The role of regret and responsibility in decision-making

Thesis submitted for the degree of Doctor of Philosophy

by

Antoinette Laura Nicolle

UCL 2010

This work was carried out under the primary supervision of Professor Raymond J. Dolan and secondary supervision of Professor Jon Driver. The work was funded by a Prize Studentship Award from the Brain Research Trust.

I, Antoinette Nicolle, 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

Regret is a cognitively mediated, multifaceted emotion engendered by thoughts of how things might have been better had we behaved differently. The causes, experience and behavioural impact of regret have been widely studied by psychologists. However, research into the neural basis of regret has been motivated primarily by economic approaches, which often reduce regret to such a simplistic construct that it loses many of its interesting qualities. This thesis attempts to build a bridge between recent functional imaging studies of regret and a more established psychological literature that addresses the subjective content and motivational impacts of regret. The thesis aims to provide a deeper understanding of the experience of regret, the factors necessary for it to be elicited, and its behavioural impact. Using functional imaging, I also provide new insights into the neural mechanisms underlying each of these levels.

In the first two studies, I provide evidence for a key role of responsibility in the experience and neuronal representation of regret, and in the efficacy of learning and decision-making more generally. In three further studies, I explore the immediate motivational impact of the experience of regret, and contrast findings with conventional models that address the impact of anticipated regret on choice. Specifically, I provide evidence that experienced regret encourages decision inertia, a bias to repeat, rather than avoid, a previous choice. These studies indicate that conventional models of the experiential content of regret, and its motivational effect, traditionally employed by economists and cognitive neuroscientists alike, do not provide a full description of behavioural responses to regret. I go on to consider multiple motivational effects of regret, including those (not always beneficial) responses through which individuals tend to manage and regulate aversive emotions.

Contents

Acknowledgements

The last three years would have been significantly more difficult without the intellectual, technical, financial and emotional support I have received from my supervisors, mentors, colleagues, friends and family. Many thanks go to Ray Dolan and Jon Driver for giving me the freedom to explore my interests, and the guidance to focus them. I have sincere respect for my mentor, Dominik Bach, whose teaching played a considerable role in making me the researcher I have become. It was also a great pleasure to work with Chris Frith, who has the gift either of being interested in all areas of research, or at least of beautifully feigning interest. For whichever one it was in my case, I thank him. Huge gratitude goes out to the rest of my colleagues at the FIL, especially Steve Fleming, Mkael Symmonds and Wako Yoshida, with whom I worked closely; Karl Friston and the other principles, whose insight and inspiration will no doubt stick with me wherever I go; Sarah Ruiz, whose brief involvement in data collection for my last experiment was so valuable; Marcia, Marina, Peter, Gareth, Holly, Isabelle, Ric, Chris and Rachael who always went out of their way to help me; David, Jan, Sheila and Kristian, who somehow managed to teach me to use an MRI scanner, and who were so warm and supportive in the process; and to my wonderful friends inside and outside the FIL.

None of this would have been possible without the financial support received from the Brain Research Trust. I hope to have lived up to their expectations of me. Finally, there are no words to describe how my parents and my brother have supported me. They are the reason I have reached where I am now and for that I dedicate everything I do in the future to them. It may surprise those who have known me these last three years but……………… Non, je ne regrette rien.

“Make the most of your regrets; never smother your sorrow, but tend and cherish it till it comes to have a separate and integral interest. To regret deeply is to live afresh.”

Thoreau (1893). *The writings of Henry David Thoreau.* Houghton, Mifflin and company. p.260

Chapter 1. Introduction

1.1 What is regret?

Regret is, above all else, an aversive experience that we are motivated to avoid. So unpleasant is the feeling of regret that people will pay money to avoid receiving information likely to induce it (Bell, 1982; Larrick & Boles, 1995), even if this risks reduced learning from mistakes and greater regret in the long-term (Reb & Connolly, 2009). The Dutch Postcode Lottery takes full advantage of our aversion to regret-inducing information. Here, non-participating individuals will discover if they would have won had they played, and must contend with the possibility of being the only person on their street not to have shared a jackpot. The anticipated regret this elicits boosts participation in the lottery (Zeelenberg & Pieters, 2004). However, as painful as they may be, our regrets can guide us to make better decisions, and help us atone for past mistakes. As posited in the quotation opening this thesis, we may benefit hugely from making the most of our regrets though, in order to take full advantage of them, we must first accept (and face) the aversive feeling they elicit.

For psychologists, economists, cognitive neuroscientists and consumer decision theorists, regret is of special interest for empirical study on several grounds. Firstly, its genesis and the content of its experience depends upon a complex set of cognitive processes, drawing on memory, causal inference, inductive reasoning, social and personal norms and beliefs. As such, regret has been termed an “unusually cognitively-laden or cognitively-determined emotion", while at the same time being "typically loaded with *feeling* and therefore qualifies as a true emotion" (Gilovich & Medvec, 1995, p. 379). Moreover, regret is thought to have a special involvement in learning and decision-making. I begin this thesis with an introduction to the unique antecedents, experiential content and motivational impact of regret.

The critical precursor of regret is knowledge that something might have been better had we acted differently. Such cognition is a form of counterfactual thinking, by which an experienced reality is compared against one that 'might have been' in a different state of the world. Counterfactual thinking requires a level of cognitive maturity, developing only between the ages of 5-7 years (Guttentag & Ferrell, 2004; Weisberg & Beck, 2010). Roese & Olson (1995) propose a two-stage model of counterfactual comparison, whereby the various possible states of the world are first activated (or made *available*) for consideration, and second the *content* of the comparison with the actual state of the world is determined. The authors argue that the mutability of the environment and the motivation of the individual will determine both the availability and content of counterfactual thought. While counterfactual thinking is a *cognitive* process, emotions can be a consequence of such cognition. Which emotion is induced depends upon the content and structure of the counterfactual, with other factors also influencing its intensity.

Firstly, counterfactual comparisons vary in their *direction*. While upward counterfactuals are thoughts that something could have been *better* and tend to induce negative emotions, downward counterfactuals are thoughts that something could have been *worse* and tend to elicit positive emotions. Upward counterfactuals are more likely to be generated when an individual perceives themselves as losing (Markman, Gavanski, Sherman, & McMullen, 1993). Although resulting in lower outcome satisfaction, these upward counterfactuals have strong motivational implications. Indeed, more upward counterfactuals are generated when individuals expect to be able to use such information to improve future circumstances (Markman et al., 1993). This also invokes the notion that the availability and content of counterfactuals depend upon the sense of control one has over the environment, such that upward counterfactuals are deemed less useful when future steps will not lead to a better outcome (Roese & Olson, 1995b). Conversely, downward counterfactuals have a more affective role, enhancing current feelings of esteem and wellbeing. However, they also have less (or a different) impact on subsequent behaviour.

Upward counterfactual thinking alone is not sufficient to induce regret. Counterfactual thoughts also vary in the *object of comparison*. Outcomes that could have been better from the same choice, but under a different state of the world, tend to invite *within-option* counterfactual thoughts (i.e. comparing an experienced outcome against what was expected from the same choice). This form of counterfactual thinking tends to induce feelings of disappointment, rather than regret, since our own action is irrelevant. Outcomes that could have been better from a different choice, on the other hand, tend to invoke *between-option* counterfactual comparisons, and induce feelings of regret (Roese & Olson, 1995a; Zeelenberg, van Dijk, Manstead, & van der Pligt, 2000). The result of these distinct counterfactual activations is that regret and disappointment differ in their experiential content and motivational effect (Zeelenberg, Van Dijk, Manstead, & Van der Pligt, 1998). I will further discuss this emotional specificity of regret in section 1.2 below.

The emphasis on our bad decisions, or between-option upwards counterfactuals, invokes the notion that regret depends upon feeling *responsible*, while disappointment does not. Responsibility may amplify feelings of regret induced from such counterfactual thinking, while also increasing the tendency to consider between-option, rather than within-option, comparisons. Indeed, evidence points to regret being dependent on a personal sense of blame or responsibility (Frijda, Kuipers, & Ter Schure, 1989; Gilovich & Medvec, 1994, study 4; Zeelenberg, van Dijk, & Manstead, 1998, 2000; Zeelenberg et al., 1998). Zeelenberg and colleagues found that, when judging the regret experienced by students who switched courses only to find that the new course was worse than the original one, participants tended to rate a student as experiencing greater regret if they themselves chose to switch, compared to if a computer randomly reassigned them. Whether responsibility is necessary for regret or whether it only amplifies the feeling of regret has been heavily debated (see Connolly, Ordónez, & Coughlan, 1997; Ordónez & Connolly, 2000; Zeelenberg, van Dijk, & Manstead, 1998, 2000 for a thorough review and discussion). In an attempt to resolve this debate, Connolly & Zeelenberg (2002) proposed two core components of regret as 1) the upwards, between-option counterfactual comparison, which they term 'outcome regret' and 2) an intense feeling of responsibility and self-blame, which is based on the justifiability of the decision or decision process (see also Connolly & Butler, 2006; Connolly & Reb, 2005). 'Outcome regret' has been the basis for computing regret in economic theories (Bell, 1982; Loomes & Sugden, 1982), and in patient and neuroimaging studies (e.g. Camille et al., 2004; Coricelli et al., 2005). Both experienced and anticipated outcome regret have been shown to depend upon feedback of a better between-option alternative (e.g. Boles & Messick, 1995; Inman, Dyer, & Jia, 1997; Ritov & Baron, 1995). The behavioural impact of regret also depends on such feedback, as shown by the aforementioned Dutch Postcode Lottery example describe previously (Zeelenberg, Beattie, Van der Pligt, & de Vries, 1996; Zeelenberg, 1999).

A study of the neural basis of regret (Coricelli et al., 2005) has addressed the factor of responsibility, showing that ventral striatum responses to absolute losses and gains depend on agency over choice. However, it is unknown whether this effect extends to responses to outcomes that are better or worse than what might have been (i.e. gains and losses relative to a counterfactual reference point). This would be a better indicator of the role of responsibility in how the brain processes regret. Moreover, feelings of self-blame for a bad outcome involves more than simple agency over a choice, depending also on how justified we believe we were in a decision in relation to personal and social norms (Connolly & Zeelenberg, 2002). Subtle variations in the internal attribution of responsibility will likely result in noticeable differences in the intensity of regret. The ability to adapt our behaviour appropriately after past mistakes may also depend on our understanding (and accepting) to what extent our actions have directly contributed towards the bad outcome. In other words, responsibility may determine our capacity to learn the consequences of our actions. Risk preference, loss-aversion, or effort in a decision may also be moderated by the responsibility an individual accepts for its outcome. One of the major aims of this thesis is to provide a better understanding of this self-blame component of regret.

Upward, between-option counterfactuals are arguably necessary for the feeling of regret to be elicited. As a result, much research has assumed that these are also *sufficient*, and only rarely do they explicitly include other factors, such as self-blame, in the equation. Gilovich & Medvec (1995) write that,

"Unfortunately, economic theorists have defined or operationalized regret so narrowly that the applicability of their work is more limited than perhaps it could be." (Gilovich & Medvec, 1995, p. 380)

The psychological literature, on the other hand, supplies other factors that can influence the intensity of regret and, in some cases, could conceivably avoid it altogether. These include distance in value between the actual reality and the better possible alternative (*relative value*); the number of antecedent events that would need to be mutated in order to reach the better alternative (*closeness*); how easily one could mutate these antecedents (*controllability*); the degree to which one can be held *responsible*, or accountable, for the antecedent decisions; and the degree to which one believes the antecedent decisions were *justifiable* at the time they were made. Other factors influencing the saliency of the decision process may also modify the intensity of regret, including external context, and social and personal norms and beliefs. In particular, much of the economic literature ignores the path by which a decision is made, placing the emphasis instead on the decision outcome. Psychologists have, however, drawn an important distinction between regret for the decision itself and regret for the decision process (e.g. Connolly & Zeelenberg, 2002; Inman & Zeelenberg, 2002), a distinction that bears similarities to that of outcome regret versus self-blame regret described previously. Indeed, Connolly & Zeelenberg (2002) argue that many findings in regret research can be explained by a Decision Justification Theory of Regret (DJT). For example, regret stemming from actions is generally rated higher than regret stemming from a failure to act (Baron & Ritov, 1994; Feldman, Miyamoto, & Loftus, 1999; Kahneman & Tversky, 1982; Landman, 1987; Tsiros & Mittal, 2000; Zeelenberg, van den Bos, van Dijk, & Pieters, 2002), and this effect may be driven by actions being deemed less justifiable and more causal than inactions. In keeping with DJT, however, Connolly & Zeelenberg (2002) provide evidence that inactions can be associated with the greater regret if previous experience deems them to be less justifiable (see also Inman & Zeelenberg, 2002).

While the simplicity of the economic definition of regret does allow for elegant experimental design, it is important to bridge the gap between the economic and psychological approaches, taking the benefits of each to help towards a unified conceptualisation of both the experience and behavioural impact of regret. Moreover, research into the neural basis of regret has primarily been motivated by the economic literature and lacks significant integration with the more established psychological literature.

1.2 The specificity of regret

Theories of the specificity of emotions are grounded in that the unique experiential content and behavioural impact of a particular emotion is determined by the particular appraisals that elicit it. For example, fear is elicited by the anticipation of approaching danger and motivates avoidant (i.e. flight) behaviour, while anger is distinct in that it is elicited by current and unjustified threat and motivates offensive (i.e. fight) behaviour. In accord with such a cognitive appraisal approach to emotion (e.g. Scherer, Schorr, & Johnstone, 2001), regret's dependency on the specific set of antecedent cognitive processes allows us to separate it from other aversive emotions – even those of a similarly counterfactual nature. Moreover, the distinct precursors of regret invoke a notion that its motivational impact should also be unique, since the goals encouraged by its experience will not be the same as those encouraged by, for example, disappointment. Zeelenberg & Pieters (2007) promote this emotional specificity of regret (see also Pieters & Zeelenberg, 2007). They write that:

"regret is distinct from related other specific emotions such as anger, disappointment, envy, guilt, sadness and shame, and from general negative affect on the basis of its appraisals, experiential content, and behavioral consequences". (Zeelenberg & Pieters, 2007, p. 7)

An elegant illustration of the specificity of regret involves a comparison against disappointment, an aversive emotion also dependent on counterfactual thinking. The two emotions differ profoundly in their experiential content and motivational effect. Zeelenberg et al. (1998) show that recalled autobiographical experiences of regret and disappointment differ most acutely in their associated action tendencies and "emotivational goals"[1]. While disappointment is associated with feelings of powerlessness, a tendency to avoid the situation and to do nothing, regret is associated with feelings of self-blame, a drive to correct one's mistake, and a tendency to "ruminate and focus on past events" (p.228). Thus, the clearest difference between these two emotions lies in the sense of reproach or self-blame elicited by a high sense of responsibility and the appraisal of between-option counterfactual information in the case of regret. While disappointment does depend upon counterfactual thinking, it is a *within-option* comparison, comparing an obtained outcome with a better one that was expected from the same action. As such, it lacks the personal responsibility component, since the focus is on the alternative *outcome* rather than on the alternative *choice*. Zeelenberg & Pieters (1999) extended these findings to everyday consumer decision-making. Disappointment with a purchase (when subjects found the obtained service to be worse than expected) was

[1] "Emotivational goals" is a term first used by Roseman (1984) to refer to the motives and goals that accompany our emotions.

associated with assigning responsibility to others, as shown in an observed desire to complain to a service provider. On the other hand, regret experienced when a consumer realised they had made the wrong choice (after comparing the obtained service with that of other service providers) was associated with a desire to switch service provider and avoidance in sharing the experience with others. These findings further emphasise the importance of a sense of responsibility, and ensuing self-blame, in the experience of regret.

More recently, evidence has emerged for an anatomical and/or functional dissociation of regret and disappointment in the brain. Camille et al. (2004) found, using a risky two-choice financial gambling paradigm inspired by Mellers, Schwartz, & Ritov (1999), that patients with selective lesions of the orbitofrontal cortex (OFC) were unable to experience regret in the way that normal controls did, as shown by affective ratings in response to gamble outcomes. Furthermore, they were unable to modify their future choices to avoid further instances of regret. On the other hand, they were able to experience normal levels of disappointment for a loss and satisfaction for a win, suggesting that the impairment caused by OFC damage was specific to regret. A follow-up study used fMRI as a further test of this dissociation (Coricelli et al., 2005). Using the same paradigm on healthy individuals, they found that activity in medial OFC, dorsal anterior cingulate cortex (ACC) and the amygdala varied with the magnitude of regret. On the other hand, middle temporal gyrus and dorsal brainstem showed activity that correlated with the magnitude of disappointment.

Other aversive emotions also show interesting distinctions from regret. Although envy is similar to regret, in that it also relies upon upward between-option counterfactual comparison, it does not depend on our having made a bad choice, but can be induced simply by observing that our state of affairs is worse than that of others. Anger depends upon our being adversely affected by the actions of others, with responsibility therefore being transferred externally. While guilt does contain the self-blame component, the focus here is on the consequences of our actions for another person, rather than for the self. In the case of shame, the comparison process relies on a reflection on one's character or reputation as opposed to one's behaviour. Mandel (2003) also emphasised the importance of responsibility in regret, showing that intensity of regret for autobiographical memories correlated only with other emotions that are associated with self-blame, for example shame and guilt, but not with anger or distrust for which responsibility is attributed externally.

The emotional specificity of regret is of particular importance when attempting to understand its role in learning and decision-making. A *"feeling-is-for-doing"* perspective states that emotions are first and foremost for motivating goal-directed behaviour, and that each particular emotion motivates behaviour in a unique way (Zeelenberg & Pieters, 2006). Moreover, these distinct motivations stem from the different appraisals that generate the emotions, and from the unique experience of the emotion itself. In the next section I will further discuss the functional (or preparative) role of regret in decision-making.

1.3 The preparative role of regret

"Never, never waste a minute on regret. It's a waste of time."

Harry S. Truman (1884-1972) (in Landman, 1993)

Emotions have, in the past, been considered obstructive for optimal, rational decision-making (e.g. see discussion by Solomon, 1993). A similar sentiment is shown by Harry Truman in the quotation above. However, a prevailing view now is that experiencing and anticipating emotions can be beneficial for learning and for guiding and motivating behaviour. This is especially the case for regret, as it involves recognition that better results that could have arisen (and could arise in the future) from exploring alternative behaviours. According to March (1978), all decisions involve predictions of how future outcomes will make us feel. Particularly in situations of uncertainty, the ability to anticipate future emotions probabilistically may provide a valuable source of information when selecting future behaviours. As written by Loomes & Sugden, two of the forerunners in regret theories of decision-making:

"...if an individual does experience such feelings, we cannot see how he can be deemed irrational for consistently taking those feelings into account". (Loomes & Sugden, 1982, p. 820)

The functional, or preparative, role of upward counterfactual thinking has been discussed in the work of Neal Roese (e.g. Epstude & Roese, 2008; Roese, 1994). Roese proposes a "problem solver" role of counterfactual thinking, such that it is activated by problems and triggers behavioural adaptation to solve them. Roese (1994) finds that upward counterfactual thinking encourages intentions to do better next time, resulting also in improved performance in a second chance at a task. This was not the case for downward counterfactual thinking, which does not have this preparative role. Markman et al. (1993) found that an increased number of upward counterfactuals were generated after failure at a game when participants expected to play again, than when they did not. This suggests that upward counterfactual thinking may be perceived as more useful when there is opportunity to improve future behaviour. Upward counterfactuals are also perceived as less useful when the outcomes are perceived as uncontrollable (Roese & Olson, 1995b). In such cases, individuals are more likely to construct downward counterfactuals (associated with feelings of relief) which have the different aim of improving mood. An example of this is evident in cancer patients or assault victims who may use downward counterfactuals as a way of coping with their traumatic, uncontrollable experience.

Regret avoidant behaviour is apparent, for example, in consumer choice (Inman & Zeelenberg, 2002; Tsiros & Mittal, 2000), sexual choice (Richard, Van der Pligt, & de Vries, 1996), and health related choice (Lechner, de Vries, & Offermans, 1997). For health related decisions, Lechner et al. (1997) found that those women who reported anticipating the regret they might feel from not getting a breast cancer screening were more likely to get the screen than those who did not report anticipated regret. A loss in the ability to experience regret, as shown by patients with OFC lesions, is associated with poor decision-making compared to normal controls (Camille et al., 2004). Patients with Parkinson's disease and schizophrenia also show impaired generation of counterfactuals, which may in part explain the decision-making deficits they exhibit (Hooker, Roese, Park, Sledge, & Penn, 2000; McNamara, Durso, Brown, & Lynch, 2003). The evidence, therefore, suggests that regret is not a *"waste of time"*, but rather plays an important role in decision-making.

Much effort has been put into advancing the descriptive appeal (and predictive validity) of economic models of decision-making – especially that of Expected Utility Theory (EUT). Although the concept of *value* is integral to EUT, it does not take into account many subjectivities attached to value, utility and probability. As such, people's actual decisions often violate the key axioms of EUT (Von Neumann, Morgenstern, Kuhn, & Rubinstein, 1947). The Allais paradox (Allais, 1953) provides an ideal example. Here choices are observed between the following two pairs of gambles.

Pair 1

1a) **£2,400 for certain**

1b) **£2,400 with a 66% chance, £2,500 with a 33% chance or £0 with a 1% chance**

Pair 2

2a) **£2,400 with a 34% chance or £0 with a 66% chance**

2b) **£2,500 with a 33% chance or £0 with a 67% chance**

It is commonly observed that people tend to choose 1a over 1b, but choose 2b over 2a. This shows a systematic violation of the independence axiom of EUT, whereby common consequences of each gamble within a choice pair should be disregarded when making a choice. This can be grasped better by deconstructing each gamble pair, as follows:

<u>Pair 1</u>

1a) ***£2,400 with a 66% chance*** **or £2,400 with a 34% chance**

1b) ***£2,400 with a 66% chance*****, £2,500 with a 33% chance or £0 with a 1% chance**

<u>Pair 2</u>

2a) **£2,400 with a 34% chance or** ***£0 with a 66% chance***

2b) **£2,500 with a 33% chance,** ***£0 with a 66% chance*** **or £0 with a 1% chance**

By removing the common consequences from within each pair (shown in red italics), the two choice pairs are now equivalent and EUT cannot explain the preference reversal observed in people's actual choices.

One possible source of this intransitivity in choice may be that the value of each option depends upon its counterfactual context. That is, an apparently good option will be valued less when in the context of an apparently better option, and valued more when in the context of an apparently worse option. Kahneman & Tversky (1979) were among the first to consider the impact of counterfactuals on choice preference, by introducing a psychological reference point to the preference function. While this reference point was originally intended to reflect comparison against the status-quo, any other comparative reference point can be important here, including what might have been from having made a different choice. Through this, the Allais paradox can be explained by the anticipation of regret. Given that the individual expects to receive feedback of both the chosen and the unchosen outcomes, they are likely to anticipate a between-option counterfactual comparison (with associated self-blame if the result is unpleasant). The preference for 1a can be explained by the decision-maker expecting to feel intense regret if they chose 1b and received the 1% chance of £0. On the other hand, in the second choice pair both gambles have a high probability of resulting in £0, making this anticipated regret less salient. The decision-maker is then motivated more by the anticipated regret of missing the larger gain in 2b (see Loomes & Sugden, 1983). When considering regret in this way, the preference reversal seems less surprising.

In a further example, Bell (1982) described how regret can explain a tendency for people to show both a risk-seeking preference in gambling behaviour, and a simultaneous risk-averse preference when taking out car insurance. This form of preference reversal is also difficult for EUT to account for. However, the risk-seeking preference may be explained by greater anticipated regret for missing a gamble jackpot than for losing a small bet. At the same time, less regret is anticipated for unnecessarily paying a small amount on car insurance than for paying the higher costs of car repairs, so predicting risk-aversion in such decisions (for details see Bell, 1982, pp. 970-972).

1.3.1 Regret Theory

Since the 1950s, researchers have been incorporating regret into models of choice (e.g. Bell, 1982; Loomes & Sugden, 1982; Mellers, Schwartz, Ho, & Ritov, 1997; Savage, 1951), an approach Connolly & Butler (2006) refer to as the "Modified Expected Utility Tradition". These models generally apply variants on a regret-minimax principle, whereby people make choices that aim to **min**imise their **max**imum possible regret. In other words, they assume that the typical decision-maker is regret-averse. Hart & Mas-Colell (2000) use the term "regret-matching" to describe a strategy of switching to a new action with a probability proportional to the regret anticipated for the current action. A turning point came when economists Bell (1982) and Loomes & Sugden (1982) incorporated regret into the utility function, such that both the possible regret *and* the expected utility of a decision are important, rather than assuming that only one of the two influences choice, as the EU theorists and the minimax theorist had done. This was termed Regret Theory.

Regret theory works as follows. Where x_{is} is the outcome of choosing the action A_i in state s, and x_{ks} is the outcome of choosing action A_k in state s, anticipated regret decreases the utility of A_i when x_{is} is expected to be worse than x_{ks}. Similarly, anticipated relief increases the utility of A_i when x_{is} is expected to be better than x_{ks}. The *modified expected utility* of A_i, given that A_k is the alternative, is then written as m_{is}^{k}, as shown in Equation 1:

Equation 1

$$m_{is}^{k} = c(x_{is}) + R[c(x_{is}) - c(x_{ks})]$$

Every choice option has a “choiceless” utility function $c(.)$, which is the utility of an option if its consequences were to occur without the individual having chosen it. $R[.]$ is the regret/relief function, whereby the difference between the choiceless utilities of what is and of what might have been influences the modified expected utility. The function is strictly increasing (i.e. more regret always leads to less utility) and concave (i.e. it has a marginally diminishing nature such that each additional unit of regret has less impact on utility than the unit before).

1.3.2 Decision Affect Theory

Barbara Mellers and colleagues later developed Decision Affect Theory, which attempts to explain choice preferences as a function of regret, disappointment *and* expected value. The heart of this theory is that decisions are influenced by the predicted emotional effect of our actions (Mellers et al., 1999; Mellers, 2000; Mellers et al., 1997).

The earliest version of this theory incorporated only a disappointment function into computations of expected pleasure with an outcome (Mellers et al., 1997). This comprised a within-option counterfactual comparison, with possible outcomes also weighted by their relative probability of occurring. Weighting emotions by their probability has the effect that surprising outcomes have relatively greater impact on behaviour than unsurprising outcomes. The authors found that this model accounted well for people's judged expected pleasure, for outcomes in a task with a single two-outcome gamble (i.e. with no choice component). In a later extension of the theory, Mellers et al. (1999) asked participants to choose between two gambles, like those shown in Figure 1-1.

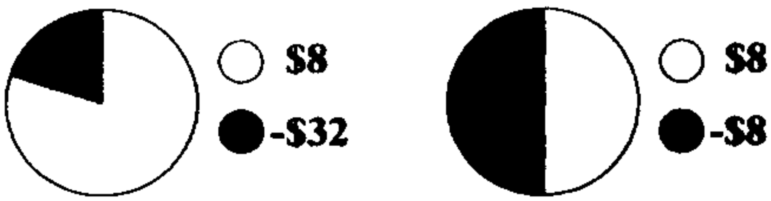

Figure 1-1 – Example gamble pair used by Mellers et al. (1999). For each gamble, outcome probabilities were indicated by wedge size, shown with corresponding monetary outcomes in black and white.

In this task, the authors varied outcome feedback in order to manipulate the object available for counterfactual comparison (see introduction the various forms of counterfactual thinking in section 1.1 above). In a *partial* feedback condition, participants observed only the obtained and unobtained outcomes of the chosen gamble, therefore allowing only for within-option comparisons. Here, the expected emotion (R) of receiving outcome A, given that B is the unobtained outcome of the same gamble, was assumed to be computed according to Equation 2.

Equation 2

$$R_A = J_R[u_A + d(u_A - u_B)(1 - s_A)]$$

Here, J_R is a function whereby the emotional response to outcome A is dependent on the utility of outcome A (i.e. u_A) and its difference from the utility of unobtained outcome B (i.e. u_B). This expected disappointment (or d) is weighted by the inverse probability of receiving outcome A (i.e. the surprise). With this partial feedback, the *subjective expected pleasure* of the full gamble becomes the summed expected emotions of each outcome, weighted by the probability that they will occur (i.e. $s_A R_A + s_B R_B$).

In a *complete* feedback condition, foregone outcomes of the unchosen gamble were also observed, thus allowing also for a between-option comparison. In this condition, the expected pleasure of receiving outcome A, given that B is the unobtained outcome of the same gamble, and that C is the outcome that would have been obtained from the unchosen gamble is computed according to Equation 3.

Equation 3

$$R_{A(C)} = J_R[u_A + d(u_A - u_B)(1 - s_A) + r(u_A - u_c)(1 - s_A s_C)]$$

Here, the pleasure of receiving outcome A is dependent *also* on the difference in utility of outcome A and that of foregone outcome C (the expected regret or r). Expected regret is weighted by the inverse of the probability that *both* outcomes A and C will occur. With complete feedback, the *subjective expected pleasure* of choosing gamble 1 then depends upon both a within- and between-option comparison of possible outcomes, and is calculated as in Equation 4.

Equation 4

$$s_A s_C R_{A(C)} + s_A s_D R_{A(D)} + s_B s_C R_{B(C)} + s_B s_C R_{B(D)}$$

Fitting the predictions of the theory, Mellers et al. (1999) observed that people's emotional responses to outcomes depend upon a mixture of the obtained outcome, what would have been obtained from the alternative outcome of the *same* gamble, and would have been obtained from the *alternative* gamble, as illustrated in Figure 1-2. The authors found that Decision Affect Theory also explains people's choices between gambles well, such that people tend to make choices that maximise their subjective expected pleasure.

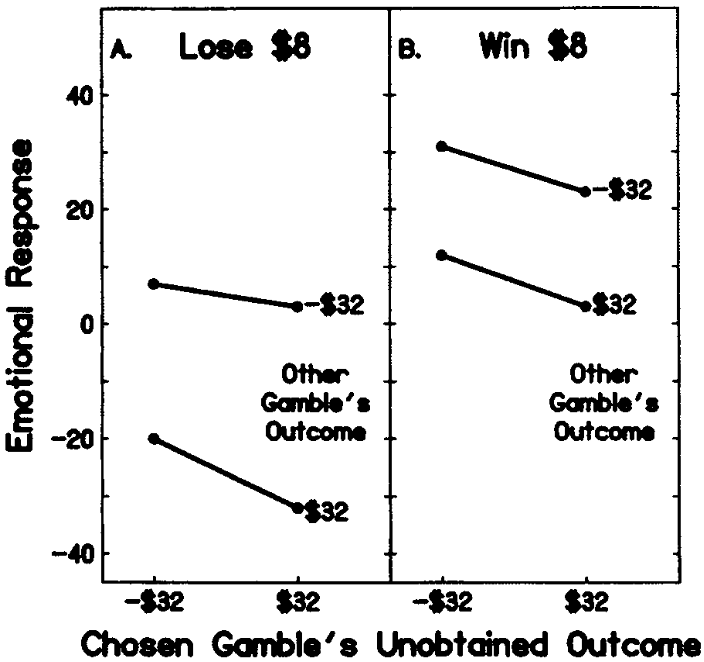

Figure 1-2 – Normalised emotional response expressed for the obtained outcome, when in the presence of better and worse outcomes that might have been obtained from the chosen and unchosen gamble. The slopes of the curves show the effect of the unobtained outcome of the chosen gamble (disappointment effect), while the distance between the curves show the effect of the unobtained outcome of the alternative outcome (regret effect). Taken from Mellers et al. (1999).

1.3.3 Impact of experienced regret

While the models described previously address the impact of *anticipated* regret on choice, it is less clear whether past experiences of regret have the same effect. Two key questions come out of this problem. Firstly, how does learning from experience of regret fit into these models? Secondly, is the impact of experienced regret always to encourage regret-avoidant future behaviour?

Concerning the first question, important advances have come from reinforcement learning approaches. In general, learning models assume that reinforcement signals elicited from experience are used to modify future behaviour towards our goals. A handful of studies have addressed the role of experienced regret in learning, whereby learning is influenced by the difference between a received outcome and one that, as realised post-hoc, would have been better from a different behaviour. Those models incorporating some form of this operational regret with high predictive success include Camerer & Ho's (1999) Experience Weighted Attraction Learning, and Normalized Fictitious Play Learning (e.g. Ert & Erev, 2007). Most recently, economists Marchiori & Warglien (2008) constructed a simple parameter-free neural network that provides an elegant model of the role of regret in interactive decision-making. In this model experienced regret is included in the online adjustment of connection weights between perceived inputs and choice propensities. As such, connections towards choices that previously induced regret are weakened by an increment proportional to the size of the regret, computed as the difference between the outcome received by the individual and the best outcome received by other players in the game. This model essentially converts previous experienced regret into anticipated regret for all possible subsequent actions. Subsequent choice then reflects a prototypical regret-minimax decision strategy. The authors compared the fit of this model with the two previously mentioned, along with other well established models in game theory, and found that their model was a good predictor of behaviour in 21 different social decision games. They also found that the Normalized Fictitious Play model (Ert & Erev, 2007), which predicts a related form of exploration (i.e. trying of new choice options) in response to regret (and likewise “stickiness” after relief), performed comparably well.

It is important to note that the interactive learning games, used by Marchiori & Warglien's (2008), may not provide the purest test of the behavioural impact of experienced regret. This is because the model reflects learning from *other people’s* better choices, a scenario that may elicit emotions such as envy and shame whose effects on behaviour may not be comparable to regret. As highlighted in a commentary by Cohen (2008), Marchiori & Warglien’s model is found to predict quicker learning than is actually observed in some games, therefore suggesting that the model does not account for some of the inconsistencies shown in real social learning.

This otherwise promising approach to understanding the behavioural impact of experienced regret has also been taken up by neuroeconomists studying how fictive (i.e. what might have been) learning signals are represented in the brain. This work also goes further in considering the role of fictive signals in private (as opposed to interactive) learning scenarios. Montague, King-Casas, & Cohen (2006) suggest that fictive signals provide individuals with a “cheaper way to learn about the world” (p.425), since individuals are able to learn simultaneously about the value of actions both taken and not taken. Fictive errors (i.e. when the fictive state is better than the actual state) also appear to have a directional impact on behaviour. In an investment decision task, Lohrenz, McCabe, Camerer, & Montague (2007) found that greater fictive error drove participants more to change their next choice. This study will be discussed further in the following section on the neurobiological basis of regret.

It is yet unclear how learning from past experiences of regret should be included in Regret Theory, and how these experiences influence behaviour in the short- and long-term. During decisions plagued by uncertainty about the objective probabilities of future occurrences, decision-makers may be largely reliant on their emotional response to similar choices made in the past. These are likely to provide the strongest influence on behaviour in iterated decision tasks, such as that used by Lohrenz et al (2007). Some studies show that cumulative experiences of regret encourage regret-minimising behaviour (Camille et al., 2004; Coricelli et al., 2005), although trial-by-trial effects on behaviour are not addressed. Some have shown that subjective desire to change a decision is greatest after regrettable trials (Chua, Gonzalez, Taylor, Welsh, & Israel Liberzon, 2009). However, others have shown in actual behaviour (as opposed to intentions) a greater tendency to switch choices after rejoice than after regret (Liu et al., 2007, data reanalysed by Sommer et al., 2009), suggesting that the immediate behavioural effect of regret may be more complex than simple avoidance.

From this perspective, experienced regret may have multiple impacts on behaviour. In particular, experienced regret may influence behaviour independent of any anticipated regret elicited (or learnt) from it. Perhaps the most significant difference between anticipated and experienced regret is that the former involves the prediction of a future emotional state, while the latter involves an immediate emotional experience. Of particular interest for this thesis is the notion that emotional experiences can influence behaviour by both moderating one's expectations of future outcomes (an indirect impact equivalent to anticipated future events) or by directly influencing our goals, arousal and attention (Loewenstein & Lerner, 2003). It is conceivable that our goals when in the midst of an intense emotional experience need not be equivalent to those when in a neutral emotional state. In particular, while anticipated regret motivates us to adjust behaviour so as to avoid the possibility of regret being elicited, current experiences of regret are unpleasant and our strongest motivation may be to regulate or reduce this current experience rather than to plan for future avoidance. The possibility that regret-regulatory goals are different under anticipated and experienced regret, and motivate distinct behavioural responses, will be discussed further in section 1.5 of this introduction.

1.4 The neurobiological basis of regret

1.4.1 Overview of the literature

The prefrontal cortex (PFC) and basal ganglia are prime candidates for involvement in counterfactual thinking and regret, given their established roles in reward processing, reasoning, planning and decision-making. Of particular importance are findings that both systems code rewards relative to the context in which they are presented, as shown through single cell recording in macaque monkeys (Hosokawa, Kato, Inoue, & Mikami, 2007; Tremblay & Schultz, 1999) and in human fMRI studies (Elliott, Agnew, & Deakin, 2008; Fujiwara, Tobler, Taira, Iijima, & Tsutsui, 2008; Nieuwenhuis et al., 2005).

The earliest studies considering how rewards signals in the brain depend on their *counterfactual* context employed a task inspired by Mellers et al. (1999). In the first of these studies, participants passively (i.e. without choice) received outcomes from good, bad and intermediate gamble wheels, like those shown in Figure 1-3 (Breiter, Aharon, Kahneman, Dale, & Shizgal, 2001). The authors concluded that nucleus accumbens, sublenticular and extended amygdala activity during *neutral* outcomes (i.e. $0) was stronger for bad spinners than for good spinners. This provided evidence that coding of rewards in these regions depended upon within-option counterfactual information.

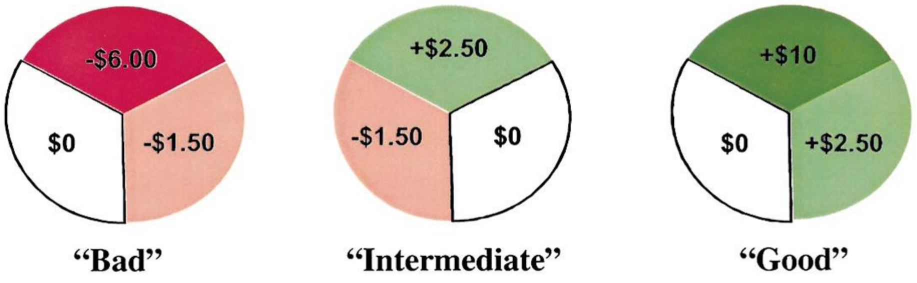

Figure 1-3 – Exemplar gamble types from Breiter et al. (2001). These authors compared the neutral, $0, outcome in the context of bad, intermediate and good gamble types.

The orbitofrontal cortex (OFC) has received the most attention with respect to regret and is widely shown to manifest context dependent responses to reward (see previously cited references on relative reward, and also Ursu & Carter, 2005, on counterfactual context). Patients with OFC lesions often have difficulty in anticipating consequences of their actions, which may account for poor decision-making and impulsive behaviour (Bechara, Tranel, & Damasio, 2000; Damasio, 1994; Mobini et al., 2002). Camille et al. (2004) used a similar Mellers-inspired task to that used in Breiter et al. (2001), but this time OFC lesioned patients made choices <u>between gamble pairs</u>. In any given pair, one of the gambles was likely to induce more regret (i.e. it is anticipated to generate a greater negative discrepancy between its outcome and that of the unchosen gamble). In addition, the authors compared trials in which participants were shown complete feedback (where the outcome of the unchosen gamble was shown and could be compared to the obtained outcome of the chosen gamble) or partial feedback (where only the obtained and unobtained outcomes of the chosen gamble were shown). Regret/relief was assumed to be experienced only on complete feedback trials. The authors found that OFC lesion patients were less likely to experience regret and were also less likely to use anticipated regret to guide their future choices in the full feedback condition, compared to non-OFC prefrontal lesion patients and healthy controls. Conversely, the patients were unimpaired in experiencing disappointment (elicited by a purely within-option comparison in partial feedback trials), and in making choices that minimised disappointment.

A follow-up fMRI study, using the same task, further explored the role of the OFC in this gambling task (Coricelli et al., 2005). In trials with complete feedback, healthy individuals showed responses in both lateral and medial OFC, anterior cingulate, and hippocampus that reflected a discrepancy between chosen and unchosen outcomes. Cumulative experience of regret was associated with enhanced activity in right somatomotor cortex, left inferior parietal lobule, left medial OFC, and left amygdala, while immediately preceding regret was associated with enhanced right dorsolateral prefrontal cortex, right lateral OFC and right inferior parietal lobule activity at subsequent choice. Moreover, both subjective pleasantness ratings and skin conductance responses reflected negative valuation and increased arousal with outcomes that were worse than what might have been from the other gamble choice. Medial OFC involvement in degree of regret was evident in a task involving avoidance of electric shocks, as opposed to reward-based motivators, while various other regions are found to be involved in degree of rejoicing, including bilateral anterior insula, bilateral striatum, right rostral anterior cingulate, left hippocampus and superior frontal cortex (Chandrasekhar, Capra, Moore, Noussair, & Berns, 2008). Moreover, these authors also showed that activity associated with both regret and rejoicing was amplified by the surprisingness of the outcome, supporting this assumption in Decision Affect Theory.

In a similar task, Chua, Gonzalez, Taylor, Welsh, & Liberzon (2009) found only lateral OFC involvement in regret compared to disappointment (along with bilateral anterior insula, right superior frontal gyrus, left middle frontal gyrus and right occipital gyrus), while medial OFC was rather involved in rejoicing. The possible different roles of medial and lateral OFC in regret is discussed in a review paper by Sommer, Peters, Gläscher, & Büchel (2009). These authors propose that the medial OFC is important in the emotional experience of regret, while activity in the lateral OFC is more likely to be involved in our behavioural response to such experiences.

In the studies discussed previously, participants are required to choose between a new pair of gambles on each trial, and only this form of task appears to recruit OFC in regret. A second class of task design requires participants to decide upon a level of risk to take on repeated decisions. In contrast to those tasks involving choices between gamble pairs, here no new information on outcome probability is made available on each trial to guide participants' choices. Rather choice is assumed to be guided by past experiences. For example, Lohrenz, McCabe, Camerer, & Montague (2007) used a task designed to explore the neural underpinnings of tracking fictive errors. On each trial, participants placed a bet on a dynamic stock market. After outcome feedback, participants were able to compare the outcome of their bet and that of the best bet they could have placed on that trial, which would be a 100% bet for wins, and a 0% bet for losses. The authors contrasted these fictive error signals in the brain against standard temporal difference (TD) error signals, computed as the difference between the experienced payoff and the expected payoff assumed from the participant's bet on that trial. For example, if a participant bets 60% on a trial in which the market then increases, the TD error is then the difference between the 60% bet and the fractional market increase, while the fictive error would be the difference between the best possible bet, i.e. 100%, and the actual 60% bet allocated. This task did not recruit the OFC, but rather striatum activity reflected these fictive error signals. A recent study, using single cell recording in monkeys, has shown cells in the anterior cingulate cortex (ACC) also respond to fictive errors (Hayden, Pearson, & Platt, 2009).

The striatum is known to be important in both the passive processing of value and in motivating goal-directed behaviour. In an early model of how counterfactual thinking might be implemented in the brain, Baird & Fugelsang (2004) proposed that the striatum may be vital for both the generation of possible counterfactuals and their use in guiding subsequent behaviour. They also suggested that the ACC may be important in checking the suitability of a counterfactual "idea", evaluating it and checking it for "errors and/or incongruities". In keeping with this model, patients with deficits to basal ganglia and PFC appear to have difficulty in appropriately considering the consequences of their actions, which may be associated with an inability to generate possible counterfactuals and to compare them against reality. Sommer et al (2008) also suggest the striatal involvement in some (but not all) regret-related tasks is due to its particular importance in the processing of fictive error signals which are vital teaching signals for guiding behaviour.

A recent study by Clark, Lawrence, Astley-Jones, & Gray (2009) explored how "near miss" outcomes are processed in the brain. These outcomes are objectively a loss, but where a win was especially close. They are, therefore, optimal for eliciting upward counterfactual thinking and regret, due to the relative ease with which the better counterfactual alternative outcome of a full win can be brought to mind (Kahneman & Tversky, 1982a). "Near miss" outcomes in gambling may have a similar conditioning effect on future behaviour as does a full win (Reid, 1986). Indeed, Clark et al. found that near misses increase the desire to continue gambling, when resulting from participant's own choices (but not from a computer-enforced choice), despite also being rated as more aversive. This finding is in keeping with the concept that upward counterfactual comparisons, though painful, are a powerful motivator for behaviour. Intriguingly, the authors also found that near misses were associated with similar neuronal responses as full wins, eliciting activity in bilateral ventral striatum and insula. Rostral ACC was also recruited by near misses, but only when participants were personally responsible for the choice. This response to a near miss is unlikely to reflect increased subjective value of the outcome, since participants reported them to be aversive experiences. They may instead be associated with the behavioural response such outcomes encourage, as the authors themselves suggest. The increased desire to continue gambling after near misses may be associated with higher sensitivity to anticipated future wins, and so this increased reward-related activity may also reflect the anticipated pleasure of future outcomes, rather than the current aversiveness of the outcome.

A key role for the striatum in such responses suggests that dopamine may be involved. McNamara et al. (2003) found that patients with Parkinson's disease had marked deficits in spontaneously generating counterfactuals and on counterfactual inference tests, where participants must form counterfactually-derived inferences about hypothetical scenarios. This supports a role of dopamine in counterfactual thinking, as Parkinson's disease is characterised by a loss of dopamine cells in the substantia nigra and ventral tegmental area. The patients were also impaired in standard tests of frontal lobe function (e.g. Stroop task and Tower of London), suggesting that disruption in dopamine connections to striatum, amygdala and PFC may result in the impaired counterfactual reasoning. Individuals with schizophrenia show similar cognitive deficits (Hooker et al., 2000; Roese, Park, Smallman, & Gibson, 2008). While deficits in counterfactual thinking with Parkinson's disease may be due to the *loss* of dopaminergic cells, the deficit in schizophrenia may be associated with excessive levels of dopamine (cf. the dopamine theory of schizophrenia, e.g. Bell, 1965; Carlsson, 1978, 1995). Dopamine may also be involved in encouraging repetitive gambling behaviour. Greater dopamine increases during gambling are found in compulsive, compared to non-compulsive, gamblers (Meyer et al., 2004), and evidence suggests a link between dopamine agonist treatment of Parkinson's disease and the onset of problematic gambling (and other forms of compulsive behaviour) in some patients (Molina et al., 2000; Weintraub et al., 2006). It is conceivable that near misses and regret may be an important precursor to such behaviour, through a dopaminergic mechanism and neural activity found by Clark et al. (2009), though such a theory is clearly in its early stages of development.

1.4.2 Synthesis and future directions

Patient studies and functional neuroimaging investigations have provided considerable advances in our understanding of regret. My review of studies exploring the neurobiological basis of regret shows consistency in brain regions involved. The OFC is the most often activated in regret. Moreover, while medial OFC involvement is the most consistent, there is apparently some role for lateral OFC. Striatum, ACC, angular gyrus/inferior parietal cortex, anterior insula and amygdala also appear to be important. In rejoicing, on the other hand, the striatum is most often implicated, though middle frontal cortex and OFC are also important. These consistencies provide a useful guide for restricting neurobiological hypotheses in the fMRI studies performed as part of this thesis.

Neuroimaging studies of regret have been largely inspired by the economic approaches, and so tend to reduce the construct of regret, along with its behavioural impact, to a simplistic form that it can be in danger of losing many of its interesting qualities. To date, this strand of research lacks significant integration with a more established psychological literature, leaving many areas open for further study. For example, how regret depends upon external context, as well as internal beliefs and regulatory strategies, is often disregarded. The role of agency in modulating brain activity associated with choice outcomes is highlighted by findings that ventral striatal responses to <u>absolute</u> gains and losses depend on agency over the causal choice (Coricelli et al., 2005). However, it is still unclear whether an agency modulation also extends to response to outcomes that are <u>relative</u> to a counterfactual reference point. Self-blame for a regrettable outcome also involves more than just agency over a choice, depending also on social norms and decision justifiability (Connolly & Zeelenberg, 2002). Moreover, Sommer et al's (2009) suggestion that the OFC may be important in tasks the allow for anticipation of possible regret, while striatum may code fictive error signals vital for learning from experienced regret, may be key to developing an understanding of the possibly dissociable roles of experienced and anticipated regret in decision-making, a distinction rarely made in the literature to date.

