Experimental methods: Eliciting risk preferences. *Journal of Economic Behavior and Organization*, *87*, 43–51. doi: 10.1016/j.jebo.2012.12.023

Chewning, E. G., & Harrell, A. M. (1990). The effect of information load on decision makers' cue utilization levels and decision quality in a financial distress decision task. *Accounting, Organizations and Society*, *15*, 527–542. doi: 10.1016/0361-3682(90)90033-Q

Chiu, Y.-C., & Lin, C.-H. (2007). Is deck C an advantageous deck in the Iowa Gambling Task? *Behavioral and Brain Functions*, *3*, 37. doi: 10.1186/1744-9081-3-37

Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*, 933–942. doi: 10.1098/rstb.2007.2098

Dai, J., Kerestes, R., Upton, D. J., Busemeyer, J. R., & Stout, J. C. (2015). An improved cognitive model of the Iowa and Soochow Gambling Tasks with regard to model fitting performance and tests of parameter consistency. *Frontiers in Psychology*, *6*, 229. doi: 10.3389/fpsyg.2015.00229

Dam, G., & Körding, K. (2009). Exploration and exploitation during sequential search. *Cognitive Science*, *33*, 530–541. doi: 10.1111/j.1551-6709.2009.01021.x

Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. (2006). Cortical substrates for exploratory decisions in humans. *Nature*, *441*, 876–879. doi: 10.1038/nature04766

Dayan, P., & Niv, Y. (2008). Reinforcement learning: The good, the bad and the ugly. *Current Opinion in Neurobiology*, *18*, 185–196. doi: 10.1016/j.conb.2008.08.003

Diehl, E., & Sterman, J. D. (1995). Effects of feedback complexity on dynamic decision making. *Organizational Behavior and Human Decision Processes*, *62*, 198–215. doi: 10.1006/obhd.1995.1043

Dunn, B. D., Dalgleish, T., & Lawrence, A. D. (2006). The somatic marker hypothesis: A critical evaluation. *Neuroscience and Biobehavioral Reviews*, *30*, 239–271. doi: 10.1016/j.neubiorev.2005.07.001

Dutt, V., & Gonzalez, C. (2012). The role of inertia in modeling decisions from experience with instance-based learning. *Frontiers in Psychology*, *3*, 177. doi: 10.3389/fpsyg.2012.00177

Edwards, K. (1996). Prospect theory: A literature review. *International Review of Financial Analysis*, *5*, 19–38. doi: 10.1016/S1057-5219(96)90004-6

Eppler, M. J., & Mengis, J. (2004). The concept of information overload: A review of the literature from organization science, accounting, MIS, and related disciplines. *The Information Society*, *20*, 325–344. doi: 10.1080/01972240490507974

Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, *112*, 912–931. doi: 10.1037/0033-295X.112.4.912

Erev, I., Ert, E., Plonsky, O., Cohen, D., & Cohen, O. (2017). From anomalies to forecasts: Toward a descriptive model of decisions under risk, under ambiguity, and from experience. *Psychological Review*, *124*, 369–409. doi: 10.1037/rev0000062

Erev, I., Ert, E., Roth, A. E., Haruvy, E., Herzog, S. M., Hau, R., … Lebiere, C. (2010). A choice prediction competition: Choices from experience and from description. *Journal of Behavioral Decision Making*, *23*, 15–47. doi: 10.1002/bdm.683

Erev, I., Ert, E., & Yechiam, E. (2008). Loss aversion, diminishing sensitivity, and the effect of experience on repeated decisions. *Journal of Behavioral Decision Making*, *21*, 575–597. doi: 10.1002/bdm.602

Erev, I., Glozman, I., & Hertwig, R. (2008). What impacts the impact of rare events. *Journal of Risk and Uncertainty*, *36*, 153–177. doi: 10.1007/s11166-

008-9035-z

Erev, I., & Haruvy, E. (2005). Generality, repetition, and the role of descriptive learning models. *Journal of Mathematical Psychology*, *49*, 357–371. doi: 10.1016/j.jmp.2005.06.009

Erev, I., & Haruvy, E. (2016). Learning and the economics of small decisions. In J. Kagel & A. Roth (Eds.), *The handbook of experimental economics* (Vol. 2, pp. 638–716). Princeton, NJ: Princeton University Press.

