

Interdependence of Utility and Probability Estimates:
The Role of Social Power in Distinguishing Theories

Laura de Molière

University College London



Thesis submitted for the degree of Doctor of Philosophy

October 2014

Declaration

I, Laura de Molière, 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.

Signed

Date

Acknowledgements

Above all, I would like to wholeheartedly thank my supervisor Adam Harris for taking me on board (or “inheriting” me, as he calls it). I am extremely grateful having been supervised by somebody so exceptionally helpful, patient, and supportive. You truly are what I always thought science would be like – critical, logical, challenging, genuine, and fun.

Thank you for everything Adam.

I would also like to acknowledge Ruud Custers and Joe Devlin for being supportive in moments where it mattered the most.

My sincerest gratitude to the fantastic group of people I shared my office(s) with over the past 3 years, in particular Marleen Gillebaart, Tiffanie Ong and Keith O’Brien – thank you for making the time at UCL (and beyond) so very enjoyable.

I would also like to thank Klaus Fiedler for starting my passion for psychological research, and Florian Kutzner, my partner in crime, for being something like an academic big brother to me.

I am very thankful for the invaluable emotional support I received from my friends. Above all, thank you Laura Loos for always being my source of calm and reason, and Kim Chisholm for making my PhD journey one that was filled with sharing joy, professional rants and mutual support.

Danke Papa und Mama, ohne euch und eure buchstäblich grenzenlose Unterstützung, euer Verständnis für meine Leidenschaften, und euren bedingungslosen Stolz wäre das alles nicht möglich gewesen. Ihr seid die Besten. Danke auch dir, Madita, für das immer offene Ohr.

Finally, thank you Anna for being so loving, patient and encouraging, and for always making me the happiest version of myself. Thank you!

Thesis Summary

This thesis aims to identify mechanisms underlying the impact of negative utility on probability estimates (Harris, Corner, & Hahn, 2009). In particular, we investigated evidence for arousal misattribution (Vosgerau, 2010) by examining a moderation by interoceptive ability. Furthermore, we investigated an account of loss function asymmetry (e.g., Weber, 1994) by examining the role of decision-control.

Important to the current aims, social power is related to both interoception and control. Power increases interoceptive awareness (Moeini-Jazani, Knöpferle, de Molière, Gatti & Warlop, 2014) and one's personal sense of control (Fast, Gruenfeld, Sivanathan, & Galinsky, 2009). Thus, we examined whether power moderates the relationship of negative utility and probability estimates. Furthermore, we aimed to observe whether this potential moderation by social power occurs due to the powerful's greater interoceptive awareness, or due to their greater sense of control, supporting an account of arousal misattribution and loss function asymmetries, respectively.

Chapter 2 provides some preliminary evidence that self-reported interoception moderates the impact of negative utility on probability across 2 experiments. However, when measuring interoceptive awareness objectively, no evidence for a moderating role was found. Furthermore, assessing arousal by means of galvanic skin responses provided also no evidence for arousal misattribution as a mechanism. Next, in Chapter 3 we examined Vosgerau's (2010) original findings demonstrating arousal misattribution. Across 4 experiments, we are unable to replicate Vosgerau's (2010) results.

Subsequently, we investigated the role of loss function asymmetries in Chapter 4. Across 4 studies we provide evidence for the notion that powerful, but not powerless

individuals assign higher probabilities to negative events. Manipulating the controllability of the event, we provide evidence for loss function asymmetries as a mechanism underlying the impact of social power on the relationship of negative utility and probability estimates.

Having demonstrated that powerful individuals can be more sensitive to negative information, Chapter 5 shows across 3 experiments that the powerful act more on affordances of negative affective states. The powerful, who have been shown to be more approach oriented than the powerless (Keltner, Gruenfeld, & Anderson, 2003) become more avoidant under negative affective states than the powerless.

In sum, this thesis provides evidence for loss function asymmetries underlying the interdependence of utility and probability, and demonstrates that power can lead to heightened sensitivity for negative (affective) information.

Contents

LIST OF FIGURES	9
LIST OF TABLES.....	11
CHAPTER 1 THEORETICAL BACKGROUND	12
INFLUENCE OF UTILITY ON SUBJECTIVE PROBABILITY ESTIMATES.....	14
CHAPTER OVERVIEW	18
INTERDEPENDENCE OF UTILITY AND PROBABILITY	20
SOCIAL POWER.....	30
MECHANISMS UNDERLYING THE INTERDEPENDENCE OF (NEGATIVE) UTILITY AND PROBABILITY	35
SUMMARY AND OVERVIEW OF THE EMPIRICAL CHAPTERS OF THE THESIS.....	64
CHAPTER 2 THE IMPACT OF INTEROCEPTION ON THE INTERDEPENDENCE OF NEGATIVE UTILITY AND PROBABILITY ESTIMATES	67
CHAPTER OVERVIEW	67
INTRODUCTION.....	67
EXPERIMENT 1.....	82
EXPERIMENT 2.....	104
EXPERIMENT 3.....	117
CHAPTER DISCUSSION	138
CHAPTER 3 REPLICATING THE STAKE-LIKELIHOOD HYPOTHESIS	141
CHAPTER OVERVIEW	141
INTRODUCTION.....	141
EXPERIMENT 4.....	145
EXPERIMENT 5.....	153
EXPERIMENT 6.....	156

BAYESIAN META-ANALYSIS ACROSS EXPERIMENTS 4, 5 AND 6	158
EXPERIMENT 7.....	160
CHAPTER DISCUSSION	165
CHAPTER 4 SOCIAL POWER AND LOSS FUNCTION ASYMMETRIES.....	170
CHAPTER OVERVIEW	170
INTRODUCTION.....	171
EXPERIMENT 8.....	185
EXPERIMENT 9.....	194
EXPERIMENT 10.....	205
EXPERIMENT 11.....	212
CHAPTER DISCUSSION	231
CHAPTER 5 JOINT IMPACT OF POWER AND NEGATIVE AFFECT ON	
APPROACH – AVOIDANCE MOTIVATION	238
CHAPTER OVERVIEW	238
INTRODUCTION.....	239
EXPERIMENT 12.....	247
EXPERIMENT 13.....	255
EXPERIMENT 14.....	259
CHAPTER DISCUSSION	266
CHAPTER 6 GENERAL DISCUSSION	274
SUMMARY	274
MECHANISMS UNDERLYING THE INTERDEPENDENCE OF UTILITY AND PROBABILITY	277
THE ROLE OF SOCIAL POWER	281
FURTHER RESEARCH.....	283
CONCLUSION	287

REFERENCES 289

List of Figures

Figure 1.1. The process of making and reporting a probability estimate.....	17
Figure 1.2. Relationship of utility and subjective probability (Slovic, 1966).....	24
Figure 1.3. Probability estimates for Lench’s (2009) Study 4.....	41
Figure 1.4. Probability estimates in Lench’s (2009) Study 4 reclassified.....	42
Figure 1.5. Asymmetric loss functions.....	51
Figure 1.6. Model of power, control, and asymmetric loss functions.....	56
Figure 1.7. Model of power, interoceptive awareness, and arousal.....	64
Figure 2.1. Mediation model of the Stake Likelihood Hypothesis.....	68
Figure 2.2. Visual probability display from Harris et al.'s (2009) Experiment 5.....	86
Figure 2.3. Probability estimates across scales in Experiment 1.....	93
Figure 2.4. Probability estimates across scales and scenarios in Experiment 1.....	94
Figure 2.5. Probability estimates (USAF scenario, CBP scale) in Experiment 1.....	96
Figure 2.6. Probability estimates (USAF scenario, BAQ scale) in Experiment 1.....	97
Figure 2.7. Probability estimates (USAF scenario, Maia-E scale) in Experiment 1.....	98
Figure 2.8. Probability estimates (USAF scenario, Maia-BL scale) in Experiment 1.....	99
Figure 2.9. Probability estimates (dice scenario, CBP scale) in Experiment 1.....	100
Figure 2.10. Probability estimates (dice scenario, BAQ scale) in Experiment 1.....	102
Figure 2.11. Probability estimates (dice scenario, MAIA-E scale) in Experiment 1.....	103
Figure 2.12. Probability estimates (dice scenario, MAIA-BL scale) in Experiment 1.....	106
Figure 2.13. Probability display in the Container scenario of Experiment 2.....	110
Figure 2.14. Probability estimates in the RAF scenario in Experiment 2.....	111
Figure 2.15. Probability estimates in the dice scenario (IS assessed after the scenario) in Experiment 2.....	113
Figure 2.16. Probability estimates in the dice scenario (IS assessed before the scenario) in Experiment 2.....	114
Figure 2.17. Probability estimates in the container scenario in Experiment 2.....	118
Figure 2.18. Phasic and tonic components of skin conductance.....	122
Figure 2.19. Display of the finger transducer used in Experiment 3.....	124

Figure 2.20. Probability display of the negative condition of the Roulette scenario.....	127
Figure 2.21. Frequency distribution of accuracy scores in heartbeat counting task in Experiment 3.....	130
Figure 2.22. Moderated mediation model assessed in Experiment 3.....	133
Figure 2.23. Probability estimates by utility and IS in Experiment 3.....	135
Figure 2.24. Probability estimates by utility and sense of power in Experiment 3.....	136
Figure 3.1. Probability estimates by utility and order in Experiment 4.....	149
Figure 3.2. Arousal ratings by utility and order in Experiment 4.....	151
Figure 3.3. Grey and pink paper used in Experiment 7.....	163
Figure 3.4. Probability estimates by paper colour in Experiment 7.....	164
Figure 4.1. Mediation model based on loss function asymmetries.....	171
Figure 4.2. Preference for a riskier plan in Anderson and Galinsky's (2006) Experiment 3.....	180
Figure 4.3. Probability estimates across scenarios in Experiment 8.....	190
Figure 4.4. Probability estimates in the dice scenario in Experiment 8.....	191
Figure 4.5. Probability estimates in the RAF scenario in Experiment 8.....	193
Figure 4.6. Probability estimates by utility and sense of power in Experiment 9.....	203
Figure 4.7. Probability estimates by utility and manipulated power in Experiment 9.....	201
Figure 4.8. Probability estimates across self- and other-relevance in Experiment 10.....	210
Figure 4.9. Probability estimates by self-relevance in Experiment 10.....	211
Figure 4.10. Visual probability displays used In Experiment 11.....	219
Figure 4.11. Probability estimates for Map A in Experiment 11.....	225
Figure 4.12. Probability estimates for Map B in Experiment 11.....	223
Figure 4.13. Probability estimates for the difference score in Experiment 11.....	227
Figure 5.1. Threat bias by power and mood in Experiment 12.....	254
Figure 5.2. BIS means by power and mood in Experiment 13.....	259
Figure 5.3. Graphical display of the office in the manager condition in Experiment 14.....	261
Figure 5.4. Graphical display of the office in the subordinate condition in Experiment 14.....	262
Figure 5.5. BIS means by power and mood in Experiment 14.....	265
Figure 5.6. BAS means by power and mood in Experiment 14.....	266

Figure 6.1. Potential relationships between mechanisms.....	280
Figure 6.2. The process of making and reporting a probability estimate.....	286

List of Tables

Table 1.1 Expected utilities of different courses of action.....	14
Table 2.1 Dimensions of interoception	71
Table 2.2 Dimensions of interoceptive sensibility.....	80
Table 2.3 Correlations between the four interoceptive sensibility scales in Experiment 1	88
Table 2.4 Orthogonal contrasts tested in Experiment 1	92
Table 2.5 Orthogonal contrasts tested in Experiment 2	109
Table 2.6 Orthogonal contrasts tested in Experiment 3	134
Table 3.1 Results of Experiments 4, 5 and 6.....	159
Table 3.2 Probability estimation questions asked in Experiment 7	162
Table 4.1 Types of control and their definitions	174
Table 4.2 Orthogonal contrasts tested in Experiment 8	189
Table 5.1 Control manipulation check in Experiment 12	251
Table 5.2 Mood manipulation check in Experiment 12.....	252

Chapter 1 Theoretical Background

Pessimism is, in brief, playing the sure game. You cannot lose at it; you may gain. It is the only view of life in which you can never be disappointed. Having reckoned what to do in the worst possible circumstances, when better arise, as they may, life becomes child's play.

-Thomas Hardy (1902)

Living in a world dominated by uncertainty, we are confronted with the situation of not knowing whether any given proposition is true on a daily basis. For instance, imagine an office worker leaving her house in the morning, looking up to the sky and pondering about the chances of rain. In order to provide an estimate, she will take several factors into account. Are there any clouds? Am I in London? One question that has recently gained attention is whether her estimate will also be determined by how bad it would be if it were to rain. Will she estimate the chance of rain to be higher when she has an important meeting and does not want to turn up soaking wet, compared to when she is sitting inside and does not mind the weather? And, as will be investigated in the current thesis, will the probability estimate differ depending on whether the office worker is a manager or a subordinate?

How likely we estimate the occurrence of an event to be governs concrete choices and decisions about the event in question. For instance, the office worker's belief about the likelihood of rain will determine her choice to go back home to take an umbrella to work. Normatively, when determining a course of action it has been proposed that we should combine the probability of an event with its utility. The probability of an event is the result of a quantification of uncertainty, varying between the two extremes of truth or falsehood, assigned the values of "0" and "1", respectively.

The utility of an event is defined as “a measure of the consequence’s *subjective* value in relation to other consequences” (Fishburn, 1981, pp. 141), and as such constitutes the subjective “goodness” or the “badness” of an event. Notably, when basing potential courses of actions on the likelihoods of future outcomes, these likelihoods are only of importance if the outcomes have utilities attached to them – otherwise rendering the course of action irrelevant.

Probably the most fundamental normative model is the Subjective Expected Utility theory (SEU, Savage, 1954), which shall be illustrated here in more detail. The SEU proposes that in order to make the best possible judgment under uncertainty, individuals should combine the utility of an outcome with its associated probability and take action in the direction of the event with the greatest expected utility (EU). For example, imagine Alex wants to make a decision as to whether to take the bus on her way home, or to walk. Alex estimates the likelihood that the bus is on time to be 70%, and the likelihood that the bus is delayed to be 30 %. Additionally, Alex assigns a utility of +10 to walking, a utility of +40 to taking the bus should it be on time, and a utility of -100 of taking the bus should the bus be delayed. Table 1.1 shows the possible outcomes and their associated utilities (a positive direction indicates a positive utility and a negative direction indicates a negative utility). Following the SEU (Savage, 1954), Alex should now weight the utilities by their probabilities to achieve an EU for each of the possible actions in order to make a decision which rationally informs action.

Table 1.1

Outcome utilities of possible actions combined with possible events and expected utilities corresponding to possible actions.

		Possible Outcomes		Expected Utilities
		<i>Bus on Time (70%)</i>	<i>Bus Delayed (30%)</i>	
Possible Actions	<i>Bus</i>	+40	-100	EU = (+40 x 70%) + (-100 x 30%) = -2
	<i>Walking</i>	+10	+10	EU = (+10 x 70%) + (+10 x 30%) = +10

Since the EU of walking (+10) is greater than the EU of taking the bus (-2), Alex should walk, and not take the bus. Whilst this normative model is seemingly straightforward, in real life individuals are not always able to adhere to normative models of decision-making (see Kahneman & Tversky, 1982). An aspect of central importance to the present thesis is that estimates of probabilities are often distorted by the decision context.

Influence of Utility on Subjective Probability Estimates

In everyday life most probabilities of events are not certain and cannot be assessed with calculations. Whilst we know that the possibility of throwing a six with a regular dice in any single throw is $\frac{1}{6} \cong 17\%$, we cannot know the precise probability of a bus being delayed, whether it will rain today, or whether the San Antonio Spurs will win the NBA play-offs next season. Thus, when making a decision about a future event, not only the assigned utilities, but also the associated probabilities are often subjective. In this thesis, we take the Bayesian approach to probability and assume that “the degree

of belief in the truth of a particular proposition” can be represented as a subjective probability (Ayton & Wright, 1994, p. 164). Furthermore, note that this thesis is observing only single-event probabilities. That is, types of probability estimates that inform the likelihood of one specific event. This, however, does not mean that individuals do not also make use of frequency data, which can inform prior beliefs in a Bayesian fashion. For example, if the bus operator states that 9 out of 10 busses are on time, then one should make use of this frequency data to estimate the likelihood that one’s specific bus will be late (see also the “principal principle”, Pettigrew, 2012). On the other hand, a Bayesian approach allows also for a subjective degree of belief about novel events where no prior frequency data is available.

One fundamental research question that has arisen in recent years is whether, and if so how, the utility of an event biases subjective probability estimates. That is, do we assign higher or lower probabilities when we *want* vs. *don’t want* an event to occur?

Identifying the circumstances under which such a bias could occur is of importance, as normative theories of decision-making assume independence of utility and probability. Whilst it is not a normative requirement of the framework of SEU that utility and probability are orthogonal to one another, Edwards (1962) states that “it is very difficult to see how the model could be applied to real decisions unless some such assumptions were made” (see also Edwards & von Winterfeldt, 1986). Moreover, most descriptive models which examine decision-making under risk (e.g. Prospect Theory, Kahneman & Tversky, 1979) are further developments of the SEU model, involving utilities and probabilities as components. Therefore, the question of the ‘if’ and the ‘how’ of an interdependence of utility and probability is fundamental.

This question is not only fundamental for theoretical underpinnings, but also with regard to the real-life consequences that occur when one misestimates probability estimates for events with utilities attached to them.

Negative and positive events differ with regard to the urgency of a correct estimate: Probabilities for negative events are undoubtedly probabilities one should get right, as they come with more drastic consequences should they occur in comparison to the non-occurrence of a positive event. Whilst foregoing a reward due to a misestimate can indeed lead to regret (Weber, 1994), underestimating the occurrence of a negative event can potentially lead to death. For example, imagine a parent failing to vaccinate their child due to underestimating the likelihood of their child catching a disease, or for someone to underestimate the likelihood of having a disease and to not seek medical advice despite having symptoms. It is this greater importance of understanding the nature of the relationship between utility and probability estimates in the negative domain that motivated the emphasis of this thesis to examine negative events in particular.

However, before presenting the evidence for a potential bias of utility on probability estimates in more detail, we would like to define the type of bias we examine in the current thesis both theoretically and empirically. The process of making and reporting a probability estimate can be broken down in three phases, all of which could potentially be biased by the utility of the event. First, evidence is accumulated and selected, then the estimate is produced, and subsequently reported (see Figure 1.1).



Figure 1.1. The process of making and reporting a probability estimate (as in Harris, Corner, & Hahn, 2009).

There is ample evidence that utility can bias the evidence accumulation stage. For example, Dai, Wertenbroch, and Brendl (2008) demonstrated that the utility of an event is used as a proxy for base rate estimates, reflecting real-life relationships of utility and base-rates (in that the more positively valenced an event is, the less frequent it is, Bock, 1968; Hirshleifer, Glazer, & Hirshleifer, 2006; Pleskac & Hertwig, 2014; see also Mandel, 2008). More, utility has been shown to impact the evidence selection stage through memory processes, where participants were more likely to remember that desirable information comes from a reliable, accurate source, than from an unreliable, less accurate source (Gordon, Franklin, & Beck, 2005). Furthermore, previous research has demonstrated that the salience of an event, which in turn is often impacted by its desirability, can equally impact the evidence-accumulation stage (Bar-Hillel, Budescu, & Amar, 2008). However, in the current research we aim to investigate theories and their mechanisms that are concerned with an *inherent* bias at the stage of the *internal estimate*. That is, we aim to examine the mechanisms of an interdependence of utility and probability as a feature of human cognition that does not depend on (biased) evidence selection.

Furthermore, note that we examine the evidence for a potential impact of negative and positive utility on probability estimates separately from one another. This distinction derives from the assumption that rather than classifying affect as a

unidimensional variable from positive to negative, evidence points towards the bidimensional and independent character of positive and negative affective states (see e.g. Berscheid, 1983; Bradburn, 1969; Isen, 1984; Watson, Clark, & Tellegen, 1988; Watson & Tellegen, 1985; Zevon & Tellegen, 1982; for a discussion see Taylor, 1991). That is, the presence of negative affect does not necessarily constitute the absence of positive affect (or vice versa). For example, Watson, Clark, & Tellegen (1988, also Watson, & Tellegen, 1985; Diener, Emmons, Larsen, & Griffin, 1985; Warr, Barter, & Brownbridge, 1983), developing the “Positive and Negative Affect” (PANAS) scales, showed that positive and negative affect were not correlated (but see Green, Goldman, & Salovey, 1993; Russell, 1980). Furthermore, Costa and McCrae (1980) demonstrated that distinct personality variables are related to positive and negative affect: neuroticism was related only to negative but not to positive affect, extraversion was related to positive, but not to negative affect. Consequently, it would be likely that, if there is reason to believe that the interdependence exists in both the positive and negative domain, different mechanisms underlie these relationships. Hence, this thesis will compare estimates for negative events to estimates for neutral, rather than to positive events.

Chapter Overview

Having provided some background to the constructs of utility, probability, normative models and the nature of a potential interdependence, we now aim to discuss the empirical evidence for the interdependence of utility and probability. First, we present and discuss the current evidence for the existence of such an interdependence in the positive domain, namely for the wishful thinking bias – the bias to assign a higher

probabilities to positive compared to neutral future events, as well as to review evidence for the existence of comparative unrealistic optimism. Note that as this thesis' focus is on the negative domain, the presentation of the literature on the evidence for the existence of a wishful thinking bias and unrealistic optimism is necessarily selective. However, pre-empting the results of our review, the evidence for an impact of positive utility on probability estimates is not convincing.

Next, we present evidence for the existence of a negativity bias, showing that individuals appear to assign higher probabilities to severely negative, compared to mildly negative or neutral events.

Following, we introduce the construct of social power and introduce theories describing mechanisms through which social power is hypothesized to impact behaviour and cognition. We then discuss the (conflicting) predictions of these theories on the impact of social power on the relationship between negative utility and probability estimates.