1.5 Recurrent themes of this thesis

Regret is what we *feel* when we realise we should have done something differently. By developing a fuller understanding of the situational factors and internal appraisals that encourage such a feeling, we can then better control and/or manipulate these conditions in the empirical study of regret. Such understanding is important for predicting the conditions in which regret is elicited in everyday life. For the purpose of this thesis, “regret-related” conditions are considered as those externally manipulated factors that fit the major specifications for inducing the upward between-option counterfactual comparison deemed *necessary* for regret to be elicited. Inspired by the work of Mellers and others, I generate these regret-related conditions by revealing to participants the higher value monetary outcomes that would have been received had they behaved differently (and similarly relief-related conditions are induced when participants made the better choice). This experimental manipulation is widely considered to induce a form of operationalised outcome regret without recourse to subjective ratings of experience (Bell, 1982; Ert & Erev, 2007; Hart & Mas-Colell, 2000; Loomes & Sugden, 1982; Marchiori & Warglien, 2008). Moreover, numerous studies show that such feedback is strongly associated with subjective ratings of regret, general negative affect, and a desire to have acted differently, in iterated choice games (e.g. Camille et al., 2004; Coricelli et al., 2005; Inman et al., 1997; Mellers et al., 1997). While I do not see such externally controlled conditions as *always sufficient* for inducing the full emotional and cognitive experience of regret, here I will refer to them as regret-related (or regrettable) in order to clarify the states in which participants are placed in the tasks I describe.

To make this operationalisation more descriptive of regret, and more predictive of its behavioural impact, a primary aim of this thesis is to understand other factors necessary for its experience, or at least important in determining its intensity. The first two studies of this thesis test whether responsibility should be added as a necessary precursor of regret, as well as addressing the importance of responsibility in learning and decision-making more generally.

Theme 1: This thesis addresses the necessary precursors currently missing from conventional economic models of regret, with a special focus on responsibility.

A second aim of this thesis is to provide a deeper understanding of the way healthy individuals respond to regrettable experiences during private decision-making. The feeling and subjective experience of emotions are, to many researchers, of great importance. However, others place greater emphasis on a behavioural response. This functionalist approach considers any particular emotion as characterised by both the external and psychological events causing it *and* its effect on behaviour. For example, while fear and anger can both be elicited by a threatening external event, fear is associated with potential threat and tends to encourage a flight response, while anger is associated with current and unjustified threat and will more often elicit a fight response. Regret is also thought to have a powerful impact on behaviour, as discussed in section 1.3 above, although further empirical work is needed to clarify the specific direction of this impact and distinguish it from anticipated regret.

This distinction between *anticipated* and *experienced* emotion is necessary in studying the behavioural effect of regret. Anderson (2003) had noted that:

"the vast majority of studies support the conclusion that emotional goals influence decision avoidance but that postdecisional emotions are infrequently measured" (p.142).

The possibly dissociable roles of experienced and anticipated regret in decision-making has been discussed in section 1.3 above. Critically, 'Regret Theory' describes only the role of anticipated regret in decision-making, and the neuroimaging literature also follows in this approach. As expected by its aversive nature, anticipated regret encourages regret-averse behaviour, i.e. decisions that minimise the likelihood of its future experience. However it is less well understood how people respond behaviourally to the actual experience of regret. As considered in section 1.3 of this introduction, experienced regret may not encourage the same regret-avoidant behaviour thought to be associated with anticipated regret.

Theme 2: This thesis considers the possibility that the experience of regret may have a distinct behavioural impact from that associated with its anticipation.

One possible reason for this difference is that experienced regret is associated with intense and current negative affect, while anticipated regret is driven by a fear of possible future regret. While anticipated regret motivates us to adjust behaviour so as to avoid the possibility of regret being elicited (in a way that resembles a *flight* response), experiences of regret are unpleasant and our strongest motivation may be to regulate or reduce this current experience rather than to plan future avoidance. It is known that individuals use a variety of cognitive and behavioural strategies to regulate current feelings of regret (Pieters & Zeelenberg, 2007; Roese & Olson, 2007; Zeelenberg & Pieters, 2007), the behavioural effects of which may trade-off against a desire to avoid repeating the regrettable choice in the future. In their Model of Regret-Regulation 1.0, Zeelenberg & Pieters (2007) advise that the regulation of anticipated regret involves attempts to prevent regret from occurring, for example by making regret-avoidant decisions or by avoiding information that can elicit regret (e.g. feedback about the outcomes of alternative options). When it comes to regulating experienced regret, however, the priorities of the decision-maker seem to change, as shown in Table 1-1. The authors write that decision-makers are,

"motivated to avoid regret from happening and when it happens they engage in ameliorative behaviors (e.g., reverse the decision or undo the consequences). When this is not possible, they manage, deny, or suppress this experience in one of many possible ways" (p.3)

Roese & Olson (2007) add to this by highlighting a difference between behavioural and cognitive responses to regret (or indeed to any emotion). They argue that undoing or reversing the decision in a behavioural response to experienced regret is likely the “primary, pivotal, and default response” (Roese, Summerville, & Fessel, 2007, p. 27). However, all other responses in part II of Table 1-1 relate to cognitive regulatory responses, which may prevail when the regrettable action cannot be undone or reversed. Such responses, Roese et al suggest, can be thought of more generally as forms of cognitive dissonance reduction strategies, associated with a ubiquitous bias to protect self-esteem.

Theme 3: This thesis assumes that individuals use various behavioural and cognitive strategies to regulate the aversive experience of regret.

Table 1-1 – Taken from Zeelenberg & Pieters (2007) and showing that regret regulation strategies differ in the case of anticipated (I) and experienced (II) regret.

Regret Regulation Strategies

I. Prevent future regret
  1. *Decision-focused*
    a. Increase decision quality
    b. Increase decision justifiability
    c. Transfer decision responsibility
    d. Delay or avoid decision
  2. *Alternative-focused*
    a. Ensure decision reversibility
    b. Avoid feedback about forgone alternatives
  3. *Feeling-focused*
    a. Anticipate regret

II. Manage current regret
  1. *Decision-focused*
    a. Undo decision
    b. Justify decision
    c. Deny responsibility for the decision
  2. *Alternative-focused*
    a. Reverse decision (switch to alternative)
    b. Re-appraise quality of alternative
  3. *Feeling-focused*
    a. Psychological repair work
    b. Suppress or deny regret

1.6 Structure of this thesis

The five studies described in this thesis are motivated by two broad aims. Firstly, I aim to provide a deeper understanding of the experiential *content* of regret, and the external and psychological *antecedents* that are necessary for it to be elicited. Secondly, I aim to shed light on the *motivational* impact of regret, by addressing the effects of regrettable events on immediately subsequent decision-making. For the latter, I explicitly examine the effect of real experiences (with real financial incentives) in games in which individuals make repeated free decisions. This approach contrasts with the bulk of experimental studies on regret which have relied on hypothetical scenarios, anticipated or imagined emotional responses, or social (as opposed to private) decisions, and is assumed to capture more realistic behaviour preferences. Investigating the neural mechanisms underlying the experience of – and the decision strategies motivated by – regret was expected to be particularly informative for both of these aims.

The role of responsibility in the experience of regret is debated, as discussed above. Moreover, it is unclear whether regret-related brain activity is a function of subtle differences in the degree of responsibility a decision-maker exerts over a regrettable outcome. In **Chapter 3**, I present the results of an experiment designed to clarify the role of responsibility in the experience (and neural representation) of regret. Here, participants made decisions under varying levels of objective responsibility, and I tested how responsibility modulates regret-related brain responses, as well as trial-by-trial subjective ratings of regret.

It is also unclear what the role of responsibility is in our ability to learn and implement optimal decisions more generally. Specifically, does being responsible for the outcomes of our actions help or hinder our learning and decision-making? To address this, in **Chapter 4** I compared the efficacy of value learning by active trial-and-error against vicarious learning through observation of the outcomes of actions performed by others. The extant literature is ambiguous as to which of these modes of learning should be more effective, since controlled comparisons of operant and observational learning are rare. Here, I contrast human operant and observational value learning, assessing implicit and explicit measures of learning from positive and negative reinforcement. I also used a model-based analysis to address possible differences in the mechanisms underlying each form of learning. Based on a hypothesised relationship between regret and responsibility, I particularly compared models that take into account the outcomes of unchosen options (i.e. which allow for between-option counterfactual comparisons to be made) with those that do not take into account this fictive information.

With its seemingly important role in motivating behaviour, is surprising that trial-by-trial influences of regrettable outcomes on subsequent choice have not been widely studied. As discussed in section 1.3 above, economic models have addressed the impact of *anticipated* regret on choice, i.e. under conditions of a knowable risk of future regret from different choice alternatives. When the risk of future regret is uncertain across choice options, past experiences may act as a powerful guide to behaviour. It is less understood, however, whether past regret encourages the regret-avoidant behaviour typically associated with anticipated regret.

Using a task in which two choice options do not differ in probability of future regret, in **Chapter 5** I aimed to tease out the immediate behavioural impact of regrettable outcomes. To stay close to the designs of earlier economic and reinforcement learning literature, here I reduce the construct of regret to those components that are best understood and most easily manipulated experimentally. These comprise a) the upward between-option counterfactual comparison, and b) responsibility for the choice. Surprisingly, results indicated that after a regrettable outcome individuals did not tend to avoid the regret-related choice on a subsequent trial, but rather appeared to *repeat it*. This behaviour fit poorly to a typical regret-minimax model of choice that was tested against the data. It may, however, be better explained in relation to some of the regulatory strategies associated with minimising experienced regret, as in part II of Table 1-1. Specifically, the observed choice repetition could instead be associated with either an attempt to make up for the previous mistake, or with a form of decision inertia elicited by experienced regret.

The behaviour observed in Chapter 5's study cannot be readily explained by conventional models of regret-aversion. In Chapters 6 and 7, I address this intriguing finding further, using fMRI to explore neuronal mechanisms underlying a bias to repeat previously bad choices. I approach this in two different ways. In **Chapter 6**, I study neuronal responses associated with outcome-type and subsequent choice within the same gambling task used in Chapter 5. For **Chapter 7**, I designed a task to address the possibility that experienced regret may drive a bias towards decision inertia. Others have provided behavioural evidence suggesting that regret may be higher when errors arise from rejection rather than acceptance of a status-quo, and that this could encourage the emergence of a status-quo bias on subsequent decisions (Baron & Ritov, 1994; Feldman et al., 1999; Kahneman & Tversky, 1982; Landman, 1987; Tsiros & Mittal, 2000; Zeelenberg et al., 2002). Such a bias could explain the previous observation that regrettable outcomes encourage choice repetition, rather than avoidance. Motivated by the purely behavioural literature connecting regret and a status-quo bias, I acquired fMRI data during a difficult perceptual decision task with a trial-to-trial intrinsic status-quo and explicit signalling of outcomes (error or correct). I examined the neural mechanisms underlying such a bias by linking choice behaviour, and neuronal activity at choice, to antecedent error processing and regret.

In **Chapter 8**, I discuss the implications of these findings for standard models of the antecedents, experience and behavioural impact of regret. General implications and future directions are also considered. In the next chapter (**Chapter 2**) I provide an introduction to the methodology used in this thesis, and how it has allowed me to address issues and concerns described above.

Chapter 2. Introduction to the methodology

2.1 Overview

This thesis relies on a combination of behavioural methods, functional brain imaging, skin conductance recordings and computational modelling of learning and decision-making, to address the questions and hypotheses discussed in Chapter 1. In all five experiments described, healthy participants complete iterated, private (i.e. not interactive) decision-making tasks that I developed to address a particular research question of interest (or modified from tasks used in previous literature). For the purpose of clarity and coherence, in this thesis these tasks will be described in full within the relevant chapters ahead. The statistical methods used, along with those functional imaging and computational modelling procedures specific to each experiment, will also be described within the individual chapters. Below, I provide a brief introduction to these methods and discuss what they add to the study of emotion and decision-making. I also take this opportunity to provide the details of those procedures that apply to multiple experiments.

2.2 Participants

All participants were right-handed with normal or corrected vision, and no history of neurological or psychiatric disorder, according to self-report. Each gave informed consent, according to procedures approved by the UCL Research Ethics Committee. Number of participants, mean ages and gender ratios will be provided for each experiment separately. Participants were an opportunity sample and, for the purpose of generalisation of findings, were assumed to represent a random sample of the population.

2.3 Behavioural measures and analysis

The experiments used iterated, private decision-making tasks, in which outcome feedback was presented trial-by-trial to encourage participants to evaluate the quality of their choices. Choice quality evaluation was also promoted by introducing financial incentives for participants. In all experiments participants were paid according to their task performance, in addition to a reimbursement for their time. Tasks were designed to capture changes in behaviour in response to previous outcomes, but with the specifics dependent on the particular study. These tasks were either uniquely designed, or inspired by previous studies (Camille et al., 2004; Coricelli et al., 2005; Fleming, Thomas, & Dolan, 2010; Lohrenz et al., 2007; Mellers, 2000; Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006). For use both inside and outside of the fMRI scanner, visual stimuli were generated and presented to participants through the Cogent 2000 Toolbox (www.vislab.ucl.ac.uk/Cogent/) for MATLAB. Responses were recorded from the computer keyboard (in behavioural studies) or 5-finger optical button-box (inside the scanner).

Analyses of behaviour and ratings were performed using SPSS 17.0. Dependent measures comprised choice propensities or ratings of subjective experience. Independent measures were the outcome experienced on the previous trial (which due to the self-blame nature of regret was critically dependent on participant's own previous choices), as well as the particular controlled task manipulations described in depth within the relevant chapters. In all experiments, the primary statistical analyses compared means across conditions in a within-subject design, using repeated-measures ANOVAs and paired t-tests. Here I assumed equal variance across conditions. I tested this assumption with Mauchly's test of sphericity, and degrees of freedom were corrected (using a Greenhouse-Geisser correction) wherever this assumption did not hold. In Chapter 3, I also tested a series of hierarchical multiple linear regression models to assess the influence of responsibility on the experience of regret. These will be explained in the behavioural analysis section of Chapter 3.

2.4 Functional magnetic resonance imaging (fMRI)

fMRI provides an elegant way of defining the neural mechanisms underlying human behaviour and subjective experience non-invasively in healthy individuals. The ability to use event-related designs in fMRI is ideal for the study of iterated decision-making. This allows for relatively good spatial resolution and 3D localisation of brain activity, although its temporal resolution (compared to M/EEG or single unit recording) is limited by the delayed and dispersed nature of the haemodynamic response function (HRF). Here I will provide a brief introduction to fMRI, followed by a description of its use in this thesis.

Magnetic resonance imaging (MRI) measures the radiofrequency energy released by hydrogen atoms as they relax from a high energy state (elicited by a radiofrequency (RF) pulse) to a resting state along the longitudinal axis of the scanner magnet (B0). This relaxation (termed T1 relaxation) increases the strength of the MR signal in the B0 direction, and its speed differs depending on tissue type. This allows for signals from grey matter, white matter and cerebrospinal fluid (CSF) to be differentiated. A second form of relaxation in the signal occurs when the spin of neighbouring molecules diphase from each other after the RF pulse forces them into spins that are all in phase. As these spins diphase, there is a decrease in MR signal, the speed of which is also unique to the tissue type. This is termed T2 relaxation.

Inhomogeneities in the magnetic field also cause the T2 signal to decay at a faster rate (the actual decaying signal being termed T2*). One variable that influences the speed of the T2* decay is the composition of the blood supplied to the particular brain region. Specifically, deoxygenated blood is paramagnetic (Pauling, 1977). This results in greater inhomogeneities in the local magnetic field, meaning that the T2* signal is relatively faster to decay in regions with a higher ratio of deoxygenated to oxygenated blood. The BOLD (Blood Oxegenation Level Dependent) response measured in fMRI is therefore associated with an increased ratio of oxygenated to deoxygenated blood. Put simply, since active brain regions require energy from oxygenated blood, local changes in this ratio allow us to make an inference about an increase in the underlying neural activity of that brain region (Ogawa et al., 1993), after correcting for the temporally delayed and dispersed nature of the HRF. Although the BOLD contrast does not measure oxygen *usage* directly, local changes in cerebral blood supply have been linked to underlying neural activity through animal models (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001; Viswanathan & Freeman, 2007) and through comparisons to E/MEG signals measured in humans (Nangini, Tam, & Graham, 2008; Ogawa et al., 2000).

For the experiments described in Chapters 3 and 6, I scanned participants in a 3T Allegra scanner (Siemens, Erlangen, Germany) operated with its standard head transmit-receive coil. For the experiment described in Chapter 7, I scanned participants in a 3T Trio whole-body scanner (Siemens, Erlangen, Germany) operated with its standard body transmit and 12-channel head receive coil. The manufacturer's standard automatic 3D-shim procedure was performed at the beginning of each experiment. Echo-planar imaging (EPI) sequences were optimised for sensitivity in brain regions of primary interest in each experiment, the details of which are available in the methods sections of the relevant chapters. EPI magnitude images were reconstructed from the complex k-space raw data using a generalised reconstruction method based on the measured EPI k-space trajectory to minimise ghosting (Josephs, Deichmann, & Turner, 2000). EPI data acquisition was monitored on-line using a real-time reconstruction and quality assurance system (Weiskopf et al., 2007).

I acquired Fieldmaps for each subject at the start of scanning (Siemens standard double echo gradient echo fieldmap sequence, echo time = 12.46 ms, TR = 10.2 ms, matrix size = 64×64, 64 slices covering the whole head, voxel size = 3 × 3 × 3 mm). These allowed for calculation of static geometric distortions caused by susceptibility-induced field inhomogeneities, which were used to correct EPI images for both these static distortions and any changes in these distortions due to head motion (Andersson, Hutton, Ashburner, Turner, & Friston, 2001; Hutton et al., 2002). I also recorded heart rate with a pulse oximeter, along with respiratory phase and volume using a breathing belt, which were used to correct for physiological noise at the stage of data analysis. At the end of all scanning sessions, I acquired a T1-weighted anatomical scan for each participant using a Modified Driven

Equilibrium Fourier Transform (MDEFT) sequence (Uğurbil et al., 1993), with optimised parameters as described in the literature (Deichmann, Schwarzbauer, & Turner, 2004): for each volunteer, 176 sagittal partitions were acquired with an image matrix of 256×224 (Read × Phase).

Image pre-processing and data analysis were implemented using Statistical Parametric Mapping software in Matlab2009a with SPM8 for Chapters 3 and 7 (Matlab7.4 with SPM5 for Chapter 6) (Wellcome Trust Centre for Neuroimaging, Institute of Neurology, UCL, http://www.fil.ion.ucl.ac.uk/spm/). The procedure for pre-processing was identical across the three fMRI experiments described in this thesis. After discarding the first 6 volumes of each run, to allow for T1 equilibration, EPI images were realigned and unwarped using SPM (Andersson et al., 2001). This corrected the images for head movement through rigid-body realignment (taking into account translation, rotation, zoom and shear), as well as geometric distortions caused by susceptibility-induced field inhomogeneities. The latter was performed utilising the Fieldmaps processed for each participant using the FieldMap toolbox implemented in SPM (Hutton, Deichmann, Turner, & Andersson, 2004).

Each participant's structural image was then co-registered to the mean of the motion-corrected functional images using a 12-parameter affine transformation, and was segmented into grey matter, white matter and CSF according to the standard procedure in SPM (Ashburner & Friston, 2005). The spatial normalisation parameters resulting from the previous step were then applied to the functional images, to remove individual differences in brain structure and place all images onto a standardised anatomical space. This step in image pre-processing allows for statistical averaging across participants, comparison across participants, and comparisons against the findings of other studies. These spatially normalised images were then smoothed using an 8mm FWHM Gaussian kernel.

Procedures for modelling and inference were similar across the fMRI experiments of this thesis, but with the particular details depending on the contrasts that addressed the research question at hand. The general linear model (GLM) allows for estimation of the parameters that best explain the spatially and temporally continuous data collected from the 3D brain over time. The model convolves a psychological variable (or multiple variables) by a canonical HRF, in order to allow for the delayed and dispersed form of the BOLD response. In the case of all the fMRI-based studies of this thesis, the temporal derivative of the HRF is also included, in order to allow for a small amount of variance in the onset of the response. Correlation between psychological regressors is avoided as this causes problems for interpretation of effects. Movement-related effects and physiological noise are also factored out from the effects of interest. This is achieved by entering the motion parameters defined by the realignment procedure as 6 regressors of no interest, along with 17 additional regressors of cardiac phase (10 regressors), respiratory phase (6 regressors) and respiratory volume (1 regressor).

In each experiment, I implemented a group-level random-effects analysis using one-sample t-tests on the contrast images obtained from each contrast of interest for each participant. Statistical parametric maps were formed in order to locate brain regions for which the measured BOLD response can be significantly explained given the model and its fitted parameters. Family-wise error (FWE) correction was used to correct the statistical threshold for multiple comparisons (since a mass-univariate approach is used by SPM, i.e. comparing many voxels against the null hypothesis of no effects). This correction also deals with spatial correlations (i.e. non-independence) between neighbouring voxels due to spatial processing, smoothing and the spatially extended nature of the HRF (i.e. Random-Field Theory). For pre-defined regions of interest (ROIs; described within the relevant chapters), I report activity that is significant at a voxel-level, FWE corrected threshold of $p<0.05$. For completeness, I also report any activity that survives whole-brain cluster-wise corrected significance of $p<0.05$. All reported activity had a voxel-level uncorrected significance of at least $p<0.001$. Anatomical ROIs were all defined through the WFU PickAtlas in SPM (Maldjian, Laurienti, Kraft, & Burdette, 2003; Tzourio-Mazoyer et al., 2002).

2.5 Physiological measures and analysis

In the 5th experiment (Chapter 7), skin conductance responses were collected, to address task-related arousal indexed in activity of the sympathetic nervous system. Such responses reflect change in the electrical resistance of the skin, a parameter known to be sensitive the emotional state of the individual (e.g. Greenwald, Cook, & Lang, 1989; Manning & Melchiori, 1974). In my experiment, these recordings were collected, as described by Bach, Flandin, Friston, & Dolan (2009), using 8 mm Ag/AgCl cup electrodes and 0.5%NaCl gel on thenar/hypothenar of the non-dominant hand. Constant voltage was provided by a custom-build coupler, and data was recorded using a 1401 signal converter and spike software (Cambridge Electronics Design). Data were z-transformed and analysed using dynamic causal modelling (Bach, Flandin, Friston, & Dolan, 2010).

2.6 Computational modelling of learning and decision-making

Two of my experiments use a model-based approach to assess mechanisms underlying learning and decision-making in the relevant task. Here, this approach is guided both by the modified expected utility literature described in section 1.3 of the Introduction and by classical models of operant learning (e.g. see Dayan & Abbott, 2001, chap. 9).

A model-based approach allows for building of a generative model that can replicate a particular effect, and so provides a way of understanding causal mechanisms in human, animal or inorganic patterns of behaviour. However, such an approach is heavily biased by assumptions about both the direction of an effect and about the mechanism itself (i.e. since one can only *compare* models, the resulting inference depends upon which models are compared). Model-free approaches, such as comparison of mean choice propensities after various previous outcome conditions (as described in section 2.3 above), provide an appropriate test of whether an effect exists as well as the nature of the particular effect. Model-free approaches are ***sensitive*** to various effects that may not have been predicted, though they do not provide an understanding of the mechanism underlying the observed effect (i.e. they are not ***specific*** when it comes to understanding the effect). A model-based approach, on the other hand, is not sensitive to effects that are not been predicted, and so is relatively more biased by our assumptions.[2]

[2] This is not to say that a model-free approach is completely unbiased, as any statistical test is biased by the type of test that is performed and how it is interpreted compared to a null hypothesis (which is still a form of model-comparison). It does, however, require fewer assumptions.

Figure 2-1 illustrates this trade-off between sensitivity and specificity in model-free and model-based approaches. One way to improve the inferences made from a model-based approach is to guide our assumptions by the findings of a prior model-free analysis. Refining the assumptions in this way should allow for our inferences to approach a more optimal level, as indicated at the top right of the chart in Figure 2-1. As such, I combine *both* approaches to reach my conclusions in Chapters 4 and 5 of this thesis.

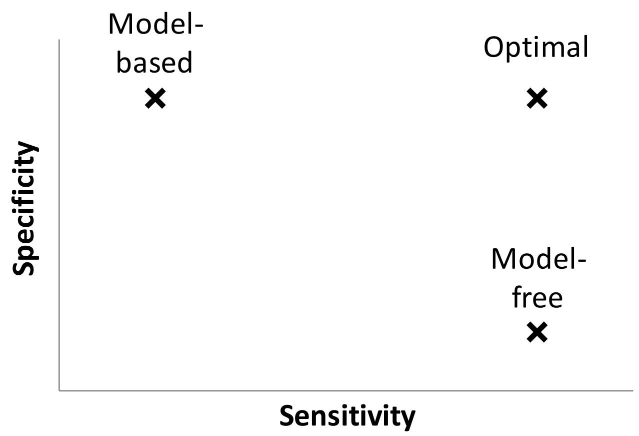

Figure 2-1 – Graphical illustration of the sensitivity-specificity trade-off encountered when empirically exploring the mechanisms underlying learning and behaviour. The optimal state is to maximise both ones sensitivity for an effect and one's specificity for capturing its underlying mechanism.

The models I test in this thesis each involve a valuation stage, in which the values of various choice options are updated based on how one's experience diverges from one's prior expectations. This is followed by a decision stage, in which values are transformed into choice propensity. Predictions from each model are fit to participants' actual choices, optimising the model's free parameters to maximise the likelihood of the model given the choices. This is realised through a standard gradient descent method (fminsearch in MATLAB). Negative log evidence (i.e. likelihoods) are penalised for the complexity of the model (i.e. the number of free parameters) by calculating the Bayesian information criterion (BIC) (Schwarz, 1978). Smaller BIC scores indicate a better fit to the data, and model comparisons were performed in order to infer the underlying mechanism that best explains observed choices.

This chapter was originally published in a shorter form in Nicolle, Bach, Frith, & Dolan, (2010). Amygdala involvement in self-blame regret. *Social Neuroscience, in press*.

Chapter 3. How important is self-blame in the experience of regret?

3.1 Introduction

"Liberty means responsibility. That is why most men dread it."

George Bernard Shaw (1903)

As discussed in Chapter 1, outcomes that could have been better from the same choice tend to induce feelings of disappointment, while outcomes that could have been better from a different choice induce feelings of regret (Roese & Olson, 1995a). This focus on the *different choice*, or the between-option counterfactual comparison, invokes a prediction that regret depends upon self-blame, whereas disappointment is associated more with a sense of bad luck and/or powerlessness (Zeelenberg et al., 1998). While some have suggested that "outcome regret" can be experienced without a feeling of self-blame (Connolly & Zeelenberg, 2002; Simonson, 1992), a dominant view is that self-blame or responsibility is important, if not necessary, for the experience of regret (see Connolly, Ordónez, & Coughlan, 1997; Ordónez & Connolly, 2000; Zeelenberg, van Dijk, & Manstead, 1998, 2000 for a thorough review and discussion).

The role of agency in modulating brain activity associated with choice outcomes has been highlighted by findings that ventral striatum responses to absolute losses and gains depend on agency (Coricelli et al., 2005). In fact in Chapter 6 of this thesis, I extend this finding by showing that agency also modulates ventral striatal responses to outcomes that are *relatively* better or worse than what might have been from a different choice. However, self-blame involves more than just agency over a choice, depending also on social norms and decision justifiability (Connolly & Zeelenberg, 2002). For example, we may more easily justify a bad decision, thus reducing feelings of self-blame, if we know that others would have made a similar choice. Similarly, if others actually played a part in the bad decision, self-blame regret can be reduced by transferring responsibility to them.

The likely importance of self-blame in regret led me to predict that regret-related responses (both psychological and neurobiological) would be modulated by the degree to which a decision-maker feels responsible for an outcome. Thus, I designed an experiment where participants experienced outcomes of 'played' and 'unplayed' gambles under various levels of responsibility (a task modified from Mellers, Schwartz, & Ritov, 1999, e.g. as in Chapter 1, Fig 1.1). This paradigm created situations where participants' experienced sense of responsibility for outcomes was systematically varied. This was realised by informing participants that each of their played gambles would depend on their own choice along with the votes of varying numbers of additional individuals. Here, I predicted that participants' subjective sense of responsibility for gamble outcomes would decrease as a function of greater numbers of these “other voters”.

Responsibility might be expected to be expressed in brain regions implicated in agency and motor control, including the insula and angular gyrus (Farrer et al., 2003; Farrer & Frith, 2002), although such regions have previously been implicated only in *being* responsible, as opposed to *feeling* responsible. In relation to my central question, I predicted regret-related brain activity would be modulated by the degree to which an outcome of a played gamble is worse than that of the unplayed gamble in regions previously implicated in regret (including OFC, amygdala, hippocampus, ACC, insula and striatum, see Chapter 1, section 1.4), but only under higher levels of responsibility.

3.2 Methods

3.2.1 Participants

I recruited 18 participants (10 female) for the experiment. Participant age ranged between 19 and 30 yrs (mean = 23.67 yrs). Due to incomplete behavioural data collection, one participant was removed from the behavioural analysis, but included in the imaging analysis. A second participant was removed from the imaging analysis due to scanner malfunction, but was included in the behavioural analysis. A further participant was excluded as an outlier, after showing a correlation between regret-related outcomes and subjective ratings of regret that was negative and 2.5 standard deviations from the mean.

3.2.2 Experimental procedure

The aim of this study was to explore how subjective and neuronal responses to regret-related outcomes are modulated by responsibility. I used a task based on that of Mellers et al. (1999). Participants were instructed to choose between two “wheel-of-fortune” gambles on each trial, each incorporating a win and loss outcome with differing probabilities (25%, 50% or 75%) (e.g. see Figure 3-1). There were 24 different gamble pairs, allowing for 4 different pairs per probability combination (e.g. 25% win against 50% win). Points allocated to the possible winning and losing outcomes were such that the two gambles in the pair were of equal expected value (EV) (i.e. probability of win × magnitude of possible win). In order to enhance feelings of skill in the game, 2 further pairs types made up 7% of trials and included one gamble of a clearly higher EV than the other. Details of the gamble pairs used are available in Appendix A.

Participants were told that their choice would count as one vote towards which gamble was played. In mini-blocks of 5 trials, participants played in a group alongside 0, 2, 4, 6 or 8 other voters, who they believed to have performed the task in advance, but who were not real. After selecting their preferred gamble, the gamble receiving the highest number of votes from the group (including their own vote) was ‘played’. When there were more than 0 other voters, there was a chance that the played gamble would not be congruent with the participant’s own choice. The likelihood of participants’ chosen gamble being played varied probabilistically as a function of the number of voters, such that increased number of voters meant an increased chance that their gamble was not played, and 0 other voters meant that their chosen gamble was always played. If the played gamble was *incongruent* with the participant’s own gamble choice, the lowest sense of responsibility was predicted, as participants were not the agents of the choice. If the played gamble was *congruent* with the participant's own gamble choice, however, participants' subjective sense of responsibility for gamble outcomes was predicted to decrease with increased numbers of "other voters", even though participants were still the agents of the choice.

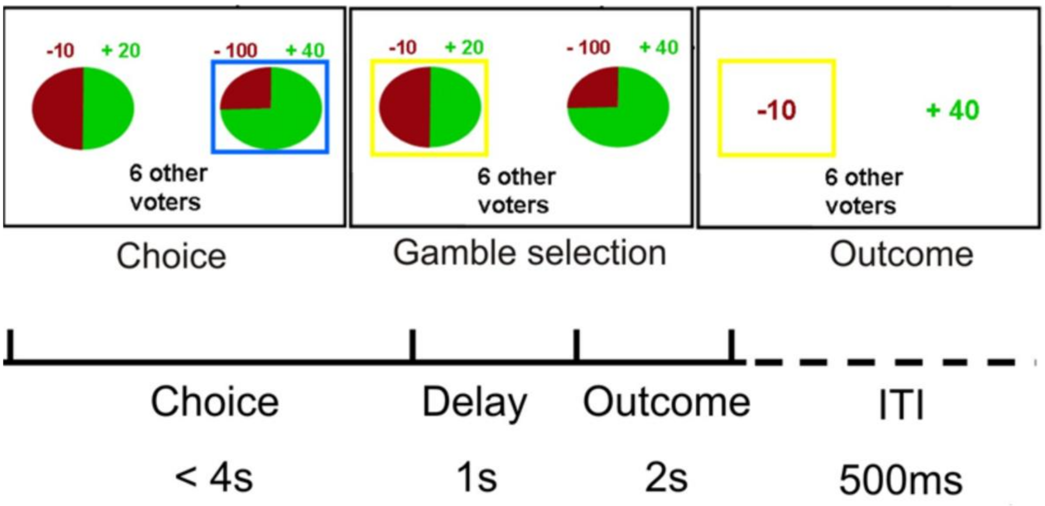

Figure 3-1 – An exemplar trial timeline. Trials included a choice phase in which participants select their choice of gamble, with choice indicated by a blue box. For each gamble, outcome probabilities were indicated by wedge size, with corresponding outcomes written above in number of points. Number of other voters was also shown, but was constant for mini-blocks of 5 trials. A gamble selection stage followed, in which the gamble receiving the highest number of votes was indicated with a yellow box. This gamble was then played, with its outcome affecting participants' winnings whether chosen by them or not. Next, the outcome of both the played and unplayed gambles was revealed.

After a 2 second delay, participants were then shown the points outcome of both the played and unplayed gamble. The outcome of played gamble, whether congruent or incongruent with their own gamble choice, determined participants' payment for the experiment. Participants received 50p for each percentage won of the maximum points they could have won in their game. This encouraged participants to compare the received outcome with what might have been under the unplayed gamble, on a trial-by-trial basis. Participants were assumed to treat all trials as having an equal impact on their financial gain.

Participants performed both a behavioural session and a scanning session, each occurring on separate days (order counterbalanced across participants). The behavioural session was optimised to gather trial-by-trial subjective ratings of regret and responsibility for different outcome types. Here, participants played 320 trials in four sessions and, after every trial, provided two subjective ratings on a 100-point horizontal visual analogue scale (this was not practical in the scanning session). The starting position of the slider on each rating scale was random between the extremes of 1 and 100, in order to avoid anchoring effects. The rating questions comprised either a memory probe asking "How much bigger or smaller was the received gamble outcome than what might have been received from the unplayed gamble?" or a rating of subjective negative feeling comprising the following probes, "How positive or negative do you feel about the trial outcome?", "How much regret do you feel for outcome?" and "How responsible do you feel for the outcome?" It was randomly decided which 2 of these 4 ratings would be presented on each trial. It was explained to participants that these ratings related to their response to the *outcome* of each trial.

Only the memory test, and not the other ratings, was used during the scanning session (in 10% of trials) in order to enhance the tendency to think counterfactually. In the scanning session, participants also played 320 trials in four 8 minute games.

3.2.3 Imaging acquisition and processing

Participants were scanned with a single-shot gradient-echo EPI sequence optimised to reduce BOLD sensitivity losses in the orbitofrontal cortex due to susceptibility artifacts, using a combination of increased spatial resolution in the readout direction and reduced echo time (Weiskopf et al., 2007). Imaging parameters were as follows: 48 oblique transverse slices tilted by 30^o, resolution of 1.5mm in the readout direction and 3mm in the PE direction, slice thickness = 2 mm, gap between slices = 1 mm, repetition time TR = 3.12 s, $\alpha = 90^o$, echo time TE = 25 ms, bandwidth (BW) = 1,953 Hz/pixel, phase-encoding (PE) direction anterior-posterior, field of view (FOV) = 192×192 mm^2, matrix size 128×64, z-shim gradient pre-pulse moment = -1.4 mT/m×ms.

Image preprocessing and data analysis were implemented using Statistical Parametric Mapping software in Matlab2009a (SPM8; Wellcome Trust Centre for Neuroimaging, at UCL), as described in Chapter 2, section 2.4.

3.2.4 Analysis of subjective ratings

I operationalised “regret-related outcomes” as those where the outcome of the played gamble showed a negative discrepancy with (i.e. is worse than) that of the unplayed gamble. Critically this outcome discrepancy relates to a between-option counterfactual comparison. The design incorporated a continuous variable of this between-gamble negative outcome discrepancy (i.e. received outcome – foregone outcome). I predicted that subjective ratings of regret (as well as negative affect) would increase with increasingly negative outcome discrepancy. I also predicted that subjective ratings of responsibility would decrease with increasing number of other voters. Finally, I predicted that the relationship between subjective regret and negative outcome discrepancy would depend upon both subjective and objective measures of responsibility. I used linear regressions to test for continuous relationships between outcome discrepancy and subjective ratings of regret, since the independent variables were on a continuous scale. To test the relationship between subjective responsibility and number of voters and choice congruence, I used a repeated-measures ANOVA, allowing for the effects of these discrete factors of objective responsibility to be addressed.

To assess the main hypothesis that regret depends on choice responsibility, and the degree thereof, I tested a multiple regression model with subjective ratings of regret as the dependent variable. The model entered as independent variables the outcome discrepancy (a continuous variable from extreme positive to extreme negative) along with the interaction of this discrepancy with a measure of responsibility (outcome discrepancy × responsibility). For the latter, three distinct measures of responsibility were tested comprising subjective ratings of responsibility, choice congruency, and the number of other voters. For its descriptive value, I also performed an ANOVA, similar to that described above in the case of subjective responsibility, to test the relationship between subjective regret and number of voters and choice congruence for negatively discrepant outcomes only (i.e. only for those outcomes objectively thought to induce regret, rather than relief).

As multiple subjective ratings were obtained from each participant, along parametric continuums of regret, responsibility and negative feeling, these ratings were standardised for each participant, to avoid anchoring effects, and the standardised regression coefficients for all analyses were calculated on an individual subject basis. I report mean standardised regression coefficients from the between-subjects level of a hierarchical linear model. Finally, with a one-sample t-test, I tested whether participants' answers on the post-trial memory tests showed performance significantly above chance. This allowed for verification that participants were taking into account both the outcomes received from the chosen and the alternative outcome of the unchosen gamble on each and every trial.

3.2.5 Imaging analysis

For the fMRI analysis, I used a two variable parametric design, with one factor for the number of other voters (0, 2, 4, 6 and 8 other voters), comprising the manipulation of responsibility, and one factor for level of outcome discrepancy, which ranged from -200 to +200 points. This factor was transformed such that positive numbers were regret-related (i.e. negative outcome discrepancy). An additional two level factor was expected to influence sense of responsibility, and comprised whether or not a participant's choice of gamble was congruent or incongruent with the gamble selected by the majority vote.

For each participant, an event-related GLM included 9 regressors of interest. These comprised one regressor for the onsets of trial outcomes at each level of responsibility, separated out as 0, 2, 4, 6 and 8 other voters. Trials in which there were more than 0 other voters were further segregated as a function of whether participants' choices were congruent, or incongruent, with the gamble selected by the majority vote. In trials with 0 other voters, participants' choices were always congruent with the gamble selected, as participants were the only voter. Each of these 9 regressors was parametrically modulated by a mean-corrected regressor of outcome discrepancy. Positive values of this parametric regressor were regret-related outcomes (i.e. negatively discrepant). Onsets were modelled with stick-functions at the time at which participants received outcome feedback, convolved with a canonical HRF and its temporal derivative.

I generated statistical parametric maps from my contrasts of interest, including the main effects of responsibility. Here, I were particularly interested in probing activity within the regions of interest (ROIs) reported by Farrer & Frith (2002), as involved in decreased (the angular gyrus) and increased motor control (the insula). As mentioned above, while these regions are implicated in simple motor responsibility (or agency), they may also be involved in processing variations in blame based upon how easily we can transfer outcome responsibility externally (of key importance for decision justification models of regret, e.g. Connolly & Zeelenberg, 2002).

Next I tested the main effect of negative outcome discrepancy (as a linear parametric effect) where I was interested in activity that increased with greater levels of negative outcome discrepancy independent of level of responsibility, as well as (using a conjunction analysis) brain activity that showed significantly increased response to negative outcome discrepancy across all levels of responsibility. I tested for these effects within anatomical ROIs in regions previously implicated in the experience of regret and regret induced decision bias, including OFC, anterior cingulate cortex (ACC), amygdala, hippocampus, insula, and striatum each defined anatomically and bilaterally. To address the hypothesis that regret-related responses are modulated by responsibility I compared the response to increasingly negative outcome discrepancy under full responsibility (i.e. 0 other voters) with that when other voters were present. Critically this contrast was performed only on trials where the participant's own chosen gamble was played (i.e. congruent choice). Additionally, I compared response to increasingly negative outcome discrepancy under congruent choice versus incongruent choice.

3.3 Behavioural results

A one sample t-test confirmed that participants had above chance memory for whether the outcome received from played gambles were better or worse than that which might have been received from the unplayed gamble, as determined by post-trial memory questions ($t(15) = 15.09$, $p < 0.001$, mean percent correct = 91.3%). This showed that participants were aware of both the received and the unplayed alternative outcomes on each trial. Mean earnings were £12.20 (*SD* £2.64) in the behavioural session, and £22.43 (*SD* £2.67) in the scanning session.

The predicted linear effect of number of voters was found on subjective ratings of responsibility ($F(1,15) = 49.5$, $p < 0.001$), reflecting an increased sense of responsibility with decreased number of other voters. This linear effect was also significant within congruent choice trials alone ($F(1,15) = 24.0$, $p < 0.001$), i.e. when participants had chosen the played gamble, but not within incongruent choice trials ($F(1,15) = 0.6$, ns). There were also significant quadratic ($F(1,15) = 9.1$, $p < 0.01$) and cubic ($F(1,15) = 5.6$, $p < 0.05$) effects of voters in the congruent choice condition, indicating that the influence of number of voters on subjective responsibility is not purely linear. Furthermore, participants showed significantly higher ratings of responsibility for congruent than incongruent choice trials ($F(1,15) = 46.9$, $p < 0.001$) (Figure 3-2-a, blue line).

Single-subject standardised regression coefficients, taken forward to a between-subject one-sample t-test, showed that increasingly negative outcome discrepancy significantly predicted greater subjective ratings of regret (mean $R = 0.52$, $t(15) = 8.04$, $p < 0.001$). This showed a successful manipulation of subjective regret by outcome discrepancy. More negatively discrepant outcomes also significantly predicted increased general negative affect (mean $R = 0.71$, $t(15) = 14.19$, $p < 0.001$).

A multiple regression analysis indicated that increased responsibility amplified the tendency for participants to report high subjective feelings of regret for outcomes with increasingly negative discrepancy. That is, outcomes were rated as more regretful when they were *both* more negative than what would have been from the alternative gamble *and* when participants felt more responsible. The three regression models used indicated that this effect was significant for all of the three measures of responsibility; namely, decreased number of voters ($t(15) = 2.82$, $p < 0.02$), increased subjective ratings of responsibility ($t(15) = 3.83$, $p < 0.002$), and choice congruency ($t(15) = 3.29$, $p < 0.005$).

I show the regression coefficients for the three regression models in Figure 3-2-c. These regression models were *not* restricted to negatively discrepant outcomes (i.e. they include the full spectrum of outcome discrepancy), and were performed using continuous functions of subjective responsibility and outcome discrepancy. To further illustrate the direction of this effect, I show subjective regret for just negatively discrepant outcomes under the different levels of objective responsibility in Figure 3-2-a (red line), and under a median split of high and low *subjective* responsibility in Figure 3-2-b. In an ANOVA, I found a linear effect of number of voters on subjective ratings of regret for negatively discrepant outcomes ($F(1,15) = 18.7$, $p < 0.001$), reflected in increased regret with decreased number of other voters. The linear effect was significant within congruent choice trials alone ($F(1,14) = 6.8$, $p < 0.05$), i.e. when participants had chosen the played gamble, but not within incongruent choice trials ($F(1,14) = 0.3$, ns). Participants also showed significantly higher ratings of regret for congruent than incongruent choice trials ($F(1,15) = 30.7$, *p* < 0.001) (Figure 3-2-a, red line), and for high compared to low subjective responsibility ($t(13) = 4.5$, $p < 0.001$) (Figure 3-2-b).

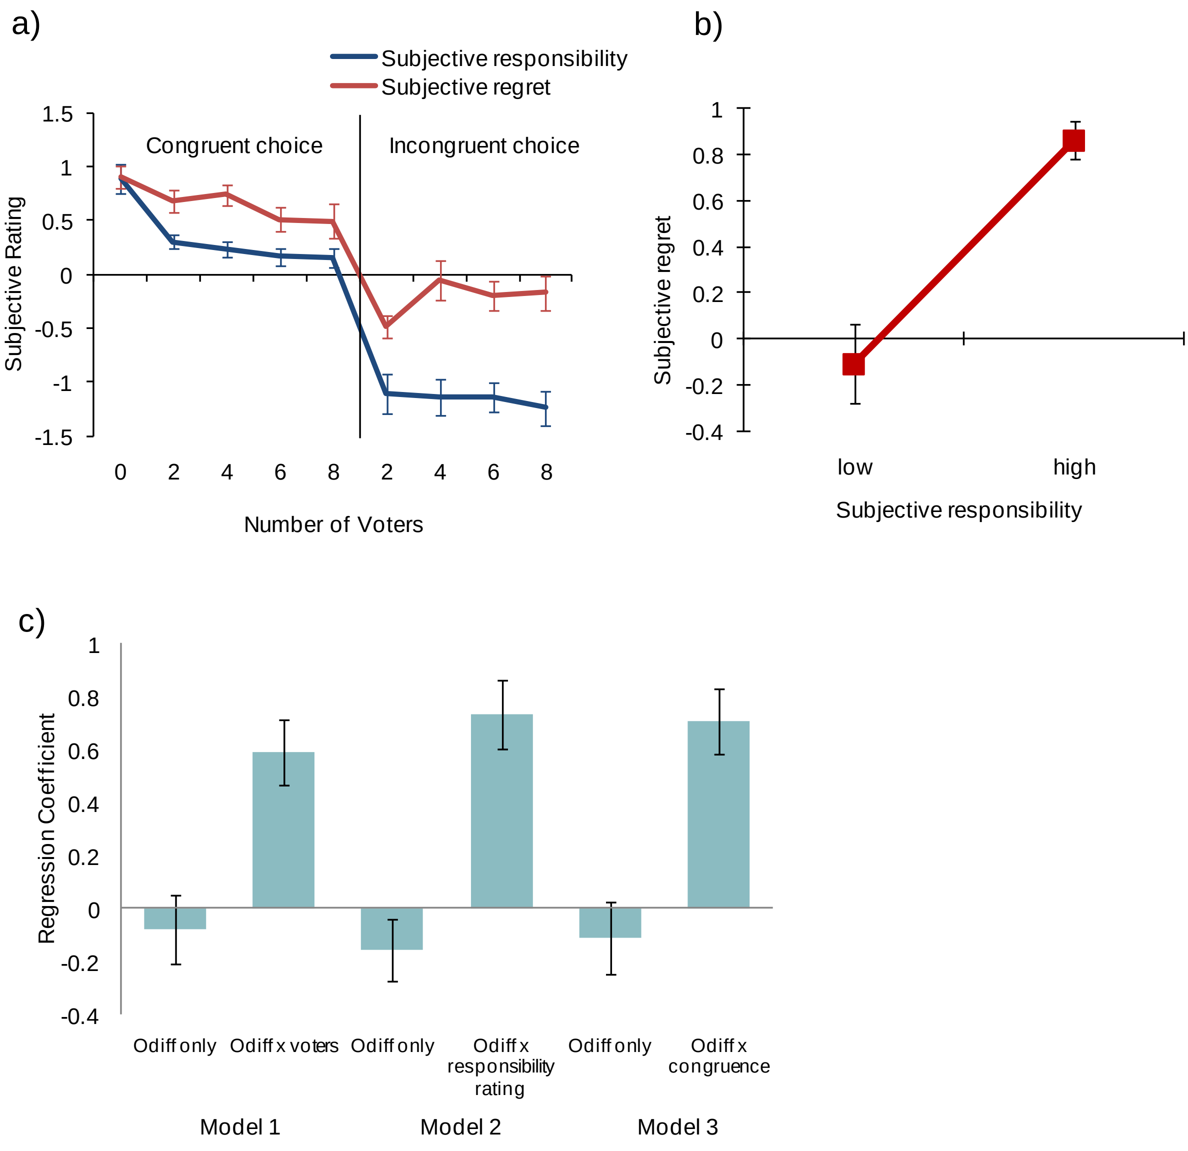

Figure 3-2 – Panel a) shows gradually decreasing mean normalised ratings of subjective responsibility (blue line) and of subjective regret (red line) under congruent and incongruent choices with 0, 2, 4, 6 and 8 other voters, otherwise considered objective responsibility. For clarity, subjective regret ratings here are shown only for negatively discrepant outcomes, i.e. where the outcome of the played gamble was worse than the outcome of the unplayed gamble. Panel b) shows the mean normalised subjective ratings of regret, displayed as a median splot of low and high subjective ratings of responsibility, for negatively discrepant outcomes. Panel c) shows the mean regression coefficients for the three regression models used, each indicating that the predictive value of outcome discrepancy (Odiff) on subjective ratings of regret is greatly enhanced by multiplying this Odiff by either of the three measures of responsibility (number of voters, subjective rating of responsibility, or choice congruence). Error bars show the standard error of the mean across participants.