Ert, E., & Erev, I. (2007). Replicated alternatives and the role of confusion, chasing, and regret in decisions from experience. *Journal of Behavioral Decision Making*, *20*, 305–322. doi: 10.1002/bdm.556

Ert, E., Erev, I., & Roth, A. E. (2011). A choice prediction competition for social preferences in simple extensive form games: An introduction. *Games*, *2*, 257–276. doi: 10.3390/g2030257

Falk, R., & Konold, C. (1997). Making sense of randomness: Implicit encoding as a basis for judgment. *Psychological Review*, *104*, 301–318. doi: 10.1037/0033-295X.104.2.301

Fantino, E., & Navarro, A. (2012). Description-experience gaps: Assessments in other choice paradigms. *Journal of Behavioral Decision Making*, *25*, 303–314. doi: 10.1002/bdm.737

Fasolo, B., Hertwig, R., Huber, M., & Ludwig, M. (2009). Size, entropy, and density: What is the difference that makes the difference between small and large real-world assortments? *Psychology & Marketing*, *26*, 254–279. doi: 10.1002/mar.20272

Fennema, H., & Wakker, P. P. (1997). Original and cumulative prospect theory: a discussion of empirical differences. *Journal of Behavioral Decision Making*, *10*, 53–64.

Ferraro, R., Shiv, B., & Bettman, J. R. (2005). Let us eat and drink, for tomorrow we shall die: Effects of mortality salience and self-esteem on self-regulation in consumer choice. *Journal of Consumer Research*, *32*, 65–75. doi: 10.1086/429601

Fox, C. R., & Hadar, L. (2006). Decisions from experience= sampling error+ prospect theory: Reconsidering Hertwig, Barron, Weber & Erev (2004). *Judgment and Decision Making*, *1*, 159–161.

Frey, R., Mata, R., & Hertwig, R. (2015). The role of cognitive abilities in decisions from experience: Age differences emerge as a function of choice set size. *Cognition*, *142*, 60–80. doi: 10.1016/j.cognition.2015.05.004

Fridberg, D. J., Queller, S., Ahn, W.-Y., Kim, W., Bishara, A. J., Busemeyer, J. R., ... Stout, J. C. (2010). Cognitive mechanisms underlying risky decision-making in chronic cannabis users. *Journal of Mathematical Psychology*, *54*, 28–38. doi: 10.1016/j.jmp.2009.10.002

Friedman, M., & Savage, L. (1948). The utility analysis of choices involving risk. *The Journal of Political Economy*, *56*, 279–304.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, *102*, 684–704. doi: 10.1037/0033-295X.102.4.684

Ginley, M. K., Whelan, J. P., Keating, H. A., & Meyers, A. W. (2016). Gambling warning messages : The impact of winning and losing on message reception across a gambling session gambling warning messages. *Psychology of Addictive Behaviors*, *30*, 931–938. doi: 10.1037/adb0000212

Glasgow, R. E., Green, L. W., Klesges, L. M., Abrams, D. B., Fisher, E. B., Goldstein, M. G., ... Tracy Orleans, C. (2006, apr). External validity: We need to do more. *Annals of Behavioral Medicine*, *31*, 105–108.

Glöckner, A., Fiedler, S., Hochman, G., Ayal, S., & Hilbig, B. E. (2012). Processing differences between descriptions and experience: A comparative analysis using eye-tracking and physiological measures. *Frontiers in Psychology*, *3*, 173. doi: 10.3389/fpsyg.2012.00173

Gonzalez, C. (2005). Decision support for real-time, dynamic decision-making tasks. *Organizational Behavior and Human Decision Processes*, *96*, 142–154. doi: 10.1016/j.obhdp.2004.11.002

Grosskopf, B., Erev, I., & Yechiam, E. (2006). Foregone with the wind: Indirect payoff information and its implications for choice. *International Journal of Game Theory*, *34*, 285–302. doi: 10.1007/s00182-006-0015-8

Güney, S., & Newell, B. R. (2015). Overcoming ambiguity aversion through experience. *Journal of Behavioral Decision Making*, *28*, 188–199. doi: 10.1002/bdm.1840

Gupta, A., Smith, K., & Shalley, C. (2006). The interplay between exploration and exploitation. *Academy of Management Journal*, *49*, 693–706.