Afterwards, we discuss 5 state-of-the-art theories on potential mechanisms for a negativity bias, and describe how social power relates to these theories. In particular, Construal Level Theory (Wakslak, 2009), Approach-Avoidance Motivation (Lench, 2009), Imaginability (Bilgin, 2012), Loss Function Asymmetries (Harris et al. 2009), and Arousal Misattribution (Vosgerau, 2010) will be discussed. As a first investigation into the impact of social power on probability estimates, this thesis focuses concretely only on arousal misattribution and loss function asymmetries, and we discuss empirical approaches of testing these accounts. In particular, we introduce a theoretically driven moderator for each of the two accounts. Providing evidence showing that social power is

related to both moderators, we then describe the theoretical setting for utilising social power as a means to distinguish between an account of loss function asymmetries and arousal misattribution. Specifically, we create two models that incorporate different processes through which social power could impact the relationship of negative utility and probability estimates.

Interdependence of Utility and Probability

Does Wishful Thinking Exist? By definition, wishful thinking exists when “the desirability of an outcome leads to an increase in the extent to which it is expected to occur” (Krizan & Windschitl, 2007, p. 96). A truism in folk psychology, there also appear to be seemingly numerous indications from field study research suggesting that wishful thinking really is a well-established empirical phenomenon (Babad, 1995; Babad & Katz, 1991; Granberg & Brent, 1983; Massey, Simmons, Amor, 2011; Simmons & Massey, 2010). On the other hand, the wishful thinking effect has been described as “elusive” (Bar-Hillel & Budescu, 1995) and the evidence for its existence as “surprisingly thin” (Krizan & Windschitl, 2007, p.95). Illustratively, direct tests of the wishful thinking phenomenon have reported conflicting results: some found evidence for wishful thinking (Price, 2000; Pruitt & Hoge, 1965), others were unable to demonstrate any effects of outcome utility (Bar-Hillel & Budescu, 1995; Erev & Cohen, 1990) while still others even found the opposite, a pessimism bias (Bilgin, 2012; Dai, Wertenbroch, & Brendl, 2008; Harris, Corner & Hahn, 2009; Mandel, 2008; Vosgerau, 2010). Thus, despite a large amount of research, to date there is little agreement as to whether or not probability estimates are truly systematically biased

towards positive events (see Krizan & Windschitl, 2007, for a review; Krizan & Windschitl, 2009, for a meta-analysis).

A vast amount of research exploring wishful thinking effects has been conducted outside the laboratory in ‘real world’ experiments. In particular, research findings in the domain of sports gave some reason to believe that probability estimates are indeed biased for desirable events. In the typical set-up of these experiments, participants report greater likelihoods of their own sports-teams winning (Babad & Katz, 1991; Babad, 1987), and this bias persists even when accuracy is incentivised (Massey, Simmons & Amor, 2011; Simmons & Massey, 2010). However, the same effects are demonstrated simply when making one team more salient than another - when told that “we are particularly interested in team X”, participants assigned higher likelihoods of this team winning in Bar-Hillel, Budescu and Amar’s (2008) experiments. Likewise, a salience effect might also underlie the results achieved by Massey et al. (2011) (and comparable findings in other domains such as wishful thinking about political outcomes, Krizan, Miller, & Johar, 2010). Thus, studies outside controlled settings can be well explained as “an unbiased evaluation of a biased body of evidence” (Bar-Hillel & Budescu, 1995, p.100), and as such do not constitute evidence for a bias of the “internal estimate” (see Figure 1.1).

However, also literature investigating the wishful thinking bias under more controlled laboratory settings is subject to criticism. For instance, laboratory research demonstrating that when participants are divided into teams they assign higher likelihoods to members of their own team winning in a competitive game (Price, 2000) can be explained by motivational factors such as the protection of the group’s positive

self-image (Jourden & Heath, 1996; see Hinkle & Schopler, 1986, for a review), and therefore do not constitute evidence for wishful thinking.

Nevertheless, some studies in laboratory settings have avoided the above mentioned confounds. Instead, early studies typically used choice paradigms. For instance, the classic ‘marked-card’ paradigm (Marks, 1951) asks participants to guess whether they will draw a marked or a non-marked card from a deck of cards, where drawing a marked card can either be a monetary gain or a loss. Typically, individuals are not “rational” in responding but attribute more than 50% to the probability that they will win (e.g. Irwin, 1953). However, Windschitl, Smith, Rose and Krizan (2009) demonstrated that guessing a favourable card does not imply that participants hold genuinely biased estimates of probability, and although their guesses *appear optimistic*, participants *maintain realism* in their subjectively held probability estimates.

Other studies have faced issues that accompany the manipulation of utility by means of monetary incentives. For instance, Pruitt and Hoge (1965) told participants that they could either win or lose money based on an outcome whose probability they were asked to estimate. Participants assigned greater probabilities to the outcome when they could win, as compared to when they could lose money. However, as argued in Harris et al. (2009), participants in the loss condition might have perceived it as unlikely that they would leave the experiment having to *pay* the experimenter, and might have therefore assigned a lower estimate to the negative outcome compared to participants’ estimates in the win condition.

More recent experiments can be criticised for the conceptualisation of affect. As briefly mentioned earlier, affect has often been conceptualised as a unidimensional

variable from negative to positive states. According to this conceptualisation, wishful thinking is conceptualized as higher estimates for positive than for negative outcomes (see left panel of Figure 1.2). For instance, in one of the most extensive research programmes in the domain of wishful thinking, Bar-Hillel and Budescu (1995) compared positive to negative outcome probabilities, and participants did not assign higher likelihoods to positive compared to negative events. However, evidence points towards affect as a bidimensional variable (e.g. Berscheid, 1983; Bradburn, 1969; Isen, 1984; Watson, Clark, & Tellegen, 1988; Watson & Tellegen, 1985; Zevon & Tellegen, 1982). Harris (2009) stresses the notion to instead investigate positive and negative utilities separately from one another, and consequently to compare the estimates for negative or positive events to neutral events instead of to each other. For instance, it is possible that both a wishful thinking bias and a negativity bias exist. However, this would not necessarily promote a difference between positive and negative events, but a significant increase of likelihood judgments from both positive and negative events to the neutral condition (see right panel of Figure 1.2).

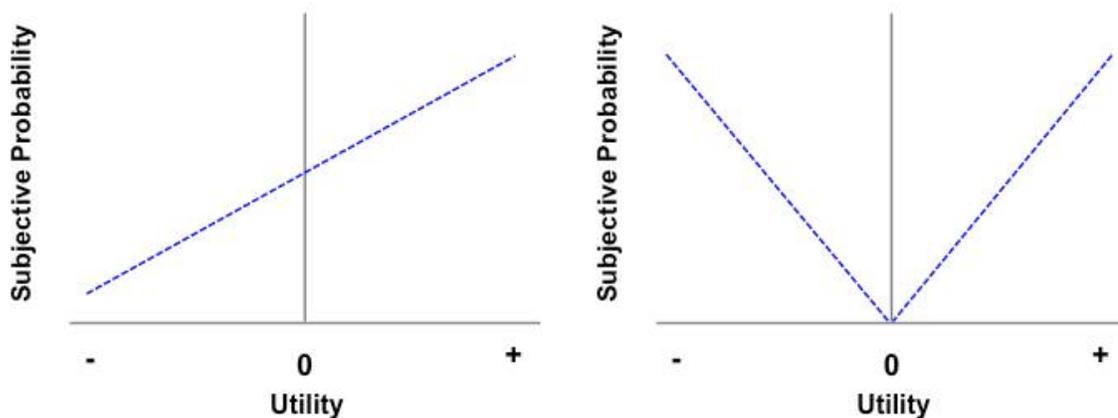


Figure 1.2. The left panel illustrates the relationship of utility and subjective probability from a general optimism theory, whereas the right panel illustrates the relationship for an “it can happen to me” account (Slovic, 1966, p. 23, reproduced with the permission of the author).

Recent work by Vosgerau (2010) included a neutral, a positive and a negative condition, potentially allowing for detecting effects in line with the right hand panel of Figure 1.2. In this experiment, participants gave higher probability estimates if they had something to gain or to lose compared to when they did not. Specifically, Vosgerau (2010) collapsed across positive and negative conditions and compared these with the neutral outcome condition. However, for reasons explained above, this does not constitute evidence for a general wishful thinking bias. By collapsing across the positive and negative utility conditions, it is impossible to know whether the difference to the neutral condition is driven by both the negative and the positive outcome conditions. It might indeed be possible that only the negative condition, in which participants gave statistically higher estimates than in the positive condition, was driving these differences. However, although Vosgerau (2010) did not statistically assess these

differences, participants in the positive condition gave higher probability estimates than participants in the neutral condition ($M_{\text{positive}}=19.77$ (16.07) vs. $M_{\text{neutral}}=13.44$ (10.43)).

More recently, de Molière, Harris, Quantmeyer and Hahn (2014) systematically went “on the hunt” for wishful thinking. To overcome the criticisms mentioned above, objective representations of probabilities were presented in minimal experimental paradigms. Individuals were presented with either positive or neutral outcome scenarios and asked to estimate the probability of the outcomes occurring. However, across ten experiments, including experiments which included potential moderators predicted to set up favourable conditions for observing such effects, de Molière et al. (2014) were unable to find any evidence for the existence of a wishful thinking bias. This conclusion was strengthened by their Bayesian meta-analysis across those ten experiments demonstrating that a null hypothesis was *twenty-four* times more likely than the evidence for a wishful thinking hypothesis, approaching “very strong evidence” (Rouder et al. 2009, p.228) in favour of a null over a wishful thinking hypothesis.

Comparative Unrealistic Optimism. Whilst the above research examined wishful thinking, that is, tested for greater likelihood estimates for positive compared to neutral events, we should at this stage also mention research ostensibly demonstrating that individuals are ‘unrealistically optimistic’ about their chances of experiencing a negative life event. Comparative unrealistic optimism (UO, Weinstein, 1980) is a phenomenon whereby people appear to perceive their chances of experiencing a positive life event (such as graduating on top of the class) as greater than the average person’s, and their chances of experiencing a negative life event (such as getting cancer) as lower. UO is typically assessed by asking participants to rate their likelihood of experiencing

positive or negative life events *compared to the average person's likelihood*. That is, participants are asked to indicate whether their chances of experiencing this future event are smaller, the same, or greater compared to the average person's chances. Therefore, UO is assessed on a group- rather than on an individual level: As it is impossible to know the true probability of a particular person experiencing such events (e.g. one's chance to get cancer is determined by multiple environmental and genetic factors), the logic behind this method is that on a group level, the difference between the average person's risk and average of the individual ratings should be zero, if unbiased. The common finding is that individuals rate their risk to experience a negative life event as smaller than the average person's, implying UO (e.g. Campbell, Greenauer, Macaluso, & End, 2007; Price, Pentecost, & Voth, 2002; Weinstein, 1980, 1982, 1984; Weinstein & Klein, 1995; Weinstein, Marcus, & Moser, 2004). However, Harris and Hahn (2011) demonstrated that the above described response patterns might in fact be a result of statistical artefacts (scale attenuation, minority undersampling and regression to the mean). Harris and Hahn (2011), using the above describe comparison method, simulated what would be classified as unrealistic optimism with perfectly rational (unbiased) agents for negative events with a base-rate of less than 50% (the majority of events in the typical versions of this paradigm). Furthermore, recent data has demonstrated that individuals also rate their chances of experiencing a rare positive event (with a prevalence of less than 50%) as lower than the average person's (Chambers, Windschitl, & Suls, 2003; Harris, 2009; Kruger & Burrus, 2004). These findings are in line with a statistical artefact account – and cast some doubt over the interpretation of human

ratings as optimistic in the comparison method (but see Shepperd, Klein, Waters, & Weinstein, 2013).

Another method used more recently to assess an optimism bias is the “Update Method” (Sharot, Korn, & Dolan, 2011). Participants are asked to estimate their chances of experiencing a negative life event, and after having made their estimate, are presented with the true base-rate of experiencing the event. Afterwards, participants are asked to re-estimate their chances, having now received the average person’s likelihood. The typical finding is that participants update their beliefs more for ‘desirable information’ (when the true base-rate was lower than the participants’ estimate) than for ‘undesirable information’ (when the true base-rate was higher than the participants’ estimate) (e.g. Sharot, Korn, & Dolan, 2011; see also Sharot et al. 2012a; Sharot et al., 2012b). However, Harris, Shah, Catmur, Bird and Hahn (2013) showed that rational agents can produce similarly biased updating patterns – due to a misclassification of what is ‘desirable’ and ‘undesirable’ information in the belief updating method: Participants’ overall estimate of their own risk (before receiving the true base-rate) should have been a combination of both an estimate of their own personal risk and their estimate of the average person’s risk. Therefore, upon receiving an update of the base-rate, participants should not necessarily change the estimate of their own risk, but their estimate of the base-rate. Harris et al. (2013) conclude that without knowing the participants’ estimate of the base-rate for the average person, it is not possible to classify base-rate information as ‘desirable’ and ‘undesirable’.

Summing up, the evidence for the existence of a general wishful thinking bias and unrealistic optimism remains, at best, elusive. Next, we will present evidence for an interdependence of utility and probability in the negative domain.

Evidence for a Negativity Bias. Evidence for the existence of a negativity bias was first provided from research on verbal probability estimates. Whilst probabilities are inherently numerical, in everyday life we often use verbal probability expressions to describe the likelihoods of future events (such as “likely”, “unlikely”, etc.). How these subjective verbal probability expressions are interpreted, however, depends largely on the context. For instance, the impact of negative utility on the perception of verbal probability expressions was examined by Weber and Hilton (1990). When statistically controlling for the event frequency, Weber and Hilton (1990) demonstrated that individuals assigned higher numerical probabilities to verbal probability expressions when the event in question (e.g. getting the flu) was severely negative compared to neutral (see also Verplanken, 1997). Avoiding the potential confound of real-life event frequencies, Harris and Corner (2011) created novel, fictional scenarios. Across three experiments, Harris and Corner (2011) demonstrated that verbal probability expressions were assigned a greater subjective likelihood when describing the likelihood of more severely negative events compared to mildly negative events.

More recently, studies have demonstrated a negativity bias also outside the domain of verbal probability estimates. For instance, Risen and Gilovich (2007) found that individuals judged the likelihood of a lottery ticket winning to be higher if the ticket would be of an enemy (negative) rather than a friend (positive). In a similar vein, Bilgin (2012) told participants to imagine that there was a 70% chance that they would have to

move into a new office that was either better or worse than their current office. When participants imagined moving into the worse office, they reported a higher subjective probability of moving than when they imagined moving into a better office.

However, central to the current investigation are the studies conducted by Harris et al. (2009), whose research is informative both with regard to methodologies employable to investigate a bias of the internal estimate as well as to potential moderators of the effect (as shall be addressed later). In their research, Harris et al. (2009) asked participants to estimate the likelihood of either neutral or negative events. Crucial to the design, participants were presented with an unambiguous, objective representation of probability (e.g. a visual cell matrix with varying proportion of yellow and grey cells). For example, in their Experiment 1 the cell matrix represented an orchard and each cell represented an apple tree. Participants in the negative condition were asked to estimate that a child playing in an orchard would eat a poisonous apple, as represented by a yellow cell in the matrix (the grey cells representing harmless apples). Participants in the neutral condition on the other hand read that the yellow cells represented a sour tasting, but harmless fruit. Harris et al. (2009) found that participants gave higher outcome probabilities when estimating the likelihood of negative as compared to neutral events. As such, these studies demonstrate the impact of negative utility on probability under controlled conditions and forego some of the confounds and criticism raised in previous research, as discussed in the section on wishful thinking above.

To conclude, the evidence for the existence of an impact of negative utility appears to be strongly supported, and recent research has started to examine the

mechanisms that could underlie this effect. However, before introducing these mechanisms in more detail, we now turn to the construct of social power as a potential moderator of the impact of negative utility on probability estimates. Examining this potential moderation is of interest for two reasons: First, the predictions made by theories on social power with regard to the direction of a moderation are conflicting. Secondly, social power is introduced at this stage as we aim to later link it to some of the to-be-introduced mechanisms. Since social power is related to some of the above proposed mechanisms, it can help to both dissociate between theories on social power, and theories on the impact of negative utility on probability

Social Power

“The fundamental concept in social science is Power, in the same sense that Energy is the fundamental concept in physics... The laws of social dynamics are laws which can only be stated in terms of power”

Russell, 1938, p.10

Social power is a fundamental, and maybe the most basic, force in social relationships (Fiske, 1993). We define power as one’s potential to influence others in psychologically meaningful ways (French & Raven, 1959) by giving reward or administering punishment (Fiske & Dépret, 1996; Keltner et al., 2003)¹. These asymmetries in control over valued resources in social relationships (Thibaut & Kelley, 1959) lead to fundamental differences in the way individuals on one or the other end of this asymmetric distribution act, feel, and think (Keltner et al. 2003).

¹ We define power as a construct distinct from social status. Status is defined as the position of respect and prestige an individual in a social hierarchy has judged to have by other members of the group (e.g., Ridgeway & Walker, 1995). Unlike social power, where the control over individuals is objective in terms of the resources an individual can or cannot give (Keltner et al. 2003), social status is primarily subjective (Foa, 1971; Podolny, 1993).

Those high in social hierarchies have the freedom to act independently, and as a result, feel in control (Fast, Gruenfeld, Sivanathan & Galinsky, 2009). Consequently, power has been associated with enhanced well-being (Kifer, Heller, Perunovic, & Galinsky, 2013), health (Scheepers, de Wit, Ellemers & Sassenberg, 2012), wealth (Domhoff, 1998) and greater happiness (Berdahl & Martorana, 2006; Gonzaga, Keltner, & Ward, 2008). Powerful individuals pursue goals with greater efficiency and persistency (Guinote, 2007b), are better negotiators (Magee, Galinsky & Gruenfeld, 2007) and are more action oriented (Galinsky, Gruenfeld & Magee, 2003). On the other hand, individuals low in social power are dependent upon individuals high in social power (Emerson, 1962) and are acutely aware of the constraining environment, particularly of the persistent threat of losing the benevolence of the powerful (e.g. Anderson & Berdahl, 2002; Fiske, 1993). As a result, powerlessness has been associated with reduced executive functioning (Smith, Jostmann, Galinsky, van Dijk, 2008), decreased authenticity (Kraus, Chen & Keltner, 2011), and decreased accessibility of goal-relevant constructs during goal-striving (Slabu & Guinote, 2010).

Moreover, social power affects judgment and decision-making. For example, powerful individuals have been demonstrated to be more overconfident in their decision-making: Fast, Sivanathan, Mayer and Galinsky (2011) showed that overconfidence led the powerful to provide smaller 95% confidence intervals when making estimates about NHL players' performance over a year, compared to the powerless. As a consequence of the decreased confidence interval, the powerful's predictions were less accurate than that of powerless or control participants. Additionally, Inesi (2010) demonstrated that power asymmetries lead to a differential

valuation of decision outcomes. Inesi (2010) showed that the powerful value anticipated gains more, and anticipated losses less than the powerless. For example, after being primed with power or powerlessness, participants were told that they were entered into a lottery for a £25 Costa Coffee [EAT] voucher. Next, they were informed that they might have the option to exchange their voucher should they win for a £40 voucher for EAT [Costa], a £40 voucher for Costa [EAT], or to keep the £25 voucher. Powerful individuals were more likely to exchange the voucher for a different retailer, and it was argued that the powerful were less loss averse than the powerless. What are the concrete mechanisms that might underlie these differences in behaviour, cognition and affect of powerful compared to powerless individuals?

Theories of Social Power. Whilst we will discuss the theories on mechanisms underlying the effects of social power in more depth in Chapters 4 and 5, we will next provide a brief overview of the most fundamental frameworks of social power. This overview will describe the predictions of different theories on the impact of negative utility on probability for powerful and powerless individuals and help positioning the rationale of the thesis in light of these predictions.

Probably the most fundamental framework of social power (1289 citations on Google scholar, 08.10.2014) is the “*Approach-Avoidance Theory of Power*” (Keltner, Gruenfeld & Anderson, 2003). Keltner et al. (2003) argue that enhanced power is associated with living in reward rich environments, freedom to act without consequences and striving autonomously for goals. Powerless individuals on the other hand have decreased access to resources and live in an environment dominated by punishment and threat. According to Keltner et al. (2003), these asymmetrical

differences in exposure toward reward and threat lead to different motivational styles: the powerful are hypothesized to be more approach oriented and the powerless to be more avoidance oriented. The notion that the powerful are more approach oriented compared to the powerless has found ample support in recent research. For example, in line with greater approach motivation, the powerful negotiate more (Magee et al. 2007), are more action oriented (Galinsky et al. 2003) and experience approach related affective states such as happiness and anger (Van Kleef, De Dreu, Pietroni, & Manstead, 2006). Moreover, approach and avoidance motivation have been strongly associated with risk seeking and risk averse behaviour, respectively (e.g. Gray, 1990; Higgins, 1998). As a result, this account predicts greater risk taking of the powerful compared to the powerless (Keltner et al. 2003). Due to their greater avoidance motivation, which is associated with greater attention paid to threatening information (e.g. Gray, 1990), the approach-avoidance theory predicts that powerless individuals should estimate the probability of a negative event to be greater than the powerful.

On the other hand, power increases attentional focus on internal states (Guinote, 2007a, 2010). Due to fewer constraints, the powerful enact internal states and automatic responses in line with situated experiential information (*Situated Focus Theory of Power* (SFTP), Guinote, 2007a). For example, the powerful's but not the powerless' amount of food eaten is predicted by actual hunger levels, and the powerful show greater reliance on subjective experiences accompanying thought processes (e.g. ease of retrieval, Weick & Guinote, 2008). These results give some evidence for the powerful's greater reliance on experiential information in the areas of bodily feelings (hunger) and thought processes. The SFTP predicts that powerful individuals pay greater attention to internal

states, and act on information relevant to the situation. Therefore, according to the SFTP, if the utility of a likelihood judgment is accessible and an integral part of the decision, the powerful will be more likely to incorporate it in their decision-making process. In addition, as will be described in more detail in a later section, one proposed mechanism underlying the impact of negative utility on probability estimates is arousal misattribution. As has been proposed by Vosgerau (2010), arousal that is created by negative utility is misattributed to the likelihood of the occurrence of the event, in that more negative (and more arousing) events are estimated to be more likely than less negative (and less arousing) events. If this were the case, according to the SFTP the powerful will demonstrate a greater negativity bias than the powerless due to their greater reliance on experiential information.