3.4 fMRI results

3.4.1 Manipulation of responsibility

I examined the effect of linearly decreasing responsibility, following the pattern shown in Figure 3-2-a. While no areas showed a significant effect of increasing responsibility, I found significant effects in superior frontal cortex (MNI 51, 26, 31), brainstem (MNI 0, -34, -26) and left insula (MNI -45, 14, -8) for decreasing responsibility. In the a priori ROIs, right angular gyrus activity (within a 20mm sphere radius of the coordinates reported by Farrer & Frith, 2002) significantly increased with decreasing responsibility (MNI 57, -43, 34), as did bilateral insula (MNI right 30, 20, -11 and left -45, 14, -8). No regions dissociated, at the time of outcome, between played gambles that were congruent or incongruent with the participant's own choice, at a whole brain corrected level. However, right insula activity showed such an effect within the a priori ROI of bilateral insula (MNI 30, 20, -14).

3.4.2 Manipulation of regret

Averaged across all levels of responsibility I found significantly increased activity as outcomes became linearly more negatively discrepant in left angular gyrus and lateral OFC (Figure 3-3-a). Activity in other a priori regret-related ROIs (in the amygdala, hippocampus, ACC and insula) did not show this effect. Crucially, no regions showed responsibility-invariant responses to negatively discrepant outcomes, as evident in a conjunction analysis, providing evidence against a neural representation of what a purely outcome-based form of regret. Instead, regret-related responses in angular gyrus and lateral OFC appeared to be dependent upon the level of responsibility. However, it is important to note that the reduced power inherent in such a conjunction analysis means that I cannot reject the null hypothesis that such responsibility-invariant responses do exist.

To address whether activity associated with negative outcome discrepancy is modulated by level of responsibility, I restricted the analysis to congruent choice trials, i.e. where the participant’s chosen gamble was played. This was since negatively discrepant outcomes received in incongruent choice trials were associated with participants actually having been the agents of the post-hoc *better* choice. Within a priori anatomical ROIs implicated in regret, I found no regions showing an entirely linear enhanced response to more negatively discrepant outcomes by decreased number of other voters (i.e. increased responsibility). However, in keeping with the marked step-like decrease in rated responsibility from 0 to 2 or more other voters (shown in Figure 3-2-a), I found that left amygdala activity was enhanced for more negatively discrepant outcomes during full responsibility trials, and not when there were any number of other voters (Figure 3-3-b and c) (within anatomically specified bilateral ROI of the amygdala). A linear interaction of outcome discrepancy and responsibility (within congruent choice trials) was seen in right amygdala activity but only in participants who showed a greater enhancement of subjective regret by responsibility in the separate behavioural session, as indexed by a greater difference between the first two bars of Figure 3-2-c. This effect (shown in Figure 3-3-d) was also seen in left amygdala at a trend level of significance.

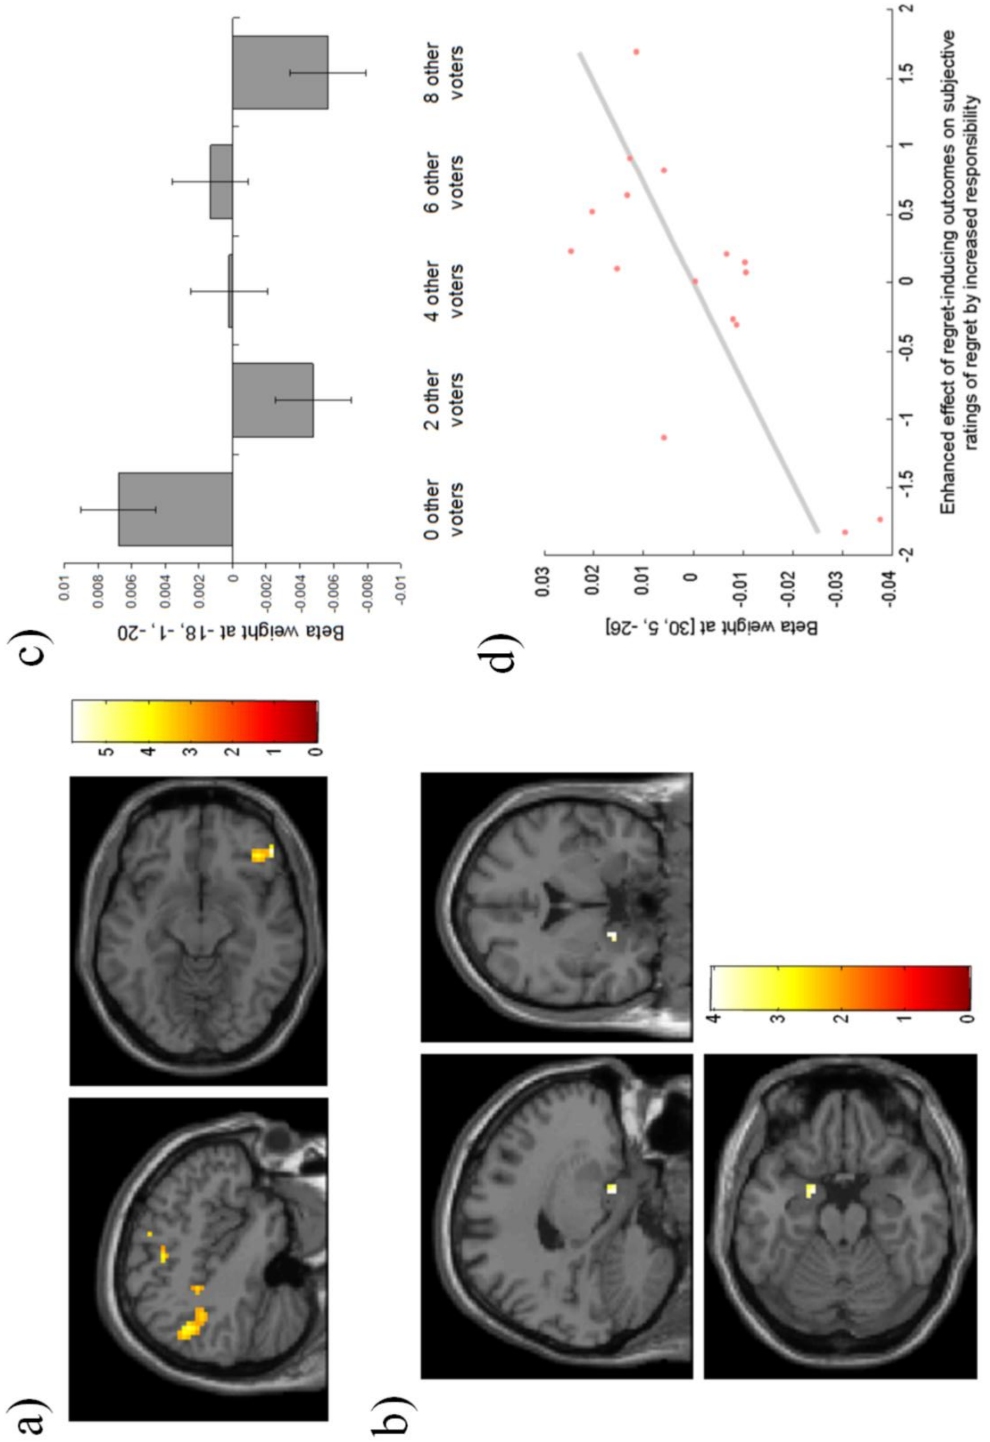

Figure 3-3 – Group SPM data, thresholded at p<0.005 for display purposes, shown on a normalised canonical template brain. a) Activity in angular gyrus (peak at -42, -64, 25) and lateral OFC (peak at 51, 38, -11) associated with the average effect of linearly increasing negative outcome discrepancy. b) Enhanced activity in amygdala (peak -18, -1, -20) associated with increasingly negatively outcome discrepancy during full responsibility (0 other voters) compared to other congruent choice trials in the presence of other voters. Plotted in c) are the beta weights showing the same amygdala response to increasingly negative outcome discrepancy under the different numbers of other voters. d) Right amygdala activity (peak 30, 5, -26) showing greater linear effect of increased responsibility on its response to more negatively discrepant outcomes (FWE corrected) in those participants exhibiting a greater enhancement of subjective regret ratings by increased responsibility (decreased number of voters) as indexed by a greater difference between the first two bars of Figure 3-2-c. Co-ordinates are in MNI space. Error bars are standard error of the mean.

By contrast, activity in middle cingulate cortex and angular gyrus showed an increased response to more negatively discrepant outcomes, when the played gamble was incongruent (whole brain cluster-level corrected $p < 0.05$), compared to when it was congruent with the participant's own choice. This indicates a role of middle cingulate cortex and angular gyrus in the processing of regret-related outcomes that are externally enforced. Furthermore, activity in middle cingulate cortex, left angular gyrus, and lateral OFC responsive to average effect of regret-related outcomes (Figure 3-3a), showed a greater response to more negatively discrepant outcomes on these incongruent trials compared to congruent trials (although this was only at $p < 0.005$ uncorrected level for the lateral OFC). This indicates that regret-related activity in these regions is not associated with self-blame, but perhaps rather to an external attribution of blame. This possibility will be explored further in section 3.6 below.

3.5 Discussion

Here I have shown that subjective ratings of experienced regret depend *both* on the outcome being worse than what might have been from a different action *and* on perceived responsibility. Moreover, I provide new evidence that subjective regret depends upon level of subjective responsibility even though the individual's own choice or action directly contributed to the regret-related outcome (as in the case of congruent trials). These findings indicate that regret is not just a function of being the agent of a choice but also depends upon subtle changes in sense of responsibility, or accountability, for the outcomes of our actions.

In keeping with these behavioural effects regret-related neuronal activity in the amygdala was enhanced by increased responsibility, suggesting a critical role in 'self-blame regret'. This effect was magnified in participants who displayed a greater enhancement of their subjective ratings of regret by responsibility. Interestingly, I did not find any brain regions responding to what has been termed "outcome regret", i.e. showing invariant responses to regret-related outcomes under all levels of responsibility. This suggests that, as for the psychological experience of regret, the way the brain processes regretful events may crucially depend upon a sense of responsibility.

The human amygdala is known to be important in emotional memory (Cahill, Babinsky, Markowitsch, & McGaugh, 1995; Richardson, Strange, & Dolan, 2004; Strange & Dolan, 2004) and in learning stimulus-reward associations (Büchel, Morris, Dolan, & Friston, 1998; Gottfried, O'Doherty, & Dolan, 2003; LaBar, Gatenby, Gore, LeDoux, & Phelps, 1998; Whalen et al., 2004). A critical model of the amygdala for the present study relates to its putative role in "relevance detection" (Bach et al., 2008; Sander, Grafman, & Zalla, 2003). This proposes that the amygdala is involved in focussing attentional and physiological resources towards cues that have special relevance for our safety or success. Amygdala activation is found in socially relevant situations, such as viewing untrustworthy or novel faces (Winston, Strange, O'Doherty, & Dolan, 2002; Wright et al., 2003) or following eye gaze (Kawashima et al., 1999), along with especially *self-relevant* situations, e.g. when one's own name is presented during sleep (Portas et al., 2000). Moreover, the amygdala is implicated in biasing future decisions by previous regret (Coricelli et al., 2005), emphasising its importance in goal-directed motivation of behaviour. Regret associated with high responsibility for its occurrence is an experience with particular self-relevance, while also being a strong motivator of future behaviour. Experimentally, self-blame regret is known to motivate active attempts to (or intentions to) undo unpleasant events (Zeelenberg & Pieters, 1999; Zeelenberg et al., 1998). A failure to appropriately accept responsibility for our mistakes may interfere with learning the accurate associations between our actions and their outcomes, which are vital for us to adapt future behaviour to improve our wellbeing. Specifically, self-relevant information may be passed to the amygdala from other cortical regions involved in the particular task at hand (for example, I have shown information about responsibility and between-option outcome comparisons to be associated with activity in the angular gyrus and prefrontal cortex). It may then be used by the amygdala to allocate processing resources to cortical mechanisms appropriate for motivating adaptive behaviour. In keeping with such a framework, evidence suggests that the amygdala imparts information necessary for acquiring stimulus-reward associations to the OFC, which uses this information to guide behaviour (Arana et al., 2003; Pickens et al., 2003; Schoenbaum, Gottfried, Murray, & Ramus, 2007), allowing for behavioural flexibility in accordance with our goals (Morris & Dolan, 2004).

In summary, I provide evidence that self-blame should be included as a necessary precursor to the experience of regret, alongside the upward between-option counterfactual comparison. I also show that in the amygdala – a region important in gathering personally relevant information, in forming stimulus-reward associations, and in guiding future behaviour – response to outcomes that could have been better from a different choice are enhanced by responsibility, even when a free choice has been made. This highlights a particular role of the amygdala in the self-blame component of our experiences of regret.

3.6 Further findings related to blaming others

A problem faced when studying the neural representation of self-blame regret is that people use strategies that regulate their aversive experiences. In an elegant paper by Zeelenberg & Pieters (2007, see also Pieters & Zeelenberg, 2007), the mechanisms of regret-regulation are discussed. They reason that regret is an aversive emotion that we are motivated either to avoid or, once it is unavoidably elicited, to regulate, modulate or suppress. One such regulation strategy, they argue, involves denying or deferring responsibility for the relevant action/decision. In the context of consumer decision-making, an example would be to blame a sales advisor for a bad product choice so as to reduce one's own feeling of regret. A separate stream of research into attribution biases also finds that people often attribute their own behaviour to uncontrollable, external causes, while attributing the behaviour of others to their internal disposition when the behaviour, or its outcome, is unpleasant (e.g. Jones & Harris, 1967; Storms, 1973). Cialdini, Braver, & Lewis (1974) showed that observers made more dispositional inferences about an actor when such inferences serve to improve their own self-image. As well as improving mood and self-esteem, this bias may be driven by a desire to feel in control of our environment (e.g. Miller, 1978; Miller & Norman, 1975), and to better predict the future behaviour of others (Heider, 1958; Kelley, 1971).

Such attribution bias is dependent on the outcome of the behaviour. Individuals tend to make dispositional attributions, i.e. take more personal responsibility, for their own behaviour when it results in a success, while making external attributions, i.e. transferring responsibility, when failure ensues. This effect has been associated with a self-serving bias and is widely replicated (e.g. Beckman, 1970; Johnson, Feigenbaum, & Weiby, 1964; Miller, 1976; Streufert & Streufert, 1969; Wolosin, 1973; Wortman, Costanzo, & Witt, 1973). For example, Johnson et al. (1964) found that teachers tended to accept responsibility for any improvement in their students' performance, but tended to blame the children for any lack of improvement. Beckman (1970) showed also that external observers attributed less responsibility to the teacher for successes than did the teacher themselves, suggesting that the teacher's bias was motivated by a desire to protect their esteem. This form of self-protecting, emotional regulation could be conceptualised as follows: The stronger the regret induced by a regrettable outcome, the more these regret-regulatory strategies are promoted. While self-blame would usually play a role in *enhancing* feelings of self-blame regret, these regulatory strategies act to *reduce* the experience of regret and perhaps also its neuronal representation. In Figure 3-4, I illustrate how regretful outcomes may excite regret-regulatory strategies, including denial or deferral of blame, with an aim to improve mood. Such strategies, however, inhibit a sense of blame for the regrettable outcome, which has the consequence of reducing regret.

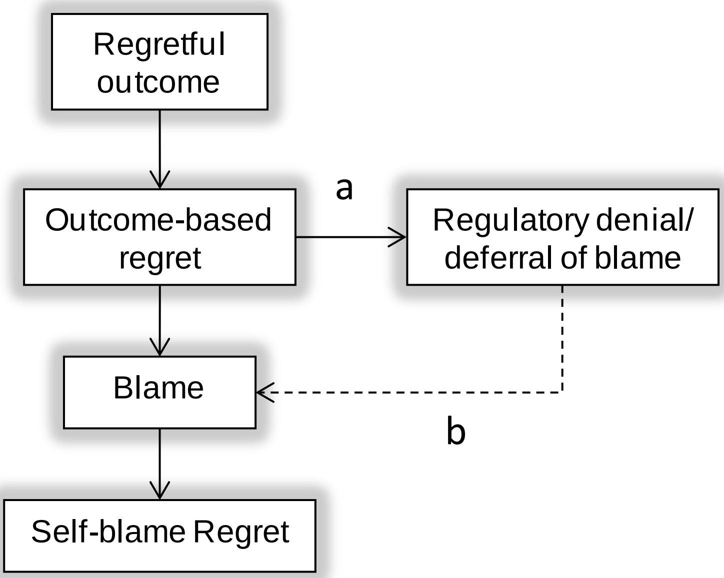

Figure 3-4 – Possible mechanism of a self-protecting attribution bias in the regulation of regret, either behaviourally or neuronally. Solid lines indicate excitatory effects, while broken lines indicate inhibition.

3.6.1 Regret reduces subjective responsibility

I also explored the current data to examine the possibility that highly regrettable outcomes may actually be associated with *decreased* subjective ratings of responsibility compared to less regrettable outcomes (in line with pathways a and b in Figure 3-4). In Figure 3-5 and Figure 3-6, the blue lines illustrate that participants indeed tend to report *lower* subjective responsibility for outcomes with *higher* negative discrepancy. A hierarchical multiple regression analysis found that both subjective regret ($t(15) = -3.914$, $p < 0.001$) and negative outcome discrepancy ($t(15) = -4.771$, $p < 0.001$) significantly diminished the tendency to claim high responsibility for trials with lower number of other voters. That is, participants were *less* likely to report feeling responsible for objectively high responsibility trials when the outcome was relatively *more* regretful.

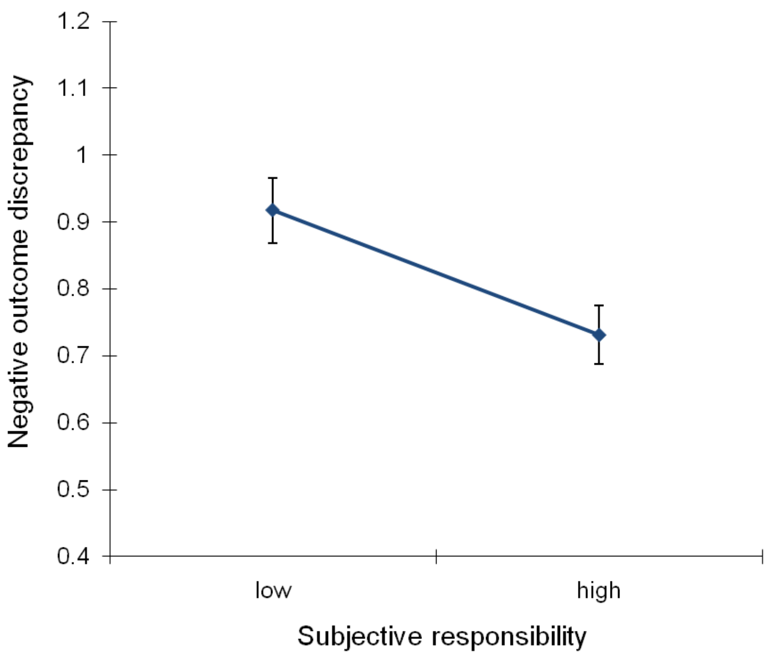

Figure 3-5 – The figure shows the mean normalised actual outcome discrepancy for negatively discrepant outcome trials only, separated according to a median split of participants' subjective ratings of responsibility. The figure shows that more negatively discrepant outcomes are associated with lower ratings of subjective responsibility. Error bars show the standard error of the mean.

3.6.2 The aversive experience of low responsibility

Higher ratings of disappointment may be associated with *lower* responsibility (Zeelenberg et al., 1998). If regret is reduced by the regulatory strategy illustrated in Figure 3-4, disappointment may then be the more salient and painful emotional response in conditions in which responsibility can be deferred. I tested whether such an effect was apparent in the subjective negative feeling ratings collected in the present study. For this, I implemented a further multiple regression model, now with *subjective negative feeling* as the dependent variable. The model entered as independent variables the outcome discrepancy along with the interaction of this discrepancy with a measure of responsibility. Analogous to the initial model with subjective regret as the dependent variable, I tested the effects of subjective ratings of responsibility, choice congruency, and the number of other voters separately. This analysis found that participants were more likely to say they felt negative about negatively discrepant outcomes when they were *less* responsible. This was significant for decreased subjective ratings of responsibility ($t(15) = 9.249$, $p < 0.001$) and choice incongruency ($t(15) = 6.179$, $p < 0.001$), but not for objective number of voters ($t(15) = 0.873$, ns). The red line in Figure 3-6 illustrates that participants felt more negative affect for *lower* subjective responsibility trials.

The regions that had shown an average effect of increasing negative outcome discrepancy described above, i.e. middle cingulate cortex (MNI -6, -4, 40) angular gyrus (MNI -42, -64, 19) and lateral OFC activity (MNI 42, 38, -8), responded more to increasing negative outcome discrepancy within incongruent choice trials compared to in congruent choice trials, as described in 3.4.2. These effects were apparent within a functional mask taken from the original average effect of negative outcomes discrepancy (with a p < 0.005 mask threshold). Within congruent trials alone, posterior cingulate cortex also showed a greater response to increasingly negative outcome discrepancy with increasing number of other voters (i.e. linearly decreased responsibility) (MNI 12, -37, 31), as did the right angular gyrus (MNI 57, -58, 34) and dorsal striatum (MNI 9, 8, 16) (see Figure 3-7). This dorsal striatal response is interesting in relation to findings that of dorsal striatum involvement in self-serving biases (Blackwood et al., 2003; Seidel et al., 2010).

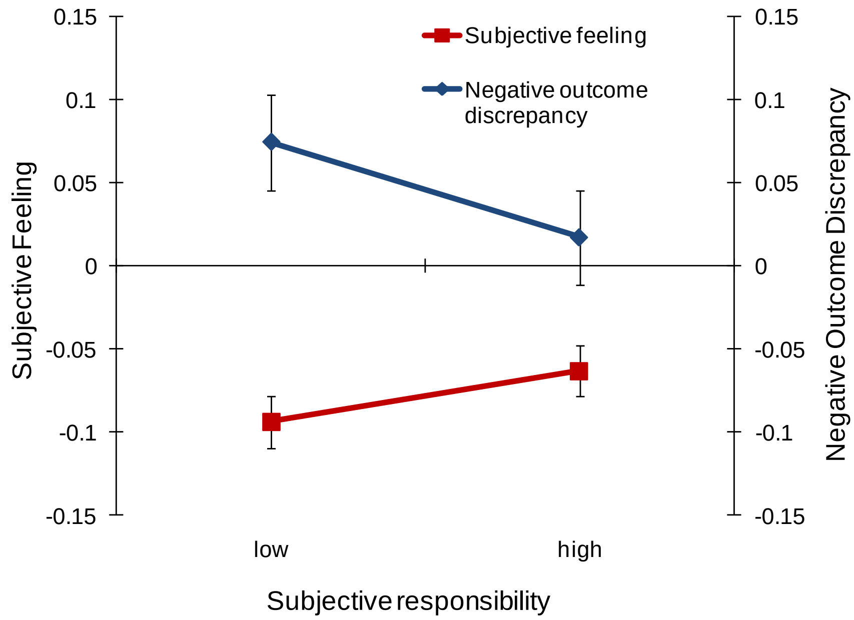

Figure 3-6- In red, the figure shows the mean normalised subjective ratings of negative feeling, for negatively discrepant outcome trials only, separated according to a median split of participants' subjective ratings of responsibility. This illustrates that participants felt more negative feeling for trials with *lower* ratings of subjective responsibility. In blue, the mean normalised actual outcome discrepancy for the same trials are presented, showing that more regretful outcomes are associated with lower subjective ratings of responsibility, as in Figure 3-5. Error bars show the standard error of the mean.

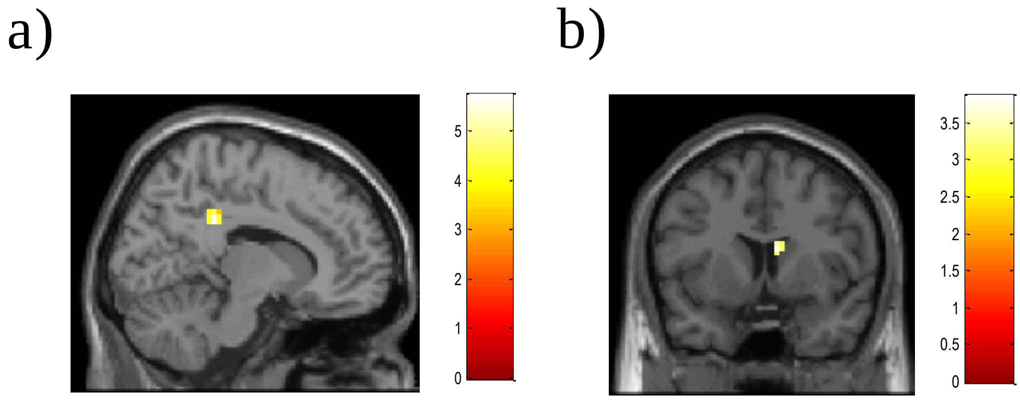

Figure 3-7- Group SPM data, thresholded at p < 0.005 for display purposes, shown on a normalised canonical template brain. Activity in posterior cingulate (peak 12, -37, 31) and caudate body (peak 9, 8, 16) associated with response to increasingly negatively discrepant outcomes, where this increase is greater with more other voters (i.e. decreasing objective responsibility). This contrast was performed on congruent trials only. Co-ordinates are in MNI space.

3.7 General discussion

In this study, subjective ratings of regret were enhanced by *increased* responsibility, with this self-blame regret recruiting activity in the amygdala as discussed in section 3.5 above. Paradoxically, additional analyses showed that individuals also actively *reduce* their feeling of responsibility for relatively worse outcomes, in keeping with the model of regret-regulation shown in Figure 3-4. This is analogous to a self-serving (self-protective) bias, and may reduce a conscious feeling of regret. When participants must make their gamble choice in the complete absence of any external force to which they could reasonably transfer blame (i.e. in the case of 0 other voters in this task), self-blame remains high, and amygdala activity reflects the personal relevance of this experience. On the other hand, as the number of other voters increases, the potential to transfer blame does so also. Under such conditions, the experience of regret and its associated amygdala response is diluted and cingulate, angular gyrus, lateral OFC and dorsal striatum are more likely to be involved.

The tendency for participants to report feeling less responsible when the outcome was relatively worse was associated with increased activity in the angular gyrus, a region previously implicated in external attributions of agency (Farrer et al., 2003; Farrer & Frith, 2002; Ruby & Decety, 2001). Similar inferior parietal/angular gyrus activation has been previously implicated in regret (Chandrasekhar et al., 2008; Coricelli et al., 2005). If activation in the angular gyrus is associated with *decreased* sense of responsibility, then previously reported activity here may actually reflect a regret-regulatory denial of responsibility for the aversive outcome, rather than the experience of regret itself.

Here, decreased responsibility was also associated with greater negative feeling, even though such outcomes elicit less regret. While in keeping with previous findings (e.g. Zeelenberg et al., 1998), these results are original here in that I show them as trial-by-trial recordings of actual emotional experience. This finding provides evidence that regret is dissociable from general negative affect, both phenomenologically and in its neural correlates. It is perhaps surprising that lateral OFC is involved in negatively discrepant outcomes *without* responsibility, since both medial and lateral OFC have previously been implicated in the experience of regret (Camille et al., 2004; Chandrasekhar, Capra, Moore, Noussair, & Berns, 2008; Chua, Gonzalez, Taylor, Welsh, & Liberzon, 2009; Coricelli et al., 2005; Liu et al., 2007). Medial-lateral differences in OFC involvement in this task may be driven by the associated feeling of responsibility. It is temping to speculate that lateral OFC response reflects anger, frustration or loss of control in participants, on the realisation that other voters produced a bad outcome for the group. The finding of strong *negative* affect associated with *low* responsibility in the task, along with reports that lateral OFC processes the unpleasantness of external stimuli (e.g. Anderson et al., 2003), supports this view.

The posterior cingulate, which shows the strongest response to bad outcomes for which the individual is *not* responsible, has been implicated in the selective allocation of attention, including cue-induced visuo-spatial bias (Small et al., 2003), a bias that may reflect a motivational shift in our attention towards the external world (Mesulam, Nobre, Kim, Parrish, & Gitelman, 2001). Frith, Friston, Liddle, & Frackowiak (1991) found decreases in posterior cingulate activity during the production of willed action relative to externally specified actions, supporting its role in externally focused attention. Tomlin et al. (2006) also found similar activity when observing a self-relevant decision of another player in a social exchange game, hinting that this external focus may be of a particularly *social* nature. This is also fitting with increased posterior cingulate activity during theory of mind tasks (Fletcher et al., 1995). While we find similar regions involved when bad outcomes were caused by externally enforced decisions, the same areas are not active for the main effect of these incongruent choices, suggesting this activity depends on evaluation of outcome valence and not simply an attention shift. Others have shown middle cingulate cortex activity involved in externally induced pain, either physical or emotional, along with induced anxiety (Eisenberger & Lieberman, 2004; Kimbrell et al., 1999; Liotti et al., 2000; Nielsen, Balslev, & Hansen, 2004; Ploghaus, Becerra, Borras, & Borsook, 2003). Increased activity in posterior cingulate has also been implicated in major depression (Ho et al., 1996), while healthy individuals also show a posterior cingulate response to threat-related words (Maddock & Buonocore, 1997). This role of the posterior or middle cingulate cortex in our response to externally induced threats, and the finding that it is also active when we are more helpless in avoiding negative outcomes, may suggest a particular importance for understanding feelings of helplessness in depression.

Surprisingly, these results suggest that self-blame regret may be a less aversive experience than the feelings of helplessness, frustration and anger associated with externally induced bad outcomes. One possible reason for this, as discussed above, is that regret-regulatory strategies (e.g. those depicted Figure 3-4) attenuate the experience and neuronal representation of self-blame regret. It is also possible that outcomes associated with self-blame are less aversive because they allow for learning and adaptive modification of future behaviour, while outcomes for which we were not responsible are not associated with this 'hope' for future change. To provide some further insight into the behavioural effects of self-blame associated outcomes, in the next chapter I question whether we actually learn better when we are actively involved in the decisions, e.g. through trial-and-error learning.

A shorter form of this study is submitted under the following reference: Nicolle, Symmonds & Dolan, (2010). Optimistic biases in observational learning of value.

Chapter 4. Does agency help or hinder learning?

4.1 Introduction

Many instances of everyday learning rely on trial-and-error. Here, we must sample between alternative actions and risk unfavourable outcomes in the early stages of learning, when action-outcome contingencies are unknown. Learning can also occur through observing the successes and failures of others, enabling us to acquire knowledge vicariously (i.e. without personal agency). The passivity involved in observational learning allows it to support "locally adaptive behaviours without incurring the costs associated with individual learning" (Boyd & Richerson, 1988, p. 30). This benefits of observational learning are ubiquitous in nature. For example, a hungry animal can avoid the energy costs incurred in active sampling of optimal feeding locations by observing actions and outcomes of conspecifics. A proliferation of customer review websites epitomises the utility of learning through the positive and negative experiences of others, so obviating our own need for expensive decisions.

Surprisingly, the efficacy of observational learning is rarely studied in the context of human value learning. Empirical evidence from animals attests to the fact that rewarded behaviour is promoted, and punished behaviour diminished, in passive observers (e.g. Bandura, 1971; Dawson & Foss, 1965; Heyes & Dawson, 1990; Mineka & Cook, 1988; Weigl & Hanson, 1980). For example, budgerigars show imitation of rewarded behaviours but a diminution of such behaviour if the observed consequences are not rewarded, suggesting that vicariously conditioned responses are goal-directed and not a mere mimicry of an observed action (Heyes & Saggerson, 2002; Heyes, 1994). However, despite these data, evidence for the effectiveness of observational learning is inconsistent. Church (1959), for example, found that rats observing lights predicting a shock to a model do not generalise these contingencies to their own risk preferences.

Several critical differences can be highlighted between vicarious and active value learning, which may lead to differences in information acquisition. One important factor is motivation, a key variable in Bandura's (1977) social learning theory, given that passive observers do not directly incur costs or benefits during learning. The previous study finds that responsibility is vital for experiences of regret, an emotion we are motivated to avoid. A preparative role of regret would suggest that superior learning may occur under conditions in which regret can be experienced, associated with the motivation to avoid future regret. Regret can also increase vigilance and carefulness, as measured by decision duration and amount of information search (Reb, 2008). By being unaffected directly by the outcomes of the actor's decisions, an observer will not anticipate possible regret, which may result in a decreased motivation to learn.

An alternative view is that emotions might impair learning by distracting us from, or "crowding out", our goals (Loewenstein, 1996). They may also bias our memory for the frequency of past events (cf. emotional biases of eyewitness testimonies, e.g. Loftus, 1996). Emotions, such as regret, may lead to a distorted (perhaps especially inflated) estimate of the likelihood of future bad outcomes. Indeed, excessive anticipation of regret has been linked to procrastination or decision avoidance (Janis & Mann, 1977). Consistent with this “dark side of emotion”, individuals with attenuated emotional responses to the outcomes of risky decisions sometimes show more advantageous decision-making (Shiv, Loewenstein, & Bechara, 2005). While regret would clearly be beneficial if it helped us learn from past mistakes, it may have unfavourable effects if it drives us towards suboptimal decisions, decision delay or avoidance, or excessive anxiety. As such, by not experiencing regret, observers might be expected to show relatively improved learning.

Operant and observational learning may also differ in how attention is allocated during learning. As discussed above, anticipated emotions may globally increase the attention focused on learning as well as the incentive to learn. Moreover, recognising a causal link between actions and consequences is vital for learning optimal behaviours, and a sense of agency may increase the salience of such information. An enhanced salience of counterfactual information, and attention to between-option comparisons, may also improve action-outcome contingency learning and is predicted to be greater when agents anticipate and experience regret. Observers may place less weight on counterfactual information, since regret/relief is not expected to occur alongside learning.

In the three experiments presented in this chapter, I make a controlled comparison between active and observational learning in the context of human probabilistic value learning. I implemented a task where individuals learnt either by active sampling (with associated reward and punishment) or by passive observation. I assessed learning efficacy as shown by goal-directed choices and individuals’ explicit estimates of value. All aspects of the tasks, save for the critical factor of self versus other choice, were matched across two modes of value learning. Specifically, differences in attention and information were controlled, as participants could track the same sequences of outcomes in both learning conditions. Moreover, between-option counterfactual comparisons could always be made and agency was the only regret-related factor manipulated. Motivation to learn was also controlled, since participants earned money according to learning performance in both conditions. It is important to note, of course, that this control did not preclude the possibility that actors and observers would have different internally generated incentives, and differences in attention, during learning.

4.2 Experiment 1

4.2.1 Participants

17 healthy participants took part in experiment 1. Participants failing to reach a criterion of 60% accuracy in either (i.e. both) session were removed from analysis, being classified as overall non-learners, i.e. considered as failing to engage sufficiently with the task. This was the case only for one participant, leaving 16 participants for the full analysis (9 female, mean age = 23.8 yrs, *SD* = 3.0).

4.2.2 Procedure

Participants completed two sessions on consecutive days. In the first (the 'actor session'), participants made choices between four abstract stimuli (letters from Agathodaimon font), presented in different pairs on each trial, while concurrently attempting to learn the probability of winning from each. Each stimulus was associated with a discrete and constant probability of winning (p{win}), and outcomes of each stimulus were drawn independently on every trial. Outcomes of both the chosen and unchosen stimuli were then shown sequentially, with yellow and red boxes indicating winning and losing outcomes, respectively. Critically, these outcomes directly influenced participant's earnings for the actor session (with £1 awarded for each chosen win from 10 randomly selected trials). Participants were instructed to choose the stimulus with the highest p{win} on every trial in order to maximise earnings.

On day 2 (the 'observer session'), participants learned the values of a novel set of stimuli. Participants were given an instruction that this time they would observe choices made previously by another participant, along with their associated outcomes. While they were not provided with any information about this other participant, they were informed that these were real choices made by a different individual in a prior session. Participants were informed that, although they could learn from the outcomes of observed choices, these outcomes would not influence their own earnings for the observer session. Unknown to them, participants actually observed the same sequence of choices they had made in their previous actor session, although now with visually novel stimuli (stimulus sets were balanced between sessions and across participants). The two sessions were, therefore, matched in terms of the information from which they learned. Observer sessions were completed on day 2 in order to reduce memory for previous choice sequence. To match for motor responses, observers indicated the observed choice on each trial with a button-press. Since learning could not be measured in these observation trials, because a free choice is not made, test trials were introduced to assess learning in both actors and observers. These comprised 9 blocks of trials (test blocks) at regular intervals throughout learning. Here, free choices were made by both actors and observers in the absence of outcome feedback (to prevent further learning). Each stimulus was presented 6 times in different pairings in each learning block, and 12 times in each test block (test pairings further explained in section 4.2.3 below).

Figure 4-1 illustrates exemplar learning and test trials and indicates the sole difference between actors and observers at the time of choice, and in that actors received the outcomes of their choices. At the end of each learning session, participants also provided explicit estimates of p{win} for each stimulus on each session. Here participants were shown each stimulus in turn and asked to explicitly write down their estimate of the probability of winning (as a percentage of trials) for the stimulus independent of its pairing. State anxiety scores were also collected, using the state component of the STAI (Spielberger, Gorsuch, & Edward, 1970), to test whether any difference in learning efficacy between the two learning sessions could be explained by a difference in state anxiety level.

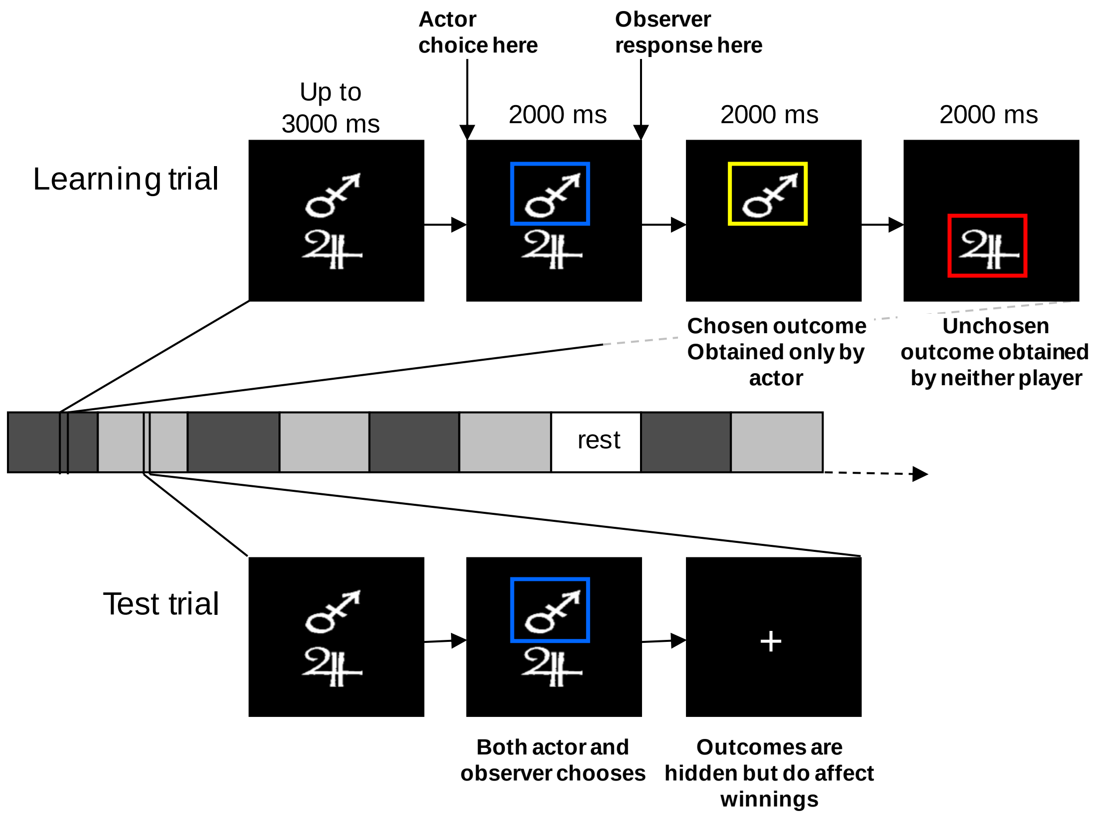

Figure 4-1 – Timeline for both actor session and observer session. Learning blocks (dark grey) and test blocks (light grey) alternate 9 times, with rests at 3 regular intervals. In learning trials, actors make a free choice between a stimulus pair, indicated by the blue box. Outcomes of the chosen and unchosen stimulus are then displayed sequentially, with a yellow box indicating a win, and red indicating no win. In observer sessions, learning trials differ only in participants' response. Here, participants wait until the blue box is shown, indicating the "other participant's" choice, and then press the button corresponding to the selected stimulus. Outcomes are presented as in the actor session. In test trials, free choices between stimulus pairs are made by both actors and observers, but outcomes are not displayed.

In the observer session, participants were paid based on the (hidden) outcomes of 10 choices from the test trials. In their actor session, earnings were based on the chosen outcomes of 5 test and 5 learning trials. This matches overall financial incentives. Full payment was given after the second session, but participants were informed that the earnings of each session were independent. Practice for both actor and observer sessions were given at the beginning of the first session.

4.2.3 Design and analysis

Analysis was restricted to test blocks where both actors and observers made measurable free choices. I used a 2 × 4 × 9 within-subject design with factors for learning session (Actor/Observer), gamble pair (80/20, 80/60, 60/40 and 40/20) and test block (1 to 9).

I measured choice accuracy for each pair, over the 9 test blocks, as the proportion of times that that participants chose the option with the highest p{win}. Probability of choosing each stimulus independent of pairings was also assessed. To eliminate differences in individual learning ability, I measured within-subject changes in choice accuracy between the two sessions. Analyses were two-tailed to test for both increases and decreases in learning against the null hypothesis of no significant change between the two learning sessions. I also tested for an effect of session on explicit estimates of p{win} for each stimulus and on state anxiety.

Reaction times (RTs) were analysed using a 2 × 2 × 9 ANOVA, with factors comprising learning sessions (Actor/Observer), size of probability discrepancy between the two gambles in the pair (80/20 versus 80/60, 60/40 and 40/20) and test block (1 to 9). I predicted an effect of probability discrepancy on RT, since 80/20 pairs were considered to allow for easier value discrimination than 80/60, 60/40 or 40/20 pairs.

4.2.4 Results

A repeated-measures ANOVA showed a significant effect of learning over the 9 test blocks ($F[8,120] = 7.72$, $p < 0.0001$), such that accuracy improved over time. This effect interacted significantly with gamble pair ($F[24,360] = 2.80$, $p <$

0.0001), with accuracy improving more steeply for 80/20 and 80/60 pair choice, than for the two remaining pairs.

I also found a main effect of the gamble pair on accuracy ($F[3,45] = 7.41$, $p < 0.001$), an effect that also interacted significantly with session ($F[3,45] = 3.76$, $p < 0.02$). Post-hoc paired t-tests showed this interaction was driven by a difference in actor and observer accuracy for the 40/20 pair alone, such that observers were significantly less accurate in such decisions ($t[15] = 3.0$, $p < 0.01$) (Figure 4-2). There was no interaction of session × gamble pair × test block, suggesting that observers' low choice accuracy for the 40/20 pair was not modulated by time.

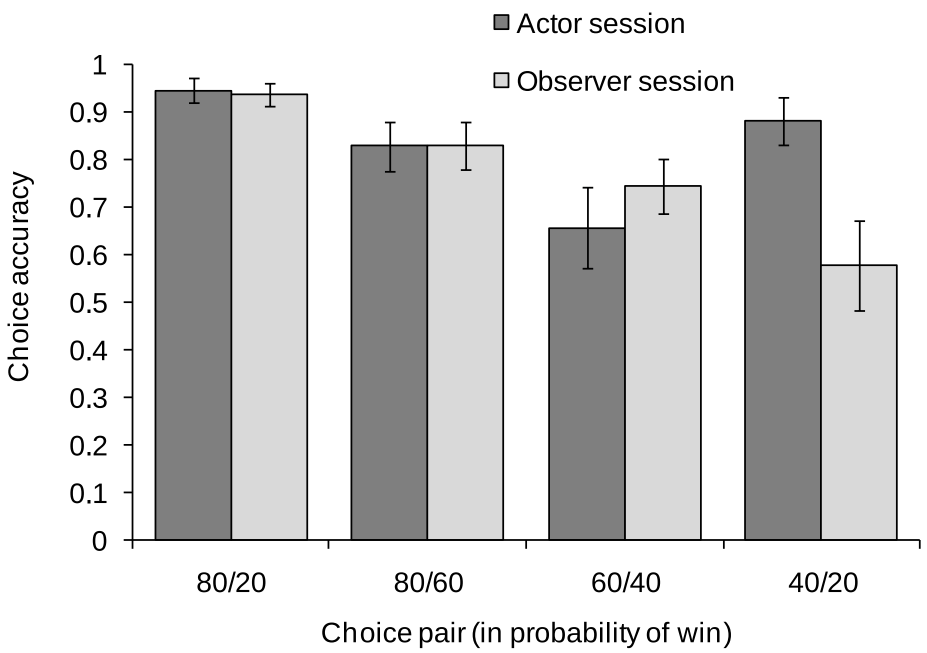

Figure 4-2 – Choice accuracy for test trial gamble pairs. Pairs are labelled according to the probability of a win for each stimulus. Choice accuracy is measured as the probability that participants chose the stimulus with the highest probability of a win. Actor and observer learning differed only for the 40/20 p{win} pair, with observers showing significantly lower accuracy compared to actors. Error bars show the standard error of the mean.

Participants' explicit estimates of stimulus p{win} also showed a specific impairment in learning in relation to lower p{win} options (Figure 4-3). A repeated-measures ANOVA showed a gamble × session interaction in estimates of p{win} ($F[3,45] = 7.29$, $p < 0.0005$), such that p{win} for the 20% win option was significantly overestimated through observation compared to action ($t(15) = 4.61$, $p < 0.005$). Observers' individual choice preference in 40/20 test choices was also strongly associated with the degree to which the 20% win gamble was overvalued when observing compared to acting ($R^2 = 0.29$, $p < 0.05$).