Hadar, L., & Fox, C. R. (2009). Information asymmetry in decision from description versus decision from experience. *Judgment and Decision Making*, *4*, 317–325.

Hansen, J., Winzeler, S., & Topolinski, S. (2010). When the death makes you smoke: A terror management perspective on the effectiveness of cigarette on-pack warnings. *Journal of Experimental Social Psychology*, *46*, 226–228. doi: 10.1016/j.jesp.2009.09.007

Hasher, L., & Zacks, R. T. (1984). Automatic processing of fundamental information: The case of frequency of occurrence. *The American Psychologist*, *39*, 1372–1388. doi: 10.1037/0003-066X.39.12.1372

Hau, R., Pleskac, T. J., & Hertwig, R. (2010). Decisions from experience and statistical probabilities: Why they trigger different choices than a priori probabilities. *Journal of Behavioral Decision Making*, *23*, 48–68. doi: 10.1002/bdm.665

Hau, R., Pleskac, T. J., Kiefer, J., & Hertwig, R. (2008). The description-experience gap in risky choice: The role of sample size and experienced probabilities. *Journal of Behavioral Decision Making*, *21*, 493–518. doi: 10.1002/bdm.598

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

Hertwig, R., & Erev, I. (2009). The description-experience gap in risky choice. *Trends in Cognitive Sciences*, *13*, 517–523. doi: 10.1016/j.tics.2009.09.004

Hertwig, R., & Pleskac, T. J. (2010). Decisions from experience: Why small samples? *Cognition*, *115*, 225–237. doi: 10.1016/j.cognition.2009.12.009

Hills, T. T., & Hertwig, R. (2010). Information search in decisions from experience: Do our patterns of sampling foreshadow our decisions? *Psychological Science*, *21*, 1787–1792. doi: 10.1177/0956797610387443

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science*, *310*, 1680–1683. doi: 10.1126/science.1115327

Huettel, S. A., Stowe, C. J., Gordon, E. M., Warner, B. T., & Platt, M. L. (2006). Neural signatures of economic preferences for risk and ambiguity. *Neuron*, *49*, 765–775. doi: 10.1016/j.neuron.2006.01.024

Hwang, M. I., & Lin, J. W. (1999). Information dimension, information overload and decision quality. *Journal of Information Science*, *25*, 213–218. doi: 10.1177/016555199423268[4]

Jacoby, J., Speller, D. E., & Berning, C. K. (1974). Brand choice behavior as a function of information load: Replication and extension. *Journal of Consumer Research*, *1*, 33–42. doi: 10.1086/208579

Jacoby, J., Speller, D. E., & Kohn, C. A. (1974). Brand choice behavior as a function of information load. *Journal of Marketing Research*, *11*, 63–69. doi: 10.2307/3150994

Jessup, R. K., Bishara, A. J., & Busemeyer, J. R. (2008). Feedback produces divergence from prospect theory in descriptive choice. *Psychological Science*, *19*, 1015–1022. doi: 10.1111/j.1467-9280.2008.02193.x

Joag, S. G., Mowen, J. C., & Gentry, J. W. (1990). Risk perception in a simulated industrial purchasing task: The effects of single versus multi-play decisions. *Journal of Behavioral Decision Making*, *3*, 91–108. doi: 10.1002/bdm.3960030203

Johnson, E. J., Häubl, G., & Keinan, A. (2007). Aspects of endowment: A query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*, 461–474. doi: 10.1037/0278-7393.33.3.461