Summing up, there is no consensus of predictions as to whether individuals high or low in social power will be more biased by negative utility when estimating probabilities. The approach-avoidance theory of power predicts that individuals low in social power should give higher probability estimates for negative events than the powerful due to greater attention to threatening and negative information and negative affective states. On the other hand, the SFTP predicts that the powerful will assign higher probability estimates due to greater attention paid to internal states and focal information accessible in the situation. Thus, focussing on negative utilities is motivated also by the fact that, by examining power as a moderator, we can distinguish between theories of social power, as the approach-avoidance and the SFTP accounts have conflicting predictions about the direction of such a moderation.

Next, we turn to five mechanisms that have been proposed to underlie the relationship of negative utility on probability estimates: The construal level theory (Wakslak & Trope, 2009), the approach-avoidance theory (Lench, 2009), an imaginability account (Bilgin, 2012), an account of Asymmetric Loss Functions (Harris et al. 2009) and arousal-misattribution (“Stake-Likelihood Hypothesis”, Vosgerau, 2010).

Social power can help dissociating between theories of the interdependence of negative utility and probability by being related to some of the mechanisms proposed to underlie this relationship. In particular, as will be discussed, social power is strongly related to the Stake-Likelihood Hypothesis and the account of Asymmetric Loss Functions. Whilst we will theoretically describe all five accounts, as this is the first research examining the relationship between social power and probability estimates for negative events, we empirically focus on the two theories that have the strongest links to social power: the Stake-Likelihood Hypothesis and the asymmetric loss function account.

Mechanisms Underlying the Interdependence of (Negative) Utility and Probability

Having provided evidence that negative utility can indeed bias probability estimates, we now turn to mechanisms that have been proposed to underlie this bias.

Construal Level Theory. Construal Level Theory (CLT, Trope & Liberman, 2003) has considerable precedent in social psychology (for a review see Trope & Liberman, 2010). The CLT proposes that individuals construe distant experiences, such as past and future events, by creating abstract, high level (abstract) mental construals of distal events or objects. Closely related, the CLT also offers a framework for the

construct of psychological distance (Trope & Liberman, 2000, 2003). Psychological distance is the subjective experience about how far an event or an object is from the current self. When an object or an event is perceived as distant, it is construed more abstractly, when an object or an event is perceived as near, it is construed more concretely. Individuals have been shown to vary in psychological distance across four major dimensions: time (how far away the object is in the present or future), spatial distance (remoteness vs. closeness in space), personal distance (high vs. low familiarity with social objects), and probability distance (how unlikely vs. likely the event is), and these four dimensions have been recently shown to be positively correlated with one another (Fiedler, Jung, Wänke & Alexopoulos, 2012; Fiedler, Jung, Wänke, Alexopoulos & de Molière, 2014). Likewise, manipulating a high or low distance of one psychological distance dimension can lead to the perception of high or low distance of another dimension, respectively (Stephan, Liberman, & Trope, 2011). Relatedly, Wakslak (2012) demonstrated that individuals assume that events with a low base-rate (and thus of greater psychological distance) are assumed to happen in more distal contexts, whereas events with a higher base-rate (low psychological distance) are judged to happen more in near contexts.

Furthermore, given the relationship between levels of construal and psychological distance, Wakslak and Trope (2009; see also Wakslak, Trope, Liberman & Alony, 2006; Wakslak, 2012) proposed that when an event is construed abstractly and is arguably high in psychological distance, the associated probability will decrease compared to when it is construed concretely and is low in psychological distance. Wakslak and Trope (2009) showed that individuals, when primed with a high vs. low

construal mind-set assigned low vs. high probability judgments to the occurrence of events, respectively. In their Study 4, Wakslak and Trope (2009) perceptually manipulated construal by means of the Navon task, where participants either focus on global or on local stimuli, manipulating abstraction or concreteness, respectively. Next, participants answered a questionnaire consisting of six statements (e.g. “Tom is waiting for the subway. How likely is the train to be late?”, “Lea is planning on going to the drugstore to buy shampoo. How likely is it that the item will be on sale?”). Individuals who adopted a high-level construal mind set gave lower probability estimates than individuals who adopted a low-level construal mind-set.

Although Wakslak (2012) does not test or discuss the impact of utility on likelihood judgments, one can infer concrete predictions from the construal level theory: since negative information is processed more concretely (lower distance) than neutral information, and positive information is processed more abstractly (greater distance) (Lieberman & Trope, 1998), estimating the likelihood of negative events should be construed on a lower level than estimating the likelihood of positive events. Thus, we can infer that the CLT account predicts that higher probabilities should be assigned to negative, and lower probabilities to positive events. However, considering the evidence for wishful thinking, where a null hypothesis (rather than any impact of positive utility on probability estimates) appears to be most strongly supported, the CLT does not appear to receive support in the positive domain. However, the prediction that events with a negative outcome utility, construed more concretely, would be assigned greater likelihoods, receives support (e.g. Bilgin, 2012; Dai et al. 2008; Harris et al. 2009; Mandel, 2008; Risen & Gilovich, 2007; Vosgerau, 2010).

Approach-Avoidance Motivation. Approach and avoidance motivation is an important factor in animal learning (Thorndike, 1935), human decision making processes (Kahneman & Tversky, 1979) and personality (Elliot & Thrash, 2002). Approach motivation is defined as “the energisation of behaviour by, or the direction of behaviour toward, positive stimuli (objects, events, possibilities)”, whereas avoidance motivation is defined as “the energisation of behaviour by, or direction of behaviour away from, negative stimuli (objects, events, possibilities)” (Elliot, 2008, p.3). In its most basic impact on human behaviour, individuals high in approach motivation are more reactive to positive cues in their environment, whereas individuals high in avoidance motivation are more reactive to negative cues (Gray, 1972). However, humans are not simply driven by these two motivational systems by taking actions towards rewards or away from punishments (Carver & White, 1994), but have developed the ability to consider the possibility for reward or punishment before the event in question happens (Gilbert & Wilson, 2008). Approach and avoidance motivation are manifested as the “vulnerability, or susceptibility, to a particular class of clues” (Carver and Scheier, 1994, p. 325), and could potentially also be related to beliefs and expectations about the frequency with which one encounters such stimuli. Moreover, the relationship between approach-avoidance motivation and affective states is reciprocal: approach-avoidance motivation makes individuals more attentive to stimuli of the compatible emotional valence, but positive and negative stimuli also signal approach and avoidance motivation, respectively (Frijda, 1986, but see Carver, 2004, for exceptions).

Given the relationship between approach-avoidance motivation and valenced stimuli, it has been proposed that the prospect of a positive or negative future event can trigger approach and avoidance motivation (Lench, 2009). In particular, the research by Lench (Lench & Bench, 2012; Lench & Darbor, 2014; Lench, 2009) has examined how these motivations can impact judgments about the likelihood of future events to occur. Aiming to provide a mechanism for optimistic and pessimistic likelihood estimates, Lench (2009) proposed the automatic optimism theory (AOT). According to the AOT, individuals estimate the occurrence of positive events as likely and of negative events as unlikely, due to affective reactions. Lench (2009) hypothesized that underlying this relationship between affect and probability are the motivations to approach and avoid positive and negative events, respectively. On one hand, positive information should lead to approach motivation, and the motivation to approach the desired outcome by recognising the possibility of the positive outcome to occur. On the other hand, negative outcomes should lead to avoidance motivation, and the motivation to avoid the negative outcome by negating the possibility of the negative outcome to occur.

In a first set of experiments, Lench (2009) established in an evaluative conditioning paradigm that participants assign higher likelihoods to events when they have been paired with positive, rather than with negative words. By pairing a target picture (e.g. a white car, Lench, 2009, Study 1) with subliminally presented images that were either positive, negative or neutral, Lench (2009) manipulated affective reactions towards the event that participants would afterwards answer likelihood questions about (here: “how likely is it that you will own a white car in the future”). Participants rated

the likelihood of owning a white car as greatest in the positive condition, followed by the neutral and least likely in the negative condition.

Next, in order to establish that approach and avoidance motivation triggered by positive and negative words is underlying this effect, in Lench's (2009) Study 4, participants' approach and avoidance motivation was manipulated orthogonally to word valence. In this experiment, first a target word (here: "garden") was paired with either a positively or negatively valenced word, in the same evaluative condition paradigm as described above. Next, approach-avoidance motivation was manipulated by means of arm flexion and tension, respectively (see Cacioppo, Priester, & Berntson, 1993). As such, the proprioceptive manipulation could either match (positive target-approach, negative target-avoidance) the target word's valence, or mismatch it (positive target-avoidance, negative target-approach). Afterwards, participants were asked to rate the likelihood that they would plant a garden.

Lench (2009) finds support for the approach-avoidance theory by demonstrating that the desirability bias (greater likelihood for the event if it was paired with a positive compared to when it was paired with a negative word) holds if there is a match, but not if there is a mismatch between word valence and proprioceptive cue (see Figure 1.3). Furthermore, Lench (2009) shows that when the target word was cued with a positive word, the likelihood was higher when participants engaged in arm flexion (approach) rather than tension, and the likelihood for the event paired with a negative word was lower when participants engaged in arm tension (avoidance) rather than flexion. According to Lench (2009), this result demonstrates that it is the motivation to approach

positive and to avoid negative events which lead to a greater likelihood assigned for positive, and the smaller likelihood assigned for more likely events.

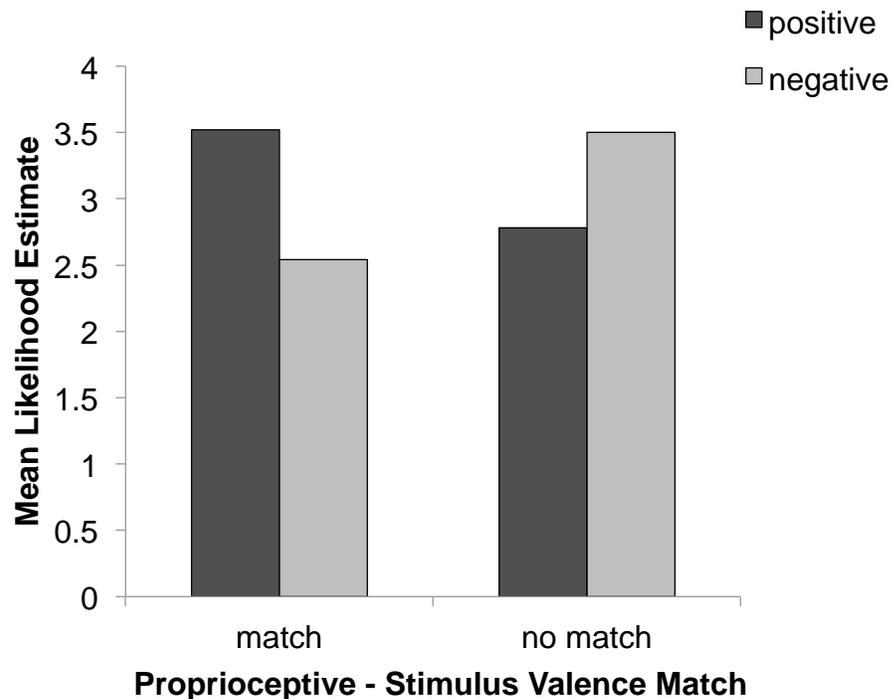


Figure 1.3. Probability estimate means for Lench's (2009) Study 4, classified according to a match between proprioception and stimulus valence.

However, we believe that these conclusions drawn by Lench (2009) are heavily confounded by the fact she used the coded *interaction* (match: positive-approach, negative-avoidance vs. no match: positive-avoidance, negative-approach) as a *factor*, which was then (falsely) analysed with a 2 (positive/negative) x 2 (match/no match) factorial ANOVA. Instead, the manipulations should be classified as utility (positive vs. negative) and motivation (approach vs. avoidance), creating a 2x2 factorial design. If one depicts the results using a factorial design, one can observe that besides a main effect of motivation (approach > avoidance), there does not appear to be an interaction between motivation and word valence (see Figure 1.4).

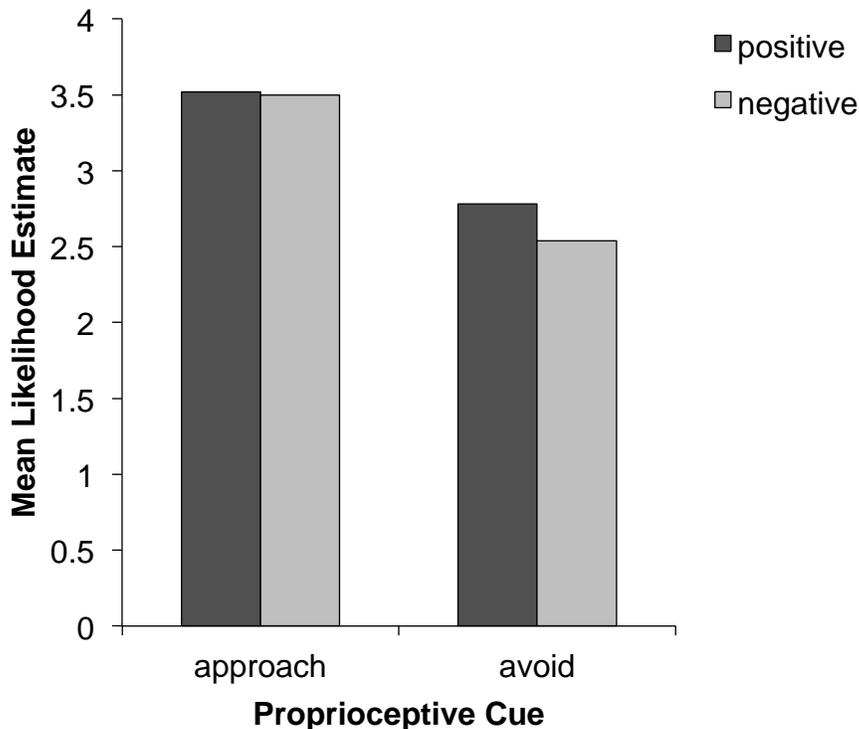


Figure 1.4. Probability estimate means in Lench's (2009) study 4, as classified according to approach-avoidance motivation.

Moreover, there does also not appear to be a noticeable effect of word valence, contradicting their earlier experiments or hinting towards a potential overwriting of the proprioceptive over the affective manipulation (as the proprioceptive manipulation was *after*, not *during* the evaluative conditioning task). To sum up, there is little direct evidence to suggest that approach-avoidance motivation underlies the findings in the evaluative conditioning experiments by Lench (2009; Lench & Ditto, 2008).

Furthermore, the conceptualisation of approach-avoidance motivation and the hypotheses by Lench (2008) are only one of several ways in which one can classify and hypothesize about the impact of approach-avoidance motivation on likelihood estimates. Rose (2009) describes 3 accounts with which incidentally (rather than integrated) approach-avoidance motivation, as in Lench's (2009) Study 4, can impact the

relationship between utility and probability estimates. The *effective action* account is most closely related to Lench's (2009) theorizing. This account proposes that since approach motivation is associated with the pursuit of desirable outcomes, likelihood estimates are inflated for positive events. Furthermore, since avoidance motivation is related to defensive denial and the aim to increase the distance between the person and the event, likelihood judgments for negative events are decreased.

In contrast, the *general outlook* account proposes that manipulating approach-avoidance should lead in fact to a generally positive or negative outlook: this would entail that approach motivation leads to individuals rating positive events as more, *and* negative events as less likely, whilst the opposite would be the case for avoidance motivation. However, given the proposition that approach-avoidance motivations are orthogonal and as such as should impact positive and negative events independently (see also Carver & White, 1994), the position favoured by us and Rose (2009) is the *compatibility-incompatibility* account (Rose, 2009). This is the account that aligns most closely with the original classification of approach-avoidance motivation in the pioneering work of Gray (1987, 1990, 1994). The compatibility-incompatibility account proposes that it is the fit between motivation and affective information that will lead to an *increase* in probability estimates, in that individuals high in avoidance will assign greater likelihoods to negative, and individuals high in approach greater likelihoods to positive future events. However, this account assumes that avoidance will not have any impact on positive or neutral, and approach no impact on negative or neutral future events. This conceptualisation of approach-avoidance motivation and its impact on likelihood estimates is in line with previous work hypothesizing that approach is related

to hope and optimism, and avoidance to fear, worry, and pessimism (Gray, 1987, 1990, 1994). Indeed, in his Experiment 2, Rose (2009) manipulated approach and avoidance motivation and measured likelihood estimates for positive and negative events.

Approach oriented individuals gave higher estimates for positive events than avoidance oriented individuals, whereas the opposite was the case for negative events. Further supporting this position is also the research by Elliot and Church (2003), who showed that the activation of the behavioural inhibition (but not approach!) system is positively correlated with defensive pessimism (the overestimation of the likelihood of negative events). Defensive pessimism has been shown to be a regulation strategy helping to avoid the negative outcome, prepare for anxiety and action (Norem, 2008), and is also related to pre-factual thinking about the event in question (see next section on imaginability). Furthermore, Carver and White (1994), who developed the BIS/BAS scales to assess individual differences in approach-avoidance motivation, found (modest) correlations between the BIS scale (measuring avoidance) and pessimism (as measured here by the Life Orientation Test, Scheier & Carver, 1985) suggesting that individuals high in BIS have a (weak) tendency to expect more negative than positive things to happen to them. Finally, Lench's (2009) results of Study 4 can potentially also be explained in light of this account. Lench (2009) asked participants to rate the likelihood that they would *plant a garden*. To our understanding, planting a garden constitutes an action that is more related to approach rather than to avoidance motives. The *compatibility-incompatibility* account would predict that approach-oriented individuals provide greater probability estimates for positive events than avoidance-oriented individuals, whilst the opposite would be hypothesized for negative events

(where individuals high in avoidance motivation would provide greater estimates than individuals high in approach motivation). Moreover, this account would also entail greater estimates of approach oriented individuals for approach related, and avoidance oriented individuals for avoidance related actions, providing a possible explanation for the main effect of approach-avoidance motivation on the estimated likelihood of planting a garden in Lench's (2009) Study 4.

Thus, the relationship between approach-avoidance motivation and the interdependence of utility and probability estimates deserves further empirical attention. However, the proposal by Lench (2009) that avoidance motivation leads to a *decrease* in probability estimates for negative events does not appear to be supported by the current evidence. If we were to make a prediction of the impact of avoidance motivation on likelihood estimates, we would take the position of Rose (2009), proposing that those high in avoidance give *higher* likelihood estimates for negative events.

Imaginability. The idea that imagining an event leads to greater probability estimates has considerable precedent. Early research by Carroll (1978) demonstrated that when participants were asked to imagine an outcome (such as Ford winning over Carter in the presidential elections), they gave higher probability estimates for the event, as compared to when participants did not imagine the outcome. In a similar vein, Gregory, Cialdini and Carpenter (1982) showed that individuals' probability estimates increased when asked to imagine the event, both for positive (e.g. winning a vacation) as well as for negative outcomes (e.g. being someday arrested for armed robbery). More specifically, Sherman, Cialdini and Reynolds (1985) demonstrated that the ease with which participants imagined the outcome of a negative event (here: contracting a

disease) to occur, mediated likelihood estimates (see also Mevissen, Meertens, Ruiter & Schaalma, 2012). Individuals who reported greater ease of imagining the outcome assigned higher likelihood judgments compared to individuals who reported difficulty in imagining the outcome. These results are most parsimoniously explained by the availability heuristic (Tversky & Kahneman, 1973), whereby the probability of an outcome is determined by cognitive availability. Since probabilities are constructed by remembering relevant information and constructing occurrences, the more available the event in question, the easier it can be constructed, leading to greater probability estimates (Kahneman & Tversky, 1982).

However, the above-mentioned research did not consider the utility of the event. Risen and Gilovich (2007) first proposed that an increased imagination for more aversive outcomes might underlie the impact of utility on probability estimates. In their critical Study 4, participants were asked to imagine that they had sold a lottery ticket to the fictitious character “Alison”, and participants were asked to estimate that the lottery ticket they had sold was actually a winning ticket. Outcome utility was manipulated by stating that Alison is a close friend, a stranger, or an enemy, from least to most negative, respectively. That participants indeed perceived the order “friend, stranger, enemy” as a linear order from least to most negative was also separately confirmed in a pilot study. Participants read the sentence “Alison wins the lottery with the ticket you exchanged” and were asked to indicate as quickly as possible whether the ending makes sense (a reaction time measure of accessibility) and to estimate the likelihood of the outcome to occur. Participants demonstrated a negativity bias, in that participants estimated the likelihood that the ticket they had sold was a winning ticket to be greater the more

aversive the outcome. Importantly, the more aversive the outcome, the more accessible was the ending as indicated by faster reaction times to the statement that Alison wins the ticket. However, in this experimental paradigm it is difficult to establish the direction of the relationship between imaginability and likelihood estimates. It remains possible that outcomes were more accessible because they were believed to be true – and not true because they were more accessible. Whilst this experiment provides more evidence for the general existence of a negativity bias, the support in favour of an imaginability account is less strong. Furthermore, an objective basis for their estimates was not provided. As utility and probability are often confounded in the real world (see Dai et al. 2008; Pleskac & Hertwig (in press); Weber & Hilton, 1990) this experimental design does not hold constant the evidence selection phase (see Figure 1.1 of this Chapter).

Bilgin (2012) further tested the hypothesis that the reason why positive outcomes are judged as less likely compared to negative outcomes is the greater propensity to imagine negative events. In Bilgin's (2012) Study 2, participants estimated the likelihood of either moving into a worse, or a better office in 2 months. Participants were informed that they had a 70% chance of moving into the office. Critically, participants in the imagination condition were asked to imagine and write down their thoughts about the outcome, whereas individuals in the control condition were not instructed to do so. Bilgin (2012) found that participants in the control condition judged the negative event to be more likely than the positive event on a 0 (not likely at all) – 11 (extremely likely) scale. However, when instructed to imagine the outcome, participants in the positive condition increased in their likelihood estimate to the same level of the negative outcome condition, whilst individuals in the negative condition did not change.

Thus, these studies suggest that positive events are judged as less likely than negative events due to negative events being more, and positive events being less vividly imagined. It is suggested that these positive-negative asymmetries are caused by greater elaboration processes employed by and greater attention directed to negative events (Baumeister et al. 2001). Note that this account is congruent with the earlier discussed CLT account. As vivid images are also more concrete, CLT would predict that the more an event is imagined, the lower the level on which it is construed, which is in line with greater probability estimates compared to more pallid, high level construal of events.