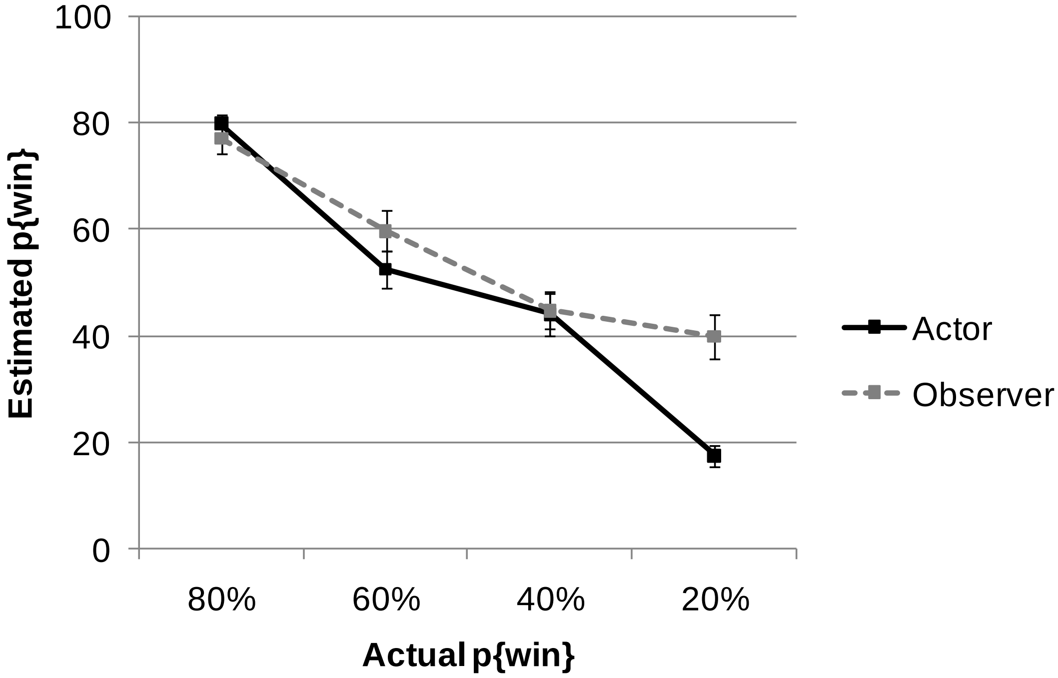

Figure 4-3 - Participants' estimated probability of a win (p{win}) for each stimulus, learned during the actor and observer sessions, plotted against the actual p{win} for each stimulus. Observers significantly overestimated the p{win} for the 20% win stimulus, compared to actors. Error bars show the standard error of the mean.

Test trial RT's were influenced by how much a gamble pair deviated in p{win}, such that participants were slower to choose between gambles with a close p{win} (80/60, 60/40, 40/20, mean = 1146 ms, SD = 54 ms) compared to distant p{win} pairs (80/20, mean = 960 ms, SD = 43 ms) ($F[1,15] = 125.81$, $p < 0.0001$).

There was also a main effect of test number, with participants becoming quicker with time (F[8,120] = 14.60, $p < 0.0001$). This effect interacted with stimulus pair such that this RT speeding was greatest for the 80/20 pair (F[8,120] = 4.31, $p < 0.0001$). There were no effects of session (mean actor = 1039 ms, SD actor = 51 ms, mean observer = 1067 ms, SD observer = 50 ms), showing that the difference found between observational and operant learning was not due to RT differences. Mean state anxiety scores for actor and observer sessions were 14.3 (*SD* = 8.11) and 13.4 (*SD* = 7.85) respectively, and did not differ significantly, indicating that the difference found between observational and operant learning was also not due to differences in state anxiety.

4.2.5 Discussion

The results from Experiment 1 show that, while value learning through trial-and-error is highly accurate, observational learning is associated with erroneous learning of low-value options. In essence, observational learners show a striking over-estimation of the likelihood of winning from the lower-value options, a fallacy leading to impaired accuracy when choosing between two low-value options. This learning difference was apparent even though monetary incentives and visual information were matched across actor and observer learning.

Two potential design weaknesses can be identified in Experiment 1. First, by yoking the sequence of actor choices to participants' subsequent observer session, to match actor and observer learning for information presented, I was not able to counterbalance session order. Since participants also learnt about novel stimuli in the second session, learning may be worse solely because the task has switched. To explicitly address these issues, I designed a second study (Experiment 2) to test for changes in learning between two actor sessions, with stimuli for each session taken from the equivalent sessions of Experiment 1. I predicted participants would show *improved* learning in the second actor session, due to generalisation of learning strategy, despite the novel stimuli.

Secondly, in Experiment 1 it is impossible to distinguish between over-valuation of low-value options versus overestimation of low probabilities. To address this, I conducted an additional experiment (Experiment 3) which reversed the framing of learning such that participants now learn in order to avoid losing, rather than to reap a reward. In so doing, options with the highest value were now associated with the lowest probability of losing, allowing for explicitly dissociation of probability and value.

4.3 Experiment 2

4.3.1 Participants

17 new participants took part in experiment 2. Again, one participant was excluded due to a failure to reach an a priori accuracy criterion. 16 participants remained (6 female, mean age = 31.2 yrs, *SD* = 10.6).

4.3.2 Procedure and analysis

Here participants performed two actor sessions on consecutive days, using the same procedure and stimuli as in Experiment 1. As in experiment 1, novel stimuli were used in the second session, and these were identical to those used in the observer session of Experiment 1 (visually *and* in associated p{win}). Explicit estimates of p{win} were also assessed after each session. While Experiment 2 used the same design as Experiment 1, critical analyses now involved the between-subject interactions in relation to findings from Experiment 1. I term Experiment 1's participants the AO group, and Experiment 2's participants the AA group.

4.3.3 Results

Choice accuracy was again measured as the probability that participants chose the stimulus with the highest p{win}. When analysing the data just within the AA group, I found a main effect of session ($F(1,15) = 6.40$, $p < 0.05$) such that accuracy was higher on the second session. There was also a main effect of gamble ($F(3,45) = 5.64$, $p < 0.005$), and of test block ($F(8,120) = 4.36$, $p < 0.001$), and an interaction between these two factors ($F(24, 360) = 1.591$, $p < 0.05$). Including a between-subject analysis against the AO participants of Experiment 1, I found a session × group interaction ($F(1,30) = 7.28$, $p < 0.02$), and a session × gamble × group interaction ($F(3,90)=3.68$, $p<0.02$), highlighting the specific impairment of observational learning for low value options shown in Experiment 1 (Figure 4-4).

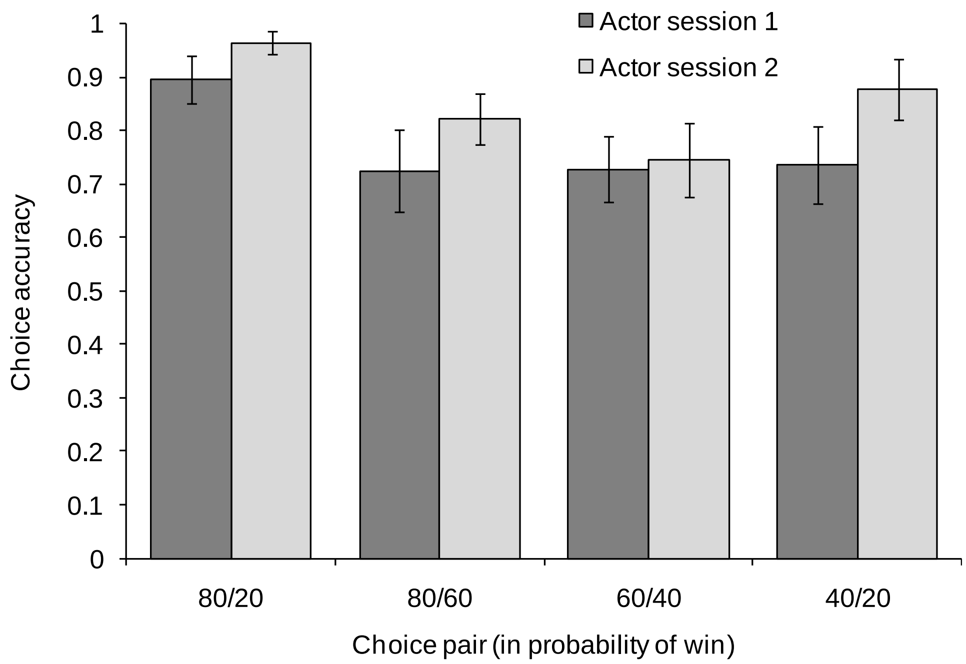

Figure 4-4 – Choice accuracy for test trial gamble pairs in Experiment 2's AA group. Pairs are labelled according to the probability of a win for each stimulus. Error bars show the standard error of the mean.

Estimated probability of a win (p{win}) for each stimulus, learned during the first and second actor sessions of Experiment 2's AA participants, is plotted in Figure 4-5 against the actual p{win} for each stimulus. There was a significant main effect of gamble ($F(3,45) = 67.87$, $p < 0.0001$) but the gamble × session interaction seen in Experiment 1 was no longer found. When comparing Experiments 1 and 2, there was a significant session × group interaction ($F(1,30) = 7.59$, $p < 0.01$) and a trend session × gamble × group interaction ($F(3,90) = 2.70$, $p = 0.051$).

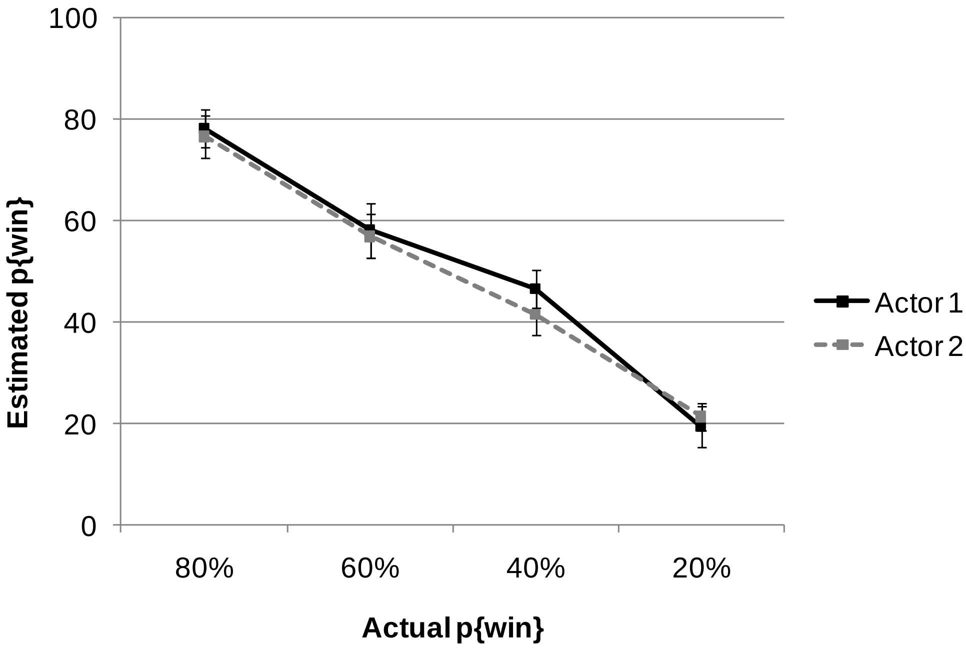

Figure 4-5 - Participants' estimated probability of a win (p{win}) for each stimulus, learned during the two actor sessions of Experiment 2's AA group, plotted against the actual p{win} for each stimulus. Error bars show the standard error of the mean.

4.3.4 Discussion

While a significant main effect of the gamble pair on accuracy is still apparent in Experiment 2, this effect no longer interacted with session. There was, however, a main effect of session, such that AA participants showed an improved accuracy from the first to the second session, and this effect interacted with participant group (i.e. AA versus AO). A session × gamble × group interaction again highlighted a specific impairment of observational learning for low value options. Explicit estimates of p{win} were also more accurate in both sessions of the AA group. The results of Experiment 2 indicate that impaired learning in the observer session of Experiment 1 cannot be attributed to a temporal order effect or to the learning of novel stimuli. The AA group actually showed improved learning in the second session, perhaps attributable to generalisation of learning strategies.

4.4 Experiment 3

4.4.1 Participants

16 new participants took part in Experiment 3 (7 female, mean age = 21.1 yrs, *SD* = 1.8).

4.4.2 Procedure and analysis

Experiment 3 utilised the same procedure and tasks (both actor and observer) as in Experiment 1, but with modified instructions and monetary incentives. Participants were now initially endowed with £10 per session. Instead of earning money from yellow boxes in the task, participants were informed that they would lose money from red boxes. In this way, the punishing power of the red boxes was assumed to attract more attention than in Experiment 1. At the end of the task, participants provided explicit estimates of the *probability of losing* (p{loss}) for each stimulus, in place of the p{win} estimates in Experiment 1.

Again, while Experiment 3 used the same design as Experiment 1, between-subject interactions with the findings from Experiment 1 were critical. I term Experiment 3’s participants the AO-loss group.

4.4.3 Results

When analysing the data just within the AO-loss group, I found main effects of session ($F(1,15) = 13.36$, $p < 0.005$), gamble ($F(3,45) = 13.98$, $p < 0.001$) and test block ($F(8,120) = 3.831$, $p < 0.001$). I also found an interaction of session × gamble

($F(3,45) = 12.15$, $p < 0.0001$), and gamble × test block ($F(24,360) = 4.47$, $p < 0.001$). The session × gamble effect was, as in Experiment 1, driven by observers showing significantly lower accuracy for the 40/20 p{win} pair compared to actors ($t(15) = 5.89$, $p < 0.0001$) (Figure 4-6). The effect of group, i.e. AO participants in Experiment 1 versus AO-loss participants in Experiment 3, interacted only with the main effects of session ($F(1,30) = 4.39$, $p < 0.05$) and of gamble ($F(3,90) = 3.36$, $p < 0.05$).

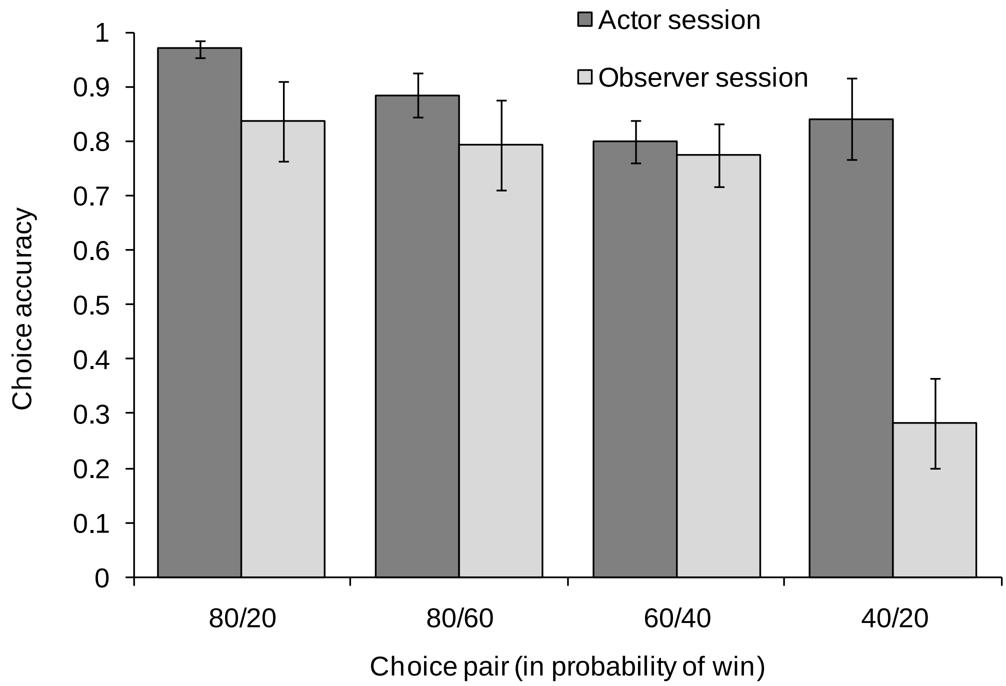

Figure 4-6 – Choice accuracy for test trial gamble pairs in Experiment 3's AO-loss group. Pairs are labelled according to the probability of a win for each stimulus. Error bars show the standard error of the mean.

Estimated probability of a loss (p{loss}) for each stimulus, learned during the actor and observer sessions of Experiment 3's AO-loss participants, is plotted in Figure 4-7 against the actual p{loss} for each stimulus. There was a significant main effect of session ($F(1,15) = 12.86$, $p < 0.005$) and of gamble ($F(3,45) = 75.85$, $p < 0.0001$), along with a gamble × session interaction ($F(3,45) = 8.87$, $p < 0.0005$). Analogous to Experiment 1, observers significantly underestimated the p{loss} for the 80% loss (20% win) stimulus, compared to actors.

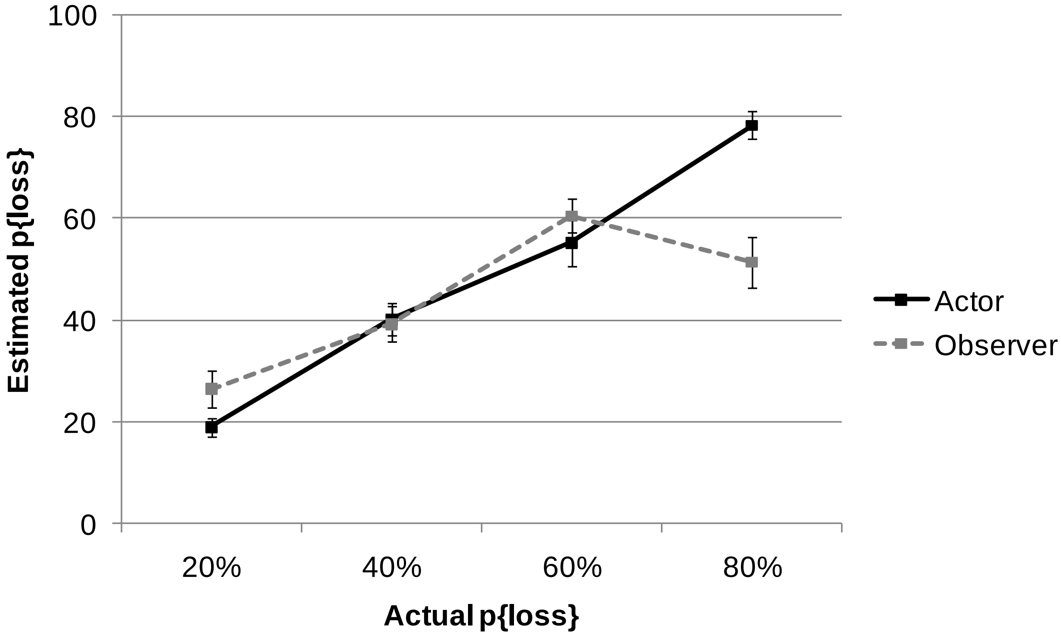

Figure 4-7 – Estimated probability of a loss (p{loss}) for each stimulus, learned during the actor and observer sessions of Experiment 3's AO-loss participants, plotted against the actual p{loss} for each stimulus. Error bars show the standard error of the mean.

With data from Experiments 1 and 3 collapsed, Figure 4-8 shows the frequency chosen of the 80%, 60%, 40% and 20% win options, in each of the 9 test blocks. Choice frequencies are shown for actor learning (filled lines) and observer learning (broken lines). Unconnected dots show the actual cumulative frequency of wins for each stimulus in the learning trials preceding each test block, which are identical for actors (diamonds) and observers (crosses). Actors can be seen to learn quickly the low value of choosing the 20% win option, avoiding this choice on the majority of trials, while observers show a higher tendency to choose this option. Actors also demonstrate an apparent overvaluation of the 40% option, likely owing to their accurate realisation that it is of a higher value than the 20% win option and resulting from an attempt to polarise the two values.

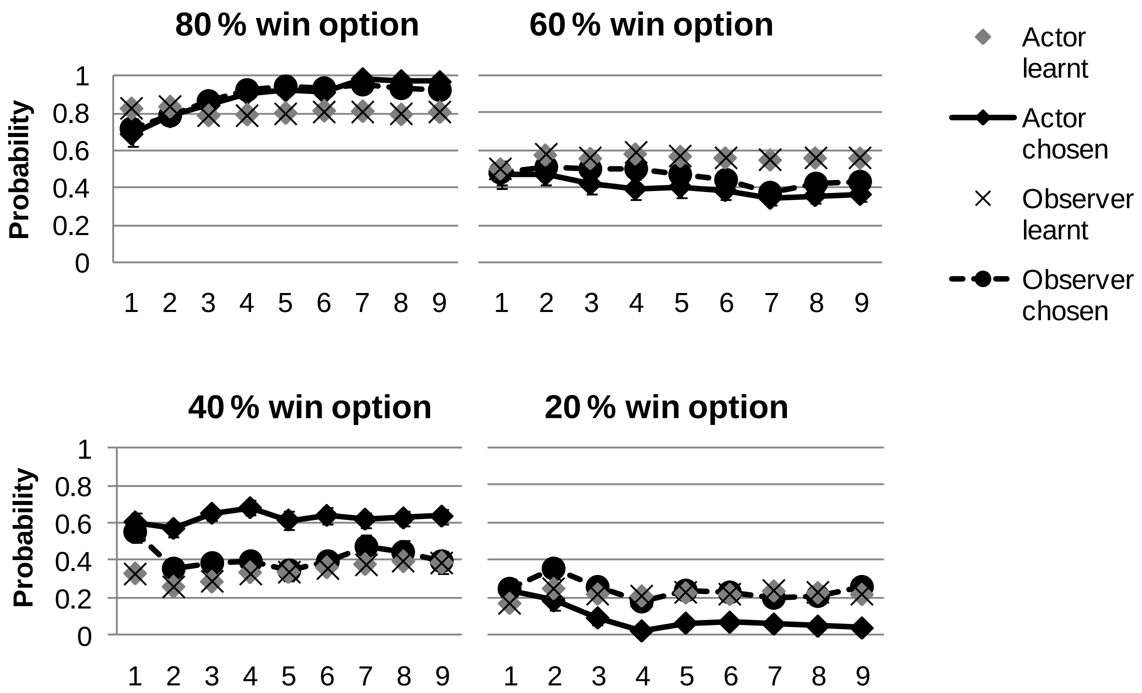

Figure 4-8 – Frequency chosen of the 80%, 60%, 40% and 20% win options, in each of the 9 test blocks across both Experiments 1 and 3. Choice frequencies are shown for actor learning (filled lines) and observer learning (broken lines). Unconnected dots show the actual cumulative frequency of wins for each stimulus in the learning trials preceding each test block, which are identical for actors (diamonds) and observers (crosses).

4.4.4 Discussion

The session × gamble interaction in choice accuracy, seen in Experiment 1, was replicated but now within the loss domain, with this effect being driven solely by observers' impaired accuracy for the lowest value 40/20 win pair (now 60/80 loss pair). Participants' explicit estimates of p{loss} for each stimulus also replicated the results of Experiment 1, supporting an over-valuing of the lowest value options (i.e. participants underestimated p{loss} for the 80% loss option) rather than an overestimation of small probabilities (participants showed high estimation accuracy for options with the lower p{loss}).

4.5 Understanding the difference at the mechanistic level

The aim of this study was to test if differences exist between operant and observational learning in the context of probabilistic value learning in healthy humans and, if so, where such differences lie. The wider goal is to understand the effect of agency (or a lack thereof) in learning. Given that I found significant differences in learning about low value options between operant and observational learning, this opens a question concerning what exactly brings about this difference at a mechanistic level. The findings of Experiment 2 indicate that the observational learning failure is not an effect of task order or task-switching. The results of the frame reversal in Experiment 3 indicate that the failure is specific to learning about low value options, not in estimating low probabilities. Moreover, the effect was not associated with differences in the amount or type of information from which participants could learn or to differences in monetary incentive, since these were matched across both learning conditions. Differences in state anxiety and in RT also cannot explain the effect. One way by which the underlying mechanisms of each form of learning may be understood is through a model-based approach.

An explanation for these findings might be that observers are less motivated by anticipated regret and do not take into account outcomes of unchosen options. In such a case, this fictive information will have less weight when updating the value of each option in a pair. Moreover, since low value options are chosen less often, they are less likely to be updated in this preferential update model. I tested this hypothesis, that observers take the outcomes of unchosen options less into account compared to actors[3], by fitting a reinforcement learning model separately to the actor and observer sessions of participants. In this model, two separate learning rates were fit for chosen and unchosen options. Although these were unconstrained in model estimation, my prediction was that actors would show no difference between their learning rates for chosen and unchosen options, while observers would show a markedly lower learning rate for unchosen options.

4.5.1 Computational model

The precise model was based on a commonly used reinforcement learning model, whereby the probability of winning for each outcome is updated trial-by-trial on the basis of outcomes experienced (or observed). These values are then converted into choice probabilities and these are then compared to the actual choices made by participants in the test blocks, in order to determine the fit of the model.

The model employs a Rescorla-Wagner update of value of stimulus i at time t (or V_t^i) (Equation 5), based on the difference between the expected reward and the reward experienced or observed for the particular stimulus on each trial (r_t^i), where $r = 0, 1$. This difference is termed the prediction error and is weighted by a free parameter, α , which determines participants' learning rate (or speed of update).

[3] Note that feedback was always given for outcomes of both chosen and unchosen options for both actors and observers. Therefore, any asymmetric effect of chosen options would indicate an attentional bias rather than an objective lack of information about the unchosen option.

The model, which I term a "choice-dependent learning" model, allows for an asymmetry between the learning rates of chosen and unchosen choice options.

Equation 5 $$V_t^i = V_{t-1}^i + \alpha(r_t^i - V_{t-1}^i)$$

Values were updated based on outcome feedback during the *learning blocks* of the task. Given the value of each stimulus at the end of each learning block, the probability of choosing stimulus i (p_i) in the following *test block* is then estimated. This is performed using the logit transform in Equation 6, where $\beta > 0$ and determines the randomness of the decision (with larger numbers indicating more random choice). This is a standard stochastic decision rule that calculates the probability of taking one of two actions according to their associated relative action values.

Equation 6 $$p_i = \frac{\exp(\beta V_i)}{\sum_{j=1,2} \exp(\beta V_j)}$$

The model was fit to choices made during the 9 test blocks of each session. I compared the fit of this choice-dependent learning model to two control models. In the first, the "choice-independent learning" model, both the chosen and unchosen stimuli of each pair are updated equivalently (with just one α). In a second control model, the learner preferentially updates only the chosen option on each trial ("chosen-only" model), while the value of the unchosen option is never updated (also with just one α). The models were fit to all 32 participants' choices from both experiment 1 (the AO group) and experiment 2 (the AO-loss group), since both groups were found to show the same behavioural effects. Each model was fit to test trial choices in actor and observer sessions of each participant separately. Note that the two control models contain only two free parameters, while the hypothesis-driven model contains three. Therefore, negative log likelihoods of each model was penalised for the complexity of the model (i.e. the number of free parameters) by calculating the Bayesian information criterion (BIC) (Schwarz, 1978).

4.5.2 Results

All models fit better to actors (choice-dependent summed *BIC* = 4692.31, choice-independent summed *BIC* = 4880.30, chosen-only summed *BIC* = 5501.45) than to observers (choice-dependent summed *BIC* = 6365.02, choice-independent summed *BIC* = 6741.90, chosen-only summed *BIC* = 6780.78). Note that smaller BIC scores indicate a better fit to the data. The relatively poor fit to observer choices overall suggests that no model fully captures mechanisms underlying observational learning in this task. This may be because observers' choices are generally more random, and so more difficult for the model to predict, an aspect apparent in the parameter comparisons outlined below. Overall, however, all models fit both actor and observer choices better than a completely random choice model (summed *BIC* = 9582.07).

The choice-dependent learning model provided the best fit to both actor and observer choices in this task. The choice accuracies predicted by this model, generated using the mean best fitting parameters, are plotted in Figure 4-9, separately for actor and observer test trial choices between each of the gamble pairs used. The model captures well the poor accuracy when observers choose between the 40% and 20% win options.

These predicted choice accuracies are generated through the model based on the probability of choosing the best option of each pair, as estimated by the softmax function. Therefore, accuracy depends on how easily the values of each option in a pair can be discriminated, along with the level of noise in the decision process (i.e. the temperature of the softmax). I found no significant difference between the temperatures of the softmax (i.e. β) of actor and observer learning (mean actor = 1.09 ± 0.14, mean observer = 1.19 ± 0.46, $t(31) = 1.13$, ns). Since actors and observers actually witnessed identical stimulus pairs and outcomes, any difference in the way options are valued must be ascribed to the learning rates. The best fitting learning rates for actor and observer's learning of chosen and unchosen options (averaged across participants) are shown in Figure 4-10. These showed a trend significant 2-way interaction ($F(1,31) = 3.54$, $p = 0.069$) such that, while actors and observers did not differ in their learning rate for chosen options ($t(31) = 0.29$, ns), observers showed a significantly lower learning rate for unchosen options compared to actors ($t(31) = 2.23$, $p < 0.05$).

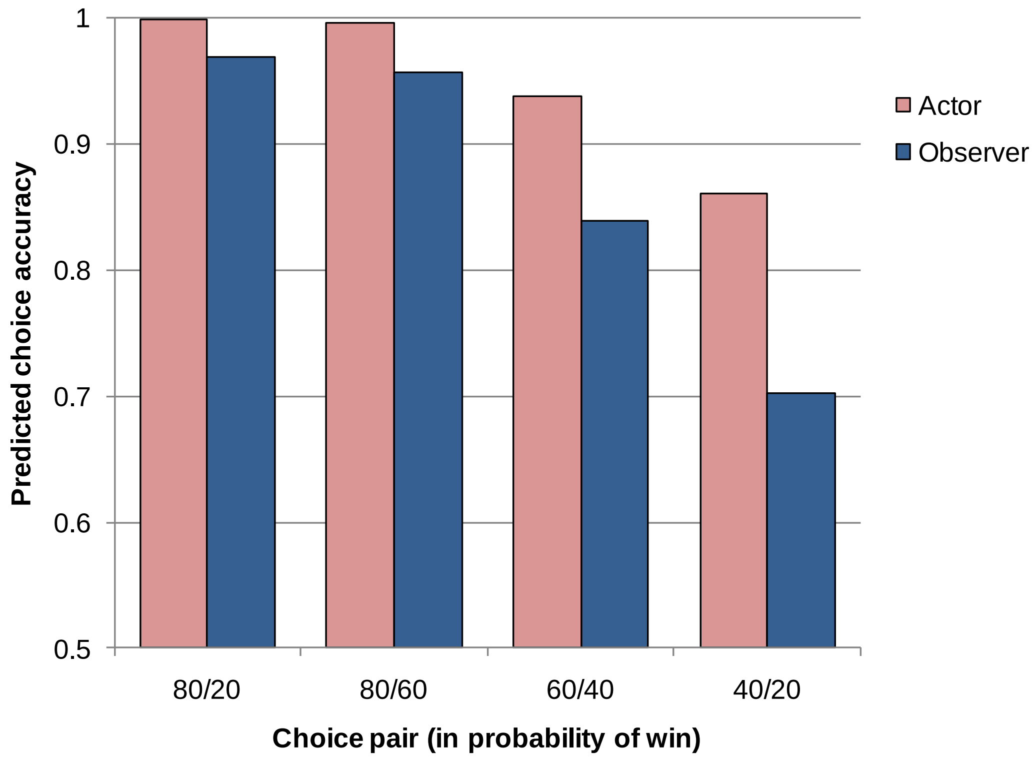

Figure 4-9 – Choice accuracy for test trial gamble pairs, as predicted by the softmax function of the choice-dependent learning model, using the mean best fitting learning rates and temperature of the softmax. Pairs are labelled according to the probability of a win for each stimulus. Choice accuracy is measured as the probability that the model chooses the stimulus with the highest probability of a win.

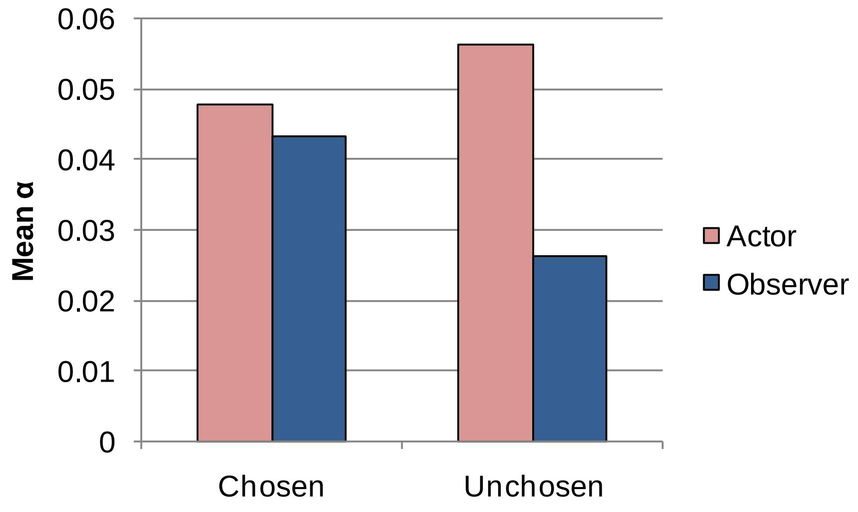

Figure 4-10 – Mean best fitting learning rates (optimised separately for chosen and unchosen stimulus options) as predicted by the choice-dependent learning model for actor and observer sessions.

Trial-by-trial stimulus values, predicted by the choice-dependent learning model were generated using the mean best fitting α and β across participants, and are shown in Figure 4-11. This figure illustrates especially close values of the 20% and 40% win options in observational learning, driven by the low learning rate associated with learning of unchosen options. The result of this is that the values of these two options are less easily discriminated by observers, leading to low accuracy when choosing between these options. The pattern of choice accuracy, value discriminability and learning rates for the two worse fitting control models (choice-independent model and chosen-only model) are available in Appendix B and do not show such effects.

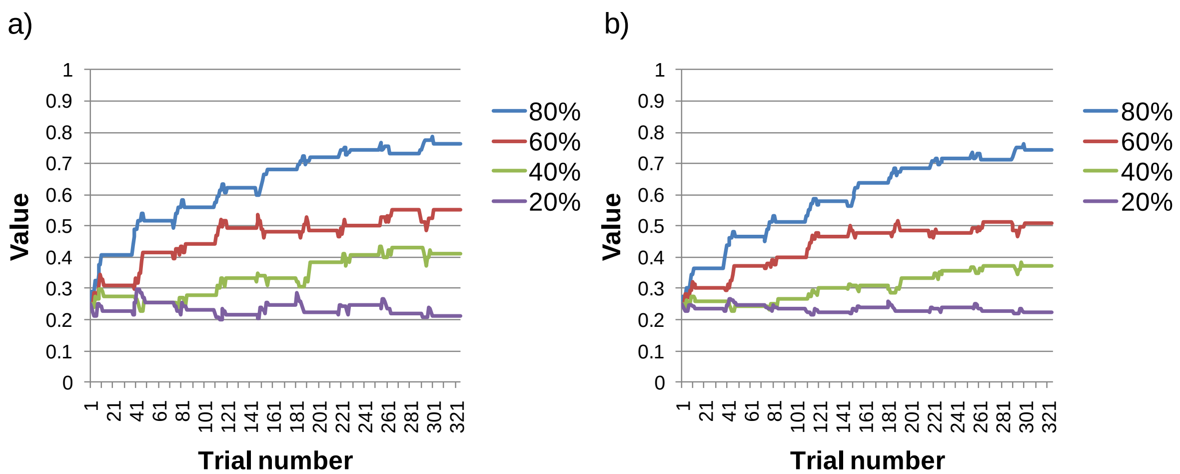

Figure 4-11- Stimulus values predicted for a) actor learning, and b) observer learning by the choice-dependent learning model, using the mean best fitting parameters across participants as the learning rates and the temperature of the softmax function. These values are shown for the four stimuli, each associated with a unique probability of winning (80%, 60%, 40% and 20%). Note that the flat portions of each line reflect the blocks of test trials. Here learning does not take place, as outcome feedback is not given.

4.5.3 Discussion

The choice-dependent learning model captures a reduced accuracy associated with observational learning of low value options, explaining it in terms of a low learning rate for unchosen, compared to chosen, options. These results are in keeping with a hypothesis that observers are less likely to use fictive (what-might-have-been) information in their learning of probabilistic value. This selective attention may be related to a lack of motivation that is usually engendered by possible regret. Specifically, fictive information is likely to be of less relevance for observers, for whom a lack of agency and direct impact of outcome precludes the possibility of experiencing regret.

This model does not provide a full explanation of the mechanism underlying observer learning, however, since it does not fit to observer choices as well as to actor choices. This is perhaps, in part, because observer choices are relatively noisier (as shown by a higher temperature of the softmax). Moreover, the values plotted in Figure 4-11b do not capture the particular over-valuation of the 20% win option as shown in Figure 4-3 and Figure 4-7. Future work will need to consider alternative models of observational learning in a value learning context, and to explicitly test models associated with optimistic overvaluation of low value options. Better models could also test the possibility observers do not update based on losses in the same way as do actors. Another possibility is that both actors and observers have an initially high prior for the value of options, but learning through trial-and-error corrects that prior for actors in a way that observational learning does not. Functional imaging studies might also be useful in understanding the mechanisms underlying the observational bias in learning. For example, we might predict differences in the way that reward (or fictive) prediction errors are coded in the brain during each form of learning, which may aid in explaining the inefficient learning through observation.

4.6 General discussion

Experiments 1 and 3 of this chapter both show an overvaluation of low value options during observational learning, an effect evident across implicit (i.e. choice preference) and explicit indices of subjective value. This difference was evident despite the observational and operant learning tasks being matched for visual information, as well as monetary incentive to learn. In contrast, Experiment 2 shows generally improved learning between two operant learning sessions despite the time delay and the novel stimuli. These results are intriguing since neither social learning theories nor reinforcement learning approaches explicitly predict that action-outcome contingency learning should depend upon the manner through which they are learned. Moreover, recent neuroimaging studies in humans report neuronal responses to errors (Koelewijn, van Schie, Bekkering, Oostenveld, & Jensen, 2008; van Schie, Mars, Coles, & Bekkering, 2004; Yu & Zhou, 2006) and success (Mobbs et al., 2009) observed from the behaviour of others, comparable to those seen in response to self-experienced outcomes, meaning one might predict little difference in learning. Experiment 3 also shows that the deficit in observational learning is specific to valuation of low value options, rather than an imprecision when estimating low probabilities, showing that observers are biased to (inappropriately) discount the chance they will experience negative outcomes seen to be incurred by others.

In contrast to a view that being responsible may create an “emotional clouding” effect on learning, these results suggest that agency produces no such bias during action-outcome learning. In fact, here I show that learning is *more* efficient in a context of agency, and that it is rather observational learning that exhibits a suboptimal bias. The improved learning through trial-and-error may reflect a greater vividness and self-relevance of direct experience (Helweg-Larsen, 1999; Stapel & Velthuijsen, 1996) or reflect improved recall of one’s own actions (Weinstein, 1987, see also Tversky & Kahneman's availability heuristic, 1974). Such an interpretation accords with findings that directly experienced information is given greater weight than observed information in guiding future behaviour in social games, even if both are equally informative and equally attended (Simonsohn, Karlsson, Loewenstein, & Ariely, 2008).

Actors and observers appear to differ in how attention is directed during learning. Specifically, the results of the model comparison suggest that an observer’s poor learning can be explained by an inattention to the outcomes of options not selected by the actor. This hints that fictive information, with associated agency, is important for optimal value learning. Moreover, this is in keeping with an enhanced salience of between-option counterfactual comparison and the experience of regret with agency. However, I acknowledge that this choice-dependent model does not fully capture an overestimation of low value options shown by observers. This particular finding reflects a behavioural manifestation of an optimistic bias, demonstrating a tendency to underweight the prospect of the most negative experience. A problem for interpretation here is that low value options were also those that were chosen less often during learning, and therefore value and choice are conflated. Future work will need to explicitly test learning models associated with optimistic overvaluation of low value options.

Could optimistic overvaluation in observers be connected with a lack of regret (or inattention to counterfactual information) during learning? Optimistic biases are pervasive in all forms of learning. During trial-and-error learning, however, a greater motivation to avoid bad outcomes, and to feel skilled as decision-makers (i.e. avoid regret), may drive actors to overcome this bias. Indeed, learning through direct experience is shown to lead to increased realism in estimating risk, thus reducing an optimistic bias (Burger & Palmer, 1992; Helweg-Larsen, 1999; Van der Velde, Van der Pligt, & Hooykaas, 1994; Weinstein, 1987, 1989). Passive observers, on the other hand, do not have this focus on their ability as skilled decision-makers, with their attention instead being directed to the quality of the choices made by an observed other. However, a recognised tendency for individuals to show an external attribution for failures and an internal attribution for successes (as discussed in Chapter 3) might interfere with accurate learning of action-outcome contingencies. Specifically, such an attribution bias distorts observational learning through a tendency to attribute an observed actor's failures to internal (i.e. dispositional) causes, encouraging an observer to believe they are less likely to fail or lose themselves. On the other hand, the actor's successes are perceived as externally determined, easily obtainable, and not due to any exceptional skill in the actor. Optimism often has such a socially comparative nature as when we tend to overestimate our own strengths and resources, while discounting those of others (Radcliffe & Klein, 2002). This bias is likely to be associated with the protection of self-esteem and avoidance of social anxiety (e.g. Hirsch & Mathews, 2000), coupled with a desire to be better than others (Weinstein, 1989), and can explain a tendency to feel that one is less likely to experience the negative events experienced by others (Weinstein & Lachendro, 1982).

A general conclusion, on the basis of these data, is that active involvement, along with direct experience of rewarding or punishing outcomes, leads to improved learning of the consequences of our actions, compared to learning through passive (and inconsequential) observation. The observer bias also indicates that actors and observers implement different weightings for positive and negative experiences as they sample outcomes. These findings also have implications for how we should apply learning theory to vicarious learning, as classical models assign no differences to these alternative models of learning.

This study was completed in collaboration with Ray Dolan, Dominik Bach and Wako Yoshida.

Chapter 5. The behavioural impact of experienced regret

5.1 Introduction

The previous study found that learning is more efficient when occurring through trial-and-error than through observation. In the former, causal links between our own actions and their outcomes can be made, and these can directly update value predictions important for guiding future decisions. Trial-and-error learning also allows for the potential to experience regret from mistakes, as well as relief from achievements. However, the experiment in Chapter 4 did not explicitly test behavioural responses to regret. While it is commonly assumed that our future decisions can be biased by counterfactual information, it is surprising that trial-by-trial influences of regrettable outcomes on immediately subsequent behaviour are not well understood. This was the inspiration for the next study of this thesis.

By incorporating regret into computations of expected utility, as the weighted difference between an outcome obtained from a performed action and the best possible outcome of a foregone action, economic models help to explain choices in the presence of anticipated regret (Bell, 1982; Loomes & Sugden, 1982; Savage, 1951). In such models, anticipated regret is computed based on outcome probabilities known by the decision-maker at the time of choice. It is unclear from this literature whether experiences of regret are assumed to have the same impact on behaviour as this anticipated regret. As discussed in Chapter 1, our goals and expectations consequent upon the actual *experience* of regret may not be equivalent to those when anticipating possible future regret.

A common assumption in theoretical discussions, and computational modelling, of regret is that it is a highly aversive emotion that motivates choices that avoid its future re-occurrence. Recent reinforcement learning approaches have been useful in understanding learning of regret avoidance. By calculating the level of post-decision regret on each trial, and incorporating this into the valuation of future actions, these models do well in predicting behaviour in various forms of decision task (e.g. Ert & Erev, 2007; Marchiori & Warglien, 2008). The assumption is that past experiences of regret enhance anticipated regret for the relevant actions during future decisions, which manifests as a bias away from repeating previously regretted choices. Such models are fitted to choices in tasks that require participants to learn, through experience, which choice options are associated with the lowest anticipated regret. Therefore, while they provide a first step in understanding the behavioural impact of experienced regret, they do not actually address this behavioural impact independent of differences in anticipated regret (i.e. since experienced regret is confounded with anticipated regret).

The aim of this third study was to explore the impact of regret-related outcomes on subsequent choice in repeated gambling decisions. I hypothesised that previous outcome would indeed influence subsequent choice, and that this influence would depend upon a comparison between the received and foregone outcomes, as well as whether or not the decision maker was responsible. This was, however, an exploratory study in that several behavioural effects were considered to be possible, as discussed below.

One possibility is that decision-makers would show *avoidance* of previously mistaken bet choices. This may be because experienced regret is used in updating the expected likelihood of experiencing regret from a choice in the future, and is in keeping with an indirect impact of experienced emotion described in Chapter 1 (Loewenstein & Lerner, 2003). However, there may be other direct behavioural impacts unique to experienced, over anticipated, regret. For example, Zeelenberg, Van Dijk, Manstead, & Van der Pligt (1998) showed that, compared to disappointment, regret encourages a desire to make up for a mistake and allow a second chance, although such compensatory responses need not necessarily involve typical avoidance. Connolly & Zeelenberg (2002) suggest that a common response to regret is to make more 'normal' decisions or 'safer' decisions, and thereby help an individual feel their decisions are more justified. Moreover, in a reanalysis of data gathered by Liu et al. (2007), Sommer, Peters, Gläscher, & Büchel (2009) found that decision-makers were actually *less* likely to switch decision after a regret-related, compared to a relief-related, outcome suggesting that regret may encourage choice repetition rather than avoidance. In a conceptually similar suggestion, Loftus & Loftus (1983) propose that 'chasing' behaviour (or repeated gambling) in response to near-misses provides gamblers with the opportunity to eliminate the impact of regret by potentially making up for their mistakes.

In this study, I used a simple gambling task in which one level of operational objective regret (and one of relief) could be induced on each trial. While outcomes (win or loss) were out of participants' control, their choice of monetary bet (which could either be a 50p or 10p bet) was within their control. Therefore, it was assumed that they would implicitly compare the outcome received from their bet choice with the one that might have been had they placed the alternative bet (conceptually similar to prior operationalisations of a fictive error, see Lohrenz, McCabe, Camerer, & Montague, 2007). Since my main question concerned behavioural responses to these regret-related outcomes, anticipated regret was held constant across the two gamble choices. This was achieved because either of the two possible bets (10p and 50p), made with the aim of winning the corresponding extra amount, could result in an outcome that could have been better from the alternative bet choice (in the form of 10p wins or 50p losses).

To address the motivational impact of these regret-related outcomes on subsequent choice, I compared choice propensities following each outcome type. Since regret also depends upon being responsible for the choice (as shown in Chapter 3), any behavioural impact of regret-related outcomes was assumed to depend upon agency. Therefore, a further feature of the design was that I varied agency over the bet selection. I was particularly interested in the tendency to *avoid* versus *repeat* previously regret-related choices. In this task, the 50p stake is considered the riskier choice, due to greater variance in possible outcome magnitude (Markowitz, 1952). This asymmetric risk of the bet options meant I was also able to address whether regret-related outcomes encourage either risk-seeking or risk-averse future behaviour.

Finally, I did not test subjective feelings of regret in this study, since the focus was on behaviour. Guided by the economic literature, and by the findings from Chapter 3 of this thesis, here I considered regret to depend upon the (upward) counterfactual context of a received outcome, as well as a sense of responsibility for the bad choice. In Chapter 3, I find that such outcomes do indeed elicit *feelings* of regret, while others also provide evidence in support of this (e.g. Camille et al., 2004; Coricelli et al., 2005; Inman, Dyer, & Jia, 1997; Mellers, Schwartz, Ho, & Ritov, 1997). This operationalisation of regret is also widely used in other accounts of regret and fictive error without recourse to subjective ratings (Bell, 1982; Ert & Erev, 2007; Hart & Mas-Colell, 2000; Lohrenz et al., 2007; Loomes & Sugden, 1982; Marchiori & Warglien, 2008), further validating this approach.