Johnson, E. J., & Payne, J. W. (1985). Effort and accuracy in choice. *Management Science*, *31*, 395–414. doi: 10.1287/mnsc.31.4.395

Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1996). Reinforcement learning : A survey. *Journal of Artificial Intelligence Research*, *4*, 237–285. doi: 10.1613/jair.301

Kahneman, D., Knetsch, J. L., & Thaler, R. H. (1991, feb). Anomalies: The endowment effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, *5*, 193–206. doi: 10.1257/jep.5.1.193

Kahneman, D., & Snell, J. (1992). Predicting a changing taste: Do people know what they will like? *Journal of Behavioral Decision Making*, *5*, 187–200. doi: 10.1002/bdm.3960050304

Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, *3*, 430–454. doi: 10.1016/0010-0285(72)90016-3

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

Kamal, A. H., & Peppercorn, J. M. (2013). The generalizability paradox within palliative care clinical trials. *Annals of Palliative Medicine*, *2*, 101–104. doi: 10.3978/j.issn.2224-5820.2013.02.06

Keller, K. L., & Staelin, R. (1987). Effects of quality and quantity of information on decision effectiveness. *Journal of Consumer Research*, *14*, 200–213. doi: 10.1086/209106

Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, *94*, 211–228. doi: 10.1037/0033-295X.94.2.211

Knight, F. H. (1921). *Risk, uncertainty, and profit*. Boston, MA: Houghton Mifflin.

Knox, W. B., Otto, A. R., Stone, P., & Love, B. C. (2012). The nature of belief-directed exploratory choice in human decision-making. *Frontiers in Psychology*, *2*, 398. doi: 10.3389/fpsyg.2011.00398

Konstantinidis, E., Ashby, N. J., & Gonzalez, C. (2015). Exploring complexity in decisions from experience: Same minds, same strategy. In D. C. Noelle et al. (Eds.), *Proceedings of the 37th annual conference of the cognitive science society* (pp. 1177–1182). Austin, TX: Cognitive Science Society.

Krizan, Z., & Windschitl, P. D. (2007). The influence of outcome desirability on optimism. *Psychological Bulletin*, *133*, 95–121. doi: 10.1037/0033-2909.133.1.95

Kudryavtsev, A., & Pavlodsky, J. (2012). Description-based and experience-based decisions: Individual analysis. *Judgment and Decision Making*, *7*, 316–331.

Lam, C. F., DeRue, D. S., Karam, E. P., & Hollenbeck, J. R. (2011). The impact of feedback frequency on learning and task performance: Challenging the "more is better" assumption. *Organizational Behavior and Human Decision Processes*, *116*, 217–228. doi: 10.1016/j.obhdp.2011.05.002

Laughery, K. R. (2006). Safety communications: Warnings. *Applied Ergonomics*, *37*, 467–478. doi: 10.1016/j.apergo.2006.04.020

Lea, S. E. G., McLaren, I. P. L., Dow, S. M., & Graft, D. A. (2012). The cognitive mechanisms of optimal sampling. *Behavioural Processes*, *89*, 77–85. doi: 10.1016/j.beproc.2011.10.004

Lejarraga, T. (2010). When experience is better than description: Time delays and complexity. *Journal of Behavioral Decision Making*, *23*, 100–116. doi: 10.1002/bdm.666

Lejarraga, T., & Gonzalez, C. (2011). Effects of feedback and complexity on repeated decisions from description. *Organizational Behavior and Human Decision Processes*, *116*, 286–295. doi: 10.1016/j.obhdp.2011.05.001

Lejarraga, T., Hertwig, R., & Gonzalez, C. (2012). How choice ecology influences search in decisions from experience. *Cognition*, *124*, 334–342. doi: 10.1016/j.cognition.2012.06.002

Lenth, R. V. (2016). Least-squares means: The R package lsmeans. *Journal of Statistical Software*, *69*, 1–33. doi: 10.18637/jss.v069.i01

Lerner, J. S., Small, D. A., & Loewenstein, G. (2004). Heart strings and purse strings: Carryover effects of emotions and economic decisions. *Psychological*

Science, *15*, 337–341. doi: 10.1111/j.0956-7976.2004.00679.x

Levy, I., Snell, J., Nelson, A. J., Rustichini, A., & Glimcher, P. W. (2010). Neural representation of subjective value under risk and ambiguity. *Journal of Neurophysiology*, *103*, 1036–1047. doi: 10.1152/jn.00853.2009

Lewandowsky, S., & Farrell, S. (2011). *Computational modeling in cognition: Principles and practice*. Thousand Oaks, CA: SAGE Publications.