However, the experiments by Risen and Gilovich (2007) and Bilgin (2012) lack experimental control. Whilst Risen and Gilovich (2007) did not provide any objective probability information, Bilgin (2012) explicitly told participants the base rate of the outcome, which might have led to similar effects as in Windschitl et al. (2010)'s examination of the Marks paradigm. In the Marks paradigm, participants guess the likelihood of drawing a "marked" card, associated with a monetary pay-off. The base-rate (50%) of drawing a marked card is known to participants. What is typically found is that participants provide optimistic guesses and overall believe that the likelihood is greater than 50%. However, as Windschitl et al. (2010) demonstrated, even though their guesses might appear in line with wishful thinking effects, their true beliefs remain unbiased.

Recent data from our laboratory (Quantmeyer, 2014) tested for the impact of imaginability in a minimal experimental context, avoiding some of the shortcomings of the studies discussed before. Participants were asked to imagine that they would take

part in a lottery where they could win or lose \$100 depending on the outcome of a series of dice throws. Participants were asked to imagine that if at least one 6 shows up in 4 dice throws they could win the \$100 (desirable condition) or lose the \$100 (undesirable condition). No outcome was associated with the dice roll in the neutral condition. Furthermore, participants in the imaginability condition were additionally instructed to imagine and to write down thoughts about the outcome. In this experimental set-up manipulating utility (positive / neutral / negative) and imaginability (present / absent), Quantmeyer (2014) failed to find evidence for a moderation by imaginability on probability estimates in a controlled setting, providing an objective basis for probability estimates. Whilst he successfully demonstrated that individuals assign higher probabilities to negative compared to neutral events, and whilst the means in the negative condition were congruent with an imaginability account ($M_{\text{NoImaginability}}=25.84$ (22.27) vs. $M_{\text{Imaginability}}=29.8$ (25.98)) this trend failed to reach significance (see also Jenkins, 2013, for a failure to demonstrate an effect of imaginability on the impact of negative utility on verbal probability estimates).

In conclusion, whilst the role of imaginability in probability estimates appears to be reliably demonstrated for probability estimates *in general*, its role as a mechanism for the relationship between utility and probability remains controversial.

Asymmetric Loss Functions. Whereas the previously mentioned theories are built mainly on affective processes, next we will discuss a more cognitive process - the account of asymmetric loss functions (ALFs), which focuses on the consequences of providing an inaccurate probability estimate.

When estimating probabilities, one can make two types of error: overestimating, or underestimating the likelihood of the event occurring. An ALF account explains that the costs associated with these two errors are not alike: underestimating the likelihood associated with a negative event often comes at a higher cost than overestimating it (Weber, 1994). Moreover, individuals are assumed to be sensitive to these differences in costs of making different errors. The crucial hypothesis is that the error associated with the lower cost will be overestimated. For instance, if it is more costly to underestimate a probability, individuals should be biased towards overestimating the probability. For negative events, it is often more costly to underestimate probability, and this asymmetry becomes more extreme the more aversive the event (see Figure 1.5). For instance, imagine that one makes an estimate about whether or not one has caught a dangerous virus. Overestimating one's chances and making a trip to the doctor is a comparatively small cost should one be healthy, compared to underestimating one's chances and not go to the doctor only to find oneself very sick later on.

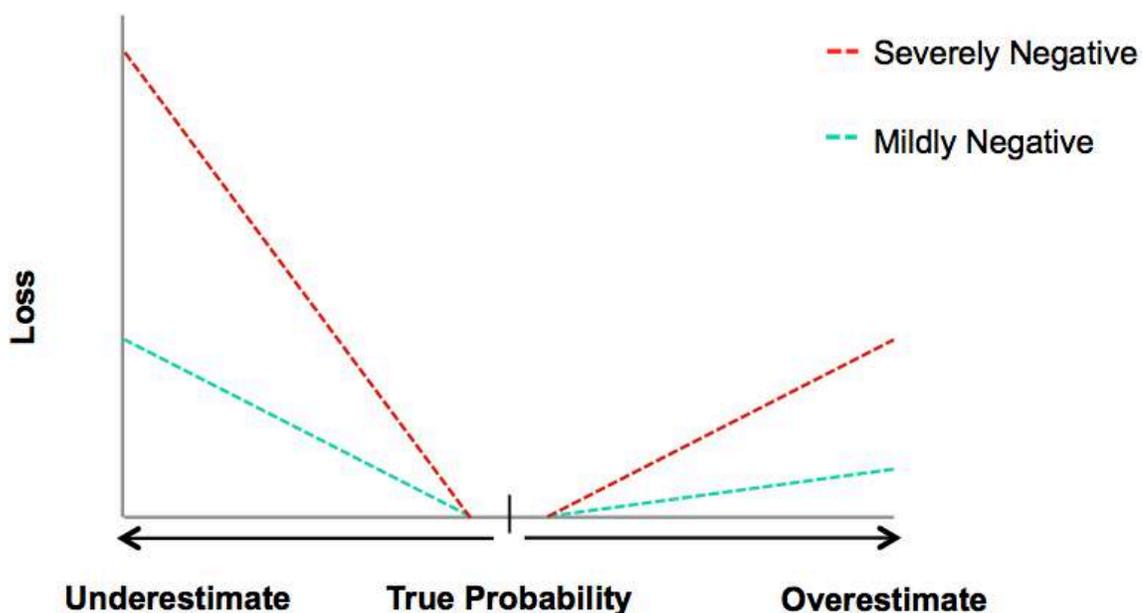


Figure 1.5. Asymmetric loss functions for severely and mildly negative events.

This theoretical account was empirically supported by Harris et al. (2009). One prediction brought forward by ALFs is that the two types of error are only then differential in loss, if a decision can be based on one's estimate, that is, if one has *decision-control*. Note that decision-control is a different type of control than *outcome-control*. If one estimates the risk for rain, taking versus not taking an umbrella based on one's estimate is equalled to decision-control, whereas being able to influence whether or not it will rain, is conceptualized as outcome control.

Importantly, without decision-control, a loss function should not exist. There should be no cost associated with over- or underestimating the occurrence of an event, if no action can be taken based on one's estimate - as there are no costs associated with an estimate for which there are no consequences. As a result, negative utility should only bias probability estimates when the estimator has decision-control, and loss functions exist.

Supporting the contention that probability estimates should only be inflated if individuals have decision-control, Harris et al. (2009) demonstrated that probability judgments for negative events were indeed not inherently biased. Instead, only for those scenarios where an element of controllability was introduced did participants give higher estimates for negative events. For example, in Experiment 4, participants were presented with an unambiguous, objective representation of probability (e.g. a visual cell matrix with varying proportion of yellow and grey cells). A cell matrix represented an orchard and each cell represented an apple tree. Participants were asked to estimate that a child playing in an orchard would eat a poisonous apple, as represented by a yellow cell in the matrix (the grey cells representing harmless apples). Decision-control was manipulated by informing participants either that there was the potential of erecting a fence, which would prevent the child from entering the orchard, or that there was no possibility to stop the girl from entering. Participants estimated the likelihood of the child eating a poisonous apple as higher if there was the potential of erecting a fence that would prevent the child from entering the orchard than if there was no possibility to prevent the girl entering the orchard (where the estimates were indistinguishable from the neutral conditions in prior experiments). These findings support the ALF account, in which there can only be costs associated with a misestimate of the probability of an event if a decision is subsequently based on this estimate, and only for these events should probability estimates be biased towards an overestimation to avoid the costs associated with an underestimation.

However, whilst the studies by Harris et al. (2009) provide some evidence for the role of decision-control and the mediation by ALFs, the studies are scenario based and

therefore without real decision-control, which might weaken their generalizability. Moreover, in their Experiment 5, Harris et al. (2009) orthogonally manipulated utility and decision-control and demonstrated a negativity bias only for the high decision-control conditions. However, the lack of manipulation checks for control and utility makes it impossible to observe in how far utility and control were manipulated truly orthogonally (the outcome might have been perceived as less bad in the no-control condition, for example). Potentially highlighting the importance of such manipulation checks is research by Siegrist and Sütterlin (2014), who demonstrated that individuals rated the same hazards as worse when man-made compared to when they were caused by nature. As can be argued, man-made hazards are more controllable than nature-made hazards, which could have led to the difference in their utility ratings.

Furthermore, Harris et al.'s (2009) scenarios all incorporate an element of communication and persuasion. Communicating one's estimate to someone else (who will experience the outcome, importantly) could have led to an inflation of probability estimates. Thus, it remains possible that not the internal estimate of the communicator was biased, but the report of the probability (see Figure 1.1) – potentially due to a feeling of responsibility toward the receiver.

A further line of research related to Harris et al. (2009) comes from Mellers, Schwarz, Ho and Ritov (1997), who showed that unexpected outcomes have a greater emotional effect compared to expected outcomes, leading individuals to be more likely to overestimate the probability of negative events to protect themselves from the emotional impact (as a self-control mechanism, Weber, 1994). This finding is in line with Shepperd et al. (2000), who showed that strategic pessimism increases the

likelihood for negative outcomes, in particular in the face of being provided with feedback about the outcome (Shepperd et al. 1996 coined the term “bracing hypothesis”). However, in the studies of Harris et al. (2009), the negativity bias only occurred for controllable outcomes. This means that if there was no potential to interfere with the risk taking process, individuals were not biased. Thus, it might not only be about anticipating “bad” news, but also about whether or not one can control in how far the events can take a turn for the better, should the prospect be of a negative nature. It is possible that the emotional impact is worse for events where one had an element of control –the emotional impact, through regret felt when underestimating the likelihood of negative events and failing to act is worse than if there was no possibility to take preventative steps (e.g. Zeelenberg, Van den Bos, Van Dijk, & Pieters, 2000). This would also be in line with the aforementioned research by Siegrist and Sütterlin (2014) showing that man-made negative events are perceived as worse than negative events caused by nature, which can be argued to constitute high and low controllability situations, respectively.

Worth mentioning, in real-life there are also boundary conditions to ALFs, where the misestimate is more costly for an overestimate, rather than an underestimate, of the true likelihood. For instance, Gigerenzer (2003) mentions certain medical situations where undergoing treatment might be more costly than foregoing treatment: Certain forms of non-progressive breast cancer are better undetected, as “it is likely that [the patient] would have led an equally long and even happier life without having undergone the rigors of treatment” (Gigerenzer, 2003, p. 66).

Social Power and Loss Function Asymmetries. Previous research has demonstrated that decision-control moderates the impact of negative utility on probability estimates by manipulating the amount of decision-control provided in the described scenario (Harris et al. 2009). Based on this first finding, we hypothesize that also the perception of how much control an individual feels over a given situation should moderate the impact of negative utility on probability estimates. In particular, in situations with low or ambiguous decision-control could personal control lead to increases in the negativity bias. One individual difference variable that strongly determines the amount of control a person feels over themselves and their environment is social power. Individuals high in social power have been shown to feel more in control, also over hard-to-control outcomes (e.g. Fast et al., 2009). Powerful individuals are both more likely to feel a greater sense of control over their own actions, as well as over other's actions. Therefore, we would expect individuals high in social power to show a greater negativity bias for scenarios where the outcome concerns themselves, as well as where they are asked to provide a probability estimate to inform someone else's actions. Importantly, this should be largely independent of the actual decision-control in a given situation for individuals high in social power, whereas those low in power should only display a negativity bias when they are explicitly provided with decision-control.

Summing up, the full proposed model based on an account of ALFs would be a mediation of the impact of negative utility on probability estimates through loss function asymmetries. However, the impact of negative utility on loss function asymmetries should be moderated by social power, due to the powerful's greater sense of control (see

Figure 1.6).

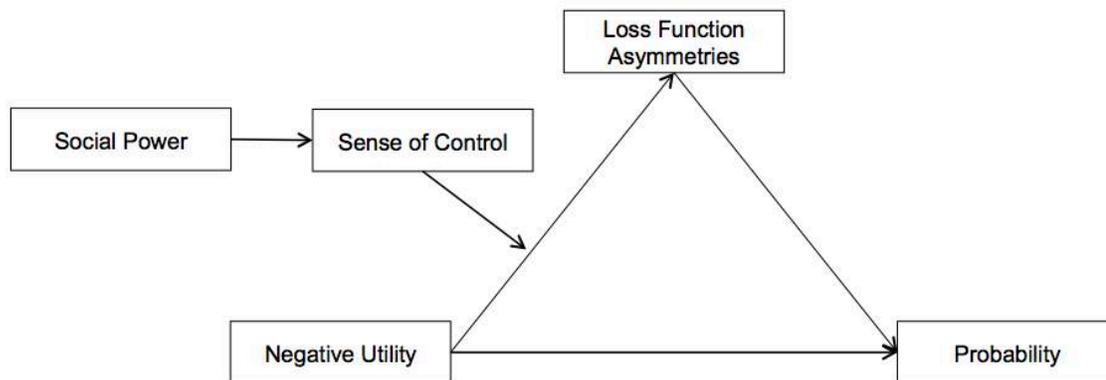


Figure 1.6. Impact of negative utility on probability via a mediation by loss function asymmetries, as well as the mediated moderation of the impact of power via sense of control on the relationship of negative utility on loss function asymmetries.

Stake Likelihood Hypothesis. With the aim of providing a mechanism both for wishful thinking and the negativity bias, Vosgerau (2010) postulated that misattribution of arousal (excitation transfer, Zillmann, 1971) exhibited by valenced (positive or negative) events leads to an overestimation of outcome probabilities. That is, the likelihood of an event to occur is judged depending on how aroused the decision maker is- the higher the arousal, the more likely the decision maker assumes the event to occur (“Stake-Likelihood Hypothesis”, SLH). Since arousal is non-specific (Schachter & Singer, 1962), and sticky (Cantor, Zillmann, & Bryant, 1975), residual arousal can influence a target despite the removal of the arousal-eliciting stimulus. In an infamous experiment by Dutton and Aron (1974) for example, male participants were more likely to accept the phone number from an attractive female confederate (and later call!), as well as to report greater sexual imagery when crossing a scary bridge, which increased arousal, compared to a non-scary, non-arousing bridge.

Vosgerau (2010) argues that the two prerequisites for an arousal transfer hold for probability estimates: first, there is evidence that having a stake in the outcome, positive

or negative, increases physiological arousal (Elliott, 1964; Schnore, 1959). Secondly, if a valenced event has a greater likelihood of occurring it should be associated with a greater level of arousal. For example, if there is a greater probability of getting into an accident, arousal levels should be higher than for lower chances of getting into an accident. Thus, greater arousal should be associated with greater probabilities of valenced events occurring, and Vosgerau (2010) postulates that the arousal from having a stake in the outcome is misattributed to the likelihood of the event occurring. The SLH was tested in four experiments, which shall be discussed in detail below.

In Study 1, Vosgerau (2010) aimed to demonstrate that greater arousal levels are associated with greater likelihoods, and manipulated arousal independently of the event. The events that were to be judged were relatively neutral in valence, but arousal was manipulated separately. Specifically, participants were presented with 9 questions probing for probability estimates (“how likely to do you think it is that Obama will be voted for?”) on either grey or bright pink paper. Vosgerau (2010) postulates that since the bright pink colour is more arousing than the flat grey paper, individuals answering on the pink paper should give higher likelihood estimates, as the arousal from the paper is misattributed to the likelihood of the event occurring. Indeed, in Vosgerau’s (2010) Experiment 1, participants that answered on pink paper provided greater probability estimates than participants that answered on grey paper.

However, some alternative explanations should be taken into account. For example, pink has been shown to worsen performance (Pellegrini & Schauss, 1980; Pellegrini, Schauss & Birk, 1980; but see Gilliam, 1991), potentially through gender based demand characteristics (Ingram & Lieberman, 1985). Since Vosgerau (2010) fails

to demonstrate the concrete mechanism through which the impact of colour might operate, this first experiment does not constitute direct evidence for the role of arousal misattribution as a mechanism.

Study 2 observed the SLH in the positive domain, where participants were asked to judge the occurrence of a 3 coming up at least once in 4 dice throws and could win \$5 if the outcome were to occur. Additionally, participants were asked how exciting it is to play the game and crucially, were asked to indicate their excitement either before or after they made the probability estimate. Thus, Study 2 made use of a classic arousal misattribution paradigm: arousal is only misattributed if participants are unaware of the actual source of arousal. If participants misattribute arousal, they should be less likely to do so once the source of the arousal is made salient and they can thus attribute their arousal to the actual source (Gorn, Pham & Sin, 2001). The data of Study 2 were indeed in line with arousal misattribution processes: individuals gave higher probability estimates when they were asked *after* providing a probability estimate of how exciting the game was. If asked before providing their estimate, participants presumably were made aware of the source of their arousal, and did not misattribute the arousal to the likelihood of the event occurring. Additional evidence was provided from the arousal ratings themselves: Arousal was greater if it was rated before the provision of probability estimates than when it was rated afterwards. Vosgerau (2010) interpreted this decrease as congruent with arousal misattribution, as participants attributed the arousal to the likelihoods rather than to having a stake in the outcome if being asked after they made probability estimates. In this experiment, Vosgerau (2010) additionally manipulated the focalism of the outcome: either participants were asked how likely it is

that they would get at least one 3, or they were asked how likely it is that they would *not* get at least one 3. The results hold for both foci (however, for the arousal ratings it is not reported whether an arousal*focus interaction was present). Nevertheless, Vosgerau (2010) does not address the alternative explanation that it might have been possible that participants did not *enjoy* providing the probability estimate. Before participants knew that they would have to provide the estimate, they might have thought that the prospect of the game sounds exciting, since they might win something. Afterwards, however, they realised it was not as exciting since they had to actually perform some math. This could explain why the ratings of excitement were lower after the probability estimate (but does not account for the decrease in probability estimates).

Following the results from Study 2, Vosgerau's (2010) Study 3 had two aims: first and foremost, it extended the SLH to the negative domain. Secondly, Vosgerau (2010) assessed in how far imaginability (see description earlier in this Chapter) might account for the effects, by including a "low imaginability" condition. In this experiment, participants judged the likelihood of a 6 coming up at least twice in 4 dice throws. Depending on their condition, participants could either win or lose a shot glass. Specifically, participants were asked to estimate the chance that they would lose (vs. win) the shot glass – a "high imaginability" condition since individuals focussed on the outcome concretely. Additionally, a "low imaginability" condition was included. In this "low imaginability" condition, instead of focussing on the outcome of winning a shot glass, participants were asked to estimate the chance that a 6 would occur at least twice in the 4 throws. Thus, participants did not focus directly on the outcome (winning a shot glass), but only on the probabilities. In this study, it was found that indeed the

conditions including a stake in the outcome were accumulatively assigned a higher probability compared to the neutral condition, and individuals in the negative condition assigned higher probabilities than participants in the positive condition. Moreover, in the “low imaginability” condition individuals still gave higher estimates for the outcome compared to the neutral condition, which was interpreted as evidence that it is not imagining an event, but rather excitation transfer, which underlies the effects. There are several noteworthy results. First, Vosgerau (2010) did not include a manipulation check for the imaginability condition. It is not shown that participants in the low imaginability condition actually imagine the outcome less. This could have been relatively easily included by asking participants “how vividly did you imagine the outcome to occur”. Secondly, Vosgerau (2010) only included a “low imaginability” condition for the win, but not for the loss condition, making this an incomplete design and allowing for the interpretation that imaginability might indeed affect positive and negative outcomes differently.

Moreover, by collapsing across the positive and negative utility conditions, it is impossible to know whether the difference to the neutral condition is driven by both the negative and the positive outcome conditions, or whether indeed only the negative condition in which participants gave statistically higher estimates than in the positive condition was driving these differences. Participants in the negative condition gave higher probability estimates ($M= 38.91$, $SD=28.92$) than participants in the positive condition ($M=19.77$, $SD=16.07$). Whilst both conditions are greater than the estimates of the neutral condition ($M =13.44$, $SD=10.43$), it is the negative condition, which has the greatest difference to the neutral condition.

Study 4 tested the SLH in a naturalistic setting – the allegiance bias in sports betting, here in football. Vosgerau (2010) asked German university students before an important football match to estimate the likelihood of the team favoured by the vast majority of students to win or to lose the match, and asked students how much they were willing to bet on the outcome. Critically, arousal was manipulated by either telling participants that they would watch the game live, or that they would watch the game with a delay (due to an apparent conflict over broadcasting rights). It was reasoned that watching a game live is more exciting and elicits greater arousal, as the results are not determined yet (Nelson, Galak, & Vosgerau, 2008). An interaction between valence (win vs. lose) and arousal was shown. Irrespective of focus, participants assigned higher probabilities when watching the game live compared to when the game was watched with delay. Additionally, when watching the game live, participants rated the chance that the team loses higher than the team winning, which was not the case when watching the game taped. Willingness to bet was marginally impacted by likelihood estimates, in that participants were willing to bet more on the team winning or losing the higher their likelihood estimates for a win or a loss, respectively. Vosgerau (2010) interprets these results as being congruent with a SLH account – presumably, arousal was only misattributed in the live condition, as individuals were less aroused when not watching the game live. However, an alternative explanation is provided by CLT, where it would be expected that events lower in time distance are construed on a lower level, and as such are also assigned a greater probability estimate.

Summing up, the current empirical evidence for the SLH thus far is subject to alternative explanations, and deserves further empirical attention.

Social power and the Stake-Likelihood Hypothesis. One way of assessing the SLH is by introducing moderators for the relationship between arousal and probability. Whilst physiological changes, such as arousal, are a necessary factor for the experience of emotions (Schachter & Singer, 1962), individual differences in the sensitivity to these changes have been shown to moderate affective responses and arousal misattribution to aversive events (Blascovitch et al. 1992; Katkin, 1985). These differences in sensitivity have been described as individual differences in interoceptive ability (body awareness), the ability to detect changes in the body, such as visceral states including hunger, pain, and temperature (Craig, 2002). Importantly, it has been demonstrated that individuals high in interoceptive awareness experience affective states, especially negative feelings, more intensely (Barrett, Quigley, Bliss-Moreau, & Aronson, 2005; Pollatos, Gramann, & Schandry, 2007). Moreover, highly interoceptive individuals report greater subjective arousal, but not valence when experiencing emotive stimuli (Feldman-Barrett et al. 2004). Logically, if the negative outcome of an event impacts arousal, which in turn leads to inflated probability judgments compared to neutral events, this bias should be increased for individuals high in interoceptive awareness.

We propose that increased social power leads to greater probability estimates for negative events by moderating the impact of arousal on probability, mediated by interoceptive ability (see Figure 1.7).