5.2 Methods

5.2.1 Participants and design

24 participants (15 female) took part in the experiment. Participants were aged 19-33yrs (mean 24.1 yrs).

Trials were ascribed to four categories, in a 2 × 2 repeated-measures design that was conditional on the outcome of the previous trial. The two factors were outcome (win/loss) and stake (high 50p/low 10p). In addition, by including a no-agency control condition (see below), I also addressed how any tendency to repeat the same choice after each outcome type interacted with agency (in a 2 × 2 × 2 design, now with the additional factor of agency versus no-agency).

5.2.2 Experimental procedure and analysis

Participants were given an initial endowment of £10 and subsequently performed a gambling task in which they placed high (50p) or low (10p) bets repeatedly on uncertain gambles. Participants were not informed of the relative probabilities of winning versus losing on each gamble, although these were in fact fixed at 50%. Participants were presented with a computer-simulated pack of red and blue cards turned face-down and were informed that the top card would be overturned on each trial. After placing their bet at the start of each trial, participants received a binary outcome of either a win (if the card was blue, as on 50% of trials) or a loss (if it was red, as on the remaining 50%). Depending on this outcome, participants either won or lost their selected bet stake. Therefore, four possible outcomes were each associated with an alternative outcome that would have been realised had participants placed the alternative bet stake. Outcomes could be 'regret-related', where the received outcome would have been better had the other bet been placed (as in a 50p loss or 10p win), or 'relief-related', where the outcome would have been worse had the other bet been placed (as in a 50p win and 10p loss). Cumulative winnings were not shown, to minimise possible 'wealth' effects (cf. Kahneman & Tversky, 1979).

The probability of a win was fixed at chance with a win being completely independent of the bet stake placed by participants. Therefore, the expected value (EV) of each stake option was, on average, zero pence thus meaning that there was no financial incentive for participants to favour either stake (other than any individual risk preferences which - as is standard - were assumed to be sufficiently stable over time and should not be influenced by preceding-trial outcome in the manner assessed here). A 50p stake can be considered the riskier choice, with risk defined as the variance in reward magnitude around the mean EV of zero pence (Markowitz, 1952). With only two choice options available, participants were always aware of the outcome that would have ensued had they placed the alternative stake. I expected that such counterfactual information would bias future choice (e.g. Mellers, Schwartz, & Ritov, 1999).

Importantly, on 3/4 of trials participants were responsible for choosing the stake, with these trials providing the ‘agency’ conditions (Figure 5-1-a). Trials in which participants were not responsible for the bet placed (this being selected by the computer instead and then executed by the participant) constituted a ‘no-agency’ condition that served to control for valence of outcome (Figure 5-1-b). This allowed for behavioural effects of agency-specific regret (i.e. outcomes worse than what would have arisen from the alternative bet choice) to be disambiguated from mere aversive outcome effects (i.e. such an outcome regardless of agency). The no-agency trials comprised the remaining 1/4 of trials on which participants were instructed to place the bet chosen by the computer.

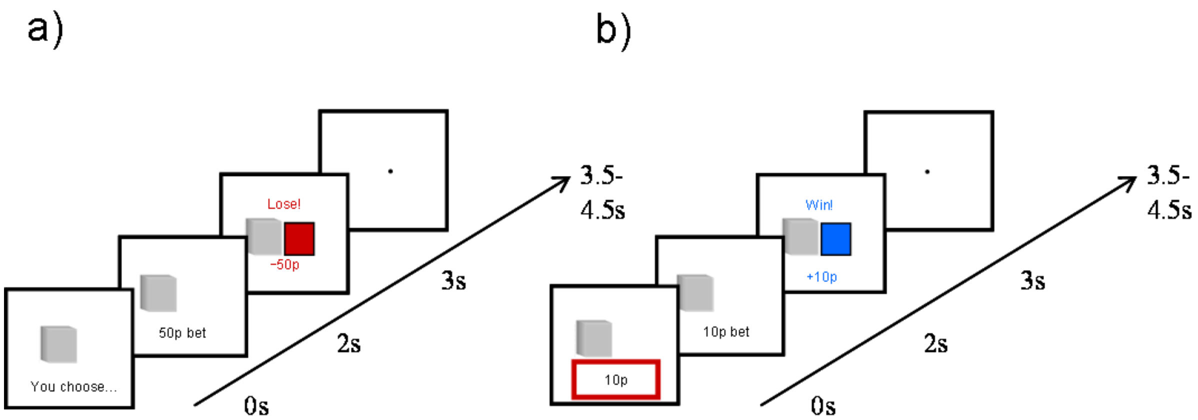

Figure 5-1 – Timeline of a) an agency trial, and b) a no agency control trial. Participants see a simulated pack of red and blue cards turned face-down and are required either to make a free-choice of betting either 50p or 10p (in an agency trial) or to place the bet selected by the computer (in a no-agency trial). The outcome of the bet is then revealed followed by an inter-trial interval.

To assess the immediate behavioural effects of experienced regret-related and relief-related outcomes, the dependent measure comprised participants’ trial-by-trial tendency to bets 50p at trial t ($p(50\text{p})_t$) contingent upon the outcome and agency of trial t-1. In particular, I tested how choices following the two regret-related outcomes (50p losses and 10p wins) differed from those following relief-related outcomes (50p wins and 10p losses), and also how any effects depended on whether or not outcomes were associated with agency. Since participants could only make a free choice on agency trials, only these trials were used to calculate the dependent measure (at trial t), and the number of these were matched for the number of preceding agency and no agency trials (120 trials each). Actual wins and losses were randomised throughout, while the overall number of trials falling into the two stake levels was choice-dependent.

Participants each played four 8-minute runs of the game. Each game included blocks in which only agency trials were played (such that dependent measure trials were always preceded by an agency trial at t-1) and blocks in which agency and no-agency trials alternated (such that dependent measure trials were always preceded by a no-agency trial at t-1). Participants were informed that the outcomes from a random selection of 100 trials, selected after the experiment, would determine their earnings for the entire experiment. Since participants did not know which trials would be selected, they were assumed to treat all trials as having an equal potential impact on their financial gain.

5.3 Results

The dependent measure was participants' trial-by-trial tendency to bet 50p at trial t contingent upon the outcome and agency of trial t-1 (i.e. $p(50\text{p})_t$). Overall participants did not show a bias towards choosing either bet (mean $p(50\text{p})_t$ = 0.46, SE = 0.03), in that they did not significantly deviate from 0.50 ($t(23)$ = -1.59, p =

0.13). There was, however, a trend level significant main effect of agency at t-1 on $p(50\text{p})_t$ (F[1,23] = 3.98, p = 0.058), such that participants preferred to opt for the safer 10p bets when they had been the agent on the preceding choice (mean $p(50\text{p})_t$ = 0.44, SE = 0.03), compared to when they had not been the agent (mean $p(50\text{p})_t$ = 0.48, SE = 0.03).

With this overall bias towards betting 10p more than 50p at trial t, it is particularly interesting that one outcome type encouraged the opposite behaviour. This exception was after 50p losses with agency (see Figure 5-2), with 19 out of 24 participants showing a higher $p(50\text{p})_t$ after the regret-related 50p loss than after the relief-related 50p win.

A Stake × Outcome interaction was found in the effect of agency outcomes, but not in no agency outcomes (stake × outcome × agency interaction, F[1,23] = 6.70, p < 0.02). This effect was such the two loss-based outcomes with agency motivated opposite bet selections on the subsequent trial, while wins did not show such difference. Relief-related 10p losses (where a 50p bet would have incurred a worse outcome) encouraged a repetition of the 10p bet, which would be in accord with a traditional regret-minimax theory whereby relief would increase the value of repeating the bet again. However, regret-related 50p losses (where a 10p bet would have incurred a better outcome) encouraged the 50p bet to be repeated subsequently, which would *not* be in keeping with models in which regret would decrease the value of repeating the same bet again.

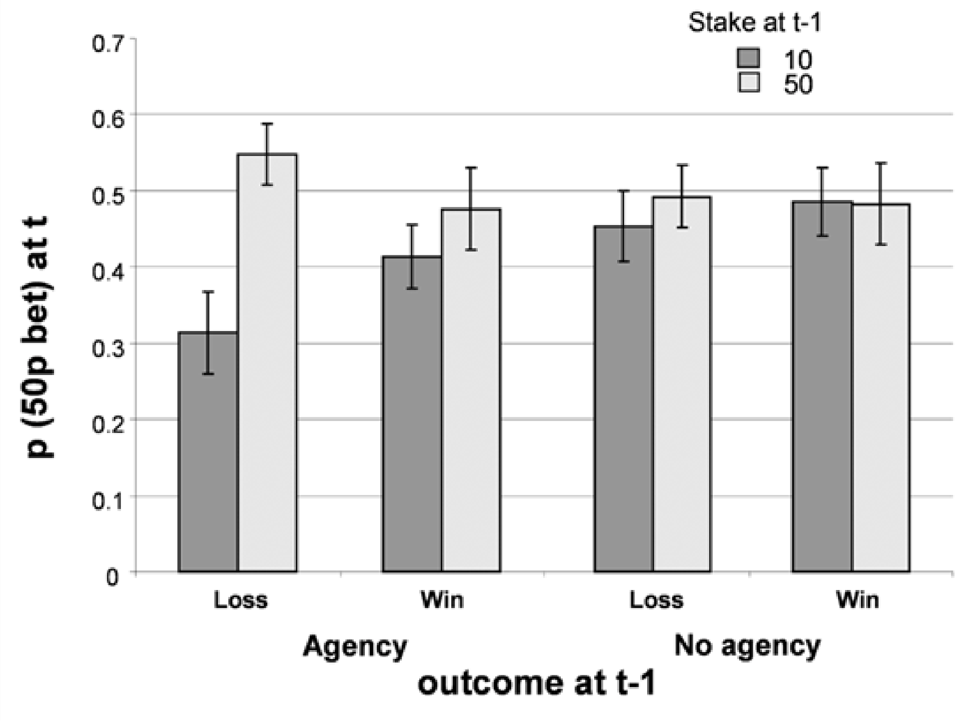

Figure 5-2 – Probability of choosing the 50p stake at trial t, after a win or loss, of the 50p or 10p stake at trial t-1. Error bars show the standard error of the mean.

One possible explanation for this effect might invoke the phenomenon of the gambler's fallacy, whereby the probability of a win is (mis)perceived to increase after a loss and decrease after a win (e.g. see Croson & Sundali, 2005). However, such an explanation would predict that both 10p and 50p losses would equally encourage the higher risk 50p bet on the next trial, since any loss should encourage the same misperception of probability. There was no main effect of outcome at t-1 ($F[1,23] = 0.07$, p = ns), and no interaction between agency and outcome ($F[1,23] = 0.01$, p = ns), suggesting that losses overall did not impact on choices in the next trial. Furthermore, gambler's fallacy alone would not predict the difference observed between agency and no agency conditions. Thus the behavioural pattern of results was not explicable solely by gambler's fallacy.

It is also important to note that the effects might be explained by decision inertia (i.e. repeat the same bet) after *both* the relief-inducing 10p loss *and* regret-inducing 50p loss, and may not be associated with regret at all. This type of behaviour could be driven by feelings of disappointment with an absolute loss as indicated by Zeelenberg et al. (1998), who reported that disappointment, in contrast to regret, is associated with a desire to avoid the situation and to do nothing. However, this line of argument would not explain the crucial interaction with agency, a core component of regret. Therefore, it is difficult to provide a clear account of why this form of decision inertia would occur here. Moreover, in the fMRI replication of the task presented in next chapter, we found no evidence for the bias to repeat 10p after a 10p loss, while the tendency to repeat 50p more after a 50p loss than after a 50p win was still evident, suggesting that the effect of 10p losses is less consistent.

5.4 Testing a 'regret-minimax' learning model

The observed tendency to repeat previously regret-related bets is contrary to expectations from those traditional learning models that update the value of a choice as a function of either standard or regret-based (or fictive) prediction-errors (e.g. Lohrenz, McCabe, Camerer, & Montague, 2007; Marchiori & Warglien, 2008; Rescorla & Wagner, 1972). On such accounts the value of a choice option should be incrementally *decreased* for negative or regrettable outcomes. Such learning models are therefore expected to provide a poor fit to the choice behaviour I observe.

I tested this prediction with a form of regret-minimax learning model, which maintains that decision-makers will avoid choices with the greatest anticipated regret. Here, similar to the neural network built by Marchiori & Warglien (2008), past *experiences* of regret act to reduce the *anticipated* value of repeating the previously regret-related bet choice. Therefore, this model assumes that experienced regret translates into anticipated regret in a linear and monotonic way, and that this increased anticipated regret then reduces the value of repeating the previously regrettable choice.

5.4.1 *The model*

At the outcome of each trial, the model employs a Rescorla-Wagner update of anticipated probability of receiving a win ($p(win)_t$) (Equation 7), based on the reward obtained on that trial ($\tilde{r}_t$), where $\tilde{r} = 0, 1$. This update is weighted by a free parameter, α, which determines participants' learning rate (or speed of update).

Equation 7

$$p(win)_t = p(win)_{t-1} + \alpha(\tilde{r}_t - p(win)_{t-1})$$

Next, the value of the choice made on the previous trial ($Q_{t+1}(i)$) is calculated as a function of the expected values of winning and losing from that action ($A(i)$) based on the outcome just obtained, as in Equation 8.

Equation 8

$$Q_{t+1}(i) = A(i)\,p(win)_t - A(i)(1 - p(win)_t)$$

The core of this regret-minimax model is that $\tilde{Q}_{t+1}(i)$ is reduced when $A(i)$ results in a regrettable outcome. This negative reinforcement is implemented by subtracting R_t from $Q_{t+1}(i)$, where R_t is the difference between the obtained and foregone outcomes of trial t (Equation 9). This term in the model is reminiscent of the modified expected utility introduced in the Chapter 1 (sections 1.3.1 and 1.3.2). Note that I use a new variable r_t for the reward obtained, where $r = -1, 1$. R_t is positive if regret is experienced, such that $\tilde{Q}_{t+1}(i)$ is reduced in the case of the two regret conditions, i.e. when there is a 10p win or a 50p loss. R_t is negative if relief is experienced, such that the value of repeating $A(i)$ on the next trial (i.e. $\tilde{Q}_{t+1}(i)$) is increased in the case of the two relief conditions, i.e. when there is a 50p win or a 10p loss. By nature of these outcomes being relief-related, the regret-minimax model predicts that they would not generate anticipated regret, thus encouraging greater desire to repeat the bet again. So where action i is taken and action j is not taken,

Equation 9 $$R_t = r_t \ (A(j) - A(i))$$

The regret-minimax model also allows for the possibility of an asymmetry between the regret of winning 10p (where a 50p bet would have been better) and the regret of losing 50p (where a 10p bet would have been better). Consequently, R_t is weighted by the parameter η when the regret follows a 10p win, and by λ when it follows a 50p loss. It was predicted that $\lambda > \eta$, but these parameters were not constrained in model estimation. So where $A(i=1)=10$ and $A(i=2)=50$, and where $c_1=\eta$ and $c_2=\lambda$,

Equation 10 $$\tilde{Q}_{t+1}(i)=Q_{t+1}(i)-c_i R_t$$

Given the value associated with making each choice option, the associated probability of making each choice is estimated through the logit transform in Equation 11, where $\beta>0$ (with larger numbers indicating more random choice). This is a standard stochastic decision rule that calculates the probability of taking one of two actions according to their relative action values.

Equation 11 $$p(i)=\frac{\exp(\beta\tilde{Q}(i))}{\sum_{j=1,2}\exp(\beta\tilde{Q}(j))}$$

5.4.2 *Fitting the model*

I fit the four free parameters (α, β, η and λ) to individual participants' choice data and adjusted these to maximise the likelihood of the choices given the model. One participant was removed from this model fitting process, leaving a total of 23 participants. This exclusion was due to the subject choosing the high-risk 50p bet less than 10% of the time, making parameter estimation problematic.

I compared this model to a default expected value (EV) model, where choices were dependent solely on the action values without any impact of R_t (this model had only the α and β free parameters, and where $Q_{t+1}(i)$ is calculated as in Equation 8). The models were only applied to agency trials to constrain their complexity and to improve parameter estimation. Choices immediately following no-agency trials were modelled as random.

5.4.3 *Results*

This regret-minimax model did not perform any better than the control EV model (summed BIC [regret minimax] = 5658.98, summed BIC [EV] = 5578.65) or a random choice default model (summed BIC [random] = 5662.37). Note that smaller BIC scores indicate a better fit to the data. The mean best fitting learning rate (α) across participants was 0.61. The temperature of the softmax (β) was 0.22. The mean regret parameters (η for 10p wins, and λ for 50p losses) were heavily affected by outliers, such that mean η = 1476.14 and mean λ = 39.08. However, the hypothesis that $\lambda > \eta$ did not hold (median λ = 0.004, median η = 0.133). These parameter fits hint at the observed behavioural response to the regret-related 50p loss condition in particular being opposite to what is predicted by the regret-minimax model, while the response to 10p wins deviates less from what is predicted by the model. This interpretation must be made with caution, however, since the model could not be fit well to actual choices.

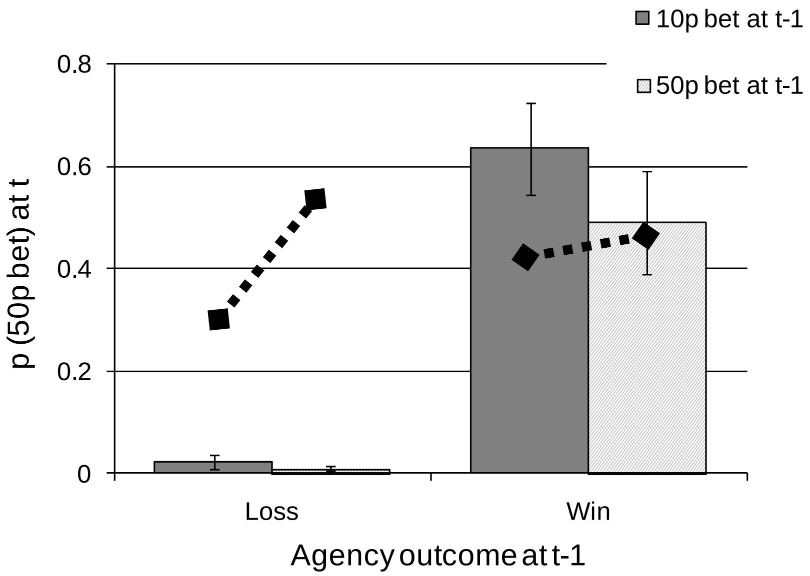

Figure 5-3 – Bars illustrate the choice propensities predicted by the regret-minimax model, showing the probability of choosing the 50p stake at trial t, after a win or loss, of the 50p or 10p stake at trial t-1. Broken lines indicate the choice propensities observed in participants' actual choices (i.e. see Figure 5-2). Error bars show the standard error of the mean.

Figure 5-3 shows the choice propensities predicted by the models, using the optimal fitting parameters for each participant. The observed behavioural response to 10p losses with agency is in general accord with the regret-minimax model, in that these relief-related outcomes tend to encourage a repetition of the relatively good 10p bet choice on the subsequent trial (although actual choices are more random than the model predicts). However, the impact of a 50p loss is most intriguing; as such outcomes appear to have the opposite effect on behaviour to that predicted by standard models of regret-minimax. The model also predicts that 10p wins encourage a slightly greater tendency to bet 50p on the next trial compared to 50p wins, due to a decreased value of betting 10p after a regrettable 10p win. Again the observed data points to an opposite tendency, although not significantly. Here the data shows a slightly decreased tendency towards the 50p bet compared to 50p wins. Both regret-related outcomes, then, appear to encourage participants to *repeat*, rather than *avoid*, the regrettable choice on the subsequent trial.

5.5 Discussion

Regret is widely regarded to exert a unique biasing influence on future behaviour (see Chapter 1, section 1.3). A general assumption of past research has been that regret motivates decision-makers to avoid actions that might engender its future occurrence. Here, however, I provide evidence that outcomes fitting the main criteria necessary for inducing regret may not always promote this typical avoidant behaviour on a subsequent choice. Rather, the findings point to a tendency to repeat the previously regret-related choice.

Previous studies of regret have assessed choices in the presence of anticipated risk of regret (e.g. Mellers et al., 1999). Since the task I used here explicitly controlled the probability of regret from each bet option, it is intriguing that the observed response to experienced regret is different to what is expected from traditional theories. This suggests that critical differences exist in the impacts of *anticipated* and *experienced* regret on future choice. Loewenstein & Lerner (2003) also make this distinction between anticipated and experienced emotions, and argue that the latter can impact behaviour in two ways. They can either indirectly affect behaviour through modifying expected consequences (and/or anticipated emotions), or they can influence behaviour directly. A direct impact may be to promote behaviour that regulates the current emotional experience rather than to prepare for future avoidance of such experiences.

One possible explanation for the regret-related choice repetition relates to an emerging view that regret invokes cognitive regulatory strategies, such that feeling better about a mistake may trade off against a more pre-potent desire to improve future behaviour (Roese & Olson, 2007; Zeelenberg & Pieters, 2007). For example, on narrowly missing a bus, we can reduce feelings of regret with self-justifying thoughts such as “I couldn’t have run any faster” or “I couldn’t have known to leave home earlier”. When an individual believes it is possible to make up for a previous mistake, future decisions may reflect this belief, either aiming for material compensation or for protection of self-esteem with a later successful decision. Repeating a previously regret-related choice, as I observer here, might provide such a way for the decision-maker to make up for the prior mistake. This view resembles a direct impact of experienced emotions, as described above.

This possibility of compensatory choice repetition after regret invokes a possible link between regret-regulatory strategies on the one hand, and the well recognised role of ‘chasing’ in problem gambling (Lesieur, 1977), on the other. In the latter, there is a continuation of gambling after a series of losses. It has also been suggested that the phenomenon of a “near miss”, which depends on comparison of an obtained outcome with a close better counterfactual outcome, may encourage chasing in the context of gambling (Reid, 1986). Moreover, a chasing strategy may provide gamblers with a potential means of reducing feelings of regret (Loftus & Loftus, 1983). This possible link with compulsive gambling with be discussed further in Chapters 6 and 8.

Rather than an active attempt to make up for the mistake, repetition of a previously regret-related outcome may also be associated with decision inertia. Notably, mistakes arising from decisions to repeat a previous course of action are sometimes associated with relatively less regret than mistakes arising from a new course of action, even if the two outcomes are objectively equally aversive. Gilovich, Medvec, & Chen (1995) found that such inflated emotional response to mistaken decisions to switch actions was also associated with stronger attempts at dissonance reduction. Furthermore, such an asymmetry may then encourage a bias towards repeating a previous action in future decisions as a form of regret-minimising strategy (e.g. Inman & Zeelenberg, 2002; Kahneman & Tversky, 1982; Ritov & Baron, 1990). Conceptually similar are findings that positive experiences can promote variety or novelty seeking (Kahn & Isen, 1993), while negative experiences may lead to a fear of novelty and may make an individual more likely to stick to familiar choices. The possibility that this form of decision inertia is promoted by experiences of regret will be followed up explicitly in Chapter 7 of this thesis.

Ritov (1996) observes that people prefer high-risk/high-gain gambles, over low-risk/low-gain gambles, but only when they expect to receive feedback of the unselected alternative. Therefore, an alternative explanation for the current findings could be that a tendency to place the 50p bet more after 50p losses with agency than after any other outcome may reflect a risk-seeking bias. Risk-seeking may arise from the anticipation of missing a large win in the future. Indeed, incorporating a regret function into a model of investment decisions predicts that a regret-averse decision-maker may become risk-seeking for fear of missing out on a large gain (Michenaud & Solnik, 2008). This is also in general accord with a prospect theoretic account of increased risk-seeking in a frame of losses, where losses here are relative to a better alternative. Similar predictions are made in a reinforcement learning context, when sequences of negative experiences tend to encourage increased risk-seeking in the short-term but more risk-neutral behaviour in the long-term (March, 1996). However, that regret-related 10p wins do not encourage the same risk-seeking behaviour creates a problem for this interpretation. It is possible that the findings reflect a mixture of risk-seeking for fear of missing the large gain that would compensate for the previous mistake, and repetition perhaps reflecting either decision inertia, or a desire to justify the previous decision. Future work will need to address which of these provide the best explanation of the results here.

Each of the above possible explanations cannot explain why the 10p win condition does not elicit identical behaviour as the 50p loss condition, when the level of operational regret associated with the two outcomes types are equivalent (i.e. the objective difference between the obtained and foregone outcomes is equal in each case). One possibility is that the impact of regret in a losing domain is greater than in a winning domain. Otherwise, this asymmetry might be driven by the nature of bet stake rather than the nature of the outcome, such that the subjective experience of regret following a risky 50p bet may be stronger than following a safer 10p bet. This may be attributable to a 50p bet being riskier, in so far as it has greater associated outcome variance, which may allow for easier construction of upward counterfactuals, stronger feelings of responsibility and greater outcome reversibility. Additionally, theories based on the role of decision justifiability in regret propose that opting for safer decisions reduces the pain of any resulting regret, since such decisions can be more easily justified by the decision-maker (Connolly & Zeelenberg, 2002). While the present study focused on behavioural effects, ratings of the subjective experience of regret during this type of task might shed light on this asymmetry.

This chapter is based largely on Nicolle, Bach, Driver and Dolan (2010). A role for the striatum in regret-related choice repetition. *Journal of Cognitive Neuroscience, early access.*

Chapter 6. A role for the striatum in regret-related choice repetition

6.1 Introduction

In the previous chapter, I suggested that the prototypical economic perspective of regret-avoidant behaviour may be more fitting with an impact of anticipated regret, and that experienced regret may be associated with different expectations or goals. In particular, one view is that regret invokes cognitive regulatory strategies, such that we mentally reconstruct an event or its antecedents in order to make ourselves feel better about mistaken choices (Pieters & Zeelenberg, 2007; Roese & Olson, 2007; Zeelenberg & Pieters, 2007). The results of the previous chapter show that decision-makers will repeat a previously regret-related choice, especially if this gives the chance of a better return - behaviour which may reflect a strategy aimed at compensating for a previous mistake. It is unclear, however, whether repetition of such choices acts to regulate the *current* experience of a regrettable outcome or whether it reflects some higher-order goal-directed behaviour that aims to regulate the *future* experience of regret.

Here I use fMRI, and the same gambling task as used in the previous chapter, to explore the neuronal mechanisms underlying such regret-related behaviour. As in the previous chapter, the task allows us to explicitly manipulate experienced regret (considered as resulting from the operationally regret-related nature of a received outcome as before), while holding anticipated outcome probabilities relatively constant. I expected to replicate a tendency for participants to repeat bets that on the immediately preceding trial had resulted in a regret-related, compared to a relief-related, outcome. As shown in the previous study, this effect was expected to be most evident after the 50p loss condition with associated agency.

For the fMRI predictions, a recent review of the relevant imaging literature by Sommer et al.(2009) is of particular importance. Here the authors suggest that the brain regions found to be involved in regret depends largely on the task used. In tasks that requiring the decision-maker to make choices between options with an associated level of *anticipated* regret, the OFC tends to be recruited (e.g. Chandrasekhar, Capra, Moore, Noussair, & Berns, 2008; Chua, Gonzalez, Taylor, Welsh, & Liberzon, 2009; Coricelli et al., 2005). On the other hand, Sommer et al argue that tasks requiring an individual to make choices based on previous *experience* (e.g. Lohrenz, McCabe, Camerer, & Montague, 2007), tend to implicate the striatum rather than the OFC. This may be because, as Sommer et al suggest, the learning signals activated by the experiences in these repeated choice tasks may depend heavily on the striatum. Since the task used here, and in the previous chapter, involves choosing between bets of differing levels of risk, we may also expect striatal involvement in the regrettable outcomes. I predicted that activity in striatum, a region that encodes both passive and action-contingent reward (O'Doherty et al., 2004), should be attenuated for regret-related outcomes compared to relief-related outcomes. Critically, given the findings of Chapter 3, I also predicted that any such reduction should be dependent upon participants being responsible for the bet selection (i.e. having agency). This is important since fictive error signals in the striatum would bear a closer relation to the psychological construct of regret if they are found to depend on agency.

If responses in the striatum act as learning signals for guiding future behaviour, then regret-related signals here may also influence resulting behaviour. To address the possibility of a specific motivational impact of (the operationally defined) regret on choice repetition, I specifically tested for BOLD responses that would distinguish choices made following a previously regret-related outcome, from choices following a previously relief-related outcome. Moreover, while ventral striatum have been widely implicated in reward learning, dorsal striatum is commonly implicated in goal-directed action (Wickens, Budd, Hyland, & Arbuthnott, 2007), and so I predicted that a similar anatomical dissociation may be found in relation to our behavioural responses to experienced regret. Since the predicted choice repetition strategy would be in apparent conflict with the relatively decreased value of the regret-related outcome, I also tested for conflict-related activity in the anterior cingulate cortex when such choices are made, a region commonly involved in the monitoring of conflict (Botvinick, Braver, Barch, Carter, & Cohen, 2001) as well as being involved in choice avoiding vs. repeating (Bush et al., 2002; Critchley et al., 2003).

6.2 Methods

6.2.1 Participants and design

I recruited 20 participants (10 female) to take part in the experiment. Participants were aged 20-38 yrs (mean = 25.6 yrs). Three participants were removed from the fMRI analysis due to faulty T1 images and resulting problems with image normalisation, but those participants were included in the behavioural analysis for completeness.

6.2.2 Experimental procedure and behavioural analysis

The same task was used as in the previous chapter's behavioural experiment. Again, trials were ascribed to four categories, in a 2 × 2 repeated-measures design that was conditional on the outcome of the previous trial for the behavioural analysis. For the fMRI analysis, onsets were modelled at the time of outcome feedback, and were separated according to outcome type. The two factors were outcome (win or loss) and stake (high 50p or low 10p). In addition, by including a no-agency control (as in the previous chapter), I could also explore how any tendency to repeat the same choice after each outcome type interacted with agency (in a 2 × 2 × 2 design, now with the additional factor of agency versus no-agency). In contrast to the procedure of the previous chapter, the no-agency trials were now randomly interleaved with the agency trials in an event-related design.

Participants each played five 8-minute sessions of the game in the scanner, each including 120 trials (2/3 agency trials). Participants were again informed that the outcomes from a random selection of 100 trials, selected after the experiment, would determine their earnings for the entire experiment. Since participants did not know which trials would be selected, they were assumed to treat all trials as having an equal potential impact on their financial gain.

The main behavioural dependent measure comprised participants' trial-by-trial tendency to repeat (versus avoid) at trial t the bet placed at trial t-1. Actual wins and losses were randomised throughout, while the overall number of trials falling into the two stake levels was choice-dependent. I also explored how any tendency to repeat the previous bet, contingent upon the outcome of the previous trial, changed from early runs (1-3) to late runs (3-5) of the game. Finally, I acquired reaction time (RT) data. Analysis of current RT to select the bet (on agency trials) was conditional on the outcome of the previous trial (which could be an agency or no-agency trial).

6.2.3 Imaging acquisition and processing

Scanning procedure was similar to that used in Chapter 3, but using a sequence of slightly lower resolution. The following imaging parameters were used: 40 oblique transverse slices, slice thickness = 2 mm, gap between slices = 1 mm, repetition time TR = 2.4 s, $\alpha = 90^o$, echo time TE = 30 ms, bandwidth BW = 3551 Hz/pixel, bandwidth in PE direction BWPE = 47.3 Hz/pixel, phase-encoding (PE) direction anterior-posterior, field of view (FOV) = 192 × 192 mm^2, matrix size 64 × 64, fat suppression. BOLD sensitivity losses in the orbitofrontal cortex and the amygdala due to susceptibility artifacts were minimised by applying a z-shim gradient moment of -0.4 mT/m*ms, a slice tilt of -30° and a positive PE gradient polarity (Weiskopf, Hutton, Josephs, & Deichmann, 2006).

Image preprocessing and data analysis were implemented using Statistical Parametric Mapping software in Matlab7.4 (SPM5; Wellcome Trust Centre for Neuroimaging, at UCL), as described in Chapter 2, section 2.4.

6.2.4 Imaging analysis

For each participant, I constructed two event-related general linear models (one to explore the response to the outcome of the current trial, and a second to test trial-to-trial effects). In the first model, 8 regressors of interest were included to allow us to assess BOLD-signal response patterns to the 8 outcome categories. These 8 outcome categories were conditionalised on the outcome of the current trial, with the three orthogonal factors of agency or no-agency; win or loss; plus 50p or 10p stake. Given a short trial length (of 4 seconds on average) I modelled trials as compound events, accounted for by one regressor onset per trial, at the time of outcome feedback. These onsets were modelled by stick-functions, and then convolved with a canonical haemodynamic response function and its temporal derivative. Motion parameters defined by the realignment procedure were entered as 6 regressors of no interest, along with 17 additional regressors of cardiac phase (10 regressors), respiratory phase (6 regressors) and respiratory volume (1 regressor).

I generated statistical parametric maps from the contrasts of interest, which included the main effects of win vs. loss, high stake vs. low stake, and agency vs. no agency along with their interactions. The interaction of critical interest was between all three factors, specifically indicating increased or decreased activity for the regret-related outcome types relative to the relief-related outcome types, on agency trials in particular. For this contrast I was particularly interested in a priori regions of interest (ROIs) within the striatum (including the caudate and putamen regularly implicated in both absolute and relative reward processing, e.g. Chandrasekhar, Capra, Moore, Noussair, & Berns, 2008; Chua, Gonzalez, Taylor, Welsh, & Liberzon, 2009; Coricelli et al., 2005; Liu et al., 2007; O'Doherty et al., 2004). The critical interaction sought was for activity greater for 50p wins than for 50p losses but conversely greater for 10p losses than for 10p wins, specifically when participants were agents of the stake choice.

A second model for each participant separated the 8 outcome-contingent regressors according to which bet-stake was chosen on the subsequent trial, giving us 16 regressors of interest and allowing us to explore differences in outcome-related responses when participants then *repeat* versus *avoid* the same choice. Again onsets were modelled with stick-functions at the time of outcome feedback, then convolved with a canonical haemodynamic response function and its temporal derivative. Since trials followed by a no-agency control trial (i.e. when there was no free choice on the subsequent trial) could not be categorised with respect to a later choice by the participant, these were included as a single regressor of no interest. I generated statistical parametric maps from the contrasts of interest. To assess the mechanism underlying the behavioural response to regret-related outcomes in the task, I tested the 2-way interaction of [‘repeat’ > ‘avoid’ after regret] > [‘repeat’ > ‘avoid’ after relief], where losing 50p and winning 10p with agency were considered operationally as regret-related, while winning 50p and losing 10p with agency were considered operationally as relief-related. This contrast allowed us to explore the brain networks involved in a tendency to repeat previous regret-related choices more than previous relief-related choices (as observed in the previous chapter’s behavioural study). Based on a possible dorsal-ventral dissociation in the roles of the striatum in reward learning and goal-directed action (Wickens et al., 2007), I constructed separate ROIs within the left and right dorsal and ventral striatum. A hypothesised involvement of the ACC in possible conflict monitoring, potentially arising since the above behavioural strategy would be in apparent conflict with the relatively decreased value of the regret-related outcome, also led us to test the same contrast within an anatomical ROI for bilateral ACC.

6.3 Behavioural results

As in the previous study, I found an increased tendency toward repeating the 50p bet after a regret-related 50p loss (when having been the agent). Here, this effect was found to reflect an early bias, which diminished significantly with increasing run number (stake × outcome × run, $F(4,72) = 3.229$, $p < 0.05$). When the outcome was associated with no-agency, neither the stake × outcome effect ($F(1,18) = 2.896$, n.s.) nor its interaction with run number ($F(4,72) = 0.567$, n.s.) was significant.

Figure 6-1 shows the probability of betting 50p at trial t, contingent on each outcome associated with agency at t-1. The early bias for repeating, compared to avoiding, a previously regret-related choice was a trend for regret-related 50p losses in runs 1 ($t(19) = 1.830$, $p = 0.083$) and significant in run 2 ($t(19) = 2.759$, $p < 0.02$). A tendency for participants to repeat 10p bets more than 50p bets was evident overall, and was found to interact significantly with agency ($F(1,18) = 26.892$, $p < 0.001$), an effect also found in the previous chapter and suggests participants are generally risk-averse when agents. Again, this was the case in all but those trials that followed a 50p loss with agency, when participants evidently preferred to repeat the 50p bet.

In this replication, a significant reaction time (RT) effect was also found (for RT to place the next agency bet) with participants being significantly quicker to respond after regret-related outcomes (that were obtained with agency), compared with relief-related outcomes (also with agency, $t(19) = 2.868$, $p < 0.01$), or compared with the equivalent outcomes with no agency ($t(19) = 2.159$, $p < 0.05$). Speeded RTs, in response to regret (Figure 6-2), accord with some previous results (Chua et al., 2009).

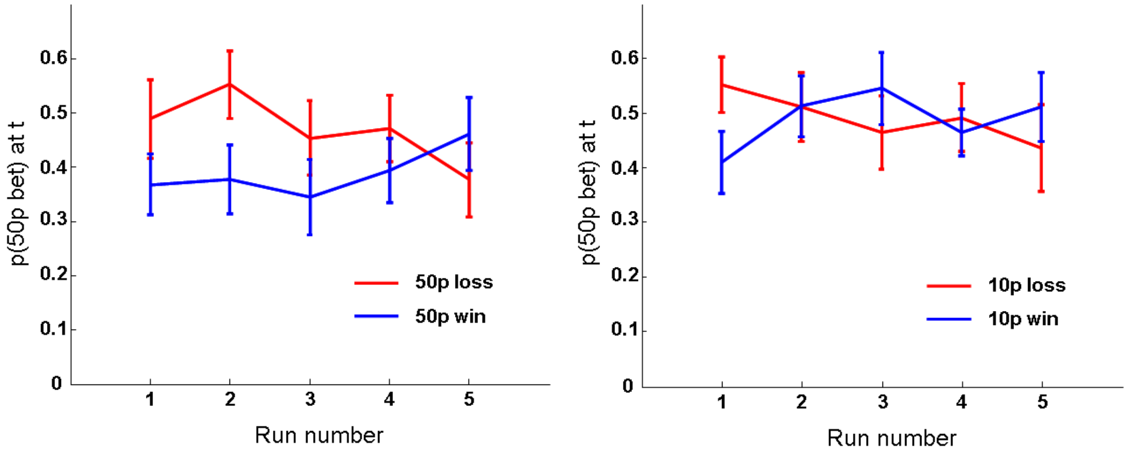

Figure 6-1 – Increased tendency to repeat, at trial t, the 50p bet after a regret-related 50p loss than after a relief-related 50p win at trial t-1. 10p wins and losses show no such difference in their effects on subsequent choice. Choice behaviour is shown for each of the 5 runs for trials associated with agency only, as no agency outcomes do not show differences in their effects on subsequent choice.

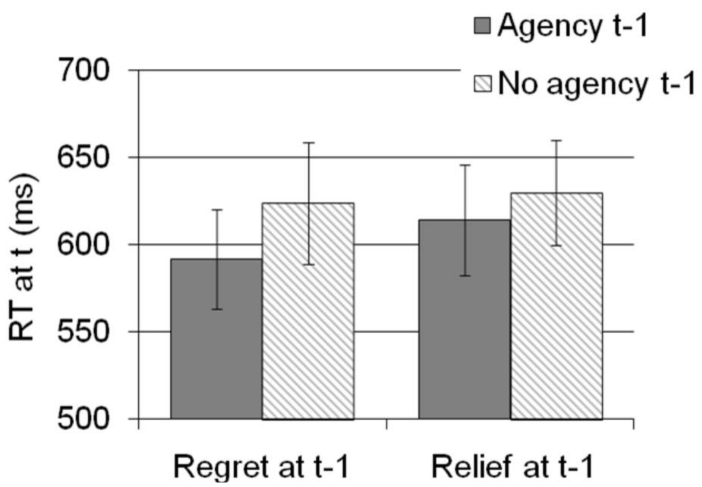

Figure 6-2 – The figure shows quickened reaction times (RT in ms) at trial t after outcomes that would have been better from the alternative choice (Regret at t-1) compared to outcomes that would have been worse from the alternative choice (Relief at t-1), but only when trial t-1 was associated with agency. Error bars show the standard error of the mean.

6.4 fMRI results

6.4.1 Main effects

The fMRI main effects are shown in Table 6-1. Increased activity for wins compared to losses overall was seen in bilateral ventral striatum (whole-brain FWE corrected at $p<0.05$); see Figure 6-3a. This pattern of increased activity was also significantly greater for 50p wins relative to 10p wins bilaterally (small-volume FWE corrected at $p<0.05$, in the whole striatum), and is consistent with previous reports of striatal responses to rewards compared to losses (Delgado, Nystrom, Fissell, Noll, & Fiez, 2000; O'Doherty et al., 2004; Fox, Trepel, & Poldrack, 2007). I found no significant interaction of this response with agency, indicating that the overall response to wins was not dependent on being responsible for the choice. No areas were significantly more active for all losses compared to all wins overall. For completeness, the main effects of agency and of stake are also presented in Table 6-1.

Table 6-1 - The table shows significant activation for the main effects of outcome, stake and agency at the time of outcome. Note that no significant voxels were found for the main effect of loss > win.

Brain regions	MNI coordinates of local maxima	Voxel number	Voxel t score
Win > Loss			
R putamen	15, 12, -3	91	7.50
L middle frontal (BA 10)	-36, 54, 6	141	6.82
Bilateral lingual sulcus (BA 30)	3, -72, 3	279	5.70
L insula (BA 13)	-27, 18, 3	32	5.40
L caudate	-12, 12, 0	31	4.38
50p > 10p			
R anterior cingulate (BA 24)	6, 24, 15	36	6.50
L substantia nigra	-3, -12, -12	57	5.25
R anterior cingulate (BA 10)	3, 51, 9	152	4.89
L inferior frontal (BA 47)	-42, 24, -15	29	4.77
R caudate	12, 3, 3	12	4.12
10p > 50p			
R parahippocampal gyrus	24,-18,-18	38	6.16
Agency > No agency			
R insula (BA 13)	33, 18, 6	15	4.27
No agency > Agency			
Bilateral precuneus (BA 7)	3, -60, 39	1471	11.70
R middle temporal gyrus	48, -54, 18	815	8.90
L middle temporal gyrus (BA 39)	-45, -60, 21	*	6.67
Bilateral lingual sulcus (BA 18)	-6, -78, -6	200	6.58
R middle temporal gyrus (BA 21)	66, -12, -12	100	6.27
R medial frontal gyrus (BA 9)	12, 51, 30	48	6.20
L middle frontal gyrus (BA 6)	-42, 6, 48	39	5.93
R hippocampus	30, -21, -18	58	5.70
L rectus gyrus (BA 11)	-3, 42, -21	138	5.62

Clusters are reported at a voxel-level significance threshold of p<0.001 uncorrected with an extent of >10 voxels

* part of the bilateral precuneus cluster

6.4.2 *Activity reflecting what-might-have-been*

A significant stake × outcome interaction reflected increased activity in bilateral ventral striatum for the two counterfactual outcomes where an outcome could have been worse (i.e. winning 50p and losing 10p) relative to when outcomes could have been better (i.e. losing 50p and winning 10p), at $p < 0.05$ FWE corrected for the whole striatum; see Figure 6-3b. This finding is consistent with the expression of a counterfactual signal in bilateral striatum, as reported previously (Chandrasekhar et al., 2008; Chua et al., 2009; Coricelli et al., 2005). The peak activity for this effect was slightly more anterior and dorsolateral within the putamen, compared to the main effect of wins. In a region of left putamen, this stake × outcome interplay was further dependent on having choice responsibility, i.e. agency (with the three-way interaction surviving FWE correction at $p < 0.05$ when using a functional ROI taken from the orthogonal 2-way interaction), thus reflecting a relief > regret difference (see Figure 6-3c). A corresponding cluster in the right putamen showed a similar effect at a lower significance level ($p < 0.002$ uncorrected, mentioned here as it points to there being no hemispheric differences in this effect). I found no regions where activity increased during the outcomes that could have been better versus worse, or showing such a pattern that interacted with agency (i.e. showing putative regret > relief effects).

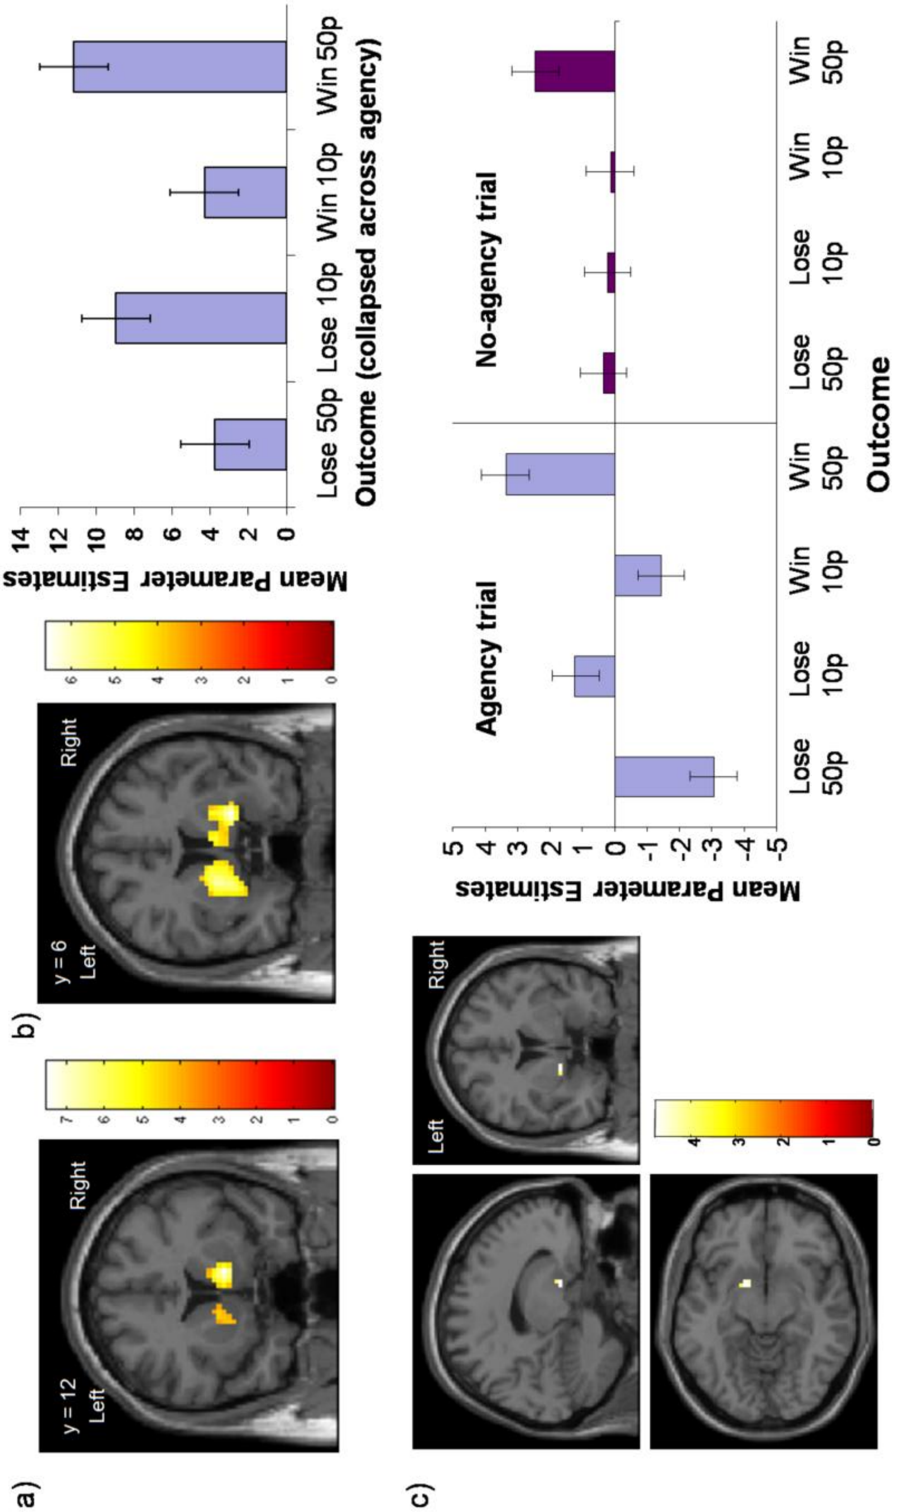

Figure 6-3 – Group SPM data, thresholded at p < 0.001 for display purposes, shown on a normalised canonical template brain. a) Activity associated with overall Win > Loss. b) Activity associated with counterfactual context, where (50p bet > 10p bet) win > (50p bet > 10p bet) loss. The upper right plot shows the mean beta values for the four outcome types (collapsed across agency) in the peak voxel at 24, 6, -9, with absolute outcome value increasing left to right on the x axis. c) SPM for the three-way interaction under which the stake x outcome interplay (as per b) was greater for agency compared to no agency trials. The lower right plots show the mean beta values for the eight trial types in the peak voxel at -15, 3, -9. Error bars show the standard error of the mean. Co-ordinates are in MNI space.