Liang, G., Konstantinidis, E., Szollosi, A., Donkin, C., & Newell, B. R. (2017). The impact of descriptions and incentives on the simultaneous underweighting and overestimation of rare events. In G. Gunzelmann, A. Howes, T. Tenbrink, & E. Davelaar (Eds.), *Proceedings of the 39th annual conference of the cognitive science society* (pp. 736–741). Austin, TX: Cognitive Science Society.

Locke, E. A. (1982). Relation of goal level to performance with a short work period and multiple goal levels. *Journal of Applied Psychology*, *67*, 512–514. doi: 10.1037/0021-9010.67.4.512

Lopes, L. L., & Oden, G. C. (1987). Distinguishing between random and nonrandom events. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 392–400. doi: 10.1037/0278-7393.13.3.392

Ludvig, E. A., & Spetch, M. L. (2011). Of black swans and tossed coins: Is the description-experience gap in risky choice limited to rare events? *PloS ONE*, *6*, e20262. doi: 10.1371/journal.pone.0020262

Lurie, N. H., & Swaminathan, J. M. (2009). Is timely information always better? The effect of feedback frequency on decision making. *Organizational Behavior and Human Decision Processes*, *108*, 315–329. doi: 10.1016/j.obhdp.2008.05.005

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2014). Remembering the best and worst of times: Memories for extreme outcomes bias risky decisions. *Psychonomic Bulletin & Review*, *21*, 629–636. doi: 10.3758/s13423-013-0542-9

Madan, C. R., Ludvig, E. A., & Spetch, M. L. (2017). The role of memory in distinguishing risky decisions from experience and description. *The Quarterly Journal of Experimental Psychology*, *70*, 2048–2059. doi: 10.1080/17470218.2016.1220608

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization Science*, *2*, 71–87. doi: 10.1287/orsc.2.1.71

Mehlhorn, K., Newell, B. R., Todd, P. M., Lee, M. D., Morgan, K., Braithwaite, V. A., … Gonzalez, C. (2015). Unpacking the exploration-exploitation trade-off: A synthesis of human and animal literatures. *Decision*, *2*, 191–215. doi: 10.1037/dec0000033

Miron-Shatz, T., Barron, G., Hanoch, Y., Gummerum, M., & Doniger, G. M. (2010). To give or not to give: Parental experience and adherence to the Food and Drug Administration warning about over-the-counter cough and cold medicine usage. *Judgment and Decision Making*, *5*, 428–436.

Morewedge, C. K., Shu, L. L., Gilbert, D. T., & Wilson, T. D. (2009). Bad riddance or good rubbish? Ownership and not loss aversion causes the endowment effect. *Journal of Experimental Social Psychology*, *45*, 947–951. doi: 10.1016/j.jesp.2009.05.014

Nevo, I., & Erev, I. (2012). On surprise, change, and the effect of recent outcomes. *Frontiers in Psychology*, *3*, 24. doi: 10.3389/fpsyg.2012.00024

Newell, B. R., Koehler, D. J., James, G., Rakow, T., & van Ravenzwaaij, D. (2013). Probability matching in risky choice: The interplay of feedback and strategy availability. *Memory & Cognition*, *41*, 329–338. doi: 10.3758/s13421-012-0268-3

Newell, B. R., & Rakow, T. (2007). The role of experience in decisions from description. *Psychonomic Bulletin & Review*, *14*, 1133–1139. doi: 10.3758/BF03193102