According to the SFTP (Guinote, 2007a), as a result of freedom of constraints, individuals high in social power engage in narrow information processing strategies and focus on selective cues accessible in the situation. Attention of the powerful is directed to primary cues relevant in the situation, such as subjective experiences. On the other

hand, powerless individuals spread their attention and attend to several sources of peripheral information (Guinote, 2007c). Thus, whilst powerful individuals are able to truly *focus* on one construct at a time, powerless individuals have a *wider* information processing strategy. The powerful are thus engaging in situational processing of experiential information and bodily cues (Guinote, 2010). Guinote (2010) demonstrated that the powerful rely more on bodily experiential information, in that only the powerful's amount of appetizing food eaten was predicted by actual hunger levels, whilst the powerless ate irrespectively. Most conclusively, however, Moeini-Jazani, Knöpferle, de Molière, Gatti and Warlop (2014) provided evidence for a main effect of social power on interoceptive ability. Following a manipulation of power, a heartbeat counting task, commonly employed to assess interoceptive ability (e.g. Dunn et al. 2010), Moeini-Jazani et al. (2014) showed that individuals high in social power were better able to feel (and count) their heartbeats compared to powerless or control participants, indicating greater interoceptive ability.

Summing up, there is evidence to believe that powerful individuals will estimate the probability of negative events to be greater than powerless individuals. This model (see Figure 1.7) states that social power leads to greater interoceptive ability, and that individuals high in interoceptive ability who feel arousal more strongly, should subsequently demonstrate greater misattribution effects of the arousal evoked from

having a stake in the outcome.

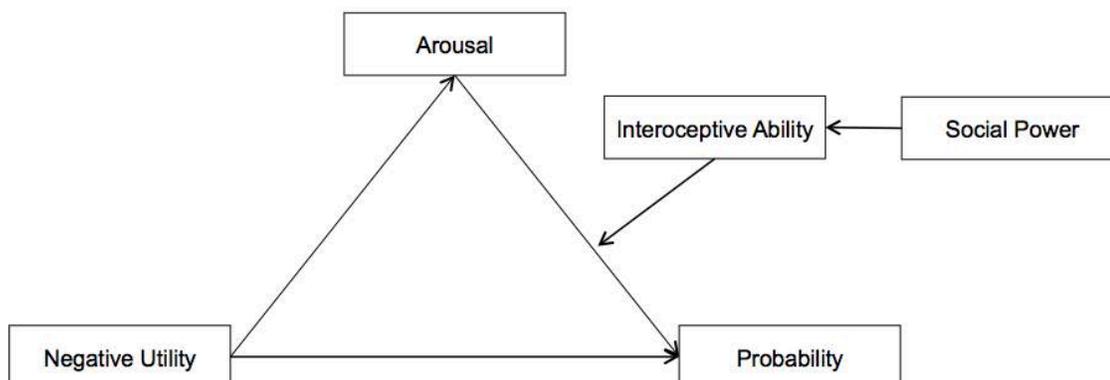


Figure 1.7. Impact of negative utility on probability via a mediation by arousal, as well as the mediated moderation of the impact of power via interoceptive ability on the relationship of arousal on probability.

Summary and Overview of the Empirical Chapters of the Thesis

In the present chapter, we first provided an overview of the evidence for an interdependence of utility and probability estimates. We discussed the evidence for an impact of positive utility on probability and conclude that this evidence does not appear to be very strong. Next, we presented the evidence for an impact of negative utility, presenting recent research showing that individuals seem to assign greater probabilities to negative compared to neutral events. We then introduced the concept of social power, and highlighted that some of the most dominant theories of social power have conflicting predictions about the impact of social power on the relationship of negative utility and probability estimates. Following a discussion of the main theories on potential mechanisms of the interdependence of negative utility and probability estimates, we demonstrated that social power is related to two of these theories, setting the rationale for the current research: On one hand, we reason that as social power increases interoceptive awareness, power could moderate the relationship between

negative utility and probability through a greater arousal misattribution of the powerful compared to the powerless. On the other hand, as social power increases one's personal sense of control, we argue that power could moderate the relationship due to the greater decision-control of the powerful.

The potential advancement of previous research from the current investigation is twofold: Our research advances previous knowledge in the domain of probability estimates, as well as to inform the literature of social power.

For one, we aim to rigorously test hypothesized mechanisms of the impact of utility on probability. In particular, we aim to identify new moderators of the negativity bias, building on previous theories. This will also allow us to identify the circumstances and individual traits that lead to the influence of negative utility on probability estimates. Secondly, we will examine more closely how individuals high and low in social power incorporate negative information when estimating probabilities of future events. Since the main theories of social power have conflicting predictions, this research can set apart these theories in the domain of probability estimates.

Having provided the theoretical background in the current chapter, Chapters 2, 3, and 4 provide an empirical investigation into mechanisms underlying the relationship between negative utility and probability estimates. Chapter 2 explores the moderating role of interoception and discusses the mechanism through which interoception can impact arousal misattribution processes. In this chapter, we provide some preliminary evidence for the impact of interoception on the relationship between utility and probability estimates, assessing interoception with a frequently used questionnaire (Miller, Murphy, & Buss, 1981). However, measuring objective interoceptive accuracy

with a heartbeat counting task (Schandry, 1981) and measuring arousal with galvanic skin responses, showed no effects on the relationship of utility on probability estimates of interoception or arousal, challenging the SLH. In Chapter 3 we therefore investigated the original effects by Vosgerau (2010) and across several replication attempts of the original findings, are unable to provide evidence for misattribution of arousal in the domain of likelihood judgments. In Chapter 4, we then examine the role of control and social power. We demonstrate that social power moderates the negativity bias, and provide some evidence that this is due to the greater decision-control of the powerful. Finally, having established that powerful individuals incorporate negative information more into their judgment in the domain of likelihood estimates, in Chapter 5 we examine the impact of negative affective states on approach and avoidance motivation in the powerful. Chapter 5 provides evidence that powerful individuals become more avoidance oriented under negative affective states than powerless individuals, supporting the situated focus theory of power in the domain of affective experiential information.

Chapter 2 The Impact of Interoception on the Interdependence of Negative Utility and Probability Estimates

Chapter Overview

As previously outlined, a way of testing the assumptions of the Stake-Likelihood Hypothesis is to examine whether the experience of arousal impacts the relationship between utility and probability estimates. Since more interoceptively aware individuals report greater intensity of arousal levels, the misattribution of arousal from negative utility onto probability estimates should be stronger than for less interoceptive individuals. As a result, we hypothesize that interoception should moderate the relationship between negative utility and probability estimates, in that more interoceptive individuals show a greater negativity bias than less interoceptive individuals. In the present chapter, we report three experiments examining this hypothesis. We examine the impact of self-reported levels of interoceptive ability on the negativity bias in Experiments 1 and 2. In Experiment 3 we assess interoception objectively with a heartbeat counting task and directly measure arousal with galvanic skin responses.

Introduction

To provide more background to the current chapter's experiments, we will first describe those aspects of the SLH that have previously been assessed, and discuss how to further examine this theory. We then create the theoretical setting for including interoception as a moderator of the effect. We define and describe interoceptive ability

and its dimensions and examine the relationship between interoceptive ability and the psychophysiological aspect of emotional experiences. Last, the evidence for the impact of interoceptive awareness on decision-making will be discussed, before providing background with regard to methodologies employed to measure interoceptive ability.

Assessing the Stake-Likelihood Hypothesis

As described in Chapter 1, the Stake-Likelihood Hypothesis (SLH) proposes that the arousal from having a stake in the outcome is misattributed to the likelihood of the outcome occurring. Vosgerau (2010) tested several of the assumed relationships in this mediation model (see Figure 2.1).

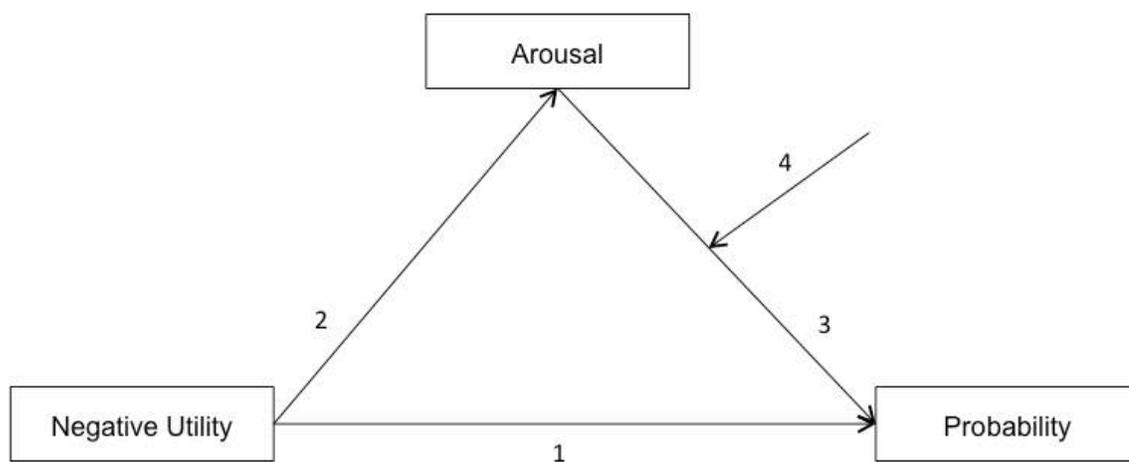


Figure 2.1. Mediation model of the stake likelihood hypothesis

Vosgerau (2010) provided some preliminary evidence for the relationship between negative utility and arousal (path 2), but arousal was not assessed explicitly. For example, in Experiment 2 Vosgerau (2010) asked participants before or after they gave a probability estimate how “exciting” the game they were playing was, and showed that when the question was asked first, participants’ likelihood estimates decreased.

Furthermore, when participants rated their excitement before providing probability estimates, arousal was rated as greater than when they rated it after providing probability estimates. Vosgerau (2010) interpreted this decrease as evidence for arousal misattribution, as participants ostensibly attributed their arousal to the true source when they were made aware of it by asking how excited they are.

Therefore, Vosgerau (2010) only indirectly assessed arousal misattribution, rather than to measure arousal as a mediator directly. Additionally, this effect was demonstrated only with a positive, but not with a negative outcome. It remains unknown in how far this arousal misattribution paradigm impacts probability estimates for events with a negative outcome utility, or in fact, probability estimates for events with a neutral outcome utility. One way of further assessing the SLH is to employ the arousal misattribution paradigm described above for negative outcomes. However, this arousal paradigm only provides indirect evidence for the process. Hence, arousal should also be assessed directly as a mediator by means of physiological measures such as galvanic skin responses (e.g. Andreassi, 2007; Boucsein, 2012; Dawson, Schell, & Courtney, 2011; Dawson, Schell, & Filion, 2007).

Next, providing evidence for path 3 (see Figure 2.1), Vosgerau (2010) demonstrated that arousal from a source other than utility (i.e. the colour of paper) is misattributed and inflates probability estimates for events where there is no stake in the outcome. However, as mentioned in Chapter 1, a manipulation of arousal from the colour of the paper is likely to be confounded (e.g. gender demand effects), and an alternative way of assessing the relationship between arousal and probability estimates

is to create arousal from an unrelated task. For instance, arousal can be subtly created by administering caffeine (Barry, Rushby, Wallace, Clarke, Johnston, & Zlojutro, 2005).

A further way of examining the SLH is through the introduction of factors that moderate the relationship between arousal and probability (path 4, see Figure 2.1). That is, manipulating the intensity of the experience of arousal could potentially impact the relationship between arousal and probability estimates. Of particular interest to the current research are individual differences in the sensitivity to physiological changes such as arousal. These differences have been described as interoceptive ability, the ability to detect changes in the body, such as visceral states including hunger, pain, and temperature (Craig, 2002), as will be defined in more detail below. Importantly, individuals that are better at distinguishing internal changes experience arousal more intensively (Dunn et al. 2010; Feldman-Barrett, Quigley, Bliss-Moreau, & Aronson, 2004; Wiens, Mezzacappa & Katkin, 2000). We propose that interoceptive ability is a potential moderator of the SLH: Individuals that perceive the arousal created by having a stake in the outcome more strongly should also misattribute more arousal, and therefore show a greater negativity bias than individuals that are less accurate in perceiving internal states.

Dimensions of Interoception

Interoception is a multidimensional concept (see Table 2.1). Interoceptive *sensibility* refers to the “dispositional tendency to be internally focused” and to “self-reported beliefs about body tendencies” (Garfinkel & Critchley, 2013, p.233).

Interoceptive *accuracy* refers to the objectively measured interoceptive ability, for example by asking participants to count their heartbeats (Garfinkel & Critchley, 2013).

Last, interoceptive *awareness* is one’s metacognitive ability to correctly report how interoceptively accurate one is (as assessed by mapping accuracy onto confidence) (Garfinkel & Critchley, 2013). Whilst interoceptive accuracy (IA) and sensibility (IS) have often been used interchangeably (Ginzburg, Tsur, Barak-Nahum, & Defrin, 2013), and whilst we will examine the moderation of negative utility on probability estimates of both dimensions, we will indicate which dimension of interoception was assessed in each of the studies when reviewing the literature below. Furthermore, an overview and discussion of the assessment of the different dimensions, and the concrete assessments utilised in the current research will be provided before the empirical section of this chapter.

Table 2.1

Dimensions of Interoception (Garfinkel, Seth, Barrett, Suzuki & Critchley, 2014) Table reproduced with permission from the author.

	Interoceptive Awareness	Interoceptive Accuracy	Interoceptive Sensibility
Definition	Metacognitive awareness of interoceptive sensitivity	Objective accuracy in detecting internal bodily sensations	Self-perceived dispositional tendency to be internally self-focused and interoceptively cognisant
Example	<i>Do you “know” whether you are accurately or inaccurately assessing your heart-timing?</i>	<i>Can you accurately report when your heart is beating?</i>	<i>To what extent do you believe you focus on and detect internal bodily sensations?</i>
Mode of Assessment	Relationship between objective performance (interoceptive accuracy) and awareness of performance	Assessed via objective tests of interoceptive accuracy	Assessed via subjective self-report measures probing perceived aptitude
Example	<i>Area under ROC curves mapping confidence onto accuracy</i>	<i>Behavioral performance accuracy during heartbeat detection / mental tracking tasks</i>	<i>Questionnaires, such as Porges Body Perception Questionnaire, or global self-report measures such as average confidence</i>

Interoception and Emotional Experiences

In order to illustrate the relationship between interoceptive ability and emotional experiences, we will first describe how emotional experiences are manifested. In particular, we will describe the role of psychophysiological processes².

One of the earliest accounts on the creation of emotional experiences goes back to James (1884, 1890) and Lange (1885, 1922). The ‘James-Lange Theory’ of emotion proposed that input from the peripheral nervous system in the form of physiological arousal (defined as the “intensity of peripheral physiological reactions”, Frijda, 2009, p. 268) is involved in creating emotional experiences. According to this account, physiological changes (in the peripheral nervous system) are experienced before the experience of the emotion. However, these peripheral changes are only related to the arousal, but not to the valence dimension of affective experiences. Therefore, following the James-Lange account of emotion, it remains unclear how emotions that have similar physiological properties lead to emotions that differ in valence. This earlier account was further developed in Schachter and Singer’s (1962) two-factor theory. The two-factor theory of emotion proposes that physiological arousal precedes cognitive appraisals, interpretations of peripheral changes in line with the given context, which together create the emotional experience and as such result in more finely tuned and distinguished subjective affective states.

What is important is that the *perception* of arousal can lead to emotional experiences. Following the two-factor model and the earlier James-Lange account of emotion, the ability to distinguish between bodily changes should moderate the

² Psychophysiological processes are defined as “the use of physiological signals to understand psychological processes”, Larsen, Berntson, Pehlmann, Ito, & Cacioppo (2008, p. 181).

relationship between peripheral changes and cognitive-affective processes. That is, the better interoceptive ability, the stronger the relationship - as peripheral changes should only be related to affective-cognitive processing if one can (accurately) detect them. Moreover, interoception should only be related to the arousal, but not to the valence dimension that together form affective experiences (see circumplex model of emotions, Russell, 1980), since peripheral changes are related to arousal, but not to valence. This is of importance to the current research: If interoceptively more accurate individuals would simply perceive negative events as worse than less accurate individuals, this could explain any moderation of the relationship between negative utility and probability estimates in terms of utility, whilst the current hypothesis relates only to the fact that interoceptively accurate individuals perceive arousal as more intense.

The above-mentioned relationship between interoception and the arousal component of affective experiences is well documented. For instance, Dunn et al. (2010) assessed interoceptive accuracy and then measured the relationship between heart rate and subjective arousal and between heart rate and subjective valence. Interoceptive accuracy moderated the relationship between heart rate and subjective arousal: interoceptively accurate individuals reported greater arousal with increasing heart rate, whilst this relationship was less pronounced for less interoceptively accurate individuals. However, there was no difference between individuals high and low in interoceptive accuracy in their relationship between heart rate and subjective valence.

In a similar vein, Wiens et al. (2000) assessed interoceptive accuracy and asked participants to rate how intense, and how pleasant or unpleasant positive and negative film clips made them feel, assessing arousal and valence, respectively. Whilst interoceptively accurate individuals reported greater intensity for both positive and negative film clips than

interoceptively less accurate individuals, there were no differences in pleasantness ratings (but see Werner, Mannhart, Del Paso, & Duschek, 2014).

Furthermore, Feldman-Barrett et al. (2004) assessed interoceptive accuracy as well as participants' *arousal focus*, defined as the "extent to which people emphasize arousal when reporting their experiences over time" (Feldman-Barrett et al. 2004, p. 1). Over a 60-day period, participants indicated the degree to which 88 emotion adjectives described their current affective state. Participants high in interoceptive accuracy rated themselves higher on adjectives that are related to feelings of activation and deactivation (related to the arousal component associated with each emotional state). However, there was no difference in the extent to which participants high or low in interoceptive accuracy reported negative or positive emotional adjectives.

Summing up, evidence is supporting the assumption that interoception is underlying the experience of arousal, and that individuals differ in the intensity with which they experience levels of arousal and the affective experiences associated with it (see also Critchley et al. 2004; Ferguson & Katkin, 1996; Herbert, Pollatos, & Schandry, 2007; Pollatos, Gramann, & Schandry, 2007; Wiens et al. 2000).

Moreover, the experience of arousal is also hypothesized to underlie decision-making in certain paradigms. If interoception causes individuals to experience arousal more intensely, then the assumption should follow that when individuals differ in their interoceptive ability and the decisions are based on the experience of arousal, they should also differ in the way they make decisions.

Interoception and Decision-Making

We will next turn to experiments demonstrating that individuals that differ in interoception also differ in the way they make use of arousal in decision-making tasks.

In order to incorporate arousal into one's decision-making, one must sufficiently experience it. Therefore, the more individuals experience arousal, the more should arousal inform their decision-making processes.

Werner, Jung, Duschek, & Schandry (2009) first assessed interoceptive accuracy and afterwards asked participants to complete the IOWA gambling task (IGT). In the IGT participants are presented with 4 decks of cards, and, across 100 trials can choose to take one card from one of the 4 decks. Each card has a pay-off as well as a loss associated with it. Two of the 4 decks yield small pay-offs and small losses, and are advantageous over the long run, whereas the other 2 decks yield large pay-offs and large losses, and are disadvantageous over the long run. As argued by Bechara, Damasio, & Damasio, 2000; Damasio, 1996), participants, over time, learn to choose from advantageous decks due to them making use of somatic responses (e.g. heart rate). In their experiment, Werner et al. (2009) demonstrated that individuals high in interoceptive accuracy chose fewer disadvantageous decks compared to individuals low in interoceptive accuracy, indicating that the former performed better due to their greater ability to perceive somatic responses.

Furthering this line of research, Dunn et al. (2010) asked participants to complete a modified version of the IGT. Measuring electrodermal activity and heart rate prior to each trial, Dunn et al. (2010) showed that individuals high in interoceptive sensitivity made more use of physiological feedback – irrespective as to whether this feedback favoured profitable or non-profitable decisions. Werner et al. (2013) supported the earlier findings by Werner et al. (2009) and Dunn et al. (2010): they demonstrated a greater activity in the right anterior insular (the brain region strongly associated with

interoceptive ability, see Critchley et al. 2004) for participants high in interoceptive sensitivity preceding disadvantageous decks. However, contrary to this previous research, Werner et al. (2013) did not show that this greater activity led to a decrease in the selection of disadvantageous decks.

Although the IGT and interferences from it has come under considerable critique (e.g. Maia & McClelland, 2004; Konstantinidis & Shanks, 2014; Steingroever et al. 2013), these results nevertheless show support for the notion that individuals high in interoceptive awareness appear to base their decision-making on physiological feedback more than individuals low in interoceptive awareness.

Whilst previous studies such as Dunn et al. (2010) drew a general connection between interoception and risky decision-making, Sokol-Hessner, Hartley, Hamilton and Phelps (2014) demonstrated more concretely that interoceptive accuracy predicts loss aversion- the tendency to be more sensitive to losses than to equal amounts of gains (Tversky & Kahneman, 1992). Physiological arousal has been shown to have direct effects on loss aversion (Sokol-Hessner, Hsu, Curley, Delgado, Camerer, & Phelps, 2009), warranting the hypothesis of interoceptive accuracy as a moderator. In their experiment, Sokol-Hessner et al. (2014) first assessed interoceptive accuracy, and then asked participants to make choices between a risky gamble and a secure option, parametrically measuring loss aversion. Here, the more interoceptively accurate participants were, the more loss averse they were, too.

Interoception and the Stake-Likelihood Hypothesis

The evidence above suggests that interoception can lead to greater reliance on arousal during decision-making processes, apparently due to the more accurate

perception of arousal. We propose that, in line with these findings, interoception should moderate the relationship between negative utility and probability estimates. As hypothesized by Vosgerau (2010), having a (positive or negative) stake in the outcome increases arousal, which should be misattributed to inform probability estimates. Individual differences in how strongly such arousal is perceived should moderate this relationship. We hypothesize that individuals that perceive the arousal evoked by having a stake in the outcome as more intense should report greater probability estimates compared to individuals that are less able to perceive the evoked arousal.