6.4.3 *Dorsal striatum reflects the tendency to repeat regret-related 50p losses*

To address what drives participants to repeat a regret-related bet choice on the immediately following trial, I divided outcomes into those where participants chose to repeat the same choice on the subsequent trial ("repeat" trials) and those where they avoided the same choice ("avoid" trials). Activity in left dorsal putamen (Figure 6-4) was greater when participants subsequently repeated a preceding regret-related choice (small-volume FWE corrected $p < 0.05$), but showed no significant difference between choices to repeat versus avoid after relief-related choices (peak MNI -24, 9, -3). I did not find any region with significantly increased activity when participants chose to avoid the previous bet after a regret-related outcome. Furthermore, no regions significantly reflected choice following outcomes associated with no agency. I found that activity in right dorsal striatum during 50p losses with agency showed a significant linear decrease from early to late runs relative to activity during 50p wins with agency (small-volume FWE corrected in the whole striatum, $p < 0.05$). This may reflect the increased behavioural tendency to repeat regret-related 50p losses in early runs (cf. Figure 6-2a).

Activity in the anterior cingulate cortex (ACC, MNI 3, 33, 24) also showed a significant interaction of choice and previous outcome (small-volume FWE corrected $p<0.05$), apparent in the same contrast that had revealed the left dorsal putamen response. This region showed increased activity associated with the subsequent choice to "repeat" a previously regret-related bet, and decreased activity associated with the subsequent choice to "repeat" a previously relief-related bet. Activity here did not differentiate regret- and relief-related outcomes when it came to decisions to "avoid" on the next trial.

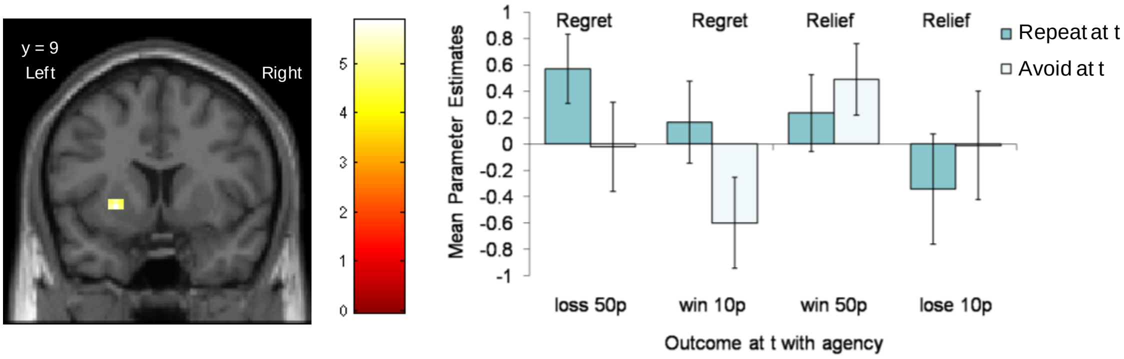

Figure 6-4 – Group SPM data thresholded at p < 0.001 for display purposes and shown on a normalised canonical template brain, showing activity for agency-related outcomes. The plot shows activity in left putamen for the contrast (repeat > avoid after regret) > (repeat > avoid after relief), with the mean beta values in the peak voxel for the left putamen at -24, 9, -3, shown for all outcome types with agency when the following choice was to repeat or avoid. Error bars show the standard error of the mean. Co-ordinates are in MNI space.

6.5 Discussion

The finding that healthy decision-makers show a behavioural tendency to repeat a previously regret-related choice is replicated here under the condition of 50p losses. Here I found this only during early runs (see Figure 6-2a), suggesting that it is an early bias (a pattern not found in the previous study perhaps because fewer trials were used). Again, one attempt to explain such behaviour invokes the gambler's fallacy, whereby the probability of a win is (mis)perceived to increase after the experience of a loss and decrease after a win (e.g. see Croson & Sundali, 2005), thus encouraging a higher risk bet after a loss. It is conceivable that such an effect might diminish over time (as for the behavioural pattern here) if participants estimated the outcome probabilities with more precision over time. But the gambler's fallacy alone would not explain the differences I observed under conditions of agency versus no-agency. Furthermore, an analysis of sequences of identical outcomes did not reveal any significant tendency to place a 50p bet more after losses than after wins when more losses were experienced in a sequence, although that would be expected from the gambler's fallacy.

The aim of this study was to investigate neural mechanisms underlying regret-related choice repetition when participants were responsible for their choice. I found increased activity in the dorsal striatum when subjects made such choices. While I did not find OFC involvement in regret-related outcomes (as might be predicted from the literature review in Chapter 1, section 1.4), this is in keeping with the suggestion by Sommer et al. (2009) that the learning signals important in a task such as this one (with these signals being elicited by *experienced*, rather than *anticipated*, regret) depend upon striatal more than OFC activity.

Previous fMRI studies show the striatum is important for processing primary rewards (e.g. O'Doherty, Deichmann, Critchley, & Dolan, 2002) as well as more abstract rewards, including money (Delgado et al., 2000), romantic love (Aron et al., 2005) and humour (Mobbs, Greicius, Abdel-Azim, Menon, & Reiss, 2003). Along with processing passive receipt of rewards (O'Doherty et al., 2004) the striatum is also implicated in encoding of action-outcome contingencies (Delgado, Miller, Inati, & Phelps, 2005; Tanaka, Balleine, & O'Doherty, 2008). Moreover, it encodes violations of our expectations for such contingencies, in the form of prediction-errors signals, which reinforcement learning models show to be central to guiding future behaviour (Barto & Sutton, 1998; Berns, McClure, Pagnoni, & Montague, 2001). While prediction errors relate to within-option counterfactual comparisons, the regret literature also implicates striatum in processing of rewards relative to between-option (i.e. choice dependent) counterfactual reference points (Chandrasekhar et al., 2008; Chua et al., 2009; Coricelli et al., 2005; Liu et al., 2007; Lohrenz et al., 2007). The finding of relatively increased activity in ventral striatum during relief-related outcomes is consistent with this behaviour-dependent coding of rewards in the striatum. Furthermore, I find a critical role for agency in this pattern of activity, as expected from Chapter 3, thus further contributing to the ongoing debate regarding the role of responsibility in regret and relief.

I found activity in ventral striatum that reflected the subjective value of the experienced outcome relative to what-might-have-been under a different choice (Figure 6-3c). However, activity in dorsal striatum was particularly involved in the behavioural response, i.e. regret-related choice repetition (Figure 6-4). Evidence for a dorsal-ventral dissociation in the roles of the striatum in reward learning and goal-directed action may be of importance here (for review see Wickens, Budd, Hyland, & Arbuthnott, 2007). Such functional-anatomical dissociation might allow independent processing of more ‘emotional’ responses versus behavioural responses to regret-related events. While not always being a conventionally regret-avoidant response, dorsal striatum responses appear to permit selection of future actions that may nonetheless bring the decision-maker toward higher-order goals, such as making up for a past mistake. This aspect of my findings is in general accord with previous work showing the dorsal striatum to be especially important in stimulus-response-reward learning, while the ventral striatum is important for stimulus-reward prediction (Delgado et al., 2005; O'Doherty et al., 2004; Tricomi, Delgado, & Fiez, 2004). O'Doherty et al. (2004) proposed that the ventral and dorsal striatum play dissociable roles in the control of our action, with the former serving a “critic” role important for passively predicting the value of future states, while the latter serves an “actor” role involved in updating stimulus-response-reward associations and reinforcing, or gating, the selection of future high value actions. Dopaminergic projections to the dorsal and ventral striatum originating from different sources (from the substantia nigra (SN) and VTA respectively) may provide a neurophysiological basis for their different roles (see review by Wickens et al., 2007). Here my findings add a new line of support for such a functional dissociation along the dorsal-ventral axis of the striatum, while also showing for the first time such dissociation in a context where rewards are relative to their counterfactual alternatives.

Activity in dorsal striatum has also been found to reflect choice-induced modifications of value. For example, Sharot, De Martino, & Dolan (2009) found increased caudate activity associated with post-choice increases in the subjective value of the selected option. They proposed that this increased activity may be associated with a desire to reduce cognitive dissonance. Their findings suggest a role for dorsal striatum in higher-order, temporally delayed, goal-directed action. Similarly, the present findings may reflect a role of the dorsal striatum in updating an increased subjective value of repeating a previously regret-related choice, perhaps motivated by a desire to make up for our mistakes. Tanaka et al. (2004, 2007) have found the dorsal striatum to be active when choosing larger delayed rewards in favour of smaller immediate rewards, supporting its role in the motivation of actions toward longer-term goals, while more ventral regions of the striatum were active when choices were more impulsive, i.e. in favour of the smaller immediate rewards.

From such previous work, I questioned whether value related signal, observed in ventral striatum, may be regulated or modified by adherence to higher-order behavioural goals. There was no evidence that decreased ventral striatal activity to regret-related outcomes was influenced by a choice to repeat or avoid on the next trial, while dorsal striatum did show this effect. Anatomical studies show ascending spirals within both a striato-nigro-striatal loop and a limbic-to-motor striatocortical loop, with the direction of both being ventromedial to dorsolateral (e.g. Draganski et al., 2008; Haber, Fudge, & McFarland, 2000; Haber & Knutson, 2009), and no evidence of direct information transmission from dorsal to ventral striatum. It is possible, however, that interactions between processing in ventral and dorsal striatum are expressed elsewhere, where anterior cingulate cortex (ACC), with its projections to ventral and dorsal striatum, is a prime candidate. Indeed, I found responses in ACC that were modulated by decisions to repeat previous bets, depending on previous outcome type. Since the ACC is often implicated in the monitoring of conflict (e.g. Botvinick et al., 2001) the observed response here may be associated with *monitoring* of conflict between value and choice. Another possibility, however, is that it regulates or *gates* value-related activity in ventral striatum, thereby facilitating choice-related activity in the dorsal striatum.

A possible alternative explanation for the findings of Studies 3 and 4 is that the 50p loss actually induces *less* regret relative to the 10p win, due to the high possible reward associated with a 50p bet affording greater justification for its selection. This may be especially the case for decision-makers who are primarily focused on possible gains, rather than possible losses, and may explain both the increased tendency to bet 50p after a 50p loss, as well as the increased activity in dorsal striatum (if this is reward-related activity). However, this explanation also predicts two other behavioural tendencies, which were not found. Firstly, it would predict that the tendency to bet 50p would be the highest overall, since betting 50p would be highly justifiable. However, participants were generally risk-averse, tending to bet the 10p more often than the 50p. Secondly, it would predict that the tendency to repeat 50p after a 50p win would be highest, since a 50p win would further defend the justifiability of a 50p bet. We actually find the tendency to bet 50p is lower after a 50p win compared to after a 50p loss. According to the first hypothesis, on the other hand, relief-related 50p wins would not generate the same desire to make up for the past mistake, thus reducing the value of the 50p bet (a tendency that appears to be associated with decreased ACC activity). Subjective ratings of emotional experience, as well as information on decision-makers' rationale for their choices, would have been useful here.

Here, I show that operational regret, under certain choice constraints, can lead to choices that appear to reflect attempts to make up for apparent mistakes, reflected in repetitive gambling behaviour. Furthermore, these findings suggest a central role for the striatum in this behaviour, in a manner that accords with current models of dorsal-ventral dissociation for striatum function. Further consideration of the role of regret-regulatory strategies in such behaviour, along with the neuronal mechanisms involved, is likely to be crucial in understanding mechanisms driving maladaptive behaviours such as gambling, as well as that seen in patient populations where compulsive gambling can sometimes be a side-effect of neuromodulatory therapy, as seen in Parkinson's disease.

This study has been submitted under the following reference: Nicolle, Fleming, Bach, Driver & Dolan, (2010). A regret-induced status-quo bias.

Chapter 7. Does regret encourage a status-quo bias?

7.1 Introduction

When faced with a complex decision, people tend to accept the status quo. Indeed, across a range of everyday decisions, such as whether to move house or trade in a car, or even whether to flip the TV channel, there is a considerable tendency to maintain the status quo and refrain from acting (Samuelson & Zeckhauser, 1988). In a recent study, Fleming, Thomas, & Dolan (2010) investigated the neural mechanisms for overcoming such a status-quo bias in a difficult perceptual task, but these data did not account for *why* the bias exists in the first place. One influential view is that a status-quo bias is associated with anticipated regret (Baron & Ritov, 1994; Kordes-de Vaal, 1996; Ritov & Baron, 1990, 1995; Tykocinski, Israel, & Pittman, 2004; Tykocinski & Pittman, 1998, 2001), such that this behaviour reflects a regret-minimising strategy. Here, I explore how asymmetric behavioural and brain responses to errors that follow choices to either reject or accept a status-quo may encourage a status-quo bias on subsequent decisions. Moreover, such a bias may account for my previous observation that regrettable outcomes encourage choice repetition, rather than avoidance.

Existing literature on counterfactual thinking and regret supports that our emotional responses to the outcomes of our actions depend not just on the nature of the outcome itself, but also on the type of behaviour that brings it about. Notably, mistakes arising from decisions to reject a status-quo apparently have an amplified emotional impact, compared to mistakes arising from a decision to accept the usual course, even when the two outcomes are objectively equally aversive. For example, Kahneman & Tversky (1982) found that the same car accident was judged to be more unpleasant subjectively if the driver had taken an unusual route home rather than their habitual route. Similarly, greater regret is judged to be incurred when value is lost on stocks that were recently switched to from a previously held stock, compared to if the failing stock was instead one that had been always held. Such asymmetries in the paths to regret have since been replicated in many different contexts in purely behavioural studies (Baron & Ritov, 1994; Feldman, Miyamoto, & Loftus, 1999; Landman, 1987; Tsiros & Mittal, 2000; Zeelenberg, van den Bos, van Dijk, & Pieters, 2002).

Several theories have been proposed to explain the asymmetric impact of status-quo rejection errors (versus status-quo acceptance errors). Rejecting a status-quo may be more salient and sometimes more costly than accepting it, such that status-quo rejection errors invite greater attention and are more easily remembered on future occasions. Of particular interest in relation to the literature on regret and counterfactual thinking is the suggestion that status-quo rejection is less “normal” (it is by definition less of a status-quo). According to Kahneman & Miller's (1986) Norm Theory, abnormal events are more mutable psychologically, which may allow for easier construction of counterfactual alternatives (which might have avoided the error). As a result, the emotional impact of status-quo rejection errors is amplified relative to otherwise equivalent status-quo acceptance errors. An action associated with status-quo rejection might also be considered as more directly causal of the outcome than the ‘inaction’ typically associated with status-quo acceptance, thereby amplifying a sense of accountability or blame for the error (Ritov & Baron, 1990; Spranca, Minsk, & Baron, 1991). With decision justifiability of key importance in theories of regret (Connolly & Zeelenberg, 2002; Inman & Zeelenberg, 2002), a broad view encompassing many of the other accounts suggested may be that the increased normality and lower potential for blame associated with status-quo acceptance makes this more justifiable and so less likely to induce regret.

An amplified psychological regret following mistaken status-quo rejection may underlie a pervasive behavioural bias towards accepting the status-quo in later decisions (Baron & Ritov, 1994; Kordes-de Vaal, 1996; Ritov & Baron, 1990, 1995; Tykocinski, Israel, & Pittman, 2004; Tykocinski & Pittman, 1998, 2001). Moreover, a status-quo bias may provide a means of regulating future experiences of regret, by improving the perceived justifiability of future choices (Zeelenberg & Pieters, 2007). Motivated by the purely behavioural literature connecting regret and the status-quo bias, and by the regret-related choice repetition found in the previous two chapters, here I examine the neural mechanisms underlying such a bias by testing an hypothesis that links choice behaviour, and neuronal activity at choice, to antecedent error processing.

Turning to the possible neural basis of this asymmetry, functional neuroimaging studies have found a critical role of medial prefrontal cortex (mPFC) and insula in error processing (Braver, Barch, Gray, Molfese, & Snyder, 2001; Carter et al., 1998; Kiehl, Liddle, & Hopfinger, 2000; Klein et al., 2007; Menon, Adleman, White, Glover, & Reiss, 2001). Interestingly, insula and anterior cingulate cortex activity has also been implicated in negative evaluation of our choices (Liu et al., 2007) as well as more specifically in the experience of regret (Chua, Gonzalez, Taylor, Welsh, & Liberzon, 2009). Given these findings, I predicted that insula and mPFC activity may also be associated with the inflated emotional response to status-quo rejection errors, as well as activity in the OFC, striatum and amygdala due to their apparent roles in regret. Furthermore, the insula is thought particularly important in predicting future consequences, for instance in terms of risk or uncertainty associated with choice (Critchley, Mathias, & Dolan, 2001; Paulus, Rogalsky, Simmons, Feinstein, & Stein, 2003; Paulus & Stein, 2006; Preuschoff, Quartz, & Bossaerts, 2008; Singer, Critchley, & Preuschoff, 2009). This leads to a further prediction that the insula may be involved not only in representing the outcomes of our choices, but also in how this information is then used to guide future behaviour. In particular, activity here is predicted to be associated with a status-quo bias being elicited in response to a prior status-quo rejection error.

Previously Fleming et al. (2010) identified a role for fronto-basal ganglia interactions in rejecting the status-quo in a difficult perceptual choice scenario, but they were unable to specify the antecedent neural mechanisms biasing choices. In the present study, I address this question directly in a similar task but now linking differences in error processing on a given trial (when explicit feedback on errors or correctness is provided) to any impact on the status-quo bias for the subsequent trial. Candidate mechanisms were predicted to lie within mPFC and insula due to their apparent involvement in error-detection and decision-making (Braver, Barch, Gray, Molfese, & Snyder, 2001; Carter et al., 1998; Kiehl, Liddle, & Hopfinger, 2000; Klein et al., 2007; Menon, Adleman, White, Glover, & Reiss, 2001). Surprisingly, previous neuroimaging studies have overlooked possible differences in the way such errors may be processed, such as a potential behavioural and neuronal asymmetry between status-quo acceptance and status-quo rejection errors.

7.2 Methods

7.2.1 Participants

Twenty healthy individuals participated in the study. Three participants were excluded after equipment failures meant their responses were not fully recorded during scanning. Thus, seventeen participants in all were included in the data analysis (8 female, mean age 23.5 yrs ± 4.8). Twenty additional participants were included in an initial manipulation-check behavioural experiment (15 female, mean age 22.8 yrs ± 4.0).

7.2.2 Experimental procedure

I modified the task reported by Fleming et al. (2010), originally designed to assess neuronal mechanisms associated with status-quo rejection, by now including explicit trial-by-trial performance feedback for errors and correct responses (Figure 7-1). In brief, this task requires a trial-by-trial perceptual decision, with participants judging whether a target ball, landing on a simulated tennis court, was "IN" (overlapping the line) or "OUT". Participants started each trial by holding the key corresponding to their choice on the previous trial (or holding IN on the very first trial). Each trial began with a central fixation cross and two peripheral lines. After a variable delay, the target ball appeared briefly (66 ms) on either the left or right of the screen. Since the side of target presentation was random, participants were instructed to fixate centrally in order to maximise performance on the task, although formal eye-tracking was not instigated as this was irrelevant to my current experimental question. Participants were then reminded of their previous choice (the key currently still pressed, either IN or OUT), and decided either to continue holding the current key to stay with their previous choice (accept status-quo) or to switch to the alternative key, thus switching their decision (reject status-quo). Participants then held the appropriate key (new if they had switched, old if they had not) until the occurrence of a trial in which they chose to switch.

In order to create a balanced design, the correct decision was to accept the status-quo on 50% of trials and to reject the status-quo on the remaining 50%. Explicit feedback was presented at the end of each trial, corresponding to a 40p gain or 40p loss for correct or incorrect decisions, respectively. In order to obtain approximately 30% errors overall, I manipulated the difficulty of the perceptual decisions by altering the distance of the target from the outside edge of the line, using a 2-up-1-down staircase procedure (Levitt, 1971) in 0.1 visual degree angle steps. Overall, each participant played 320 trials, over 4 separate sessions, while fMRI brain data were simultaneously acquired. Keys were pressed with the index or middle finger of the right hand, and key-choice contingencies (i.e. whether index or middle meant IN or OUT) were counterbalanced across sessions. Participants were paid according to the actual winnings for their two best sessions, which averaged £25.51 in the main experiment (*SD* = £2.7).

After scanning, participants gave subjective ratings of their emotional responses to errors in the task out of the scanner. These were post-hoc overall ratings, not trial-by-trial, and were completed on a 9-point Likert scale. Ratings probed relative responses to reject status-quo and accept status-quo errors along dimensions of regret, disappointment, blame, sense of "kicking yourself", and desire to "make up for the mistake" (separate Likert scales for each of these 5 dimensions). For relative regret ratings, participants were asked "Which felt more regretful – when you stayed with your previous choice and made an error, or when you switched your choice and made an error." Absolute ratings of errors overall were also gathered along the same dimensions (i.e. collapsed across accept and reject errors).

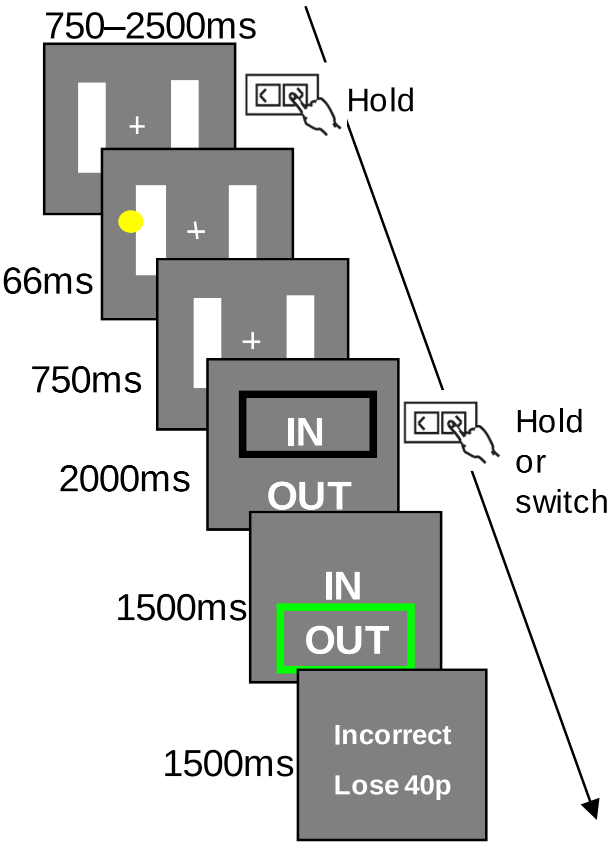

Figure 7-1- An exemplar trial timeline. Participants began each trial holding the key corresponding to their choice on the previous trial, while fixating on the central cross flanked by two tramlines. After a varied delay, a ball landed on either the left or right of the screen, at any height on the tramline. Participants were then asked to judge whether the ball landed "IN" (overlapping the line) or "OUT". Their decision was indicated by either continuing to depress their previous decision (a black box served to remind them of this status-quo), or to switch to the alternative key to reject the status-quo and switch their decision. On rejecting the status-quo, participants then held the new key until they chose to switch again on a later trial. Accuracy feedback, and associated monetary win or loss, was presented at the end of each trial.

In an initial manipulation-check behavioural experiment, a separate cohort of participants gave absolute ratings of experienced regret after *every* error trial from 320 consecutive trials, on a 9-point Likert scale, where 9 was the highest possible regret. With this additional experiment, I sought to confirm that errors on the task did indeed induce reported experiences of regret, while also aiming to provide unique evidence for a trial-by-trial account of the rejection-acceptance differences in regret (since previous studies have largely relied on anticipated relative regret for hypothetical scenarios, rather than actual personal current experiences of regret). The task used in this experiment was otherwise the same as in the scanning study. I also collected skin conductance responses (SCRs) in this additional experiment (as described in the Chapter 2, section 2.5). Each trial onset, ball appearance, decision cue, and outcome, were modelled as separate events, followed by an evoked SCR with a canonical shape (Bach, Daunizeau, Kuelzow, Friston, & Dolan, 2010).

7.2.3 Imaging acquisition and processing

Participants were scanned with a single-shot gradient-echo EPI sequence. Imaging parameters were as follows: 48 oblique transverse slices tilted by 30^o, slice thickness = 2mm with a 1mm gap between slices, repetition time TR = 3.36s, α = 90°, echo time TE = 30ms, BW_{PE} = 27 Hz/pixel, positive phase-encoding gradient polarity in an anterior-posterior direction, field of view = 192×192 mm^2, matrix size 64×64, fat suppression, z-shim gradient pre-pulse moment = -1.4 mT/m×ms.

Image preprocessing and data analysis were implemented using Statistical Parametric Mapping software in Matlab2009a (SPM8; Wellcome Trust Centre for Neuroimaging, at UCL), as described in 2.4.

7.2.4 Behavioural analyses

In our manipulation-check behavioural experiment, I tested for a difference between experienced regret ratings for rejection and acceptance errors, using a paired sample t-test. Ratings here were standardized to remove individual differences in mean regret ratings.

Post-hoc (rather than online) ratings of subjective response to errors were analysed in the scanning experiment. These included absolute ratings of regret, disappointment, blame, sense of "kicking yourself", and desire to "make up for the mistake" in response to errors overall. Relative responses to rejection and acceptance errors were collected along the same dimensions, and then any mean bias of each rating for one type of error over the other was tested with a one-sample t-test against the null hypothesis of no bias (midpoint rating of 5 on the scale).

I predicted an overall behavioural status-quo bias (tendency to accept the status-quo option), in line with what has previously been shown with this particular 'line-judgement' task (Fleming et al., 2010). I measured such bias as a tendency towards status-quo acceptance over-and-above what was the correct (optimal) choice on each trial, as assessed with a one-sample t-test. Finally, I measured how this status-quo bias may be influenced by the outcome of the preceding trial, using a 2×2 repeated-measures ANOVA on data for the current trial, with factors for outcome (correct, incorrect) and choice (accept, reject) on the previous trial.

7.2.5 Imaging analysis

In this experiment, I used fMRI to test two hypotheses in relation to the link between errors types and the status-quo bias. First, guided by the regret literature (see section 7.1), I predicted that error-related brain responses would be greater for erroneous status-quo rejections than for erroneous status-quo acceptances. Since the overall main effect of rejection error compared to acceptance error would be confounded by the motor response (changed or unchanged key), I tested for the critical interaction between outcome (error/correct feedback) and status-quo rejection/acceptance. Specifically, for the BOLD data corresponding to the feedback event, I tested the following contrast: Reject [error - correct] > Accept [error - correct]. Note that this contrast tests for a predicted increase in brain responses to errors on rejection than acceptance trials, while perfectly controlling for motor-response related activations via the subtraction of correct status-quo rejection or acceptance response. This contrast was performed within an event-related general linear model using a model that included the 4 regressors of interest: accept-error, accept-correct, reject-error, and reject-correct. Onsets were modelled with stick-functions at the time of outcome feedback, convolved with the standard canonical haemodynamic response function and its temporal derivative. Regressors of no interest comprised target ball onsets and button press, along with start of the decision phase parametrically modulated by RT if a decision to reject the status-quo was made. This modulation by RT was included to factor out activity previously shown to reflect local variations in difficulty in this task (Fleming et al, 2010). Head motion parameters defined by the realignment procedure were entered as 6 regressors of no interest, along with 17 additional regressors of cardiac phase (10 regressors), respiratory phase (6 regressors) and respiratory volume (1 regressor).

I generated statistical parametric maps from the contrasts of interest. At the time of outcome feedback, I was interested in the main effect of error corrected for multiple comparisons across the whole brain volume. For the outcome-related 2-way interaction, Reject [error - correct] > Accept [error - correct], my focus was especially on error-related activity. Accordingly, I restricted the search volume for this 2-way interaction to a functional ROI from the orthogonal main effect of error > correct, set at a threshold on $p < 0.005$ uncorrected. Within these ROIs, I report voxel-wise activity significant at a FWE corrected threshold of $p < 0.05$. For completeness, I also report activity that survives whole-brain cluster-wise corrected significance of $p < 0.05$.

My second hypothesis was that differences in processing of errors (versus correct) status-quo rejections versus acceptances would be associated with an increased behavioural status-quo bias on the subsequent trial. To explore the neural basis of such bias, I tested for two forms of a 3-way interaction whereby enhanced responsivity to status-quo rejection, compared to status-quo acceptance, errors (i.e. Reject [error - correct] > Accept [error - correct]) should be associated more with a subsequent decision to accept (rather than reject) the status-quo. I hypothesised that such an interaction could be associated either with outcome-driven responses (predicting subsequent choice) or choice-driven responses (associated with the initiation of the choice). To test these two means though which asymmetries in error processing could underlie a behavioural status-quo bias, I implemented two further event-related general linear models per participant. The first modelled onsets at the time of outcome feedback (for the 4 outcome types – accept-error, accept-correct, reject-error, and reject-correct) contingent on the choice made on the *subsequent* trial (subsequent accept or subsequent reject). The second modelled onsets at the start of the subsequent decision phase, for either status-quo acceptance or status-quo rejection on the *current* trial contingent on outcome type (i.e. one of the four outcomes described above) on the previous trial. In both models, onsets were stick functions convolved with a canonical haemodynamic response function and its temporal derivative. Motion and physiological regressors of no interest were as in the previous model. I generated statistical parametric maps from the 3-way interaction contrast of interest. For both models, I restricted my analysis to functional ROIs identified from the outcome-related effects found in the first model. The same statistical thresholds were used, as described above.

7.3 Results – Manipulation-check study

The preliminary behavioural experiment showed trial-by-trial regret ratings were higher than 5 on the 9-point scale (mean = 5.83 out of 9), indicating significant experienced regret for errors in the task. Critically, regret ratings were greater for status-quo rejection errors (mean = 6.00, z mean = 0.09) than status-quo acceptance errors (mean = 5.66, z mean = -0.06) ($t(19) = 2.21$, $p < 0.05$) in line with previous studies. Here, I uniquely extend this to actual trial-by-trial outcomes and personal experiences in the present task, rather than merely hypothetical scenarios as in previous studies. Overall, I found a significant 6.2% bias towards the status-quo ($t(19)=5.01$, $p<0.0001$). I also observed an effect of correctness in skin conductivity response, at the timepoint of the outcome, where the SCR was bigger for incorrect than for correct responses ($t(20) = 2.9$, $p < 0.01$). No effect was observed for action or the action × correctness interaction, and no effect was observed at any other time point in the trial.

7.4 Results – Main experiment

7.4.1 Post-hoc subjective ratings

As predicted, regret was again significantly greater for status-quo rejection errors than status-quo acceptance errors (mean deviation from neutral point on the scale = 1.24, $t(16) = 2.32$, $p < 0.05$). No such effect was found for ratings of disappointment (mean deviation from neutral = 0.35, $t(16) = 0.63$, p = ns), despite absolute ratings of regret and disappointment for errors being similar (mean regret = 5.29, mean disappointment = 5.35, $t(16) = -0.13$, p = ns). The only other rating that was higher for status-quo rejection errors than status-quo acceptance errors was self-blame (mean deviation neutral = 1.06, $t(16) = 2.31$, $p < 0.05$), although the feeling of "kicking yourself" was trend significant (mean deviation neutral = 0.88, $t(16) = 1.99$, $p = 0.06$).

7.4.2 fMRI data: Enhanced response to status-quo rejection errors in anterior insula and mPFC

I addressed BOLD responses corresponding to the time of outcome feedback. A main effect of error > correct was found in bilateral anterior insula (Figure 7-2a). These were FWE significant whole-brain ($p < 0.05$ whole brain cluster-corrected). No effect of error was found within the OFC. The opposite contrast, correct > error responses, revealed effects in striatum, postcentral gyrus, superior temporal cortex and superior occipital cortex, but are of less interest due to a lack of a priori hypotheses for my current purposes.

The critical interaction contrast explored whether the main effect of error (versus correct) was stronger for status-quo rejection errors than for status-quo acceptance errors, as predicted from the subjective rating results. I tested for activity fitting this pattern at the time of outcome feedback. At whole-brain cluster-corrected significance, I found activity in medial prefrontal cortex (extending into the rostral anterior cingulate) showed greater responsivity to status-quo rejection errors compared to status-quo acceptance errors, as evident in the contrast Reject [error - correct] > Accept [error - correct] (Figure 7-2b). I also tested specifically whether the error-related anterior insula activity, within a mask defined by the (orthogonal) main effect of error (extent threshold $p < 0.005$), showed any such enhancement for the case of status-quo rejection errors. I found that left anterior insula was significantly greater for status-quo rejection than to status-quo acceptance errors ($p < 0.05$ FWE corrected for the ROI) (Figure 7-2c & d). No regions showed significant effects in the inverse interaction. All effects are summarised in Table 7-1.

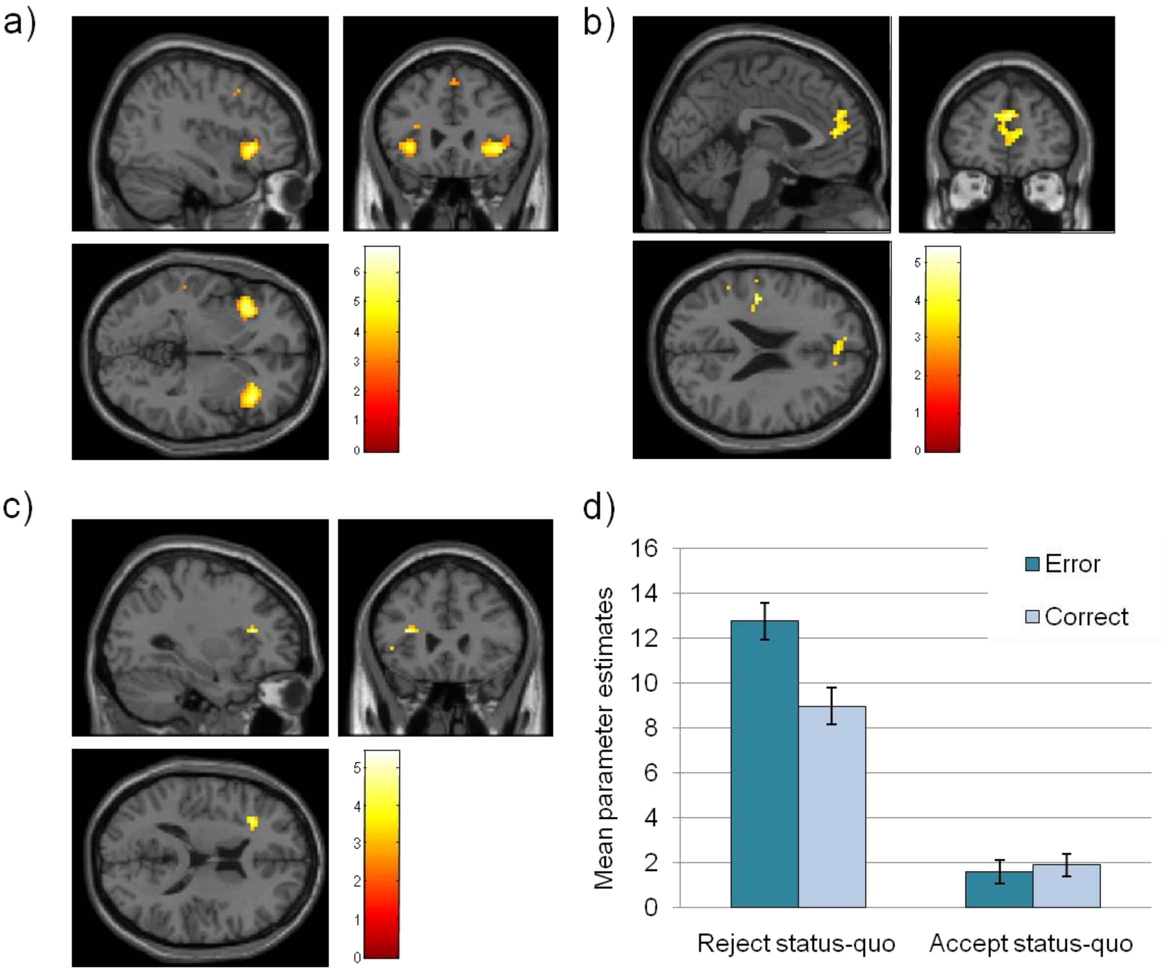

Figure 7-2 – Group SPM data showing responses at the time of outcome feedback, thresholded at p<0.005 for display purposes, shown on a normalised canonical template brain. a) While brain corrected bilateral anterior insula response to the main effect of error > correct (MNI peaks -36,17,-5 and 39,26,-2). b) Whole brain corrected activity in medial prefrontal cortex (MNI peak 3, 47, 25) reflecting the interaction of choice and outcome (Reject [error > correct] – Accept [error > correct]). c) Anterior insula activity from the main effect of error, showing the same interaction of choice and outcome and d) the plotted mean parameter estimates for the four outcome types at the peak voxel (MNI -30, 26, 16). Error bars show within-subject standard errors of the difference between error and correct responses for the two decision types.

Table 7-1 – The table shows activation for the main effects of error and correct trials, and the interaction of outcome with rejecting or accepting the status-quo. Contrasts were performed on responses modelled at the time point of trial outcome.

Brain regions	MNI coordinates of local maxima	Voxel number at p<0.001 uncorrected	Voxel t score
Error > Correct			
L anterior insula (BA 47)	-36, 17, -5	68	6.89
R anterior insula	39, 26, -2	69	6.44
Superior frontal (BA 8)	3, 32, 52	11	5.88
Correct > Error			
R postcentral	36, -31, 64	133	6.59
R superior temporal (BA 41)	57, -22, 10	121	6.42
R putamen	21, 8, -8	23	6.06
R superior occipital (BA 19)	30, -85, 25	61	5.97
R supplementary motor area	12, -16, 73	50	5.50
R paracentral lobule (BA 6)	3, -37, 61	193	5.30
R caudate	27, 2, 13	43	5.21
Reject [error > correct] – Accept [error > correct]			
Medial prefrontal cortex/anterior cingulate (BA 9)	3, 47, 25	14	4.76
L anterior insula	-30, 26, 16	3	4.62

7.4.3 Anterior insula and mPFC activity at choice predicts a status-quo bias

In the main experiment, I observed a 7.9 % status-quo bias ($t(16) = 7.59$, $p < 0.00001$). I tested whether this bias was dependent on the outcome of the previous trial, using a 2 × 2 repeated-measures ANOVA with factors for outcome (error/correct) and reject/accept status-quo on the previous trial. I found a main effect of reject/accept, such that overall participants tended to repeat the strategy used on the previous trial (i.e. to keep rejecting or keep accepting, $F(1,16) = 13.79$, $p < 0.005$). However, this effect interacted with outcome ($F(1,16) = 11.25$, $p < 0.005$), with erroneous status-quo rejection encouraging significantly greater subsequent status-quo acceptance (mean = 0.56) compared to correct status-quo rejection (mean

= 0.49) ($t(16) = 3.15$, $p < 0.01$), which was not the case for erroneous status-quo acceptance (mean = 0.59) compared to correct status-quo acceptance (mean = 0.62) ($t(16) = -1.88$, ns, Figure 7-3). Furthermore, erroneous rejection of the status-quo was the only outcome type that significantly drove a tendency towards an alternative strategy (i.e. from rejection to acceptance of the status-quo) on the subsequent trial, as shown by a one-sample t-test against no bias ($t(16) = -3.21$, $p < 0.01$).

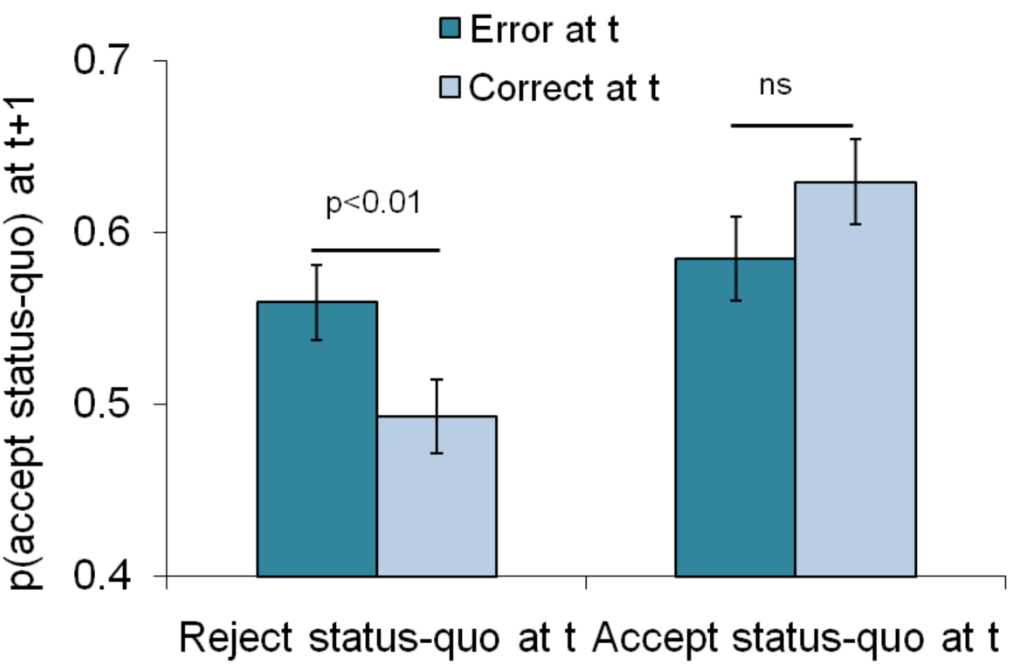

Figure 7-3 – Figure shows how the status-quo bias is contingent upon the outcome of the previous trial. Error bars show within-subject standard errors of the difference between error and correct responses for the two decision types.

At the time of choice on the subsequent trial, pooling correct and error trials, I found enhanced brain activity associated with a decision to reject the status-quo (compared to accept), in bilateral cerebellum, bilateral supramarginal gyrus, and bilateral insula ($p < 0.05$, whole-brain cluster-level FWE corrected). Decisions to accept the status-quo were associated with a trend significant cluster in rostral ACC ($p < 0.06$).

I next explored the possible link between the enhanced error-related activity following rejection errors and the behavioural bias towards subsequent status-quo acceptance. Accordingly, I designed a three-way interaction contrast which tested whether the critical 2-way interaction (Reject [error - correct] > Accept [error - correct]) was more closely associated with participants making a subsequent decision to accept rather than reject the status-quo, thereby providing a link between the neural outcome response and the subsequent behavioural status-quo bias. I found significant effects when modelling the interaction at the time of subsequent choice. Specifically, I found that the anterior insula pattern of activity shown in Figure 7-2 c-d was also expressed when participants choose to accept the status-quo, but not when they decide to reject the status-quo (contrast estimates at the peak of the previous anterior insula 2-way interaction [MNI -30, 26, 16] taken forward to a 2 × 2 × 2 ANOVA, 3-way interaction, $F(1,16) = 10.51$, $p < 0.005$). I found that mPFC showed the same pattern of activity to anterior insula, such that the enhanced responsivity to errors under status-quo rejection was also expressed at subsequent choice, but only when participants are currently choosing to accept, rather than reject, the status-quo. However, this response was not expressed within the same peak as for the 2-way outcome-related response. These choice-related responses are shown in Figure 7-4 and in Table 7-2.

When the 3-way interaction was tested at the time of preceding outcome, now categorised by subsequent choice, no significant effects were found. The time-course plots in Figure 7-4 b & d also indicate these responses appear to emerge at the time of choice on the subsequent trial. Since a 3-way interaction was not identified in these regions at the time of outcome, predicting subsequent choice, then these results point to a *choice-related* fMRI response (dependent on the outcome of the previous trial). I note, however, that a limited inter-trial interval and low temporal resolution of fMRI prevents us from unequivocally assigning modulation of the BOLD response to either previous outcome or subsequent choice. Future studies exploiting the higher temporal resolution of MEG may be useful for teasing apart the exact temporal dynamics of outcome-driven versus choice-driven responses.

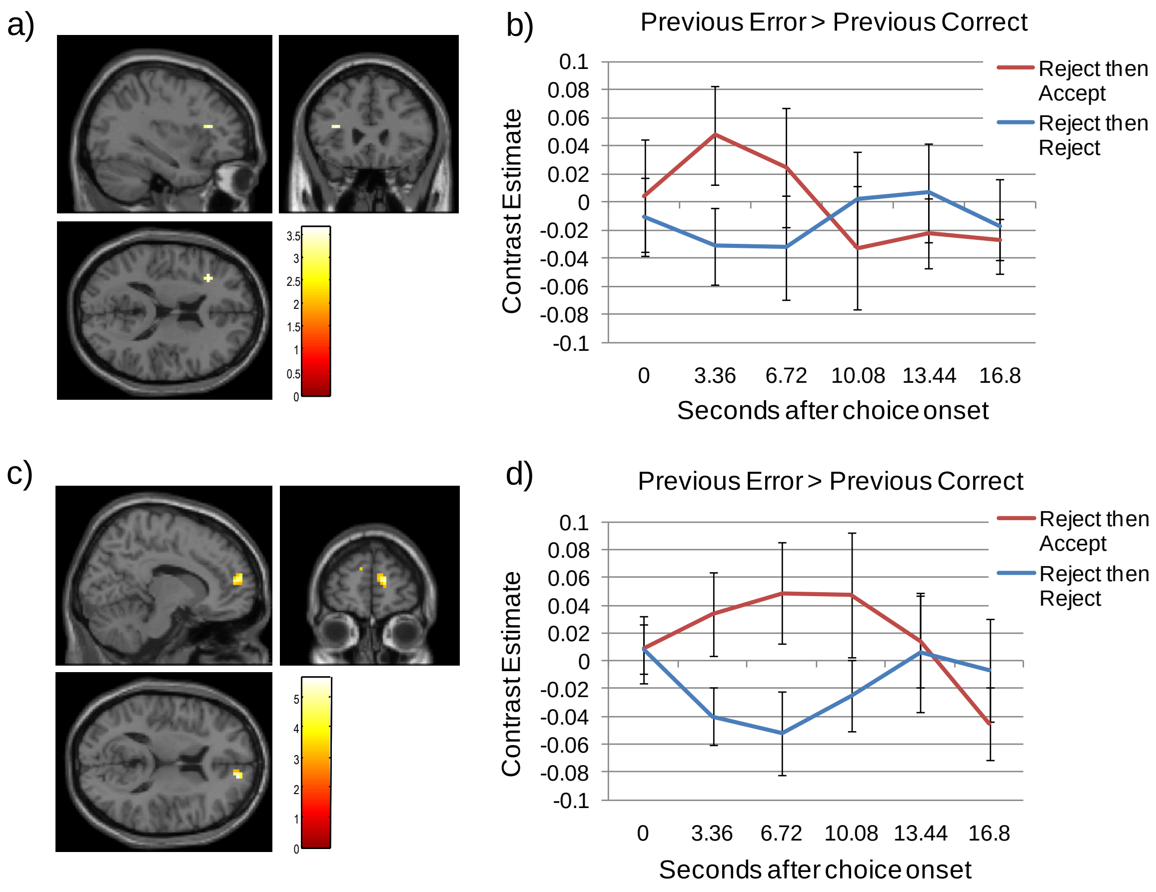

Figure 7-4 – Group SPM data showing responses at the time of subsequent choice, thresholded at p<0.005 for display purposes, shown on a normalised canonical template brain. a) and b) show anterior insula activity reflecting an interaction of previous choice and outcome with current choice (MNI peak -33, 26, 16) (i.e. (Reject [error - correct] > Accept [error - correct]) only when the current choice is to Accept). c) and d) show the same 3-way interaction within medial prefrontal cortex (peak MNI 15, 56, 13).