Niv, Y. (2009). Reinforcement learning in the brain. *Journal of Mathematical Psychology*, *53*, 139–154. doi: 10.1016/j.jmp.2008.12.005

Noguchi, T., & Hills, T. T. (2016). Experience-based decisions favor riskier alternatives in large sets. *Journal of Behavioral Decision Making*, *29*, 489–498.

doi: 10.1002/bdm.1893

Otto, A. R., Gershman, S. J., Markman, A. B., & Daw, N. D. (2013). The curse of planning: Dissecting multiple reinforcement-learning systems by taxing the central executive. *Psychological Science*, *24*, 751–761. doi: 10.1177/0956797612463080

Otto, A. R., Taylor, E. G., & Markman, A. B. (2011). There are at least two kinds of probability matching: Evidence from a secondary task. *Cognition*, *118*, 274–279. doi: 10.1016/j.cognition.2010.11.009

Paquette, L., & Kida, T. (1988). The effect of decision strategy and task complexity on decision performance. *Organizational Behavior and Human Decision Processes*, *41*, 128–142. doi: 10.1016/0749-5978(88)90051-9

Paul, S., & Nazareth, D. L. (2010). Input information complexity, perceived time pressure, and information processing in GSS-based work groups: An experimental investigation using a decision schema to alleviate information overload conditions. *Decision Support Systems*, *49*, 31–40. doi: 10.1016/j.dss.2009.12.007

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Decision Processes*, *16*, 366–387. doi: 10.1016/0030-5073(76)90022-2

Payne, J. W. (1982). Contingent decision behavior. *Psychological Bulletin*, *92*, 382–402. doi: 10.1037/0033-2909.92.2.382

Payne, J. W., & Braunstein, M. L. (1978). Risky choice: An examination of information acquisition behavior. *Memory & Cognition*, *6*, 554–561. doi: 10.3758/BF03198244

Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin*, *68*, 29–46.

Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the eye of the beholder: Divergent perceptions of bias in self versus others. *Psychological Review*, *111*, 781–799. doi: 10.1037/0033-295X.111.3.781

R Core Team. (2014). *R: A language and environment for statistical computing.* R Foundation for Statistical Computing, Vienna, Austria. Retrieved from `http://www.r-project.org/`

Radley, D. C., Finkelstein, S. N., & Stafford, R. S. (2006). Off-label prescribing among office-based physicians. *Archives of Internal Medicine*, *166*, 1021–1026. doi: 10.1001/archinte.166.9.1021

Rakow, T., Demes, K. A., & Newell, B. R. (2008). Biased samples not mode of presentation: Re-examining the apparent underweighting of rare events in experience-based choice. *Organizational Behavior and Human Decision Processes*, *106*, 168–179. doi: 10.1016/j.obhdp.2008.02.001

Rakow, T., & Miler, K. (2009). Doomed to repeat the successes of the past: History is best forgotten for repeated choices with nonstationary payoffs. *Memory & Cognition*, *37*, 985–1000. doi: 10.3758/MC.37.7.985

Rakow, T., & Newell, B. R. (2010). Degrees of uncertainty: An overview and framework for future research on experience-based choice. *Journal of Behavioral Decision Making*, *23*, 1–14. doi: 10.1002/bdm.681

Rakow, T., Vincent, C., Bull, K., & Harvey, N. (2005). Assessing the likelihood of an important clinical outcome: New insights from a comparison of clinical and actuarial judgment. *Medical Decision Making*, *25*, 262–282. doi: 10.1177/0272989X05276849

Reutskaja, E., & Hogarth, R. M. (2009). Satisfaction in choice as a function of the number of alternatives: When "goods satiate". *Psychology & Marketing*, *26*, 197–203. doi: 10.1002/mar.20268

Rick, S. (2011). Losses, gains, and brains: Neuroeconomics can help to answer open questions about loss aversion. *Journal of Consumer Psychology*, *21*, 453–463. doi: 10.1016/j.jcps.2010.04.004

Rogers, W. A., Lamson, N., & Rousseau, G. K. (2000). Warning research: An integrative perspective. *Human Factors*, *42*, 102–139. doi: 10.1518/001872000779656624

Samuelson, W., & Zeckhauser, R. J. (1988). Status quo bias in decision making.