Note, however, that it would also be possible to predict a moderation in the opposite direction. It could be possible that individuals high in interoceptive ability are more likely to be aware of the *source* of their arousal, and are consequently *less* likely to misattribute arousal. For example, Reisenzein and Gattinger (1982) measured interoceptive sensibility and later manipulated arousal by means of a brief period of exercise. Next, either a positive or a negative mood manipulation took place, after which participants were asked to rate their affective states. It was hypothesised that interoceptively sensible individuals would be less likely to misattribute the residual arousal from exercising to their current affective state. This hypothesis was supported only in the negative domain, where individuals high in interoceptive sensibility experienced less self-reported affect in line with the mood induction procedure. In addition, Gibbons, Carver, Scheier, and Hormuth (1979) demonstrated that interoceptively sensible individuals were less likely to report bodily changes that were ostensibly induced by a (placebo) pill (see also Gibbons and Gaeddert, 1984). However, our studies differ with regard to the nature of the above-described studies. Riding a bike

can evoke arousal too intense to be misattributed to a stimulus of much weaker amplitude. If this difference between stimuli is too large, one becomes aware of the source of the arousal, which stops misattribution processes (Zillman, Mody & Cantor, 1974). Our arousal manipulation is a lot more subtle than the exercise manipulation by Reizenzein and Gatteringer (1982), making the source of arousal less prominent. Furthermore, contrary to Gibbons et al. (1979), our experiments actually manipulate arousal. Therefore, if interoceptive individuals are better at perceiving when there is *no* bodily change, they should also be better at perceiving when there *is* a bodily change.

Assessing Interoception

In the current studies, we will assess both the potential moderation by interoceptive sensibility (self-report measure of one's internal focus) and interoceptive accuracy (objective measure of interoceptive ability). Whilst most of the research that we have discussed before has described the impact of interoceptive accuracy on affective experiences and decision-making processes, there is reason to assume that the impact of interoceptive sensibility is comparable. For example, research by Schnall, Haidt, Clore, & Jordan (2008) demonstrated that individuals make harsher moral judgments when having been disgusted by an irrelevant stimuli beforehand. However, only individuals high in interoceptive sensibility demonstrated this misattribution of feeling disgusted to how morally wrong certain actions (e.g. marrying your 1st degree cousin) are. Häfner (2013) demonstrated identical effects of interoceptive accuracy and sensibility on embodied cognition, supporting the notion that interoceptive sensibility is a good proxy of interoceptive accuracy.

Moreover, Critchley et al. (2004) measured both interoceptive accuracy and interoceptive sensibility and observed the neural correlates of both dimensions of interoception. Importantly, both dimensions correlated with grey matter volume in the right anterior insula, a brain region proposed to mediate attention to visceral and somatic states (however, the two dimensions were not inter-correlated, despite similar relationships with the right anterior insula). Likewise, Terasawa, Fukushima and Umeda (2013) measured neural activity whilst participants appraised both body physiology (“I have a fast pulse”) and emotions (“I am happy”). Here, the right interior insular cortex and the ventromedial prefrontal cortex were strongly activated by appraisal of both body physiology and emotions, suggesting an overlap between the two constructs.

Therefore, given that both dimensions of interoception are related to greater attention to internal states and arousal, as a first affordable step, we will examine the hypothesis that interoception moderates the impact of negative utility on probability estimates by assessing interoceptive sensibility. Finally, we will also assess the moderation by interoceptive accuracy (a hypothesis, which one might argue is more strongly supported by prior research).

Measuring Interoceptive Accuracy. As interoception is a multidimensional construct, several different methods of assessing each dimension have been proposed. The ability to perceive distension locations in the gut (Hälzl, Erasmus, & Möltner, 1996) or the effect of drugs on one’s heart rate (Khalsa, Rudrauf, Sandesara, Olshanksy & Tranel, 2009) and the ability to count one’s heartbeats (cardiac perception, Schandry et al. 1981) have all been utilised to assess interoceptive accuracy, the latter being the most frequently used method. Since the James-Lange and Schachter and Singer theories of

emotion focus primarily on feedback from the autonomic nervous system (see Reisenzein, 1995), cardiac perception constitutes a good proxy for measuring interoceptive accuracy. Indeed, Critchley et al. (2004; see also Craig, 2002; Pollatos, Schandry, Auer, & Kaufmann, 2007) demonstrated that cardiac perception is related to brain networks associated with monitoring of internal states (Craig, 2002; Critchley et al. 2004; Pollatos et al. 2007) and emotional experiences (Damasio et al. 2000). The details of the measurement of cardiac perception will be discussed in Experiment 3.

Measuring Interoceptive Sensibility. As this dimension is defined as the subjective self-report of one's attentional focus on internal bodily changes and awareness of bodily sensations, interoceptive sensibility is assessed by means of questionnaire measures. Interoceptive sensibility can be broken down further into sub-dimensions. Mehling et al. (2012) distinguishes between 7 dimensions that are displayed in Table 2.2.

Table 2.2

Dimensions of interoceptive sensibility (Mehling et al. 2012)

Dimension	Definition
Noticing	Awareness of bodily sensations
Distracting	Distracting oneself from painful sensations
Worrying	Worrying about painful or uncomfortable sensations
Attention Regulation	Control over attention to bodily sensations
Emotional Awareness	Awareness over the connection between bodily sensations and affective states
Self-Regulation	Regulation of distress by directing attention to bodily sensations
Body Listening	Directing attention to bodily sensations for insight
Trusting	Trusting one's body

Different questionnaires examine different dimensions of interoceptive sensibility. In the current studies, we are particularly interested in the dimension of

“noticing”. As this dimension assesses the awareness of bodily sensations such as one’s heartbeat, it is the most direct self-report of interoceptive accuracy. The scale most straightforwardly relating to the dimension of “noticing” is the Consciousness of Body: Private scale (CBP, Miller, Murphy, & Buss, 1981). This scale directly assesses interoceptive skills by asking participants about the noticing of physical sensations such as one’s mouth getting dry, the ability to feel one’s heartbeat, etc. This scale is therefore the scale most relevant to the current research. However, Experiment 1 is the first investigation into this subject, and for exploratory purposes we will also assess interoceptive sensibility with three other scales described below.

We also administered the Body Awareness Questionnaire (Shields, Mallory & Simon, 1989). This questionnaire has four sub-domains: The “note response or changes in body process”, “sleep-wake cycle”, “predict body reactions”, and “prediction of the onset of illness” sub-domains. These four sub-domains are not scored separately. Thus, this questionnaire does not assess “noticing” as directly as the CBP scale.

Furthermore, we also assessed two other dimensions of interoceptive sensibility that are relevant to the current research: Emotional awareness, as a noticing of how one’s body changes when one is in different emotional states, and body listening, the active scanning of bodily reactions as a top-down process. Mehling et al.’s (2012) “Multidimensional Assessment of Interoceptive Awareness” (MAIA) reports two instruments that measure emotional awareness (MAIA-E) and body listening (MAIA-BL), and both will be employed in Experiment 1.

Experiment 1

Background

The purpose of Experiment 1 was to examine the impact of interoceptive sensibility on the relationship between negative utility and probability estimates. As mentioned above, interoceptive sensibility was measured by 4 different scales in Experiment 1. We administered the CBP scale as the main scale of interest, as well as the BAQ scale as the most frequently utilised, to demonstrate generalizability across different means of assessing the “noticing” dimension. Besides, we also assessed whether the potential moderation would hold for individuals who pay attention to the physiological manifestation of affective states by administering the MAIA-E subscale of the MAIA. In addition, testing for a top-down influence on this process, participants also completed the body listening subscale (MAIA-BL). The administration of these different scales and testing for different components did not serve the concrete aim of being able to distinguish between different impacts, but rather served as a pilot study testing for the generalizability of the effect with scales where we would expect to see a moderation based on the SLH.

The paradigms of eliciting probability estimates followed previous work by Harris et al. (2009) and Vosgerau (2010). Here, participants are presented with fictional scenarios, avoiding confounds of real-world base-rates. On one hand, more negative events are less frequent than less negative events in the real world (Pleskac & Hertwig, in press; see also, Dai et al., 2008; Weber & Hilton, 1990). On the other hand, some negative events (e.g. plane accidents) are judged as more prevalent than they truly are as a result of, for example, media coverage (e.g. Slovic, Fischhoff, & Lichtenstein, 1982).

Fictional scenarios avoid both types of potential confounds. Also crucial to the experimental design, participants are supplied with an objective basis for their subjective estimates and this objective basis is identical across the utility manipulations. Any systematic difference between the estimates of probability across conditions is consequently directly attributable to the manipulation of utility.

Experiment 1 employed two different scenarios. One scenario was the unmodified scenario from Harris et al.'s (2009) Experiment 5. In Experiment 5, Harris et al. (2009) manipulated the controllability of the event, showing that participants only demonstrate a negativity bias when the event involved decision-control (see Chapter 1). Therefore, we included the scenario from the 'high controllability' condition. In this scenario, participants are asked to imagine that the Royal Air-Force (here adapted to the US Air-Force, USAF) is looking for a new training site. Furthermore, participants are told that debris and crashes are not uncommon for training sites. Participants in the negative condition are informed that the area currently favoured would involve flying over a densely populated town, whereas participants in the neutral condition are informed that the area is unpopulated wasteland. Participants are presented with a visual display that consists of a map displaying white dry land and a blue river flowing through this land, and are asked to estimate the chance that if debris were to fall it would fall onto the dry land and not into the river. Therefore, outcome utility is manipulated by describing that if the debris were to land on the white land it would kill inhabitants of the town (negative condition), or litter the environment (neutral condition).

The second scenario employed was a version of Vosgerau's (2010) Experiment 3, adapted to be more appropriate for an online experiment. In the original experiment,

participants gambled with real money: they were given \$3 and told that they would lose the money should a 6 come up on at least 2 out of 4 dice rolls. Following, participants were asked to estimate the likelihood of losing the \$3. In the negative condition of this modified version, we asked participants to imagine they would walk down the street with their friend and find \$100, and that their friend proposes that they would roll a dice four times. They were further told to imagine that if a 6 comes up on at least two of these four throws, they would lose the \$100 to their friend. No outcome was attached to the rolling of the dice in the neutral condition.

Method

Participants and Design. 233 subjects were recruited via Amazon Mechanical Turk for a reimbursement of \$0.2. After removing 36 subjects for having duplicate IP addresses (see Harris et al. 2009), the final sample consisted of 198 participants (97 female, median age=29), who took part in this experiment online. Participants were randomly assigned to a 4 (Body Awareness Questionnaire (BAQ)/Private Consciousness Scale (CBP)/Emotional Awareness (MAIA-E)/Body Listening (MAIA-BL)) x 2 (scenario: USAF/Dice) x 2 (utility: negative/neutral) mixed design, the last factor manipulated between participants.

Materials. Participants completed the CBP scale (Miller et al., 1981). Five items assess sensitivity to bodily reactions (“I am sensitive to internal bodily tensions”, “I know immediately when my mouth or throat gets dry”, “I can often feel my heart beating”, “I am quick to sense the hunger contractions in my stomach”, “I am very aware of changes in my body temperature”), and responded on a Likert scale from 1 (disagree strongly) to 6 (agree strongly) (see Appendix A). Moreover, the BAQ (Shields

et al., 1989) was administered. This questionnaire consists of 18 statements referring to the sensitivity to non-emotive bodily processes (“I notice differences in the way my body reacts to various foods”, “I notice distinct body reactions when I am fatigued”, etc.), which participants answer on a Likert scale from 1 (not at all true of me) to 7 (very true of me) (see Appendix B). In addition, participants answered two subscales of the MAIA instrument (Mehling et al., 2012). Participants completed the emotional awareness subscale (MAIA-E) (“I notice how my body changes when I am angry”, “I notice that my breathing becomes free and easy when I feel comfortable”, etc.) (see Appendix C), and the body listening subscale (MAIA-BL) (“I listen for information from my body about my emotional state”, etc.) (see Appendix D). Both MAIA subscales were answered on a Likert scale from 0 (never) to 5 (always).

USAF Scenario. For the elicitation of probability estimates, participants were presented with the visual stimuli used in Experiment 5 from Harris et al. (2009) (see Figure 2.2). The cover story was identical to the high control condition from Harris et al. (2009). Participants in the negative outcome condition (neutral outcome in brackets) read:

The USAF are in need of a new training site for their pilots. The location currently favoured would involve flying over the area pictured below, in which the white area represents a densely populated town (uninhabited wasteland) and the blue area represents the river that flows through that town. Crashes and falling plane debris are not uncommon occurrences in USAF training sites, and if falling debris were to land on a populous area (dry land), it would kill anybody beneath it (litter that area). Any debris falling from the sky during training could land in any of the grid squares in the picture below.

The USAF have asked you to use the picture below to estimate the chance that any falling debris would land on the densely populated (uninhabited) dry land. By looking at the picture below, please estimate the chance that any falling debris will land on the densely populated (uninhabited) dry land.

The visual probability display from the USAF scenario is displayed in Figure 2.2.

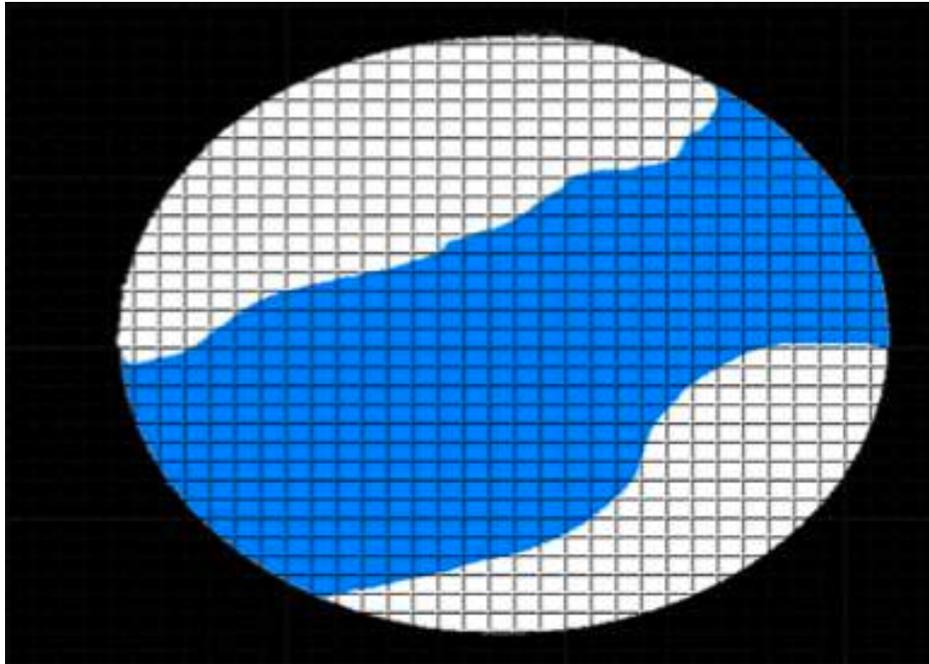


Figure 2.2. Visual probability display from Harris et al.'s (2009) Experiment 5.

Dice Scenario. In this modified version of the dice scenario from Vosgerau's (2010) Experiment 3, participants read the following text in the negative condition:

Imagine you and your friend walk together down the street and you find \$100. You picked it up, so the \$100 is in your pocket. However, your friend proposes the following:

Your friend will roll a six-sided dice 4 times. If a 6 comes up on at least 2 of these throws, your friend will get the \$100 and you will lose the \$100. Otherwise you can keep it.

What do you think is the chance that a 6 would turn up on at least 2 out of 4 throws, so you would lose the money?

In the neutral condition, on the other hand, participants read:

Imagine you and your friend are together in a park. Your friend gets out a dice from their pocket and tells you that they will roll the dice four times.

What do you think would the probability be like for a 6 to turn up on at least 2 out of 4 throws?

Estimates for both scenarios were provided on slider scales anchored at 0 and 100 on the other end, and participants could move the slider in increments of 1. Participants were informed that 0 means that the outcome is “absolutely impossible” and 100 that the outcome is “absolutely certain”. Participants saw the value of the position of the slider, but the scale was not anchored further.

Procedure. Participants were informed that the study consists of 2 parts, the first one being a pilot study for a future experiment testing three different questionnaires (MAIA-E and MAIA-BL were presented on one page), and the second part being a study on risk perception. Participants were provided with the four scales (the order of which was counterbalanced, but the MAIA-E and MAIA-BL scales were presented together) and asked to complete them independently of one another. Before completing the two scenarios in a counterbalanced order, participants were informed that we are only interested in their estimates rather than in calculations. This was done to avoid participants looking up the formula for the binomial distribution with which one can calculate the probability of the dice scenario, or counting the squares in the USAF scenario. Following, a manipulation check asked participants to estimate how “bad” the negative outcomes (debris falling down, friend rolling 2x a 6) would be on a 7-point Likert scale. Afterwards, participants were thanked and debriefed.

Results

Descriptive Statistics. For the consciousness of private scale (CBP) the mean was 4.18, the variance was .85 and the range from 1.6 to 6 (on a scale from 1 to 6). For

the body awareness questionnaire (BAQ) the mean was 4.5, the variance was .73 and answers ranged from 1.35 to 6.65 (on a scale from 1 to 7). The emotion-information subscale from the Multidimensional Assessment of Interoceptive Awareness (MAIA-E) had a mean of 3.25 with a variance of 1.18, whilst the spread was satisfactory from 0 to 5 (on a scale from 0-5). For the body information subscale (MAIA-BL) the mean was slightly lower with 2.39 with a variance of 1.63 and a range from 0 to 5 (also on a scale from 0 to 5). Cronbach's alpha was satisfactory for CBP, BAQ, MAIA-E, MAIA-IB, with .76, .85, .88, and .88, respectively. See Table 2.3 for correlations between the scales (showing medium to strong correlations).

Table 2.3

Correlations between the four scales MAIA-E, MAIA-BL, CBP and BAQ.

	CBP	BAQ	MAIA-E	MAIA-BL
CBP		.70**	.61**	.53**
BAQ	.70**		.61**	.54**
MAIA-E	.61**	.61**		.64**
MAIA-BL	.53**	.54**	.64**	

** . Correlation is significant at the 0.01 level.

Manipulation Checks

USAF Utility. An ANOVA with utility as a between subjects factor was run on the answer to the question “how bad would it be if debris would be dropped onto the dry land”? Based on a 7-point scale from 1 (not at all bad) to 7 (very bad), participants in the negative condition indicated that debris falling on dry land would be worse ($M=6.25$, $SD=.86$) compared to those in the neutral condition ($M=4.35$, $SD=1.79$), and this difference was significant $F(1,196)=88.67$, $p<.001$, $\eta^2=.31$.

To test for potential interactions between utility and interoceptive sensibility (IS), four 2 (utility: negative/neutral) x 2 (IS: high/low) ANOVAs with the four IS questionnaires were run. This was important in order to rule out the possibility that individuals high in IS simply perceive the outcome as worse than individuals low in IS.

In order to determine the moderating role of IS, we chose to perform median-splits on the 4 questionnaires (see Schnall et al. 2008). Conducting a median split enabled us to also assess concrete hypotheses by means of orthogonal contrasts (as discussed later).

Including CBP did not lead to a significant interaction $F(1,194) = 2.59, p = .11, \eta^2 = .01$, however, a significant main effect of CBP was found, $F(1,194) = 5.25, p = .02, \eta^2 = .03$. Participants low in CBP rated both negative and neutral events as worse ($M = 5.52, SD = 1.52$) than participants high in CBP ($M = 5.04, SD = 1.84$). Similarly, for the BAQ there was no interaction, $F(1,194) = 2.22, p = .14, \eta^2 = .01$, but a marginal main effect of BAQ, $F(1,194) = 3.64, p = .06, \eta^2 = .02$. Again, those scoring low on the BAQ overall indicated the outcome to be worse ($M = 5.57, SD = 1.51$) compared to those scoring high ($M = 4.99, SD = 1.84$). For the MAIA-E, the interaction was not significant, $F < 1$, and neither was the main effect of MAIA-E $F(1,194) = 1.24, p = .27, \eta^2 = .01$. For the MAIA-BL, neither the interaction nor the main effect were significant, both $Fs < 1$.

Thus, two out of the four questionnaires showed a main effect of IS, in that individuals low in IS perceive *both* neutral and negative events as worse than individuals high in IS. Importantly, however, individuals high in IS did not perceive the negative outcome as worse than individuals low in IS. As Harris et al. (2009) demonstrated, severely negative events are assigned a greater probability than mildly

negative events. Therefore, if those high in IS would simply perceive the utility of the event as worse, this difference could potentially explain differences in probability estimates for negative events.

Dice Utility. An ANOVA with utility as a between subjects factor was run on the answer to the question “how bad would it be if your friend would roll at least two times a 6”? Again, participants answered on a 7-point scale from 1 (not bad at all) to 7 (very bad). Participants in the negative condition rated the outcome as worse ($M=4.69$, $SD=1.86$) compared to those in the neutral condition ($M=1.58$, $SD=2.16$), and this difference was significant, $F(1,195)=208.57$, $p<.001$, $\eta^2=.52$.

As before, four ANOVAs were run including the four IS measures to examine the impact of IS on the utility manipulation check. For the CBP and BAQ scales there were neither main effects, nor an interaction, all F s <1 . For MAIA-E, there was no significant interaction, $F(1,193)=1.48$, $p=.23$, $\eta^2=.01$ ³, but a marginally significant main effect of MAIA-E, $F(1,193)=2.84$, $p=.09$, $\eta^2=.01$. Participants low in MAIA-E indicated the outcomes to be worse ($M=3.47$, $SD=2.2$) compared to those high in MAIA-E ($M=2.76$, $SD=2.08$). Including MAIA-BL revealed a significant interaction, $F(1,193)=4.63$, $p=.03$, $\eta^2=.02$, but no significant main effect. Upon inspection of the interaction, it became evident that there was a difference between participants scoring high and low on the scale for negative events only: those low in MAIA-BL rated those events as more negative ($M=5.02$, $SD=1.7$) compared to those high in MAIA-BL ($M=4.33$, $SD=1.94$), and this difference was significant $F(1,193)=5.2$, $p=.02$. This was not the case for neutral events, $F<1$.

³ One participant did not complete this manipulation check, accounting for the difference in degrees of freedom.

In sum, two of the four scales impacted the perception of utility, individuals low in MAIA-E perceived the scenarios to be worse, and individuals low in MAIA-BL perceived the outcome as worse in the negative condition. However, only the opposite finding, if individuals high in IS would have perceived the outcome as worse would have been a reason for concern, as this could potentially explain why those high in IS would assign greater probability estimates to negative events, if the data were to follow the predicted pattern.