Table 7-2 - The table shows activation for the main effects of reject and accept at choice, and the interaction of these choices with the choice and outcome on the previous trial. Contrasts were performed on responses modelled at the time point of subsequent trial choice.

Brain regions	MNI coordinates of local maxima	Voxel number	Voxel t score
Reject > Accept			
R cerebellum	24, -52, -23	94	10.17
L postcentral	-48, -25, 52	190	9.97
L supramarginal	-57, -34, 28	9	8.81
R supplementary motor area (BA 6)	15, 2, 67	4	8.80
R insula	39, 5, 1	41	8.72
R supramarginal (BA 40)	45, -34, 40	44	8.66
L insula	-42, 11, -8	2	8.64
L cerebellum	-27, -52, -26	32	8.56
Posterior cingulate	-12, -28, 46	7	8.43
Middle cingulate (BA 32)	-9, 14, 34	5	8.41
Accept > Reject			
Rostral anterior cingulate	6, 35, -11	30	4.82
Ventral striatum	0, 8, -8	6	4.19
Reject [error > correct] – Accept [error > correct] in subsequent accept – reject			
Medial prefrontal cortex/ACC	15, 53, 13	30	5.90
L anterior insula	-30, 26, 13	-	3.66

7.5 Discussion

While a status-quo bias was previously found to arise for difficult perceptual decisions (Fleming et al., 2010), it was unclear *why*, in such situations of uncertainty participants are driven to stay with a status-quo, rather than change to a different option. Here, I hypothesised that erroneous rejection of the status-quo may lead to greater regret than erroneous acceptance of the status-quo, both psychologically and neurobiologically, and that such an asymmetry can be a key driver of a status-quo bias. Thus greater regret after status-quo rejection may reduce its justifiability and appeal, and encourage future decisions that align with a status-quo option.

In keeping with this perspective, here I show erroneous status-quo rejection is associated with stronger feelings of regret than erroneous status-quo acceptance, with anterior insula and medial prefrontal cortex (mPFC) BOLD signals also showing an inflated response specifically for erroneous status-quo rejections (as opposed to erroneous status-quo acceptance). Both regions have been previously implicated in error-processing (e.g. Carter et al., 1998; Klein et al., 2007; Menon et al., 2001; Ullsperger & Von Cramon, 2004), but hitherto without any consideration of status-quo rejection versus status-quo acceptance.

A particular role of the insula in interoceptive awareness (Critchley, Mathias, & Dolan, 2002; Critchley, Wiens, Rotshtein, Öhman, & Dolan, 2004), and more generally subjective feeling states (Craig, 2002), has shaped theories of insula involvement towards the detection of external threat stimuli or loss, and for awareness of resulting physiological arousal. Errors that people become aware of are associated with stronger insula activity compared to those of which we are unaware, in keeping with an increased autonomic response to the former errors (Klein et al., 2007). I did not find evidence for increased physiological arousal, as assessed with SCRs during status-quo rejection versus status-quo acceptance errors. Although this null outcome should be interpreted with caution, I did find significant effects for error versus correct feedback overall, just no selectivity for status-quo acceptance versus rejection. This hints that the insula activation I highlight is likely to represent more than just increased physiological arousal. Apart from a suggested role in error-processing, the insula (and mPFC) is more generally implicated in reflection on one's own (or others') performance (Bengtsson, Lau, & Passingham, 2009; Kelley et al., 2002; Mitchell, Banaji, & MacRae, 2005; Ochsner et al., 2004). Despite the psychological literature consistently showing that errors stemming from decisions to reject a status-quo are not experienced as equivalent to those stemming from a decision to accept a status-quo, to my knowledge the present study is the first to address error-processing in the brain in terms of such differences.

Participants also showed a clear overall bias towards accepting the status-quo. Previous literature suggests that past experience of higher regret for status-quo rejection errors encourages higher anticipated regret for the prospect of a similar status-quo rejection error in the future (Baron & Ritov, 1994; Kordes-de Vaal, 1996; Ritov & Baron, 1990, 1995; Tykocinski, Israel, & Pittman, 2004; Tykocinski & Pittman, 1998, 2001). Amplified emotional and neuronal responses to erroneous status-quo rejection, compared to acceptance, could lead decision-makers to consider accepting the status-quo as the more justifiable choice, thus a subsequent status-quo bias (see Connolly & Zeelenberg, 2002; Inman & Zeelenberg, 2002; Zeelenberg et al., 2002). Here, I found such a decision bias to be associated with activity in the insula and mPFC, regions which also show an inflated response to status-quo rejection errors at the time of such a preceding outcome.

Existing evidence suggests that the insula is involved in processing information necessary for learning and goal-directed behaviour (e.g. risk and uncertainty, Huettel, Stowe, Gordon, Warner, & Platt, 2006; Paulus, Rogalsky, Simmons, Feinstein, & Stein, 2003; Preuschoff, Quartz, & Bossaerts, 2008). The insula is critically involved in the anticipation of future aversive outcomes, which is important for avoidance learning (Ploghaus et al., 1999). It is also an important component of the neurocircuitry for stress, fear and anxiety, showing heightened activity during experience and anticipation of aversive events in people suffering from anxiety disorders (Shin & Liberzon, 2009). Its involvement in anxiety disorders is suggested to be associated with inappropriately high prediction of future aversive events (Paulus & Stein, 2006). In healthy decision-making, the insula is particularly found to be involved in post-error behavioural modification and response slowing (Clark et al., 2008; Hester, Barre, Mattingley, Foxe, & Garavan, 2007; Kuhnen & Knutson, 2005; O'Doherty, Critchley, Deichmann, & Dolan, 2003; Paulus et al., 2003). Singer, Critchley, & Preuschoff (2009) provide a functional model of the insula, that postulates a role in integrating internal physiological and external sensory information, along with their associated uncertainty, contextual information and individual preferences, to motivate adaptive subsequent behaviour. The results from this chapter are in keeping with a role of the insula in post-error behavioural adaptation, in this case encouraging a status-quo bias in response to previously amplified emotional (and neuronal) responses to erroneous status-quo rejection in particular.

In this study, I have considered the status-quo to be the previous decision, in accord with literature on the functional role of regret in consumer decision-making (Inman & Zeelenberg, 2002; Tsiros & Mittal, 2000). In such a case, an inaction bias is implied, since remaining with the current decision involves inaction. Under certain other conditions, however, status-quos can involve action (e.g. a goalkeeper may have a stronger bias to dive to save the ball than to remain stationary). It is currently unclear whether an inaction bias and a status-quo bias are independent, or are rather driven by the same underlying cause (Anderson, 2003; Baron & Ritov, 1994; Ritov & Baron, 1992; Schweitzer, 1994). In particular, the likelihood of regret may be the driving force behind *both* biases, so the two could be considered as a unitary construct for the purpose of this study. Unfortunately, here and in Fleming et al (2010), these two possible forms of bias (status-quo and inaction) are synonymous by design. It remains possible, however, that an inaction bias arises primarily or exclusively at the level of the motor response, with the effort needed to act or switch behaviour a salient feature driving the asymmetric impact of erroneous decisions to accept and reject the status-quo. Status-quo biases, on the other hand, may arise at the more abstract level of the decision, with features such as the normality and perceived causality of the behaviour playing a greater role. Indeed, the fact that I found a tendency for participants to repeat their previous form of response (i.e. accept or reject), hints at the presence of another form of bias, perhaps at a higher hierarchical level than the bias to repeat the previous IN or OUT response. This form of bias (or perseveration in choice) appears to be encouraged when the same strategy succeeded in the past, as we found it to be strongest after correct trials (see Figure 7-3). More broadly, I note that perseveration is increasingly incorporated into formal models of learning and decision-making (e.g. Schonberg, Daw, Joel, & O'Doherty, 2007). Future research could address possible differences in the way the various forms of bias – inaction, status-quo, and strategy perseveration – are processed in the brain.

In conclusion, I suggest that inflated emotional and neurobiological responses to status-quo rejection errors (as compared with status-quo acceptance errors) may be a key contributor to the emergence of a status-quo bias. In support of this, I show enhanced error-related responses in anterior insula and mPFC, as well as enhanced subjective feeling of regret, for status-quo rejection errors. The observed asymmetries in neural and emotional responses predict a bias towards the status-quo in subsequent decision-making. This regret-induced status-quo bias may also underlie the choice repetition effects observed in the previous two chapters.

Chapter 8. General discussion

8.1 Summary and synthesis of results

In this thesis, I have addressed the experiential *content* of regret, the external and psychological *antecedents* that are necessary for it to be elicited, and its *motivational* impact. Using fMRI, I examined the *neural mechanisms* involved at each of these levels of processing. In addition, three further themes flow through this thesis, as discussed in Chapter 1. The first concerns identifying any *necessary* precursors of regret that are currently absent from regret-based economic models, with a special focus on the decision-maker's attribution of responsibility. Secondly, I question whether the actual *experience* of regret may be associated with goals that are distinct from those associated with its anticipation, with these unique goals manifesting in behaviour that cannot be easily explained by existing models of anticipated regret. The third concerns behavioural and cognitive strategies that may be used to *regulate* the experience of regret. I suggest that emotional regulation may need to be taken into account when addressing the experience and motivational impact of any human aversive experience. Below I discuss how each of the five studies, that comprise this thesis, add to our understanding the role of regret and responsibility in decision-making.

8.1.1 *Responsibility as a necessary antecedent to regret*

The results of Study 1 (Chapter 3) indicate that self-blame should be considered a vital antecedent to the experience and neuronal expression of regret. In this study, trial-by-trial subjective ratings of regret were found to depend on subjective and objective responsibility. Using fMRI, I showed a key role for the amygdala in this self-blame-dependent regret and for the angular gyrus in attribution of responsibility. The pattern of ventral striatum activity in Study 4 (Chapter 6) also supports a view that the way the brain responds to regret-related outcomes is critically dependent on a decision-maker having a sense of agency, as does the behavioural response to such outcomes. Moreover, the results of Study 1 provide no evidence for what Connolly & Zeelenberg (2002) term “outcome regret”, i.e. showing invariant psychological or brain responses to regret-related outcomes under *all* levels of responsibility. Following Theme 1 of this thesis, these data support responsibility as a *necessary* precursor to the experience of regret, alongside a between-option upward counterfactual comparison process.

Curiously, self-blame does not necessarily make the experiences more aversive. Externally enforced outcomes may actually encourage a stronger emotional response when one feels hard done by. In contrast to previous studies reporting involvement of OFC and cingulate cortex in regret (Camille et al., 2004; Chandrasekhar, Capra, Moore, Noussair, & Berns, 2008; Chua, Gonzalez, Taylor, Welsh, & Liberzon, 2009; Coricelli et al., 2005; Liu et al., 2007), in Study 1 these regions actually show enhanced response to regret-related outcomes when participants were *not* objectively responsible. This provokes a reconsideration of the roles of these brain responses in regret, where activity may in fact be associated with externally attributed frustration, helplessness or denial of responsibility.

Study 5 (Chapter 7) also adds to our understanding of the antecedents of regret. Previous work has shown that regret depends not just on the consequences of our behaviour, but also on how easily the behaviour itself can be justified post-hoc. An example, supported by the findings in Study 5, is that errors from action (or decisions to deviate from a status-quo or norm) result in greater regret compared to errors from inaction. Decisions that deviate from a norm or status-quo may induce greater regret because they are less easily justified (Connolly & Zeelenberg, 2002). Consequently, these results encourage us to embrace perceived low justifiability of behaviour as another necessary antecedent to the experience of regret. As discussed below, low perceived justifiability may be learnt from previous experiences of this asymmetry in regret, and may encourage a future bias towards accepting the status-quo. Alternatively actions may be perceived, by the decision-maker, as more directly *causal* of its outcome than inaction, thereby amplifying a sense of accountability or blame for the error (Ritov & Baron, 1990; Spranca, Minsk, & Baron, 1991). With self-blame as an important precursor to the experience of regret, this may account for the action-inaction asymmetry in regret.

8.1.2 Responsibility improves learning of action-outcome contingencies

Study 2 (Chapter 4) was motivated by the possibility that internal attribution of responsibility may improve learning of action-outcome contingencies compared to when we are not responsible. This may be due to of a greater incentive to learn optimal future behaviours, stemming from the direct personal impact of these outcomes. An enhanced focus on between-option counterfactual comparisons, associated with regret, might also lead to more efficient learning from the outcomes of actions taken and those (fictive actions) not taken. In Study 2, I compared the efficacy of value learning by active trial-and-error against vicarious learning through observation of the outcomes of actions performed by others. Here I showed that trial-and-error learning is consistently more accurate than passive observational learning. Moreover, a model comparison suggested that this asymmetry in learning is associated with greater attention to fictive information at a mechanistic level. Although this study did not directly tap into the motivational effects of regret, it goes some way to suggest that two core components of regret (i.e. responsibility and between-option comparisons) can be beneficial for learning and decision-making. In general, the results also suggest that the descriptive and predictive value of social learning theories and reinforcement learning models may benefit from taking into account the unique mode through which action-outcome contingencies are learned, including the subtle biasing influence of social comparison.

From the design of Study 2, we cannot be sure that improved learning is directly associated with regret. Although active learning benefits from an enhanced attention to fictive information, this does not mean that learners are motivated specifically by an aim of avoiding *upward* between-option counterfactual comparison. In other words, it is not clear that active learners are motivated to minimise regret rather than to maximise relief. Future studies could develop this task design, so as to more explicitly test whether actors learn better to be *regret-averse* than observers, e.g. by providing one option on each trial which acts as a regret-minimizing choice. On the other hand, the pattern of learning deficit shown by observers does hint at a relative absence of regret-aversion. Passive observational learning was found to be associated with a biased appreciation of the likelihood of receiving a bad outcome, and this was associated with a relative inattention to fictive information. This poor observational learning manifests exclusively as optimistic over-valuation of *low value* options. In other words, observers underestimate the likelihood of experiencing negative events as observed occurring to others, a bias not present in actors' direct learning by trial-and-error. I suggest that observers' inattention to fictive information and lack of personal involvement with the outcomes of the actors' decisions means that they are not motivated by a possibility of committing later bad decisions. In contrast, active learners are motivated by anticipated regret during learning, founded in their greater attention to fictive outcomes, between-option counterfactual comparison and agency, which motivates more efficient and accurate learning, especially for low value options.

Improved learning with personal responsibility may also explain why, in Study 1, self-blame regret is seemingly a less aversive experience than feelings of helplessness, frustration or anger associated with externally induced bad outcomes. This may be because perceived control over ones future circumstances, and the potential to improve them, reduces the pain associated with a current mistake. In contrast, helplessness has been shown in animal models to interfere with avoidance learning, while also producing stress and reactive depression in both animals and humans (cf. review of learned helplessness, by Seligman, 1972).

8.1.3 The motivational impact of experienced regret

In the remaining three studies, I measured trial-by-trial influences of regret-related outcomes on subsequent behaviour, with the aim to provide a better understanding of the motivational impact of experienced regret. In Studies 3 and 4 (Chapters 5 & 6) I found a tendency for decision-makers to repeat, rather than avoid, previously regrettable choices. A role of the striatum *both* in the experience and the behavioural impact of such outcomes, shown in Study 4, is in keeping with previous suggestions that regret-related learning signals are processed in the striatum (e.g. Lohrenz et al., 2007; Sommer et al., 2009), although the present findings suggest the two processing levels involve different striatal structures. This activity may be related to previous findings of increased striatal response and increased desire to continue gambling associated with near misses (Clark, Lawrence, Astley-Jones, & Gray, 2009), and may be associated with higher sensitivity to anticipated future wins that may compensate for the previous outcome. The behaviour is moreover in keeping with suggestions that 'chasing' behaviour (or repeated gambling) in response to near-misses provides gamblers with the opportunity to eliminate the impact of regret by potentially making up for their mistakes (Loftus & Loftus, 1983).

From this literature, one explanation for the tendency to repeat previously regret-related choices is that experienced regret encourages a desire to make up for a previous mistake. If true, this explanation may be important for understanding the unique motivations associated with current experiences of regret, compared to its anticipation. While anticipated regret motivates us to adjust behaviour so as to avoid possible future regret, current experiences of regret are unpleasant and our strongest motivation may be to regulate or reduce this experience rather than to plan for future avoidance. Others have found that, while fear and anxiety (associated with anticipated emotion) promote safety-seeking behaviour, anger and sadness (associated with present threat) promote high-risk/high-reward behaviours, in keeping with a goal of "reward replacement" (Lerner & Keltner, 2000, 2000). Similarly, while fear or anticipatory anxiety in relation to possible future regret may initiate a form of *flight* response (i.e. regret-avoidance), experienced regret may encourage active attempts to reduce or undo the aversive feeling, or to repair the damage with a compensatory good choice or outcome. The latter bears closer analogy to a *fight* response. It may also reflect a form of Loewenstein & Lerner's (2003) *direct* impact of current emotional state, which putatively occurs independently of ones expectations about future outcomes and their probabilities (i.e. an *indirect* impact).

Metcalfe & Mischel's (1999) comparison of hot versus cool processing may also be important in the distinction between a direct and indirect impact of regret. The putative “hot system” is emotional, simple, reflexive, fast, stimulus-controlled and accentuated by stress. The “cool system”, on the other hand, is cognitive, complex, reflective, slow, controlled and attenuated by stress. In Studies 3 and 4, the *hot* behavioural effects of experienced regret may take precedent over the *cool* anticipation of future regret, because of the immediacy of the current experience and the ambiguity surrounding future outcome probabilities. The cool system is thought to depend on the later development of executive control mediated via the prefrontal cortex, while the hot system is though to be mediated by phylogenetically and ontogenetically older systems, involving limbic brain structures such as the amygdala and the striatum. It is possible that experienced regret, mediated by a hot processing mechanism in the striatum, influences behaviour independent of any anticipated regret learnt from it, which may be separately processed in the prefrontal cortex (especially the OFC). That the expected behavioural effects of anticipated regret on decision-making were not observed in these studies (i.e. towards regret-minimizing or regret-avoidant choices) may explain why we see involvement of the striatum, rather than the OFC.

The results of Study 5 (Chapter 7) indicate that the regret-related choice repetition found in the aforementioned studies might reflect a regret-induced status-quo bias. Rather than a *flight* or *fight* response, this behaviour may reflect a form of *freeze* response. In Study 5, I show that experienced regret is higher after an erroneous decision to reject a previous action, compared to an erroneous decision to accept it. The fMRI data reveal that anterior insula and medial prefrontal cortex show increased BOLD signal after such status-quo rejection errors, mirroring the asymmetry in the subjective experience of regret. A similar pattern of signal change was associated with accepting the status-quo on a subsequent trial, thereby linking error-related activity on the previous trial to a subsequent behavioural bias towards accepting a status-quo option. As such, these results point to the status-quo bias being driven by previous mistakes that have induced a greater intensity of experienced regret. It is unclear, however, if the behavioural bias is a direct result of experienced regret, or if it is promoted by greater anticipated regret for subsequent status-quo rejection errors. Zeelenberg & Pieters (2007) highlighted decision delay or avoidance as a means of preventing or reducing future regret, while sticking with a status-quo may also increase the justifiability of the subsequent decision (Chapter 1, Table 1.1). A regret-induced status-quo bias may be more in keeping, then, with a response to anticipated rather than experienced regret.

A further question that might be addressed in future studies of this regret-induced status-quo bias concerns the temporal pattern on the action-inaction asymmetry in experienced regret. Gilovich & Medvec (1994) found that, while actions are subject to the emotion of regret most immediately following a mistake, inactions are regretted more over a longer temporal horizon. The authors suggest that various factors decrease the pain of regretful actions over time, including behavioural steps to make up for a mistake, psychological attempts at dissonance reduction and the identification of 'silver linings'. On the other hand, factors that increase the pain of regrettable inactions over time include gradually increased confidence in ones ability to act, less salient evidence for justification of inaction compared to action, and the open-endedness of inaction-induced regret that encourages its growth with time. Future studies could explore how the regret-induced status-quo bias, and the neuronal responses associated with it, are modulated over time.

We cannot change previous outcomes, as much as we may strive to, and so it could be considered irrational to use past experiences to steer future choices. On the other hand, it is rational to take them into account if these experiences contribute to improved learning from our mistakes, have direct impact on our goals, vigilance or effort, or indirect effects on our future expectations. While Study 2 points to a beneficial impact of regret on learning, Studies 3-5 suggest that experienced regret can also promote inertia or behaviour akin to repetitive gambling which may have negative effects on our wellbeing. The different results here may be due to an increased uncertainty regarding future outcome probabilities in studies 3-5, compared to study 2.

8.1.4 Experienced regret promotes strategies of regret-regulation

The third theme of this thesis concerns the various behavioural and cognitive strategies that can be used to regulate the aversive experience of regret. Zeelenberg & Pieters (2007) suggest that one such strategy is to deny or defer responsibility for the mistaken action/decision (Chapter 1, Table 1.1, II.1.c – "Deny responsibility for the decision"). In Study 1, individuals were found to actively *reduce* their feeling of responsibility for relatively worse outcomes, in a manner akin to such a regulation strategy.

These regulatory strategies may also alter the effect of regret on future behaviour. I have suggested that regret-related choice repetition, found in Studies 3, 4 (and perhaps in a different way in Study 5) is best explained by a regulatory strategy associated with minimising experienced regret. Specifically, this behaviour could be associated with an attempt either to make up for a previous mistake (Chapter 1, Table 1.1, II.1.a – “Undo decision”), or to improve the justifiability of the choice by making it a norm or status-quo (Chapter 1, Table 1.1, II.1.b – “Justify decision”). The results of Study 5, also suggest that similar behaviour may stem from a strategy of delaying or avoiding a subsequent decision, as a means of avoiding possible future regret (Chapter 1, Table 1.1, I.1.d – “Delay or avoid decision”). If the active learning measured in Study 2 is assumed to elicit experiences of regret, then the results of this study also emphasize the benefits of regret for improving the quality of learning and decision-making (Chapter 1, Table 1.1, I.1.a – “Increase decision quality”).

Speculatively, these regulatory responses may stem from a direct impact of experienced regret, and may prevail when avoidance of future regret is not an option (i.e. when it is not possible to follow Chapter 1, Table 1.1, II.2.a – “Reverse decision (switch to alternative)”). It is possible that these strategies come from a higher cognitive level and may, to some extent, be independent of a more primitive valuation system. If this is the case, the two systems may run in parallel so that a regulated conscious feeling of regret may not interfere with an unconscious impact of regret on future choice behaviour. In keeping with this, the fMRI results from Study 4 suggest that regret-related activity (in ventral striatum) is not affected by a decision-maker's subsequent behaviour. This provides little evidence for individuals suppressing or denying regret (Chapter 1, Table 1.1, II.3.b) or reappraising the quality of the alternative option (Chapter 1, Table 1.1, II.2.b) at the level of the initial valuation.

If individuals regulate their aversive experiences of regret, an important question concerns whether regret can ever be truly measured empirically before it is actively suppressed. This is an issue that we, as regret researchers, must take into account. Future work should directly explore which features of regret are *unaffected* by behavioural and cognitive strategies, and so can be measured regardless. It is also conceivable that, by empirically measuring the regulatory strategies themselves, perhaps through the observable responses of decision-makers, we can incorporate these strategies into future models of regret and its impact on decision-making.

8.2 Implications of these findings

The findings presented in this thesis not only deepen our understanding of the experience of regret and its role in decision-making, but also raise issues for how regret can be approached in future research. By developing a fuller understanding of the situational factors and internal appraisals that elicit regret, we can then better control and/or manipulate these conditions empirically. Traditional learning models update the value of a choice as a function of either standard or regret-based (or fictive) prediction-errors (e.g. Lohrenz et al., 2007; Marchiori & Warglien, 2008; Rescorla & Wagner, 1972). On such accounts the value of a choice option should be incrementally *decreased* for negative or regrettable outcomes, and a regret-averse decision-maker should then avoid choice options that are considered most likely to induce regret. Therefore, such models currently provide a poor fit to an observed tendency to repeat previously regret-related gambles, shown in studies 3, 4 and 5 of this thesis. The descriptive and predictive value of such models may benefit from taking into account an influence of perceived responsibility and justifiability, the different roles of experienced and anticipated regret (where only the latter may be a good predictor of regret-avoidant behaviour), and an influence of regret-regulatory strategies that modify our response to *experienced* regret.

Understanding subjective responsibility and blame may have wider implications in formulating therapies for painful and debilitating life regrets. An ability to externally shift responsibility for our bad choices can reduce aversive feelings of regret (Zeelenberg & Pieters, 2007), while understanding that others have made similarly bad choices provide justification for our actions by placing us within the social norm (Connolly & Zeelenberg, 2002). On the other hand, an ability to accept responsibility for our actions can motivate adaptive future behaviours, leading to long-term improvements in quality of life. Decision-making in business, politics, the legal world, health care, and in personal relationships may also be improved by an ability to effectively anticipate levels of self-blame. Indeed, anticipating future regret has been shown to have positive influences on decisions in sexual behaviour (e.g. Richard, Van der Pligt, & de Vries, 1996), consumer choices (e.g. Inman & Zeelenberg, 2002; Tsiros & Mittal, 2000), and health related choices (e.g. Lechner, de Vries, & Offermans, 1997). Reb (2008) also found that anticipating regret can improve the quality of the decision process by increasing vigilance and care.

The relationship between regret and choice-repetition found in this thesis has potential implications for understanding the “chasing” behaviour exhibited by some problem gamblers. In a large sample of horse races, gamblers tended to go for long shot gambles (low probability of high possible gain) on the last race of the day, as if trying to break even or to make right earlier losses (McGlothlin, 1956). Losing gamblers also tended to increase their bets more than did winning gamblers. Anticipated regret is also found to significantly increase intentions to play the lottery (Sheeran & Orbell, 1999; Zeelenberg & Pieters, 2004). Such behaviour may provide gamblers with an opportunity to modulate current regret by making up for their mistakes. Evidence for a link between compulsive gambling and dopamine agonist treatment of Parkinson patients (Molina et al., 2000) supports a possibility that dopaminergic projections to striatum may play a critical role here, which is in keeping with the role of the dorsal striatum in comparable behaviour in Study 4. Moreover, the experience of “near misses” in gambling has been proposed as having a similar conditioning effect on future behaviour as experiencing a full win (Griffiths, 1990; Reid, 1986), while also being associated with strong feelings of regret (Kahneman & Tversky, 1982a). Near misses have recently been associated with increased striatum activity as well as with a simultaneously increased desire to continue gambling, despite these outcomes being rated as more unpleasant than full misses (Clark et al., 2009; see also Clark, 2010). Some gambling tools are now designed to allow for a higher than chance frequency of near misses (Reid, 1986), suggesting that casinos may apprehend the role that anticipated regret and near misses play in reinforcing continued gambling, and that they employ this for their own advantage.

This thesis takes a step in bridging psychological, economic and neuroimaging approaches to studying regret. Taken together, the findings provide some clarification as to which external and internal events cause us to experience regret and how we may respond to it, optimally or otherwise. As discussed above, the findings have important theoretical and practical implications for understanding the role of regret in promoting gambling-like behaviour in healthy or problematic gamblers or in patients receiving dopaminergic neuromodulatory therapy for Parkinson's disease. Appreciation of the role of regret and attribution of blame in optimal learning and decision-making, as well as an ability to predict how regret influences individual decisions, likely has widespread practical implications in, for example, policy development, healthcare, marketing, advertisement and investment. Furthermore, understanding the importance of regret-regulatory strategies in personal quality of life, and appreciating how such strategies can be impaired, may help in developing therapies for depression and anxiety disorders.

References

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine. *Econometrica: Journal of the Econometric Society*, *21*(4), 503–546.

Anderson, A. K., Christoff, K., Stappen, I., Panitz, D., Ghahremani, D. G., Glover, G., Gabrieli, J. D. E., et al. (2003). Dissociated neural representations of intensity and valence in human olfaction. *nature neuroscience*, *6*(2), 196–202.

Anderson, C. J. (2003). The psychology of doing nothing: Forms of decision avoidance result from reason and emotion. *Psychological Bulletin*, *129*(1), 139–166.

Andersson, J. L., Hutton, C., Ashburner, J., Turner, R., & Friston, K. (2001). Modeling geometric deformations in EPI time series. *Neuroimage*, *13*(5), 903–919.

Arana, F. S., Parkinson, J. A., Hinton, E., Holland, A. J., Owen, A. M., & Roberts, A. C. (2003). Dissociable contributions of the human amygdala and orbitofrontal cortex to incentive motivation and goal selection. *Journal of Neuroscience*, *23*(29), 9632–9638.

Aron, A., Fisher, H., Mashek, D. J., Strong, G., Li, H., & Brown, L. L. (2005). Reward, Motivation, and Emotion Systems Associated With Early-Stage Intense Romantic Love. *J Neurophysiol*, *94*(1), 327-337.

Ashburner, J., & Friston, K. J. (2005). Unified segmentation. *Neuroimage*, *26*(3), 839–851.

Bach, D. R., Daunizeau, J., Kuelzow, N., Friston, K. J., & Dolan, R. J. (2010).

Dynamic causal modelling of spontaneous fluctuations in skin conductance. *Psychophysiology*, *in press*.

Bach, D. R., Flandin, G., Friston, K. J., & Dolan, R. J. (2010). Modelling event-related skin conductance responses. *International Journal of Psychophysiology*, *in press*.

Bach, D. R., Grandjean, D., Sander, D., Herdener, M., Strik, W. K., & Seifritz, E. (2008). The effect of appraisal level on processing of emotional prosody in meaningless speech. *Neuroimage*, *42*(2), 919–927.

Bach, D. R., Flandin, G., Friston, K. J., & Dolan, R. J. (2009). Time-series analysis for rapid event-related skin conductance responses. *Journal of Neuroscience Methods*, *184*(2), 224-234.

Baird, A. A., & Fugelsang, J. A. (2004). The emergence of consequential thought: evidence from neuroscience. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *359*(1451), 1797.

Bandura, A. (1971). Vicarious and self-reinforcement processes. In *The nature of reinforcement* (pp. 228–278). New York and London: Academic Press Inc.

Bandura, A. (1977). *Social learning theory*. General Learning Press.

Baron, J., & Ritov, I. (1994). Reference points and omission bias. *Organizational Behavior and Human Decision Processes*, *59*, 475–475.

Barto, A. G., & Sutton, R. S. (1998). *Reinforcement Learning: an introduction*. MIT Press Cambridge, MA.

Bechara, A., Tranel, D., & Damasio, H. (2000). Characterization of the decision-making deficit of patients with ventromedial prefrontal cortex lesions. *Brain*, *123*(11), 2189.

Beckman, L. (1970). Effects of students' performance on teachers' and observers' attributions of causality. *Journal of Educational Psychology*, *61*(1), 76–82.

Bell, D. S. (1965). Comparison of amphetamine psychosis and schizophrenia. *The British Journal of Psychiatry*, *111*(477), 701.

Bell, D. E. (1982). Regret in Decision Making under Uncertainty. *Operations Research*, *30*(5), 961-981.

Bengtsson, S. L., Lau, H. C., & Passingham, R. E. (2009). Motivation to do well enhances responses to errors and self-monitoring. *Cerebral Cortex*, *19*, 797-804.

Berns, G. S., McClure, S. M., Pagnoni, G., & Montague, P. R. (2001). Predictability modulates human brain response to reward. *Journal of Neuroscience*, *21*(8), 2793.

Blackwood, N. J., Bentall, R. P., Ffytche, D. H., Simmons, A., Murray, R. M., & Howard, R. J. (2003). Self-responsibility and the self-serving bias: an fMRI investigation of causal attributions. *Neuroimage*, *20*(2), 1076–1085.

Boles, T. L., & Messick, D. M. (1995). A reverse outcome bias: The influence of multiple reference points on the evaluation of outcomes and decisions. *Organizational Behavior and Human Decision Processes*, *61*(3), 262–275.

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review, 108(3),* 624–652.

Boyd, R., & Richerson, P. J. (1988). An evolutionary model of social learning: the effects of spatial and temporal variation. *Social learning: Psychological and biological perspectives*, 29–48.

Braver, T. S., Barch, D. M., Gray, J. R., Molfese, D. L., & Snyder, A. (2001). Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. *Cerebral Cortex*, *11*(9), 825.

Breiter, H. C., Aharon, I., Kahneman, D., Dale, A., & Shizgal, P. (2001). Functional imaging of neural responses to expectancy and experience of monetary gains and losses. *Neuron*, *30*(2), 619–639.

Büchel, C., Morris, J., Dolan, R. J., & Friston, K. J. (1998). Brain systems mediating aversive conditioning: an event-related fMRI study. *Neuron*, *20*(5), 947–957.

Burger, J. M., & Palmer, M. L. (1992). Changes in and generalization of unrealistic optimism following experiences with stressful events: Reactions to the 1989 California earthquake. *Personality and Social Psychology Bulletin*, *18*(1), 39.

Bush, G., Vogt, B. A., Holmes, J., Dale, A. M., Greve, D., Jenike, M. A., & Rosen, B. R. (2002). Dorsal anterior cingulate cortex: a role in reward-based decision making. *Proceedings of the National Academy of Sciences*, *99*(1), 523.

Cahill, L., Babinsky, R., Markowitsch, H. J., & McGaugh, J. L. (1995). The amygdala and emotional memory. *Nature*, *377*, 295-296

Camerer, C., & Ho, T. H. (1999). Experienced-weighted attraction learning in normal form games. *Econometrica*, *67*(4), 827–874.

Camille, N., Coricelli, G., Sallet, J., Pradat-Diehl, P., Duhamel, J., & Sirigu, A. (2004). The Involvement of the Orbitofrontal Cortex in the Experience of Regret. *Science*, *304*(5674), 1167-1170.

Carlsson, A. (1995). The dopamine theory revisited. *Schizophrenia. Oxford: Blackwell Science*, 379–400.

Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior cingulate cortex, error detection, and the online monitoring of performance. *Science*, *280*(5364), 747.

Chandrasekhar, P. V., Capra, C. M., Moore, S., Noussair, C., & Berns, G. S. (2008).

Neurobiological regret and rejoice functions for aversive outcomes. *NeuroImage*, *39*(3), 1472–1484.

Chua, H. F., Gonzalez, R., Taylor, S. F., Welsh, R. C., & Liberzon, I. (2009). Decision-related loss: Regret and disappointment. *NeuroImage*, *47*(4), 2031-2040.

Church, R. M. (1959). Emotional reactions of rats to the pain of others. *Journal of Comparative and Physiological Psychology*, *52*(2), 132–134.

Cialdini, R. B., Braver, S. L., & Lewis, S. K. (1974). Attributional bias and the easily persuaded other. *Journal of Personality and Social Psychology*, *30*(5), 631–637.

Clark, L., Bechara, A., Damasio, H., Aitken, M. R. F., Sahakian, B. J., & Robbins, T. W. (2008). Differential effects of insular and ventromedial prefrontal cortex lesions on risky decision-making. *Brain*, *131*(5), 1311-1322.

Clark, L., Lawrence, A. J., Astley-Jones, F., & Gray, N. (2009). Gambling near-misses enhance motivation to gamble and recruit win-related brain circuitry. *Neuron*, *61*(3), 481–490.

Clark, L. (2010). Decision-making during gambling: an integration of cognitive and psychobiological approaches. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *365*(1538), 319-330. doi:10.1098/rstb.2009.0147

Cohen, M. D. (2008). Economics: Learning with Regret. *Science*, *319*(5866), 1052.

Connolly, T., & Butler, D. (2006). Regret in economic and psychological theories of choice. *Journal of Behavioral Decision Making*, *19*(2), 148–158.

Connolly, T., Ordónez, L. D., & Coughlan, R. (1997). Regret and responsibility in the evaluation of decision outcomes. *Organizational Behavior and Human*

Decision Processes, *70*(1), 73–85.

Connolly, T., & Reb, J. (2005). Regret in cancer-related decisions. *Health psychology(Hillsdale, N. J.)*, *24*(4), 29–34.

Connolly, T., & Zeelenberg, M. (2002). Regret in decision making. *Current directions in psychological science*, *11*(6), 212–216.

Coricelli, G., Critchley, H. D., Joffily, M., O'Doherty, J. P., Sirigu, A., & Dolan, R. J. (2005). Regret and its avoidance: a neuroimaging study of choice behavior. *Nature neuroscience*, *8*(9), 1255–1262.

Craig, A. D. (2002). How do you feel? Interoception: the sense of the physiological condition of the body. *Nat Rev Neurosci*, *3*(8), 655-666.

Critchley, H. D., Mathias, C. J., & Dolan, R. J. (2001). Neural activity in the human brain relating to uncertainty and arousal during anticipation. *Neuron*, *29*(2), 537–545.

Critchley, H. D., Mathias, C. J., & Dolan, R. J. (2002). Fear Conditioning in Humans:: The Influence of Awareness and Autonomic Arousal on Functional Neuroanatomy. *Neuron*, *33*(4), 653–663.

Critchley, H. D., Mathias, C. J., Josephs, O., O'Doherty, J., Zanini, S., Dewar, B. K., Cipolotti, L., et al. (2003). Human cingulate cortex and autonomic control: converging neuroimaging and clinical evidence. *Brain*, *126*(10), 2139.

Critchley, H. D., Wiens, S., Rotshtein, P., Öhman, A., & Dolan, R. J. (2004). Neural systems supporting interoceptive awareness. *Nature neuroscience*, *7*(2), 189–195.

Croson, R., & Sundali, J. (2005). The gambler's fallacy and the hot hand: Empirical data from casinos. *Journal of Risk and Uncertainty*, *30*(3), 195–209.

Dawson, B. V., & Foss, B. M. (1965). Observational learning in budgerigars. *Animal*

Behaviour, *13*(4), 470–474.

Dayan, P., & Abbott, L. F. (2001). *Theoretical neuroscience* (Vol. 83). Citeseer.

Deichmann, R., Schwarzbauer, C., & Turner, R. (2004). Optimisation of the 3D MDEFT sequence for anatomical brain imaging: technical implications at 1.5 and 3 T. *Neuroimage*, *21*(2), 757–767.

Delgado, M. R., Miller, M. M., Inati, S., & Phelps, E. A. (2005). An fMRI study of reward-related probability learning. *Neuroimage*, *24*(3), 862–873.

Delgado, M. R., Nystrom, L. E., Fissell, C., Noll, D. C., & Fiez, J. A. (2000). Tracking the hemodynamic responses to reward and punishment in the striatum. *Journal of neurophysiology*, *84*(6), 3072.

Draganski, B., Kherif, F., Kloppel, S., Cook, P. A., Alexander, D. C., Parker, G. J., Deichmann, R., et al. (2008). Evidence for segregated and integrative connectivity patterns in the human basal ganglia. *Journal of Neuroscience, 28*(28), 7143.

Eisenberger, N. I., & Lieberman, M. D. (2004). Why rejection hurts: a common neural alarm system for physical and social pain. *Trends in Cognitive Sciences*, *8*(7), 294–300.

Elliott, R., Agnew, Z., & Deakin, J. F. W. (2008). Medial orbitofrontal cortex codes relative rather than absolute value of financial rewards in humans. *European Journal of Neuroscience*, *27*(9), 2213–2218.

Epstude, K., & Roese, N. J. (2008). The functional theory of counterfactual thinking. *Personality and Social Psychology Review*, *12*(2), 168.

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*(3).

Farrer, C., Franck, N., Georgieff, N., Frith, C. D., Decety, J., & Jeannerod, M. (2003). Modulating the experience of agency: a positron emission tomography study. *Neuroimage*, *18*(2), 324–333.

Farrer, C., & Frith, C. D. (2002). Experiencing oneself vs another person as being the cause of an action: the neural correlates of the experience of agency. *Neuroimage*, *15*(3), 596.

Feldman, J., Miyamoto, J., & Loftus, E. F. (1999). Are actions regretted more than inactions? *Organizational Behavior and Human Decision Processes*, *78,* 232–255.

Fleming, S. M., Thomas, C. L., & Dolan, R. J. (2010). Overcoming status quo bias in the human brain. *Proceedings of the National Academy of Sciences*, *107*(13), 6005-6009.

Fletcher, P. C., Happe, F., Frith, U., Baker, S. C., Dolan, R. J., Frackowiak, R. S. J., & Frith, C. D. (1995). Other minds in the brain: A functional imaging study of "theory of mind" in story comprehension. *Cognition*, *57*(2), 109–128.

Frijda, N. H., Kuipers, P., & Ter Schure, E. (1989). Relations among emotion, appraisal, and emotional action readiness. *Journal of Personality and Social Psychology*, *57*(2), 212–228.

Frith, C. D., Friston, K., Liddle, P. F., & Frackowiak, R. S. J. (1991). Willed action and the prefrontal cortex in man: a study with PET. *Proceedings: Biological Sciences*, *244*(1311), 241–246.

Fujiwara, J., Tobler, P. N., Taira, M., Iijima, T., & Tsutsui, K. I. (2008). Personality-dependent dissociation of absolute and relative loss processing in orbitofrontal cortex. *European Journal of Neuroscience*, *27*(6), 1547.

Gilovich, T., & Medvec, V. H. (1994). The temporal profile to the experience of regret. *Journal of Personality and Social Psychology*, *67*, 357–365.

Gilovich, T., Medvec, V. H., & Chen, S. (1995). Omission, commission, and dissonance reduction: Overcoming regret in the Monty Hall problem. *Personality and Social Psychology Bulletin*, *21*, 182–190.

Gilovich, T., & Medvec, V. (1995). The experience of regret: what, when, and why. *Psychological Review*, *102*(2), 379-395.

Gottfried, J. A., O'Doherty, J., & Dolan, R. J. (2003). Encoding Predictive Reward Value in Human Amygdala and Orbitofrontal Cortex. *Science*, *301*(5636), 1104-1107.

Greenwald, M. K., Cook, E. W., & Lang, P. J. (1989). Affective judgment and psychophysiological response: Dimensional covariation in the evaluation of pictorial stimuli. *Journal of Psychophysiology*, *3*(1), 51–64.

Griffiths, M. D. (1990). The cognitive psychology of gambling. *Journal of Gambling Studies*, *6*(1), 31–42.

Guttentag, R., & Ferrell, J. (2004). Reality compared with its alternatives: Age differences in judgments of regret and relief. *Developmental Psychology*, *40*(5), 764–775.

Haber, S. N., Fudge, J. L., & McFarland, N. R. (2000). Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *Journal of Neuroscience*, *20*(6), 2369.

Haber, S. N., & Knutson, B. (2009). The Reward Circuit: Linking Primate Anatomy and Human Imaging. *Neuropsychopharmacology*, *35*(1), 4-26.

Hart, S., & Mas-Colell, A. (2000). A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, *68*(5), 1127–1150.

Hayden, B. Y., Pearson, J. M., & Platt, M. L. (2009). Fictive reward signals in the anterior cingulate cortex. *science*, *324*(5929), 948.

Heider, F. (1958). *The psychology of interpersonal relations*. Wiley New York.

Helweg-Larsen, M. (1999). (The lack of) optimistic biases in response to the 1994 Northridge earthquake: The role of personal experience. *Basic and Applied Social Psychology*, *21*(2), 119–129.

Hester, R., Barre, N., Mattingley, J. B., Foxe, J. J., & Garavan, H. (2007). Avoiding another mistake: error and posterror neural activity associated with adaptive posterror behavior change. *Cognitive, affective & behavioral neuroscience*, *7*(4), 317.

Heyes, C., & Saggerson, A. (2002). Testing for imitative and nonimitative social learning in the budgerigar using a two-object/two-action test. *Animal Behaviour*, *64*(6), 851–859.

Heyes, C. M. (1994). Social learning in animals: categories and mechanisms. *Biological Reviews of the Cambridge Philosophical Society*, *69*(2), 207–232.

Heyes, C. M., & Dawson, G. R. (1990). A demonstration of observational learning in rats using a bidirectional control. *The Quarterly Journal of Experimental Psychology Section B*, *42*(1), 59–71.

Hirsch, C. R., & Mathews, A. (2000). Impaired positive inferential bias in social phobia. *Journal of Abnormal Psychology*, *109*(4), 705–712.

Ho, A. P., Gillin, J. C., Buchsbaum, M. S., Wu, J. C., Abel, L., & Bunney, W. E. (1996). Brain glucose metabolism during non-rapid eye movement sleep in major depression. A positron emission tomography study. *Archives of General Psychiatry*, *53*(7), 645-652.

Hooker, C., Roese, N. J., Park, S., Sledge, W., & Penn, D. L. (2000). Impoverished counterfactual thinking is associated with schizophrenia. Commentaries.

Psychiatry(Washington, D. C.), *63*(4), 326–343.

Hosokawa, T., Kato, K., Inoue, M., & Mikami, A. (2007). Neurons in the macaque orbitofrontal cortex code relative preference of both rewarding and aversive outcomes. *Neuroscience research*, *57*(3), 434–445.

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*(5), 765-775.

Hutton, C., Bork, A., Josephs, O., Deichmann, R., Ashburner, J., & Turner, R. (2002). Image distortion correction in fMRI: a quantitative evaluation. *Neuroimage*, *16*(1), 217–240.

Hutton, C., Deichmann, R., Turner, R., & Andersson, J. L. R. (2004). Combined correction for geometric distortion and its interaction with head motion in fMRI. In *Proceedings of ISMRM* (Vol. 12).

Inman, J. J., & Zeelenberg, M. (2002). Regret in repeat purchase versus switching decisions: The attenuating role of decision justifiability. *Journal of Consumer Research*, *29*(1), 116–128.

Inman, J. J., Dyer, J. S., & Jia, J. (1997). A Generalized Utility Model of Disappointment and Regret Effects on Post-Choice Valuation. *Marketing science*, *16*(2), 97-111.

Janis, I. L., & Mann, L. (1977). Anticipatory regret. *Decision making: A psychological analysis of conflict, choice, and commitment*, 219–242.

Johnson, T. J., Feigenbaum, R., & Weiby, M. (1964). Some determinants and consequences of the teacher's perception of causation. *Journal of educational Psychology*, *55*(5), 237–246.

Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of*

Experimental Social Psychology, *3*(1), 1–24.

Josephs, O., Deichmann, R., & Turner, R. (2000). Trajectory measurement and generalised reconstruction in rectilinear EPI. *NeuroImage*, *11*(5; PART 2), 543–543.

Kahn, B. E., & Isen, A. M. (1993). The influence of positive affect on variety seeking among safe, enjoyable products. *Journal of Consumer Research*, *20*(2), 257.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, *93*(2), 136–153.

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

Kahneman, D., & Tversky, A. (1982). The simulation heuristic. In D. Kahneman, P. Slovic, & Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201-208). New York: Cambridge Univ Pr.

Kahneman, D., & Tversky, A. (1982). The psychology of preferences. *Scientific American*, *246*(1), 160–173.

Kawashima, R., Sugiura, M., Kato, T., Nakamura, A., Hatano, K., Ito, K., Fukuda, H., et al. (1999). The human amygdala plays an important role in gaze monitoring: A PET study. *Brain*, *122*(4), 779.

Kelley, H. H. (1971). *Attribution in social interaction*. General Learning Press.