Journal of Risk and Uncertainty, *1*, 7–59. doi: 10.1007/BF00055564

Schmidt, J., & Spreng, R. (1996). A proposed model of external consumer information search. *Journal of the Academy of Marketing Science*, *24*, 246–256. doi: 10.1177/0092070396243005

Schultz, W. (2006, jan). Behavioral theories and the neurophysiology of reward. *Annual Review of Psychology*, *57*, 87–115. doi: 10.1146/annurev.psych.56.091103.070229

Shah, A. M., & Wolford, G. (2007). Buying behavior as a function of parametric variation of number of choices. *Psychological Science*, *18*, 369–370. doi: 10.1111/j.1467-9280.2007.01906.x

Shanks, D. R., Tunney, R. J., & McCarthy, J. D. (2002). A re-examination of probability matching and rational choice. *Journal of Behavioral Decision Making*, *15*, 233–250. doi: 10.1002/bdm.413

Shiffrin, R., Lee, M. D., Kim, W., & Wagenmakers, E.-J. (2008). A survey of model evaluation approaches with a tutorial on hierarchical bayesian methods. *Cognitive Science*, *32*, 1248–1284. doi: 10.1080/03640210802414826

Shlomi, Y. (2014). Subjective integration of probabilistic information from experience and description. *Judgment and Decision Making*, *9*, 491–499.

Slovic, P. (1987). Perception of risk. *Science*, *236*, 280–285. doi: 10.1126/science.3563507

Speekenbrink, M., & Konstantinidis, E. (2014). Uncertainty and exploration in a restless bandit task. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th annual conference of the cognitive science society* (pp. 1491–1496). Austin, TX: Cognitive Science Society.

Speekenbrink, M., & Konstantinidis, E. (2015). Uncertainty and exploration in a restless bandit problem. *Topics in Cognitive Science*, *7*, 351–367. doi: 10.1111/tops.12145

Stanovich, K. E. (2003). Is probability matching smart? Associations between probabilistic choices and cognitive ability. *Memory & Cognition*, *31*, 243–251. doi: 10.3758/BF03194383

Steingroever, H., Wetzels, R., Horstmann, A., Neumann, J., & Wagenmakers, E.-J. (2013). Performance of healthy participants on the Iowa Gambling Task. *Psychological Assessment*, *25*, 180–193. doi: 10.1037/a0029929

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1–26. doi: 10.1016/j.cogpsych.2005.10.003

Stout, J. C., Rock, S. L., Campbell, M. C., Busemeyer, J. R., & Finn, P. R. (2005). Psychological processes underlying risky decisions in drug abusers. *Psychology of Addictive Behaviors*, *19*, 148–157. doi: 10.1037/0893-164X.19.2.148

Streufert, S., & Driver, M. J. (1965). Conceptual structure, information load and perceptual complexity. *Psychonomic Science*, *3*, 249–250. doi: 10.3758/BF03343120

Sutton, R., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.

Swait, J., & Adamowicz, W. (2001). The influence of task complexity on consumer choice: A latent class model of decision strategy switching. *Journal of Consumer Research*, *28*, 135–148. doi: 10.1086/321952

Thorngate, W. (1980). Efficient decision heuristics. *Behavioral Science*, *25*, 219–225. doi: 10.1002/bs.3830250306

Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, *211*, 453–458. doi: 10.1126/science.7455683

Tversky, A., & Kahneman, D. (1992). Advances in Prospect-Theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323. doi: 10.1007/Bf00122574

Ungemach, C., Chater, N., & Stewart, N. (2009). Are probabilities overweighted or underweighted when rare outcomes are experienced (rarely)? *Psychological Science*, *20*, 473–479. doi: 10.1111/j.1467-9280.2009.02319.x

Wagenaar, W. A., Hudson, P. T. W., & Reason, J. T. (1990). Cognitive failures and accidents. *Applied Cognitive Psychology*, *4*, 273–294. doi: 10.1002/acp.2350040405

Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using

Akaike weights. *Psychonomic Bulletin & Review*, *11*, 192–196. doi: 10.3758/BF03206482

Weber, E. U. (2006). Experience-based and description-based perceptions of long-term risk: Why global warming does not scare us (yet). *Climatic Change*, *77*, 103–120. doi: 10.1007/s10584-006-9060-3

Weber, E. U., & Johnson, E. J. (2009). Mindful judgment and decision making. *Annual Review of Psychology*, *60*, 53–85. doi: 10.1146/annurev.psych.60.110707.163633

Weber, E. U., Shafir, S., & Blais, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, *111*, 430–445. doi: 10.1037/0033-295X.111.2.430

Weinstein, N. D. (1989). Effects of personal experience on self-protective behavior. *Psychological Bulletin*, *105*, 31–50. doi: 10.1037/0033-2909.105.1.31

Wilson, T. D., Reinhard, D. A., Westgate, E. C., Gilbert, D. T., Ellerbeck, N., Hahn, C., ... Shaked, A. (2014). Just think: The challenges of the disengaged mind. *Science*, *345*, 75–77. doi: 10.1126/science.1250830

Wood, R. E. (1986). Task complexity: Definition of the construct. *Organizational Behavior and Human Decision Processes*, *37*, 60–82. doi: 10.1016/0749-5978(86)90044-0

Worthy, D. A., Pang, B., & Byrne, K. A. (2013). Decomposing the roles of perseveration and expected value representation in models of the Iowa Gambling Task. *Frontiers in Psychology*, *4*, 640. doi: 10.3389/fpsyg.2013.00640

Wulff, D. U., Hills, T. T., & Hertwig, R. (2015a). How short- and long-run aspirations impact search and choice in decisions from experience. *Cognition*, *144*, 29–37. doi: 10.1016/j.cognition.2015.07.006

Wulff, D. U., Hills, T. T., & Hertwig, R. (2015b). Online product reviews and the description-experience gap. *Journal of Behavioral Decision Making*, *28*, 214–223. doi: 10.1002/bdm.1841

Yaniv, I. (2004). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes*, *93*, 1–13. doi:

10.1016/j.obhdp.2003.08.002

Yechiam, E., Barron, G., & Erev, I. (2005). The role of personal experience in contributing to different patterns of response to rare terrorist attacks. *Journal of Conflict Resolution*, *49*, 430–439. doi: 10.1177/0022002704270847

Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, *12*, 387–402. doi: 10.3758/BF03193783

Yechiam, E., & Busemeyer, J. R. (2006). The effect of foregone payoffs on underweighting small probability events. *Journal of Behavioral Decision Making*, *19*, 1–16. doi: 10.1002/bdm.509

Yechiam, E., & Busemeyer, J. R. (2008). Evaluating generalizability and parameter consistency in learning models. *Games and Economic Behavior*, *63*, 370–394. doi: 10.1016/j.geb.2007.08.011

Yechiam, E., Busemeyer, J. R., Stout, J. C., & Bechara, A. (2005). Using cognitive models to map relations between neuropsychological disorders and human decision-making deficits. *Psychological Science*, *16*, 973–978. doi: 10.1111/j.1467-9280.2005.01646.x

Yechiam, E., Druyan, M., & Ert, E. (2008). Observing others' behavior and risk taking in decisions from experience. *Judgment and Decision Making*, *3*, 493–500.

Yechiam, E., Erev, I., & Barron, G. (2006). The effect of experience on using a safety device. *Safety Science*, *44*, 515–522. doi: 10.1016/j.ssci.2005.11.006

Yechiam, E., & Ert, E. (2007). Evaluating the reliance on past choices in adaptive learning models. *Journal of Mathematical Psychology*, *51*, 75–84. doi: 10.1016/j.jmp.2006.11.002

Yechiam, E., & Rakow, T. (2012). The effect of foregone outcomes on choices from experience. *Experimental Psychology*, *59*, 55–67. doi: 10.1027/1618-3169/a000126