Outlier Removal. Consistently throughout this thesis, unless otherwise indicated, we removed disproportionately influential data points. Whilst Harris et al. (2009), using a similar methodology, did not remove any outliers, we wanted to avoid that single data points would strongly impact our results. We reasoned that this is important in particular in often noisy online experiments (the predominant work of this thesis) or when working with physiological measures (Experiments 3 and 9).

Thus, the analyses reported in this thesis are *always after* these outliers have been removed (unless otherwise indicated), and all graphs or descriptive statistics are indeed on the means where outliers have already been excluded. However, should there be a difference in the results between analyses were outliers have or have not been removed, these are always indicated in footnotes. If no further footnotes follow an analysis, the analyses with or without outlier removal produce the same results.

In order to identify individuals that disproportionately influenced the results, we plotted Studentized Deleted Residuals against Cook's Distance and visually identified participants that are high on both indexes (see Cohen, Cohen, West, & Aiken, 2003). Here the cut-off points of +2 and - 2 (Fox, 1991) and 4/N (Bollen, Kenneth, & Jackman,

1990) were adopted for Studentized Deleted Residuals and Cook's distance, respectively.

Probability Estimates. For the analyses reported, we first assessed the joint impact of interoceptive sensibility and utility on probability estimates with a 2 (IS: high/low) x 2 (utility: negative/neutral) ANOVA. However, since our concrete hypothesis was that only individuals high in IS should demonstrate a negativity bias (but not that those low in IS would *decrease* probability estimates under negative affect), we did not predict a disordinal interaction between the two variables per se. Therefore, in order to test the concrete hypothesis that those high in IS would show an increase in probability estimates for negative scenarios, but not individuals low in IS, we conducted orthogonal contrasts (see Table 2.4). Concretely, we predict only contrast 1 to be significant, in line with the SLH. That is, we predict a significant difference between individuals high in IS estimating negative events compared to all other conditions, but we don't expect there to be a difference between individuals low in IS estimating negative events compared to the neutral condition, and we also did not hypothesize a difference between individuals high and low in IS in the neutral condition.

Table 2.4

Orthogonal contrasts tested in Chapter 2.

	High IS	Low IS	High IS	Low IS
	Negative	Negative	Neutral	Neutral
Contrast 1	3	-1	-1	-1
Contrast 2	0	2	-1	-1
Contrast 3	0	0	1	-1

Aggregated Analyses. First, we conducted a 2 (scenario: USAF/Dice) x 2 (utility: neutral/negative) x 2 (IS: low/high) repeated measures ANOVA. Here, the median split of the z-scored average across all four IS scales was entered as a between subjects factor (“IS”). Unsurprisingly, a main effect of scenario was present, $F(1,188)=295.76, p<.001, \eta^2=.61$. Neither the scenario*IS, nor the scenario*utility, nor the three-way interaction was significant, all $F_s<1$. There was, however, a significant main effect of utility, $F(1,188)=8.69, p<.01, \eta^2=.04$, as well as a significant main effect of IS, $F(1,188)=8.83, p<.01, \eta^2=.05$. The interaction between utility and IS was not significant, $F(1,188)=2.5, p=.12, \eta^2=.01$ (see Figure 2.3).

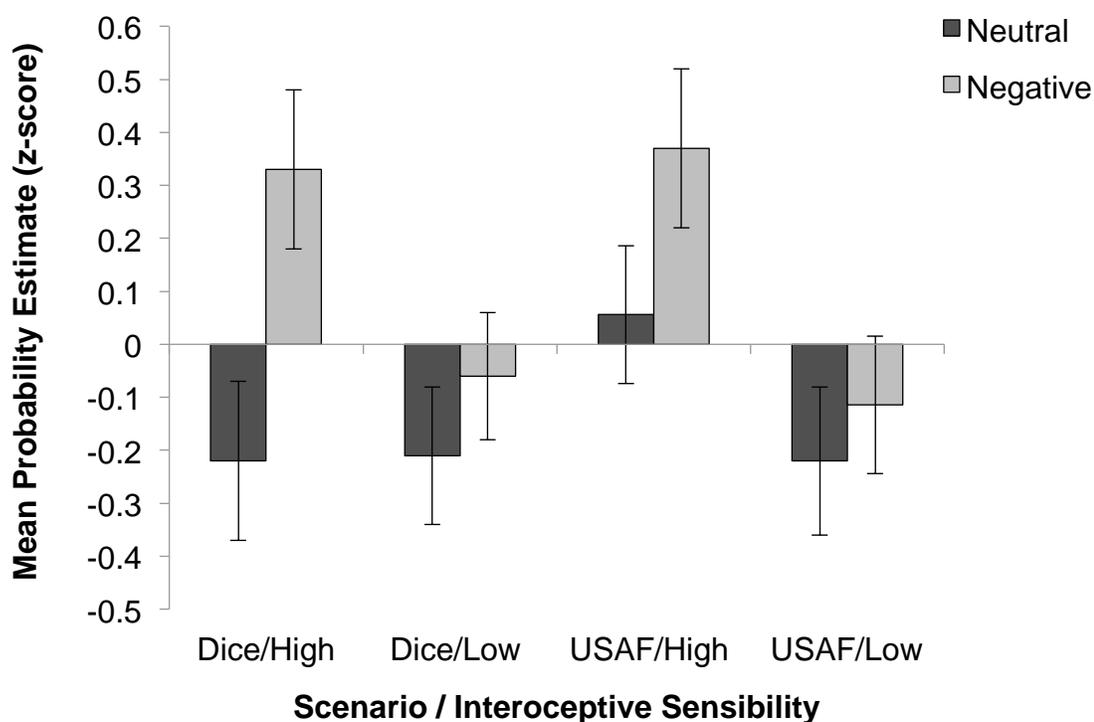


Figure 2.3. Probability estimate means for participants high and low in IS (averaged across all 4 scales) for neutral and negative conditions for both scenarios. Error bars represent +/- 1 standard error of the mean.

Next, as there were no interactions involving the scenario, we z-scored and averaged the two scenarios. This allowed us to conduct orthogonal contrasts combining both scenarios and scales (see Table 2.4).

The critical contrast 1, comparing those high in IS in the negative condition to all other conditions was significant, $t(188)=3.77, p<.001$ (see Figure 2.4). No other contrast was found significant, all $ps>.60$. An omnibus test conducted assessing whether there was any rest-variance not accounted for by the assessed contrasts was not significant, $F<1$. This means that there is no significant amount of variance that the contrasts we tested do not describe⁴.

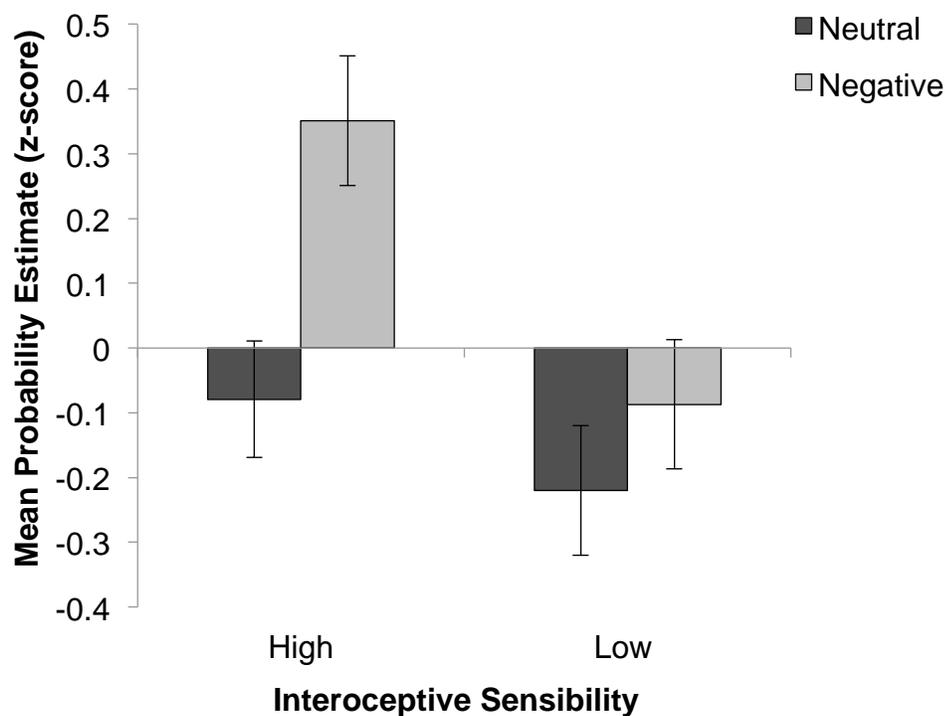


Figure 2.4. Probability estimate means for participants high and low in IS (averaged across all 4 scales) for neutral and negative conditions for both scenarios. Error bars represent +/- 1 standard error of the mean.

⁴ Note that before the removal of outliers, the main effect of utility is not significant, $F<1$, and the CBP*utility interaction is marginally significant, $F(1,194)=3.71, p=.06$. All other results remain the same.

Whilst these first analyses support our hypothesis that only individuals high in IS demonstrate a negativity bias, examining individual scenarios and scales painted a slightly less consistent picture, despite an absence of an interaction between the scenarios, utility and the scales. In the interest of being open about the present results in this new area of research, we report the individual analyses examining each scenario with each scale separately below.

USAF Scenario.

CBP (USAF Probability Estimates). Whilst the effect of utility was marginally significant, $F(1,183)=2.42$, $p=.06$, $\eta^2=.02$, there was a strong main effect of CBP, $F(1,188)=12.71$, $p<.001$, $\eta^2=.07$, and a significant utility*CBP interaction, $F(1,188)=6.58$, $p=.01$, $\eta^2=.04$. Simple effect analyses revealed that individuals high in CBP provided greater probability estimates in the negative as compared to the neutral condition, $F(1,183)=9.77$, $p<.01$, $\eta^2=.05$, whilst this was not the case for those low in CBP, $F<1$. Furthermore, there was a significant difference between those high and low in CBP on probability judgments for the negative, $F(1,183)=18.49$, $p<.001$, $\eta^2=.09$, but not for the neutral condition $F<1$ (see Figure 2.5).

Conducting the above described orthogonal contrasts, the critical contrast 1 was significant, $t(183)=4.59$, $p<.001$. Neither contrast 2, $t(183)=-.98$, $p=.33$, nor contrast 3 were significant, $t(183)=.71$, $p=.48$. There was no rest-variance not accounted for by the tested contrasts, $F<1$.⁵

⁵ Note that before the removal of outliers, the main effect of utility is not significant, $F<1$, and the CBP*utility interaction is marginally significant, $F(1,194)=3.71$, $p=.06$. All other results remain the same.

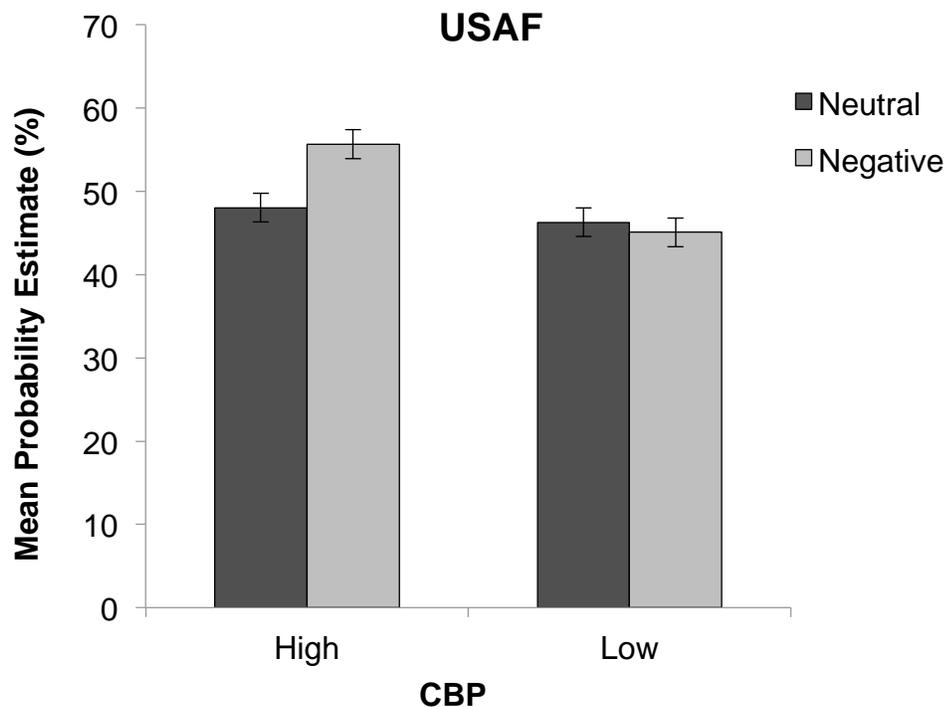


Figure 2.5. Probability estimates for participants high and low in CBP for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

BAQ (USAF). There was a marginally significant main effect of utility, $F(1,185)=3.57, p=.06, \eta^2=.02$, in that negative events were estimated to be more likely ($M=50, SD=12.13$) than neutral events ($M=46.78, SD=13.30$). Furthermore, there was a marginally significant main effect of BAQ, $F(1,185)=3.54, p=.06, \eta^2=.02$. Participants high in BAQ provided greater probability estimates ($M=49.89, SD=13.73$) than participants low in BAQ ($M=46.75, SD=11.7$). The BAQ*utility interaction was not significant, $F<1$. However, the critical contrast 1 was significant, $t(185)=2.55, p=.01$ (see Figure 2.6). No other contrasts were significant, all $ps>.45$, and no rest-variance was unaccounted for, $F<1$.⁶

⁶ Before the removal of outliers there were no main effects of utility, $F<1$, or BAQ, $F(1,194)=1.86, p=.17, \eta^2=.01$. None of the orthogonal contrasts were significant, all $ps>.21$.

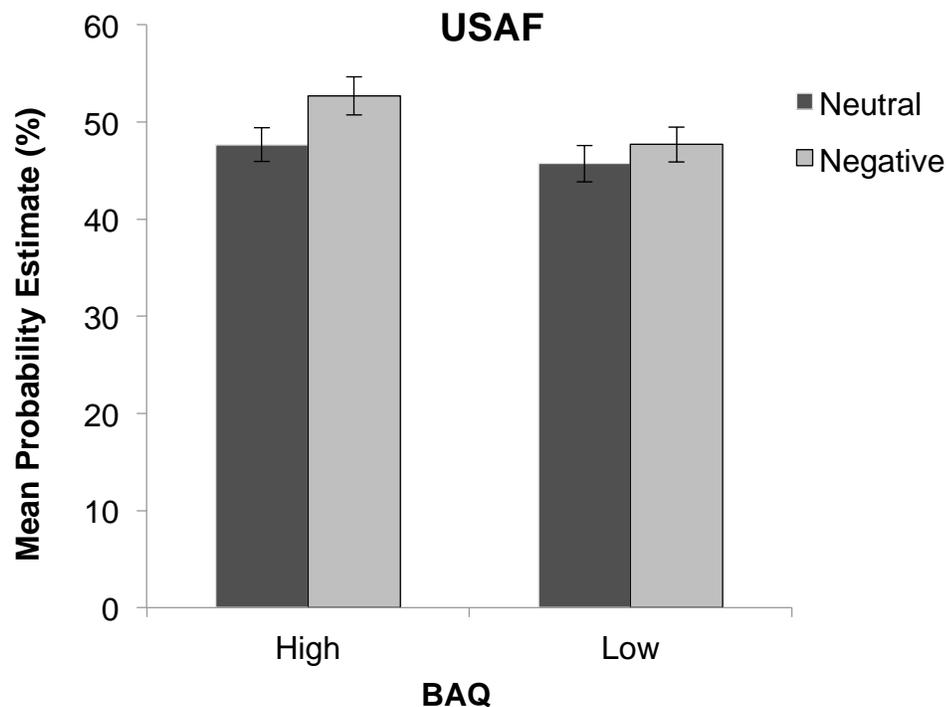


Figure 2.6. Probability estimates for participants high and low in BAQ for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

MAIA-E (USAF). There was a main effect of MAIA-E, $F(1,183)=4.13$, $p=.04$, $\eta^2=.02$. Individuals high in MAIA-E provided greater probability estimates ($M=50.35$, $SD=12.57$) than individuals low in MAIA-E ($M=46.98$, $SD=11.91$). However, the main effect of utility, $F(1,183)=2.58$, $p=.11$, $\eta^2=.01$, and the interaction, $F<1$, were non-significant (see Figure 2.7)⁷. However, the critical contrast 1 was marginally significant, $t(184)=1.91$, $p=.06$. Neither contrast 2, $t(184)=.66$, $p=.51$, nor contrast 3, $t(184)=1.58$, $p=.12$, were significant⁸. An omnibus test showed that there was no further between-condition variance left to account for, $F(2,184)=1.47$, $p=.23$.

⁷ The reason why utility is significant in one but not in other scenarios is most likely due to the fact that across different scales, different participants are identified as outliers.

⁸ Before the removal of outliers, neither the main effects of utility, nor of MAIA-E were significant, all $F_s<1$. None of the orthogonal contrasts were found significant, all $p_s>.38$.

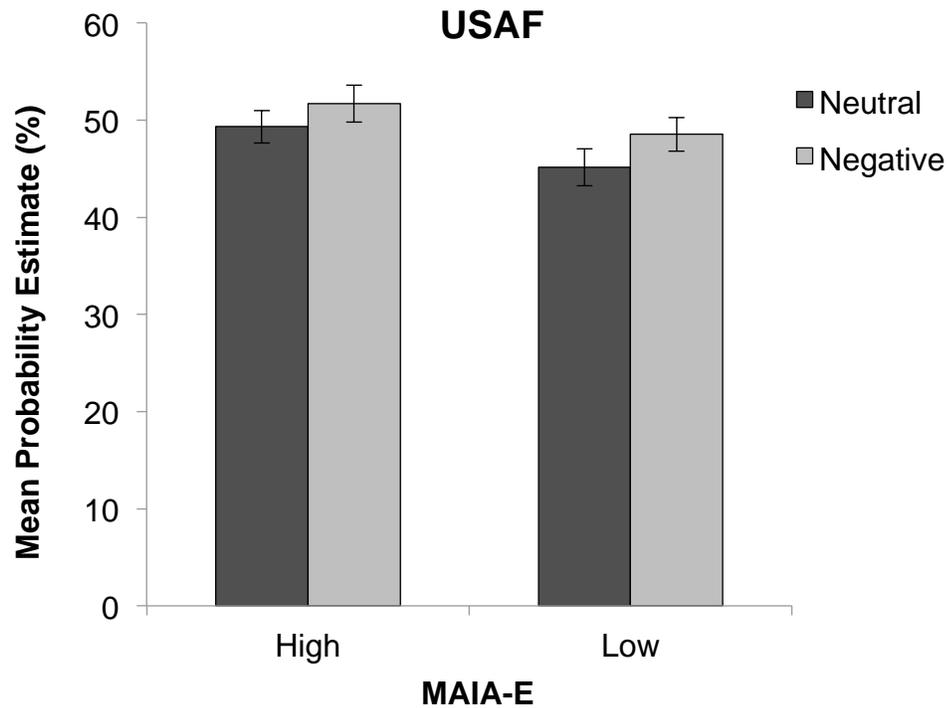


Figure 2.7. Probability estimates for participants high and low in MAIA-E for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

MAIA-BL (USAF). The main effect of MAIA-BL was marginally significant, $F(1,184)=3.44, p=.07, \eta^2=.02$, in that individuals high in MAIA-BL provided greater probability estimates ($M=50.28, SD=13.41$) compared to individuals low in MAIA-BL ($M=46.93, SD=11.62$). Neither the main effect of utility, $F(1,184)=2.53, p=.11, \eta^2=.01$, nor the MAIA-BL*utility interaction, $F<1$, were significant (see Figure 2.8).

Orthogonal contrasts showed that the critical contrast 1 was marginally significant, $t(184)=1.68, p=.09$, and contrast 2 was not significant, $t(184)=.73, p=.47$. However, contrary to our hypotheses, contrast 3 was also found marginally significant,

$t(184)=1.64, p=.10$, indicating that the MAIA-BL scale impacted probability estimates in the neutral condition, too.

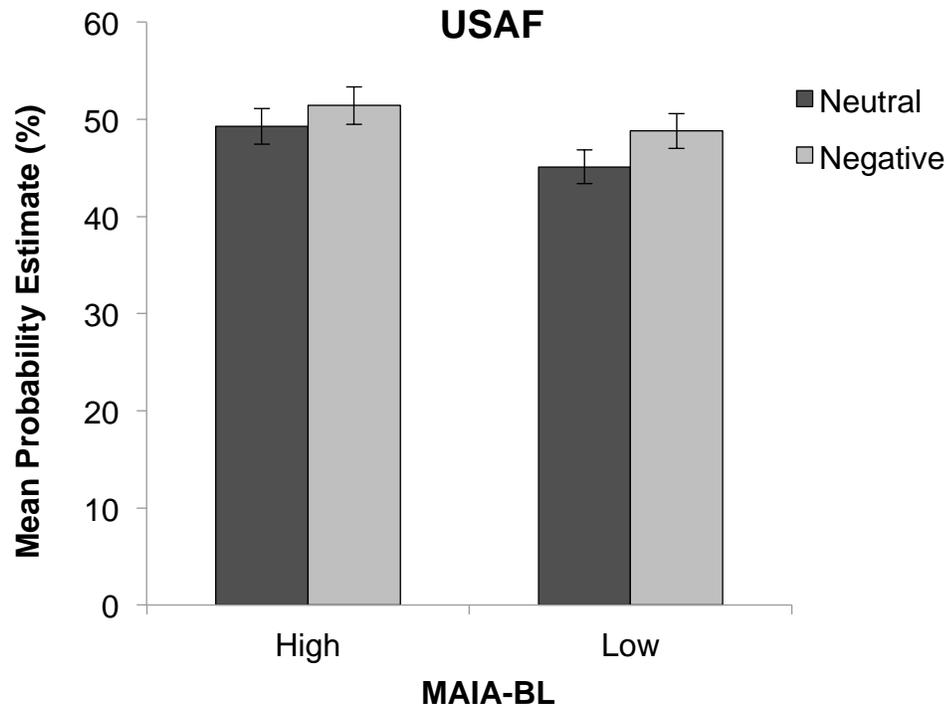


Figure 2.8. Probability estimates for participants high and low in MAIA-BL for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean

Dice Scenario.