Kelley, W. M., Macrae, C. N., Wyland, C. L., Caglar, S., Inati, S., & Heatherton, T. F. (2002). Finding the self? An event-related fMRI study. *Journal of cognitive neuroscience*, *14*(5), 785–794.

Kiehl, K. A., Liddle, P. F., & Hopfinger, J. B. (2000). Error processing and the rostral anterior cingulate: an event-related fMRI study. *Psychophysiology*,

37(02), 216–223.

Kimbrell, T. A., George, M. S., Parekh, P. I., Ketter, T. A., Podell, D. M., Danielson, A. L., Repella, J. D., et al. (1999). Regional brain activity during transient self-induced anxiety and anger in healthy adults. *Biological Psychiatry*, *46*(4), 454–465.

Klein, T. A., Endrass, T., Kathmann, N., Neumann, J., von Cramon, D. Y., & Ullsperger, M. (2007). Neural correlates of error awareness. *NeuroImage*, *34*(4), 1774-1781.

Koelewijn, T., van Schie, H. T., Bekkering, H., Oostenveld, R., & Jensen, O. (2008). Motor-cortical beta oscillations are modulated by correctness of observed action. *Neuroimage*, *40*(2), 767–775.

Kordes-de Vaal, J. H. (1996). Intention and the omission bias: Omissions perceived as nondecisions. *Acta Psychologica*, *93*(1-3), 161–172.

Kuhnen, C. M., & Knutson, B. (2005). The neural basis of financial risk taking. *Neuron*, *47*(5), 763–770.

LaBar, K. S., Gatenby, J. C., Gore, J. C., LeDoux, J. E., & Phelps, E. A. (1998). Human amygdala activation during conditioned fear acquisition and extinction: a mixed-trial fMRI study. *Neuron, 20*, 937–945.

Landman, J. (1993). *Regret. The Persistence of the Possible*. New York: Oxford University Press.

Landman, J. (1987). Regret and Elation Following Action and Inaction: Affective Responses to Positive Versus Negative Outcomes. *Pers Soc Psychol Bull*, *13*(4), 524-536.

Larrick, R. P., & Boles, T. L. (1995). Avoiding regret in decisions with feedback: A negotiation example. *Organizational Behavior and Human Decision*

Processes, *63*(1), 87–97.

Lechner, L., de Vries, H., & Offermans, N. (1997). Participation in a breast cancer screening program: influence of past behavior and determinants on future screening participation. *Preventive Medicine*, *26*(4), 473–482.

Lerner, J. S., & Keltner, D. (2000). Beyond valence: Toward a model of emotion-specific influences on judgement and choice. *Cognition & Emotion*, *14*(4), 473–493.

Lesieur, H. R. (1977). *The chase: career of the compulsive gambler*. Anchor Press.

Levitt, H. (1971). Transformed up-down methods in psychoacoustics. *Journal of the Acoustical Society of America*, *49*(2), 467–477.

Liotti, M., Mayberg, H. S., Brannan, S. K., McGinnis, S., Jerabek, P., & Fox, P. T. (2000). Differential limbic-cortical correlates of sadness and anxiety in healthy subjects: implications for affective disorders. *Biological Psychiatry*, *48*(1), 30–42.

Liu, X., Powell, D. K., Wang, H., Gold, B. T., Corbly, C. R., & Joseph, J. E. (2007). Functional dissociation in frontal and striatal areas for processing of positive and negative reward information. *Journal of Neuroscience*, *27*(17), 4587-4597.

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational behavior and human decision processes*, *65*(3), 272–292.

Loewenstein, G., & Lerner, J. S. (2003). The role of affect in decision making. *Handbook of affective science*, 619–642.

Loftus, E. F. (1996). *Eyewitness testimony*. Harvard Univ Pr.

Loftus, G. R., & Loftus, E. F. (1983). *Mind at Play; The Psychology of Video Games*. Basic Books, Inc. New York, NY, USA.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, *412*(6843), 150–157.

Lohrenz, T., McCabe, K., Camerer, C. F., & Montague, P. R. (2007). Neural signature of fictive learning signals in a sequential investment task. *Proceedings of the National Academy of Sciences*, *104*(22), 9493.

Loomes, G., & Sugden, R. (1983). A rationale for preference reversal. *The American Economic Review*, *73*(3), 428–432.

Loomes, G., & Sugden, R. (1982). Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty. *The Economic Journal*, *92*(368), 805-824.

Maddock, R. J., & Buonocore, M. H. (1997). Activation of left posterior cingulate gyrus by the auditory presentation ofthreat-related words: an fMRI study. *Psychiatry Research: Neuroimaging*, *75*(1), 1–14.

Maldjian, J. A., Laurienti, P. J., Kraft, R. A., & Burdette, J. H. (2003). An automated method for neuroanatomic and cytoarchitectonic atlas-based interrogation of fMRI data sets. *Neuroimage*, *19*(3), 1233–1239.

Mandel, D. (2003). Counterfactuals, emotions, and context. *Cognition & Emotion*, *17*(1), 139–159.

Manning, S. K., & Melchiori, M. P. (1974). Words that upset urban college students: Measured with GSRs and rating scales. *The Journal of Social Psychology*, *94*(2), 305–306.

March, J. G. (1978). Bounded rationality, ambiguity, and the engineering of choice. *The Bell Journal of Economics*, 587–608.

March, J. G. (1996). Learning to be risk averse. *Psychological Review*, *103*(2), 309–319.

Marchiori, D., & Warglien, M. (2008). Predicting human interactive learning by regret-driven neural networks. *Science*, *319*(5866), 1111.

Markman, K. D., Gavanski, I., Sherman, S. J., & McMullen, M. N. (1993). The mental simulation of better and worse possible worlds. *Journal of experimental social psychology(Print)*, *29*(1), 87–109.

Markowitz, H. (1952). Portfolio selection. *Journal of finance*, *7*(1), 77–91.

McGlothlin, W. H. (1956). Stability of choices among uncertain alternatives. *The American Journal of Psychology*, 604–615.

McNamara, P., Durso, R., Brown, A., & Lynch, A. (2003). Counterfactual cognitive deficit in persons with Parkinson's disease. *British Medical Journal*, *74*(8), 1065–1070.

Mellers, B., Schwartz, A., & Ritov, I. (1999). Emotion-based choice. *Journal of experimental psychology General*, *128*, 332–345.

Mellers, B. A. (2000). Choice and the relative pleasure of consequences. *Psychological bulletin*, *126*(6), 910–924.

Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: How we feel about risky options. *Psychological Science*, *8*, 423–429.

Menon, V., Adleman, N. E., White, C. D., Glover, G. H., & Reiss, A. L. (2001). Error-related brain activation during a Go/NoGo response inhibition task. *Human Brain Mapping*, *12*(3), 131–143.

Mesulam, M. M., Nobre, A. C., Kim, Y. H., Parrish, T. B., & Gitelman, D. R. (2001). Heterogeneity of cingulate contributions to spatial attention. *Neuroimage*, *13*(6), 1065–1072.

Metcalfe, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, *106*, 3–19.

Meyer, G., Schwertfeger, J., Exton, M. S., Janssen, O. E., Knapp, W., Stadler, M. A., Schedlowski, M., et al. (2004). Neuroendocrine response to casino gambling in problem gamblers. *Psychoneuroendocrinology*, *29*(10), 1272–1280.

Michenaud, S., & Solnik, B. (2008). Applying regret theory to investment choices: Currency hedging decisions. *Journal of International Money and Finance*, *27*(5), 677–694.

Miller, D. T. (1976). Ego involvement and attributions for success and failure. *Journal of Personality and Social Psychology*, *34*(5), 901–906.

Miller, D. T. (1978). What constitutes a self-serving attributional bias? A reply to Bradley. *Journal of Personality and Social Psychology*, *36*(11), 1221–1223.

Miller, D. T., & Norman, S. A. (1975). Actor-observer differences in perceptions of effective control. *Journal of Personality and Social Psychology*, *31*(3), 503–515.

Mineka, S., & Cook, M. (1988). Social learning and the acquisition of snake fear in monkeys. In T. Zentall & B. J. Galef (Eds.), *Social learning: Psychological and biological perspectives* (pp. 51–75). Erlbaum, Hillsdale, NJ.

Mitchell, J. P., Banaji, M. R., & MacRae, C. N. (2005). The link between social cognition and self-referential thought in the medial prefrontal cortex. *Journal of Cognitive Neuroscience*, *17*(8), 1306–1315.

Mobbs, D., Greicius, M. D., Abdel-Azim, E., Menon, V., & Reiss, A. L. (2003). Humor modulates the mesolimbic reward centers. *Neuron*, *40*(5), 1041–1048.

Mobbs, D., Yu, R., Meyer, M., Passamonti, L., Seymour, B., Calder, A. J., Schweizer, S., et al. (2009). A key role for similarity in vicarious reward. *Science*, *324*(5929), 900.

Mobini, S., Body, S., Ho, M. Y., Bradshaw, C. M., Szabadi, E., Deakin, J. F., &

Anderson, I. M. (2002). Effects of lesions of the orbitofrontal cortex on sensitivity to delayed and probabilistic reinforcement. *Psychopharmacology*, *160*(3), 290–298.

Molina, J. A., Sainz-Artiga, M. J., Fraile, A., Jimenez-Jimenez, F. J., Villanueva, C., Orti-Pareja, M., & Bermejo-P, F. (2000). Pathologic gambling in Parkinson's disease: a behavioral manifestation of pharmacologic treatment? *Movement Disorders*, *15*(5), 869–872.

Montague, P. R., King-Casas, B., & Cohen, J. D. (2006). Imaging valuation models in human choice. *Annual Review of Neuroscience*, *29*, 417-448.

Morris, J. S., & Dolan, R. J. (2004). Dissociable amygdala and orbitofrontal responses during reversal fear conditioning. *Neuroimage*, *22*(1), 372–380.

Nangini, C., Tam, F., & Graham, S. J. (2008). A novel method for integrating MEG and BOLD fMRI signals with the linear convolution model in human primary somatosensory cortex. *Human brain mapping*, *29*(1), 97–106.

Nielsen, F. A., Balslev, D., & Hansen, L. K. (2004). Mining posterior cingulate. *NeuroImage*, *22*, 14–17.

Nieuwenhuis, S., Heslenfeld, D. J., Alting von Geusau, N. J., Mars, R. B., Holroyd, C. B., & Yeung, N. (2005). Activity in human reward-sensitive brain areas is strongly context dependent. *Neuroimage*, *25*(4), 1302–1309.

Ochsner, K. N., Knierim, K., Ludlow, D. H., Hanelin, J., Ramachandran, T., Glover, G., & Mackey, S. C. (2004). Reflecting upon feelings: an fMRI study of neural systems supporting the attribution of emotion to self and other. *Journal of cognitive neuroscience*, *16*(10), 1746–1772.

O'Doherty, J., Critchley, H., Deichmann, R., & Dolan, R. J. (2003). Dissociating valence of outcome from behavioral control in human orbital and ventral prefrontal cortices. *Journal of Neuroscience*, *23*(21), 7931.

O'Doherty, J., Dayan, P., Schultz, J., Deichmann, R., Friston, K., & Dolan, R. J. (2004). *Dissociable roles of ventral and dorsal striatum in instrumental conditioning* (Vol. 304). American Association for the Advancement of Science.

O'Doherty, J. P., Deichmann, R., Critchley, H. D., & Dolan, R. J. (2002). Neural responses during anticipation of a primary taste reward. *Neuron*, *33*(5), 815–826.

Ogawa, S., Lee, T. M., Stepnoski, R., Chen, W., Zhu, X. H., & Ugurbil, K. (2000). An approach to probe some neural systems interaction by functional MRI at neural time scale down to milliseconds. *Proceedings of the National Academy of Sciences of the United States of America*, *97*(20), 11026.

Ogawa, S., Menon, R. S., Tank, D. W., Kim, S. G., Merkle, H., Ellermann, J. M., & Ugurbil, K. (1993). Functional brain mapping by blood oxygenation level-dependent contrast magnetic resonance imaging. A comparison of signal characteristics with a biophysical model. *Biophysical Journal*, *64*(3), 803–812.

Ordónez, L. D., & Connolly, T. (2000). Regret and responsibility: A reply to Zeelenberg et al.(1998). *Organizational Behavior and Human Decision Processes*, *81*(1), 132–142.

Pauling, L. (1977). Magnetic properties and structure of oxyhemoglobin. *Proceedings of the National Academy of Sciences of the United States of America*, *74*(7), 2612.

Paulus, M. P., Rogalsky, C., Simmons, A., Feinstein, J. S., & Stein, M. B. (2003). Increased activation in the right insula during risk-taking decision making is related to harm avoidance and neuroticism. *Neuroimage*, *19*(4), 1439–1448.

Paulus, M. P., & Stein, M. B. (2006). An Insular View of Anxiety. *Biological Psychiatry*, *60*(4), 383-387.

Pessiglione, M., Seymour, B., Flandin, G., Dolan, R. J., & Frith, C. D. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*(7106), 1042–1045.

Pickens, C. L., Saddoris, M. P., Setlow, B., Gallagher, M., Holland, P. C., & Schoenbaum, G. (2003). Different roles for orbitofrontal cortex and basolateral amygdala in a reinforcer devaluation task. *Journal of Neuroscience*, *23*(35), 11078–11084.

Pieters, R., & Zeelenberg, M. (2007). A theory of regret regulation 1.1. *Journal of Consumer Psychology*, *17*(1), 29–35.

Ploghaus, A., Becerra, L., Borras, C., & Borsook, D. (2003). Neural circuitry underlying pain modulation: expectation, hypnosis, placebo. *Trends in Cognitive Sciences*, *7*(5), 197–200.

Ploghaus, A., Tracey, I., Gati, J. S., Clare, S., Menon, R. S., Matthews, P. M., & Rawlins, J. N. (1999). Dissociating pain from its anticipation in the human brain. *Science*, *284*(5422), 1979.

Portas, C. M., Krakow, K., Allen, P., Josephs, O., Armony, J. L., & Frith, C. D. (2000). Auditory Processing across the Sleep-Wake Cycle:: Simultaneous EEG and fMRI Monitoring in Humans. *Neuron*, *28*(3), 991–999.

Preuschoff, K., Quartz, S. R., & Bossaerts, P. (2008). Human insula activation reflects risk prediction errors as well as risk. *Journal of Neuroscience*, *28*(11), 2745.

Radcliffe, N. M., & Klein, W. M. (2002). Dispositional, unrealistic, and comparative optimism: Differential relations with the knowledge and processing of risk information and beliefs about personal risk. *Personality and Social Psychology Bulletin*, *28*(6), 836.

Reb, J. (2008). Regret aversion and decision process quality: Effects of regret salience on decision process carefulness. *Organizational Behavior and Human Decision Processes*, *105*(2), 169–182.

Reb, J., & Connolly, T. (2009). Myopic regret avoidance: Feedback avoidance and learning in repeated decision making. *Organizational Behavior and Human Decision Processes*, *109*(2), 182-189.

Reid, R. L. (1986). The psychology of the near miss. *Journal of Gambling Studies*, *2*(1), 32–39.

Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II*. New York: Appleton-Century-Crofts.

Richard, R., Van der Pligt, J., & de Vries, N. (1996). Anticipated regret and time perspective: Changing sexual risk-taking behavior. *Journal of Behavioral Decision Making*, *9*(3).

Richardson, M. P., Strange, B. A., & Dolan, R. J. (2004). Encoding of emotional memories depends on amygdala and hippocampus and their interactions. *Nature Neuroscience*, *7*(3), 278–285.

Ritov, I. (1996). Probability of regret: Anticipation of uncertainty feedback in choice. *Organizational Behavior and Human Decision Processes, 66,* 228–36.

Ritov, I., & Baron, J. (1990). Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making*, *3*(4), 263–277.

Ritov, I., & Baron, J. (1995). Outcome knowledge, regret, and omission bias. *Organizational Behavior and Human Decision Processes*, *64*(2), 119–127.

Ritov, I., & Baron, J. (1992). Status-quo and omission biases. *Journal of Risk and Uncertainty*, *5*(1), 49-61.

Roese, N. J. (1994). The functional basis of counterfactual thinking. *Journal of Personality and Social Psychology*, *66*, 805–805.

Roese, N. J., & Olson, J. M. (1995). Counterfactual thinking: A critical overview. In N. J. Roese & J. M. Olson (Eds.), *What might have been: The social psychology of counterfactual thinking* (pp. 1-55). Lawrence Erlbaum Mahwah, NJ.

Roese, N. J., & Olson, J. M. (1995). Outcome controllability and counterfactual thinking. *Personality and Social Psychology Bulletin*, *21*(6), 620–628.

Roese, N. J., & Olson, J. M. (2007). Better, stronger, faster: Self-serving judgment, affect regulation, and the optimal vigilance hypothesis. *Perspectives on Psychological Science*, *2*(2), 124.

Roese, N. J., Park, S., Smallman, R., & Gibson, C. (2008). Schizophrenia involves impairment in the activation of intentions by counterfactual thinking. *Schizophrenia Research*, *103*(1-3), 343–344.

Roese, N. J., Summerville, A., & Fessel, F. (2007). Regret and behavior: Comment on Zeelenberg and Pieters. *Journal of consumer psychology: the official journal of the Society for Consumer Psychology*, *17*(1), 25.

Roseman, I. J. (1984). Cognitive determinants of emotion: A structural theory. *Review of Personality & Social Psychology*.

Ruby, P., & Decety, J. (2001). Effect of subjective perspective taking during simulation of action: a PET investigation of agency. *Nature Neuroscience*,

4(5), 546–550.

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

Sander, D., Grafman, J., & Zalla, T. (2003). The human amygdala: an evolved system for relevance detection. *Reviews in the Neurosciences*, *14*(4), 303–316.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical Association*, 55–67.

Scherer, K. R., Schorr, A., & Johnstone, T. (2001). *Appraisal processes in emotion: Theory, methods, research*. Oxford University Press, USA.

van Schie, H. T., Mars, R. B., Coles, M. G., & Bekkering, H. (2004). Modulation of activity in medial frontal and motor cortices during error observation. *Nature Neuroscience*, *7*(5), 549–554.

Schoenbaum, G., Gottfried, J. A., Murray, E. A., & Ramus, S. J. (2007). Linking affect to action: critical contributions of the orbitofrontal cortex. Preface. *Annals of the New York Academy of Sciences*, *1121*.

Schonberg, T., Daw, N. D., Joel, D., & O'Doherty, J. P. (2007). Reinforcement learning signals in the human striatum distinguish learners from nonlearners during reward-based decision making. *Journal of Neuroscience*, *27*(47), 12860.

Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, *6*(2), 461–464.

Schweitzer, M. (1994). Disentangling status quo and omission effects: An experimental analysis. *Organizational Behavior and Human Decision Processes*, *58*(3), 457–476.

Seidel, E. M., Eickhoff, S. B., Kellermann, T., Schneider, F., Gur, R. C., Habel, U., & Derntl, B. (2010). Who is to blame? Neural correlates of causal attribution in social situations. *Social neuroscience, 5*(4), 335–350.

Seligman, M. E. (1972). Learned helplessness. *Annual Review of Medicine*, *23*(1), 407–412.

Sharot, T., De Martino, B., & Dolan, R. J. (2009). How Choice Reveals and Shapes Expected Hedonic Outcome. *Journal of Neuroscience*, *29*(12), 3760.

Shaw, G. B. (1903). Maxims for revolutionists. *Man and Superman*, 217–29.

Sheeran, P., & Orbell, S. (1999). Augmenting the Theory of Planned Behavior: Roles for Anticipated Regret and Descriptive Norms. *Journal of Applied Social Psychology*, *29*(10), 2107-2142.

Shin, L. M., & Liberzon, I. (2010). The neurocircuitry of fear, stress, and anxiety disorders. *Neuropsychopharmacology*, *35*, 169-191.

Shiv, B., Loewenstein, G., & Bechara, A. (2005). The dark side of emotion in decision-making: When individuals with decreased emotional reactions make more advantageous decisions. *Cognitive Brain Research*, *23*(1), 85–92.

Simonsohn, U., Karlsson, N., Loewenstein, G., & Ariely, D. (2008). The tree of experience in the forest of information: Overweighing experienced relative to observed information. *Games and Economic Behavior*, *62*(1), 263–286.

Simonson, I. (1992). The Influence of Anticipating Regret and Responsibility on Purchase Decisions. *The Journal of Consumer Research*, *19*(1), 105-118.

Singer, T., Critchley, H. D., & Preuschoff, K. (2009). A common role of insula in feelings, empathy and uncertainty. *Trends in Cognitive Sciences*, *13*(8), 334-340.

Small, D. M., Gitelman, D. R., Gregory, M. D., Nobre, A. C., Parrish, T. B., &

Mesulam, M. M. (2003). The posterior cingulate and medial prefrontal cortex mediate the anticipatory allocation of spatial attention. *Neuroimage*, *18*(3), 633–641.

Solomon, R. C. (1993). *The passions: Emotions and the meaning of life*. Hackett Publishing.

Sommer, T., Peters, J., Gläscher, J., & Büchel, C. (2009). Structure–function relationships in the processing of regret in the orbitofrontal cortex. *Brain Structure and Function*, *213*(6), 535–551.

Spielberger, C. D., Gorsuch, R. L., & Edward, L. R. (1970). *STAI manual for the State-Trait Anxiety Inventory (" self-evaluation questionnaire").* Consulting Psychologists Press.

Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology*, *27*(1), 76-105.

Stapel, D. A., & Velthuijsen, A. S. (1996). " Just As If It Happened To Me": The Impact of Vivid and Self-Relevant Information on Risk Judgments. *Journal of Social and Clinical Psychology*, *15*, 102–119.

Storms, M. D. (1973). Videotape and the attribution process: Reversing actors' and observers' points of view. *Journal of Personality and Social Psychology, 27*(2), 165–175.

Strange, B. A., & Dolan, R. J. (2004). ß-Adrenergic modulation of emotional memory-evoked human amygdala and hippocampal responses. *Proceedings of the National Academy of Sciences*, *101*(31), 11454–11458.

Streufert, S., & Streufert, S. C. (1969). Effects of Conceptual Structure, Failure, and Success on Attribution of Causality and Interpersonal Attitudes. *J Personality Soc Psychol*.

Tanaka, S. C., Balleine, B. W., & O'Doherty, J. P. (2008). Calculating consequences: brain systems that encode the causal effects of actions. *Journal of Neuroscience*, *28*(26), 6750.

Tanaka, S. C., Doya, K., Okada, G., Ueda, K., Okamoto, Y., & Yamawaki, S. (2004). Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nature Neuroscience*, *7*(8), 887–893.

Tanaka, S. C., Schweighofer, N., Asahi, S., Shishida, K., Okamoto, Y., Yamawaki, S., & Doya, K. (2007). Serotonin differentially regulates short-and long-term prediction of rewards in the ventral and dorsal striatum. *PLoS ONE*, *2*(12).

Thoreau, H. D. (1893). *The writings of Henry D. Thoreau*. Boston: Houghton Mifflin Company.

Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, *315*(5811), 515.

Tomlin, D., Kayali, M. A., King-Casas, B., Anen, C., Camerer, C. F., Quartz, S. R., & Montague, P. R. (2006). *Agent-specific responses in the cingulate cortex during economic exchanges* (Vol. 312). American Association for the Advancement of Science.

Tremblay, L., & Schultz, W. (1999). Relative reward preference in primate orbitofrontal cortex. *Nature*, *398*(6729), 704–708.

Tricomi, E. M., Delgado, M. R., & Fiez, J. A. (2004). Modulation of caudate activity by action contingency. *Neuron*, *41*(2), 281–292.

Tsiros, M., & Mittal, V. (2000). Regret: A Model of Its Antecedents and Consequences in Consumer Decision Making. *Journal of Consumer Research: An Interdisciplinary Quarterly*, Journal of Consumer Research: An Interdisciplinary Quarterly, *26*(4), 401-17.

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

Tykocinski, O., Israel, R., & Pittman, T. S. (2004). Inaction Inertia in the Stock Market1. *Journal of Applied Social Psychology*, *34*(6), 1166–1175.

Tykocinski, O. E., & Pittman, T. S. (1998). The consequences of doing nothing: Inaction inertia as avoidance of anticipated counterfactual regret. *Journal of Personality and Social Psychology*, *75*, 607–616.

Tykocinski, O. E., & Pittman, T. S. (2001). Product aversion following a missed opportunity: price contrast or avoidance of anticipated regret? *Basic and Applied Social Psychology*, *23*(3), 149–156.

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., et al. (2002). Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. *Neuroimage*, *15*(1), 273–289.

Ugurbil, K., Garwood, M., Ellermann, J., Hendrich, K., Hinke, R., Hu, X., Kim, S. G., et al. (1993). Imaging at high magnetic fields: initial experiences at 4 T. *Magnetic Resonance Quarterly*, *9*(4), 259.

Ullsperger, M., & Von Cramon, D. Y. (2004). Neuroimaging of performance monitoring: error detection and beyond. *Cortex*, *40*(4-5), 593–604.

Ursu, S., & Carter, C. S. (2005). Outcome representations, counterfactual comparisons and the human orbitofrontal cortex: implications for neuroimaging studies of decision-making. *Cognitive Brain Research*, *23*(1), 51–60.

Van der Velde, F. W., Van der Pligt, J., & Hooykaas, C. (1994). Perceiving AIDS-related risk: Accuracy as a function of differences in actual risk. *Health*

Psychology, *13*(1), 25–33.

Viswanathan, A., & Freeman, R. D. (2007). Neurometabolic coupling in cerebral cortex reflects synaptic more than spiking activity. *Nature Neuroscience, 10*(10), 1308–1312.

Von Neumann, J., Morgenstern, O., Kuhn, H. W., & Rubinstein, A. (1947). *Theory of games and economic behavior*. Princeton university press Princeton, NJ.

Weigl, P. D., & Hanson, E. V. (1980). Observational learning and the feeding behavior of the red squirrel Tamiasciurus hudsonicus: the ontogeny of optimization. *Ecology*, *61*(2), 214–218.

Weinstein, N. D. (1987). Unrealistic optimism about susceptibility to health problems: Conclusions from a community-wide sample. *Journal of Behavioral Medicine*, *10*(5), 481–500.

Weinstein, N. D. (1989). Optimistic biases about personal risks. *Science, 246*(4935), 1232–1232.

Weinstein, N. D., & Lachendro, E. (1982). Egocentrism as a source of unrealistic optimism. *Personality and Social Psychology Bulletin*, *8*(2), 195.

Weintraub, D., Siderowf, A. D., Potenza, M. N., Goveas, J., Morales, K. H., Duda, J. E., Moberg, P. J., et al. (2006). Association of dopamine agonist use with impulse control disorders in Parkinson disease. *Archives of neurology*, *63*(7), 969.

Weisberg, D. P., & Beck, S. R. (2010). Children's thinking about their own and others' regret and relief. *Journal of Experimental Child Psychology*, *In Press*. Retrieved from http://www.sciencedirect.com/science/article/B6WJ9-4YNBRX6-1/2/a27f4b6aa5af59b1b46ad7815cfd9a65

Weiskopf, N., Hutton, C., Josephs, O., & Deichmann, R. (2006). Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *Neuroimage*, *33*(2), 493–504.

Weiskopf, N., Sitaram, R., Josephs, O., Veit, R., Scharnowski, F., Goebel, R., Birbaumer, N., et al. (2007). Real-time functional magnetic resonance imaging: methods and applications. *Magnetic Resonance Imaging*, *25*(6), 989–1003.

Whalen, P. J., Kagan, J., Cook, R. G., Davis, F. C., Kim, H., Polis, S., McLaren, D. G., et al. (2004). Human Amygdala Responsivity to Masked Fearful Eye Whites. *Science*, *306*(5704), 2061.

Wickens, J. R., Budd, C. S., Hyland, B. I., & Arbuthnott, G. W. (2007). Striatal contributions to reward and decision making: making sense of regional variations in a reiterated processing matrix. *Annals of the New York Academy of Sciences*, *1104*, 192.

Winston, J. S., Strange, B. A., O'Doherty, J., & Dolan, R. J. (2002). Automatic and intentional brain responses during evaluation of trustworthiness of faces. *Nature Neuroscience*, *5*(3), 277–283.

Wolosin, R. J. (1973). Effects of Cooperation and Competition on Responsibility Attribution After Success and Failure. *Journal of Experimental Social Psychology*, *9*(3), 220–35.

Wortman, C. B., Costanzo, P. R., & Witt, T. R. (1973). Effect of anticipated performance on the attributions of causality to self and others. *Journal of Personality and Social Psychology*, *27*(3), 372.

Wright, C. I., Martis, B., Schwartz, C. E., Shin, L. M., Fischer, H., McMullin, K., & Rauch, S. L. (2003). Novelty responses and differential effects of order in the amygdala, substantia innominata, and inferior temporal cortex. *Neuroimage*,

18(3), 660–669.

Yu, R., & Zhou, X. (2006). Brain responses to outcomes of one's own and other's performance in a gambling task. *Neuroreport*, *17*(16), 1747.

Zeelenberg, M., Beattie, J., Van der Pligt, J., & de Vries, N. K. (1996). Consequences of regret aversion: Effects of expected feedback on risky decision making. *Organizational Behavior and Human Decision Processes*, *65*(2), 148–158.

Zeelenberg, M., van den Bos, K., van Dijk, E., & Pieters, R. (2002). The inaction effect in the psychology of regret. *Journal of Personality and Social Psychology*, *82*(3), 314–327.

Zeelenberg, M., van Dijk, W. W., & Manstead, A. S. (1998). Reconsidering the relation between regret and responsibility. *Organizational Behavior and Human Decision Processes*, *74*(3), 254–272.

Zeelenberg, M., van Dijk, W. W., & Manstead, A. S. (2000). Regret and responsibility resolved? Evaluating Ordonez and Connolly's (2000) conclusions. *Organizational Behavior and Human Decision Processes*, *81*(1), 143–154.

Zeelenberg, M., van Dijk, W. W., Manstead, A. S., & van der Pligt, J. (2000). On bad decisions and disconfirmed expectancies: The psychology of regret and disappointment. *Cognition & Emotion*, *14*(4), 521–541.

Zeelenberg, M., van Dijk, W. W., van der Pligt, J., Manstead, A. S., van Empelen, P., & Reinderman, D. (1998). Emotional reactions to the outcomes of decisions: The role of counterfactual thought in the experience of regret and disappointment. *Organizational Behavior and Human Decision Processes*, *75*(2), 117–141.

Zeelenberg, M., & Pieters, R. (1999). Comparing service delivery to what might have been. *Journal of Service Research*, *2*(1), 86–97.

Zeelenberg, M., & Pieters, R. (2004). Consequences of regret aversion in real life: The case of the Dutch postcode lottery. *Organizational Behavior and Human Decision Processes*, *93*(2), 155–168.

Zeelenberg, M., & Pieters, R. (2006). Feeling is for doing: A pragmatic approach to the study of emotions in economic behavior. *Social psychology and economics*, 117–137.

Zeelenberg, M., & Pieters, R. (2007). A theory of regret regulation 1.0. *Journal of Consumer Psychology*, *17*(1), 3–18.

Zeelenberg, M., Van Dijk, W. W., Manstead, A. S., & Van der Pligt, J. (1998). The experience of regret and disappointment. *Cognition and Emotion*, *12*(2), 221–230.

Zeelenberg, M. (1999). Anticipated regret, expected feedback and behavioral decision making. *Journal of Behavioral Decision Making*, *12*(2), 93-106.

Appendices

Appendix A

Supplementary table to Chapter 3: Details of the gamble pairs used in the task, including the probabilities of each of the winning and losing outcomes (in points) of each gamble option. Gamble pairs 1-24 have equal expected value (i.e. [outcome 1 × probability 1] + [outcome 2 × probability 2]). Gamble pairs 25 and 26 are catch trials with noticeably different expected values.

Pair	Gamble 1				Gamble 2			
	Outcome 1	Probability 1	Outcome 2	Probability 2	Outcome 1	Probability 1	Outcome 2	Probability 2
1	80	0.5	-80	0.5	60	0.25	-20	0.75
2	10	0.5	-20	0.5	130	0.25	-50	0.75
3	70	0.5	-170	0.5	20	0.5	-120	0.5
4	30	0.5	-50	0.5	10	0.5	-30	0.5
5	50	0.5	-30	0.5	80	0.5	-60	0.5
6	150	0.5	-10	0.5	190	0.5	-50	0.5
7	160	0.5	-80	0.5	60	0.75	-20	0.25
8	40	0.5	-80	0.5	20	0.75	-140	0.25
9	70	0.5	-50	0.5	20	0.75	-20	0.25
10	200	0.5	-30	0.5	130	0.75	-50	0.25
11	50	0.5	-30	0.5	80	0.75	-200	0.25
12	50	0.5	-30	0.5	40	0.75	-80	0.25
13	120	0.25	-60	0.75	60	0.25	-40	0.75
14	10	0.25	-50	0.75	40	0.25	-60	0.75
15	100	0.25	-80	0.75	20	0.75	-200	0.25
16	80	0.25	-40	0.75	20	0.75	-100	0.25
17	140	0.25	-30	0.75	20	0.75	-5	0.25
18	20	0.75	-100	0.25	110	0.25	-50	0.75
19	60	0.75	-60	0.25	150	0.25	-10	0.75
20	60	0.75	-20	0.25	200	0.25	-10	0.75
21	100	0.75	-140	0.25	60	0.75	-20	0.25
22	100	0.75	-80	0.25	80	0.75	-20	0.25
23	10	0.75	-70	0.25	40	0.75	-160	0.25
24	80	0.75	-40	0.25	100	0.75	-100	0.25
25	150	0.5	0	0.5	0	0.5	-50	0.5
26	10	0.5	-200	0.5	200	0.75	-10	0.25

Appendix B

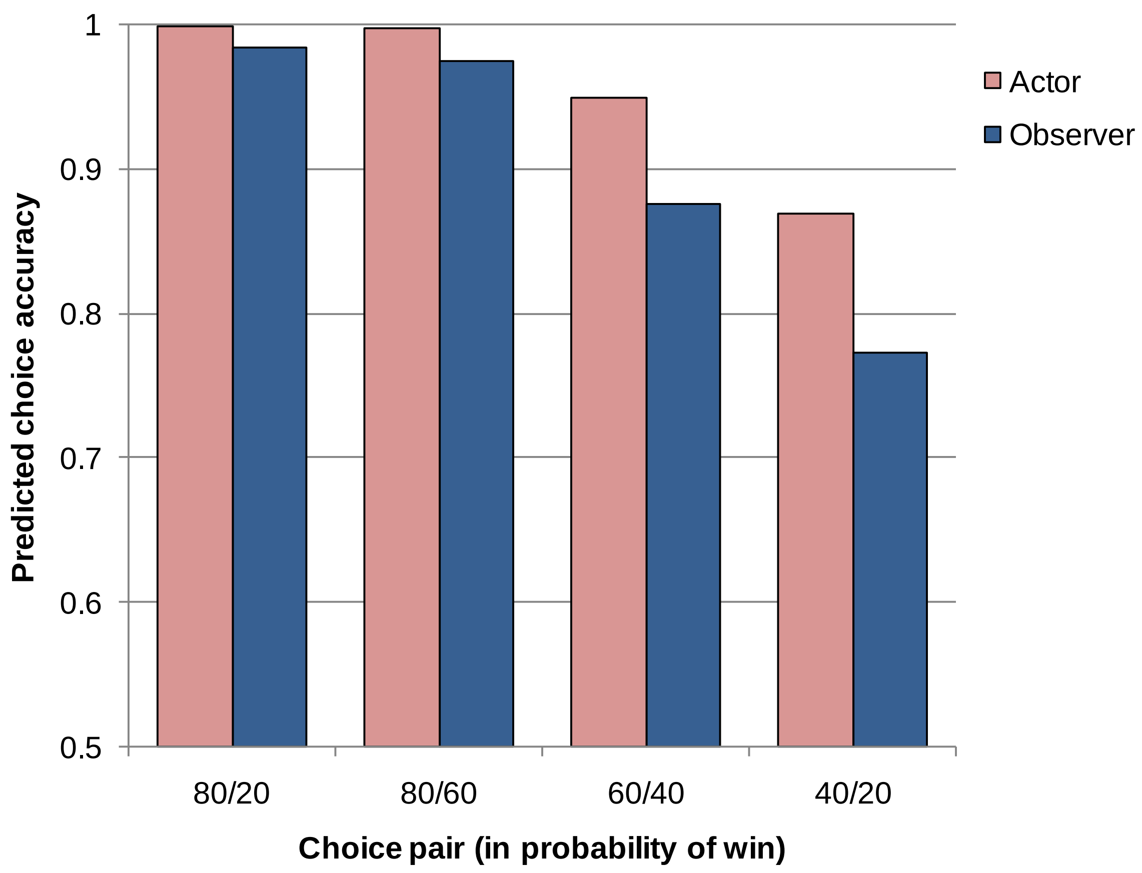

Choice accuracy for test trial gamble pairs, as predicted by the softmax function of the choice-independent learning model, using the mean best fitting learning rates and temperature of the softmax. Pairs are labelled according to the probability of a win for each stimulus. Choice accuracy is measured as the probability that the model chooses the stimulus with the highest probability of a win.

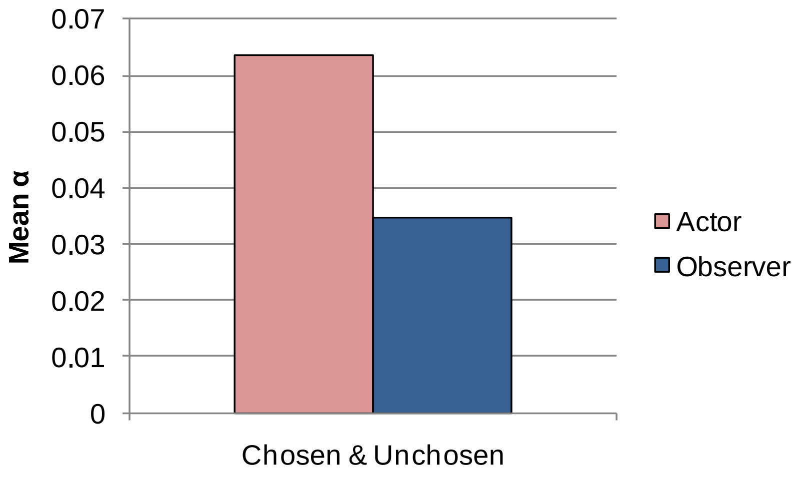

Mean best fitting learning rates (optimised separately for chosen and unchosen stimulus options) as predicted by the choice-independent learning model for actor and observer sessions. Best fitting learning rates were significantly different for actors (mean α = 0.06) and observers (mean α = 0.03) ($t(31) = 2.47$, $p < 0.02$). Best fitting temperatures of the softmax did not differ significantly between actors (mean β = 1.09) and observers (mean β = 1.13) ($t(31) = -0.81$, ns).

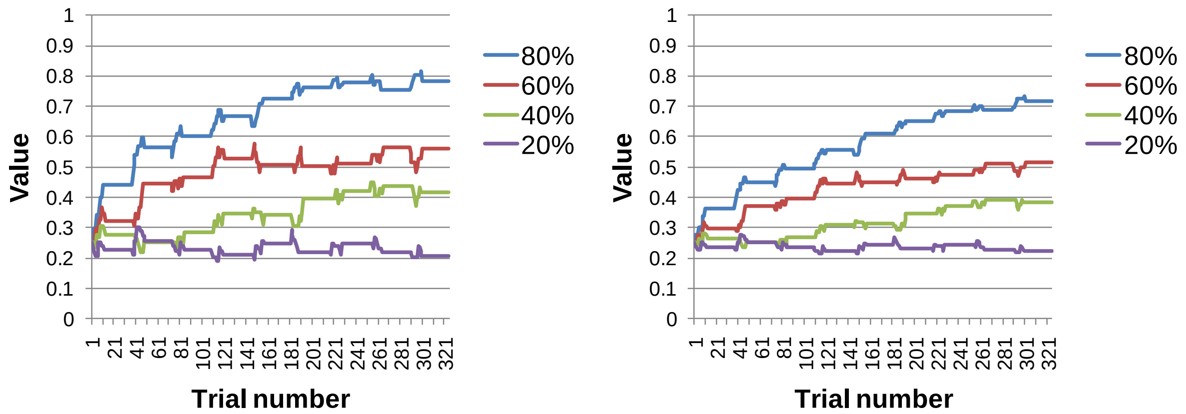

Stimulus values predicted for a) actor learning, and b) observer learning by the choice-independent learning model, using the mean best fitting parameters across participants as the learning rates and the temperature of the softmax function. These values are shown for the four stimuli, each associated with a unique probability of winning (80%, 60%, 40% and 20%). Note that the flat portions of each line reflect the blocks of test trials. Here learning does not take place, as outcome feedback is not given.

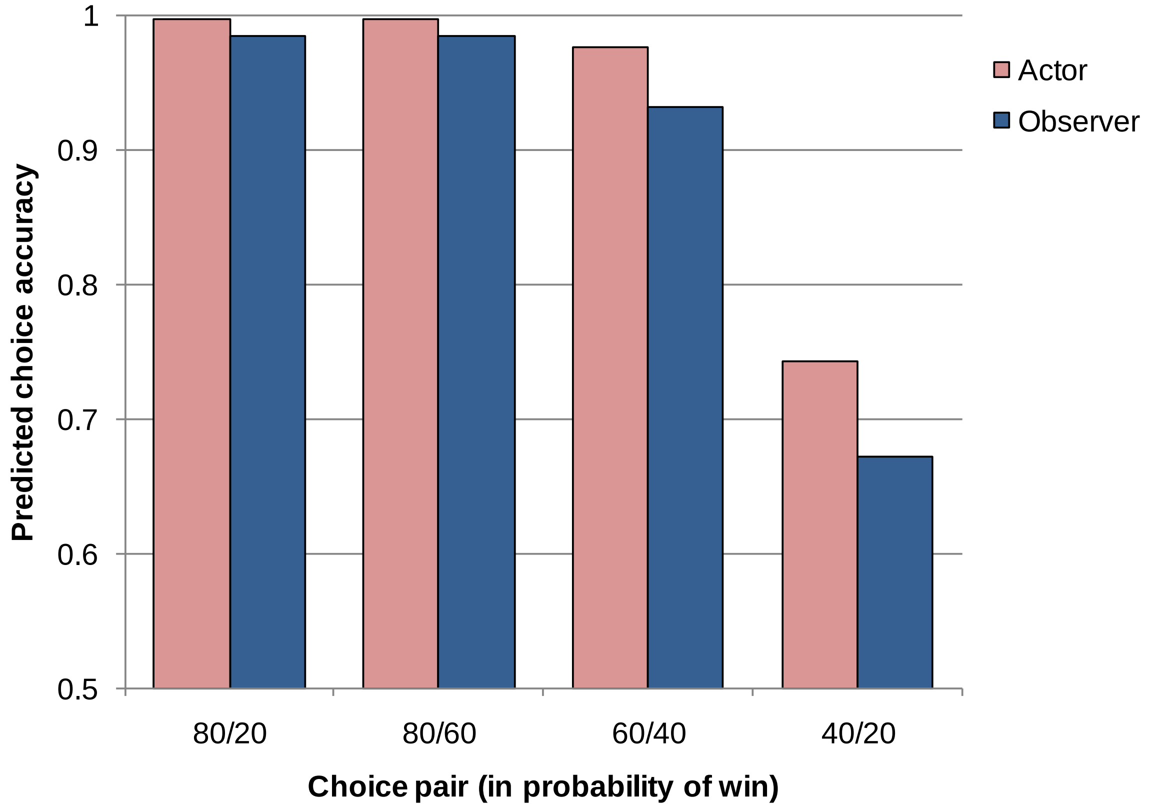

Choice accuracy for test trial gamble pairs, as predicted by the softmax function of the chosen-only learning model, using the mean best fitting learning rates and temperature of the softmax. Pairs are labelled according to the probability of a win for each stimulus. Choice accuracy is measured as the probability that the model chooses the stimulus with the highest probability of a win.

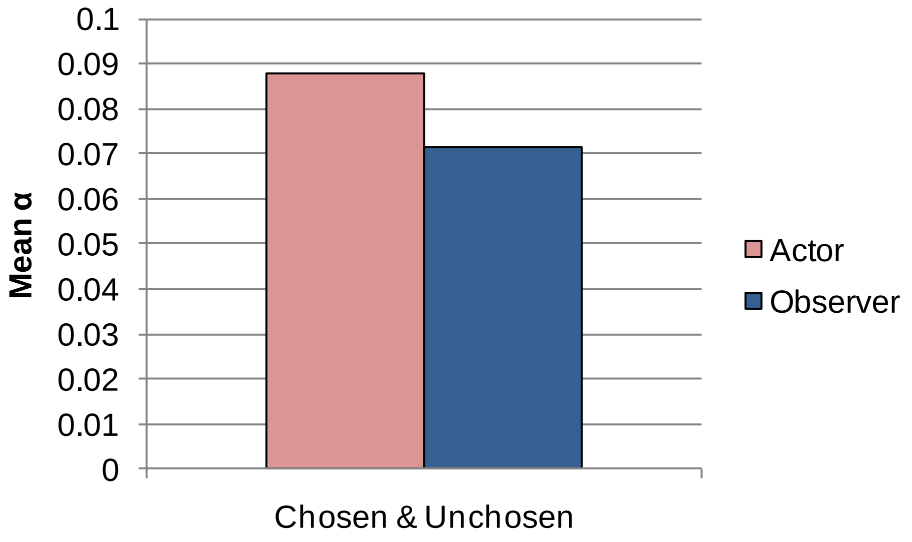

Mean best fitting learning rates (optimised separately for chosen and unchosen stimulus options) as predicted by the chosen-only learning model for actor and observer sessions. Best fitting learning rates did not significantly differ between actors (mean α = 0.09) and observers (mean α = 0.07) (*t*(31) = 0.71, ns). Best fitting temperatures of the softmax also did not differ significantly between actors (mean β = 1.11) and observers (mean β = 1.16) (*t*(31) = -1.13, ns).

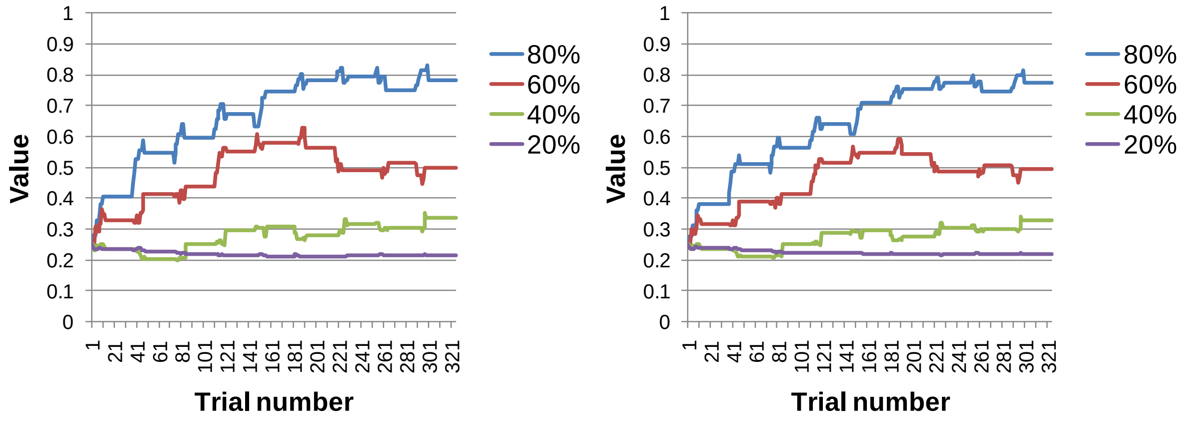

Stimulus values predicted for a) actor learning, and b) observer learning by the chosen-only learning model, using the mean best fitting parameters across participants as the learning rates and the temperature of the softmax function. These values are shown for the four stimuli, each associated with a unique probability of winning (80%, 60%, 40% and 20%). Note that the flat portions of each line reflect the blocks of test trials. Here learning does not take place, as outcome feedback is not given.