CBP (Dice). A significant main effect of condition, $F(1,184)=7.77, p<.01, \eta^2=.04$, ($M_{\text{negative}}=24.55 (14.49)$ vs. $M_{\text{neutral}}=19.26 (11.8)$), as well as a significant main effect of CBP, $F(1,184)=6.14, p=.02, \eta^2=.03$, ($M_{\text{High}}=24 (14.41)$, $M_{\text{Low}}=19.5 (11.99)$) was found. The interaction between CBP and utility was not significant, $F(1,184)=2.45, p=.12$. Orthogonal contrasts showed that only the critical contrast 1 was significant,

$t(184)=3.92, p<.001$ (see Figure 2.9). No other contrasts were significant, all $ps>.51$, and no rest-variance was left to be accounted for, $F<1$.⁹

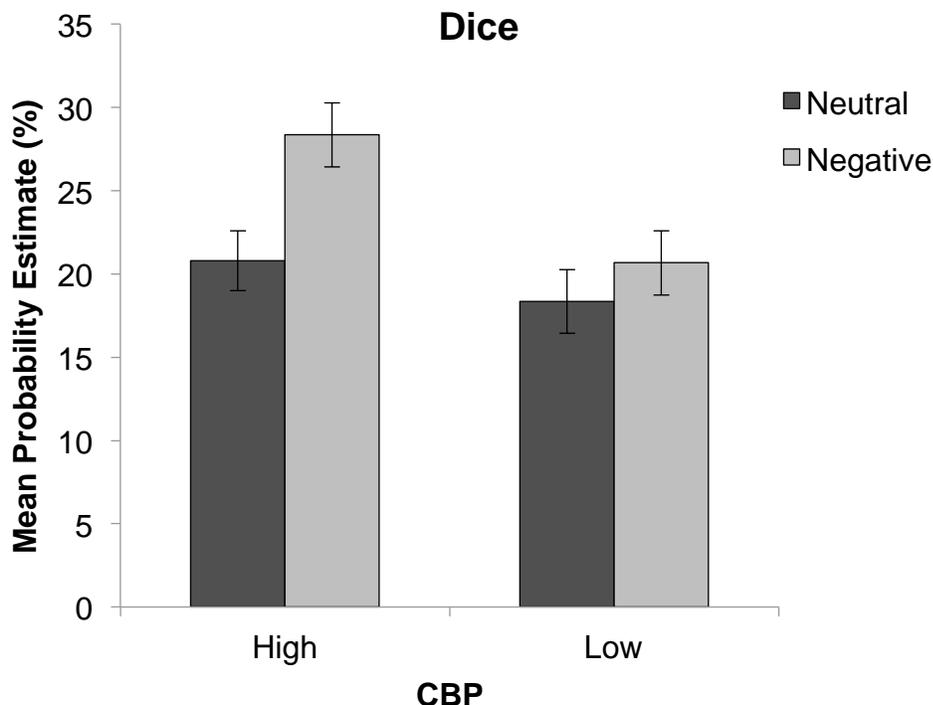


Figure 2.9. Probability estimates for participants high and low in CBP for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

BAQ (Dice). A main effect of utility, $F(1,185)=9.35, p<.01, \eta^2=.05$, and a marginally significant main effect of BAQ were found, $F(1,185)=3.21, p=.08, \eta^2=.02$, ($M_{\text{High}}=23.4 (14.57), M_{\text{Low}}=20.55 (12.5)$). The interaction between BAQ and utility was not significant, $F<1$ (see Figure 2.10).

However, orthogonal contrasts showed that the critical contrast 1 was significant, $t(185)=3.21, p<.01$. Contrast 2 was not significant, $t(185)=1.43, p=.15$, and neither was contrast 3, $t(185)=.63, p=.53$. No rest-variance was unaccounted for, $F(2,185)=1.22, p=.3$.

⁹ These results remain the same before the removal of outliers, however, the utility*CBP interaction is then marginally significant, $F(1,194)=3.13, p=.08$.

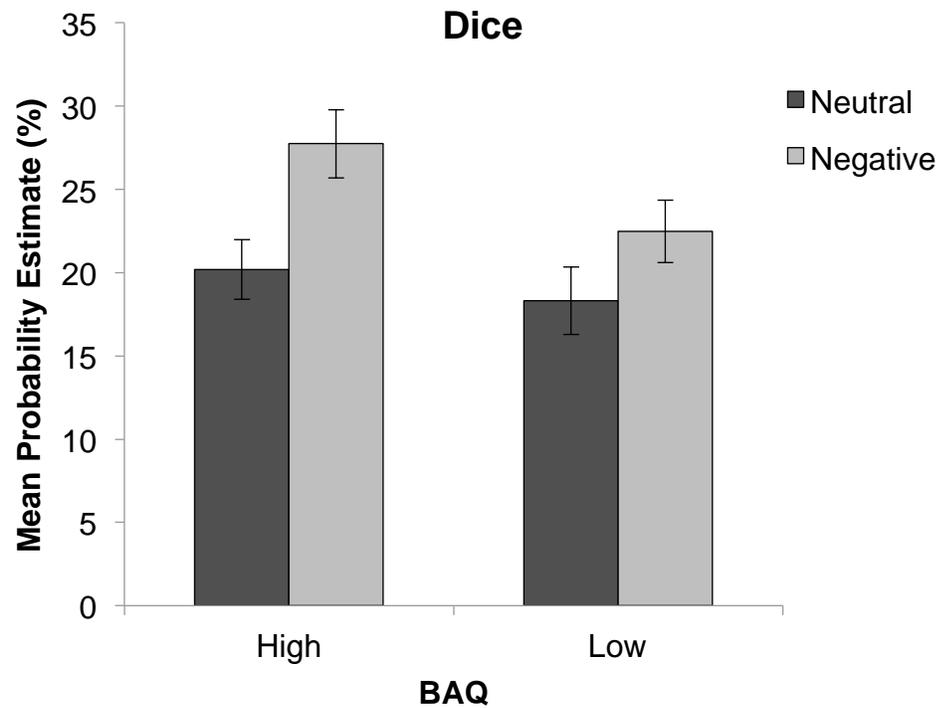


Figure 2.10. Probability estimates for participants high and low in BAQ for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean

MAIA-E (Dice). The main effect of utility was significant, $F(1,185)=7.63$, $p<.01$, $\eta^2=.04$. Participants in the negative condition estimated the probability to be higher ($M=24.55$, $SD=14.59$) compared to participants in the neutral condition ($M=19.62$, $SD=12.28$). There was also a significant main effect of MAIA-E, $F(1,185)=4.7$, $p=.03$, $\eta^2=.03$, participants high in MAIA-E rated the probability as higher ($M=23.64$, $SD=14.46$) compared to participants low in MAIA-E ($M=20.22$, $SD=12.51$). The interaction between MAIA-E and utility was not significant, $F(1,185)=2.35$, $p=.13$, $\eta^2=.01$ (see Figure 2.11).

Furthermore, the critical contrast 1 (comparing those high in IS to all other 3 conditions) was found significant, $t(185)=3.58$, $p<.001$. No other contrasts were significant, all $ps<.44$, and no rest-variance was unaccounted for.¹⁰

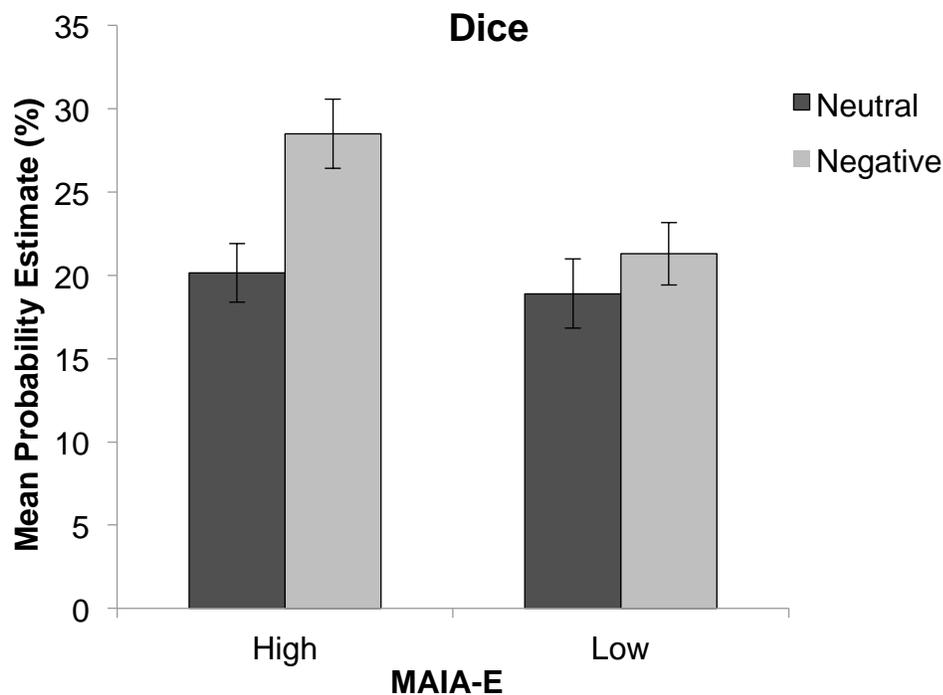


Figure 2.11. Probability estimates for participants high and low in MAIA-E for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

MAIA-BL (Dice). A significant main effect of condition was found, $F(1,186)=7.77$, $p<.01$, $\eta^2=.04$, ($M_{\text{negative}}=24.88$ (14.85), $M_{\text{neutral}}=19.62$ (12.28)), as well as a significant main effect of MAIA-BL, $F(1,186)=4.58$, $p=.04$, $\eta^2=.02$, ($M_{\text{High}}=24.17$ (15.2) vs. $M_{\text{Low}}=20.24$ (12.09)). The interaction between MAIA-BL and utility was not significant, $F<1$ (see Figure 2.12).

¹⁰ Before the removal of outliers, the MAIA-E*utility interaction was significant, $F(1,194)=4.7$, $p=.03$, $\eta^2=.02$. All other results remained the same.

Conducting orthogonal contrasts, only the critical contrast 1, comparing the high MAIA-BL negative condition to all other conditions was significant, $t(186)=3.07$, $p<.01$, all other $ps>.23$. Furthermore, an omnibus test showed that there was no rest-variance left to be accounted for, $F(2,186)=1.36$, $p=.26$.

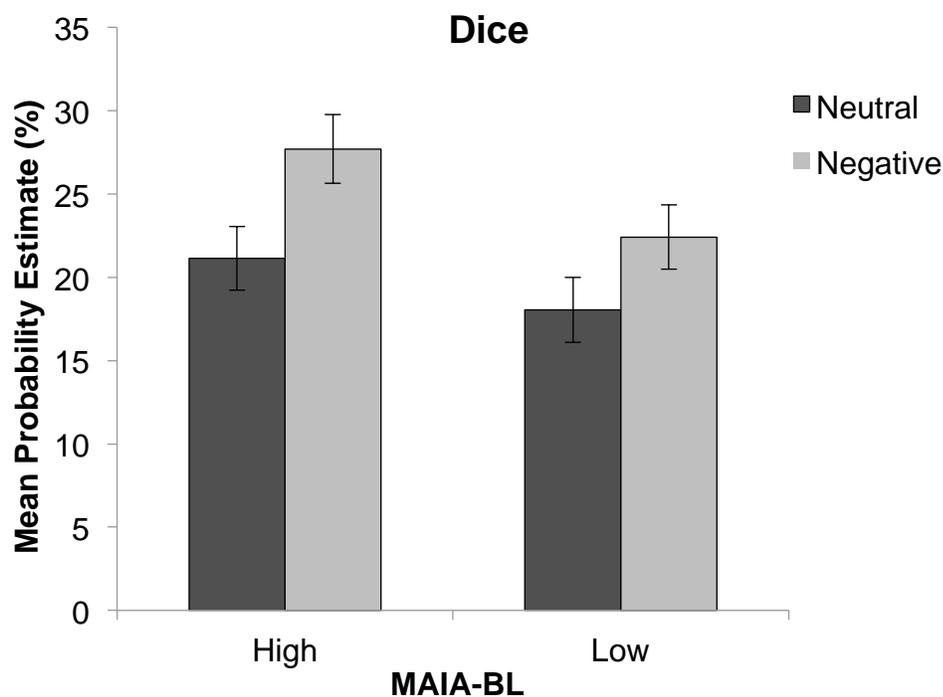


Figure 2.12. Probability estimates for participants high and low in MAIA-BL for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

Discussion

Experiment 1 provided some first evidence for a moderation by interoceptive sensibility on the impact of negative utility on probability estimates. Importantly, whilst not all scales moderated the relationship for all scenarios, in 6 out of the 8 analyses examining individual scenarios with the four scales the means were in the direction predicted by the SLH. That is, individuals high in IS, irrespectively of how IS was assessed, demonstrated a stronger negativity bias than individuals low in IS. Moreover,

when averaging across scales and IS measures, we demonstrated an overall effect in line with the hypothesis that interoceptively sensible individuals feel arousal created by aversive events more strongly, and therefore provide greater probability estimates for negative outcomes. This difference was not accounted for by differences in the perception of utility, as the manipulation checks demonstrated that those high in IS did not indicate that they experienced the negative outcome as worse than those low in IS.

Experiment 2

Background

Experiment 2 sought to replicate the results of Experiment 1 in a different sample, making several small adjustments. First, we now only included the CBP scale, as the main scale of interest measuring self-reported interoceptive accuracy. This scale is shorter than the BAQ, and measures how much individuals feel their bodily functions, rather than including items on the prediction of the onset of disease, etc.

In addition, we included a novel scenario with a different way of displaying objective probabilities. This scenario employed an adapted version of the visual display of George Louis Leclerc's "Buffon's Coin" problem (see Strick, 2007). In this problem a coin is tossed randomly at a grid of square tiles, and participants have to guess how likely it is that the coin lands entirely on a tile (rather than on a boundary). We did not employ square tiles in our experiment but instead a map of a river supplying drinking water to a large city on which a container could land (the impact area of the container constituting the "coin"). The container either contained toxic materials that would

poison the city (negative condition), or natural materials that would make the water taste slightly different (neutral condition).

Furthermore, we examined the potential effect of the order of administration of the scale. Whilst in Experiment 1 the scales were administered before the probability estimates were made, in Experiment 2 we counterbalanced the order to ensure that the moderation by IS does not depend on the order of the scale administration.

Method

Participants and Design. 744 participants from University College London (455 female, median age=22) were recruited by email announcement. Participants were entered into a lottery for 25 pounds, and completed the experiment online. The experiment employed a 2 (utility: negative/neutral) x 2 (order: CBP before/afterwards) x 3 (scenario: dice/RAF/Container) mixed design, the last factor manipulated within participants.

Materials. The USAF (here: Royal Air-Force, RAF, since the experiment was conducted in the UK) and dice scenarios were identical to Experiment 1. For the new container scenario, participants read the following in the negative condition (neutral condition in brackets):

A container is to be dropped from the air, and will land somewhere in the area depicted below, with all locations equally likely. The container contains toxic chemicals (natural, organic materials), which are fatally poisonous to humans (pose no risk to people or the environment).

Below, you see the area where the container could land. The blue lines are an underground watercourse, which supply drinking water to a large city. The red circle indicates the size of the area where toxic chemicals will be released. If this area overlaps at all with one of the water veins, the chemicals will be released into the drinking water, killing thousands of people (making it taste very slightly different with no threat to health).

By looking at the picture below, what is the chance that the container lands so that it overlaps with one of the water veins, thus poisoning the large city's drinking water and killing thousands (causing the large city's drinking water to taste very slightly different)?

The “area where the container could land” and the “red circle” shown to participants is reproduced in Figure 2.13.

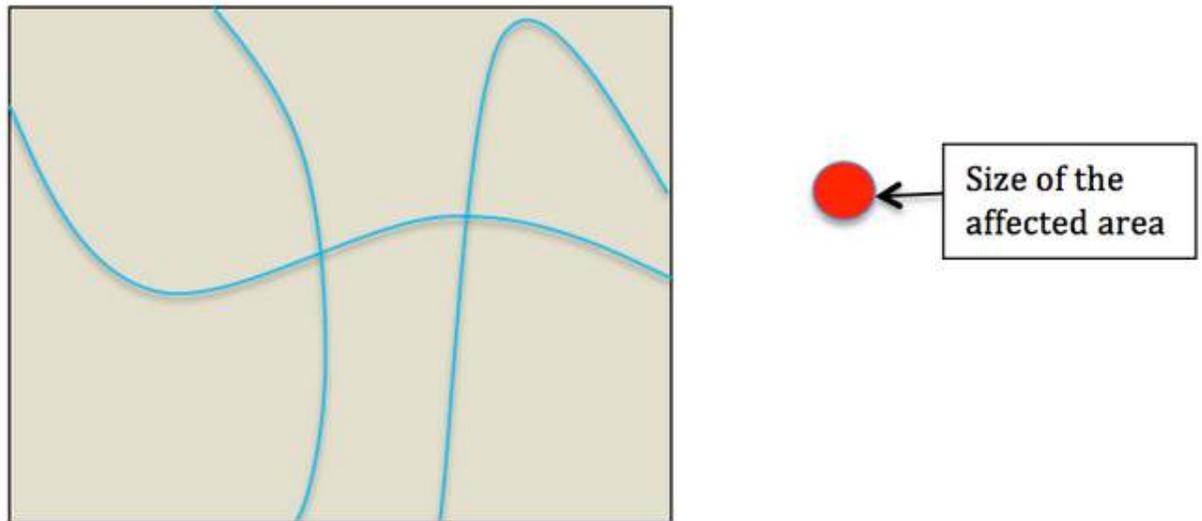


Figure 2.13. Probability display in the “Container” scenario.

Procedure. Participants either completed all three scenarios, followed by the IS scale, or the IS scale followed by all three scenarios. Afterwards, participants completed manipulation checks for utility, provided demographic information and were thanked and debriefed.

Results

The 5 items of the consciousness of body private scale (CBP, see Appendix A) were summed ($\alpha=.66$) and a median split was performed on the outcome.

Manipulation Checks.

RAF Manipulation Checks. The answer to the question “Thinking back about the RAF risk estimation task you just did: how bad would it be if debris would be dropped?” was submitted to a one-way ANOVA with utility as a between subjects factor. Those in the negative condition rated the outcome as worse ($M=5.7$, $SD=1.33$) than those in the neutral condition ($M=3.26$, $SD=1.72$), and this difference was significant, $F(1,742)=467.81$, $p<.001$, $\eta^2=.39$. Entering IS as a factor to test for possible interactions between utility and IS showed a main effect of IS: those high in IS rated the outcome on average as worse ($M=4.61$, $SD=1.9$) than those low in IS ($M=4.36$, $SD=2.01$), $F(1,740)=9.18$, $p<.01$, $\eta^2=.01$. A marginally significant interaction, $F(1,740)=3.03$, $p=.08$, $\eta^2=.004$, followed up by simple effect analyses, showed that the difference between those high and low in IS was only significant for neutral, $F(1,740)=11.38$, $p<.001$, $\eta^2=.02$, ($M_{\text{high}}=3.52$, $SD=1.79$ vs. $M_{\text{low}}=2.99$, $SD=1.62$) but not for negative outcomes, $F < 1$. The results of these simple effects are important, as any impact of IS on the relationship between negative utility and probability could otherwise be explained in terms of different ratings of utility – if those high in IS were simply to perceive the negative events as worse. Importantly, however, those high in IS only rated the *neutral*, and not the negative items as worse.

Dice Manipulation Checks. The manipulation of utility in the dice scenario was successful, with those in the negative condition rating the outcome as worse ($M=3.64$, $SD=1.75$) compared to those in the neutral condition ($m=1.61$, $SD=1.28$), $F(1,742)=325.21$, $p<.001$, $\eta^2=.31$. Entering IS as a factor showed that neither the main effect of IS, $F < 1$, nor the interaction, $F(1,742)=2.27$, $p=.13$, $\eta^2=.003$, were significant.

Container Manipulation Checks. The manipulation of utility was also successful for the container scenario, $F(1,742)=672.88, p<.001, \eta^2=.48$. Those in the negative condition rated the outcome as worse ($M=6.5, SD=1.1$) compared to those in the neutral condition ($M=3.62, SD=1.84$). Entering IS as a factor revealed a marginally significant main effect of IS, with those high in IS giving slightly higher utility ratings ($M=5.1, SD=2.02$) compared to those low in IS ($M=5.03, SD=2.15$), $F(1,740)=2.95, p=.09, \eta^2=.004$. A marginally significant interaction, $F(1,742)=3.38, p=.07, \eta^2=.01$, followed by simple effect analyses revealed that the differences between high and low IS were only significant for the neutral, $F(1,740)=6.32, p<.01, \eta^2=.01$, ($M_{\text{high}}=3.81, SD=1.8$ vs. $M_{\text{low}}=3.42, SD=1.86$), but not for the negative condition, $F<1$. As described above, the finding that individuals high in IS perceive neutral outcomes as worse would not be able to explain why they would assign greater probability estimates to negative events and therefore does not need to be taken into consideration at this point.

Probability Estimates.

Mixed ANOVA. Running an ANOVA on the full 2 (utility: negative/neutral) x 2 (order: IS first/second) x 2 (IS: high/low) x 3 (scenario: RAF/Dice/Container) design (repeated measures on the last factor) revealed a significant scenario*order*utility* IS four way interaction, $F(2,1416)=4.23, p=.02, \eta^2=.001$ (Greenhouse Geisser corrections applied), a marginally significant scenario*order*utility interaction, $F(2,1416)=2.43, p=.09, \eta^2=.003$, and a significant scenario*utility interaction, $F(2,1416)=4.42, p=.01, \eta^2=.01$. No other main effects or interactions were significant, all $ps>.14$.

In order to explore and further understand the significant four way interaction between order, utility, IS and scenario, three ANOVAs were run on the different

scenarios including utility, IS and order as a factor, as well as conducting orthogonal contrasts as before (see Table 2.5).

Table 2.5

Orthogonal contrasts tested in Experiment 2.

	High IS	Low IS	High IS	Low IS
	Negative	Negative	Neutral	Neutral
Contrast 1	3	-1	-1	-1
Contrast 2	0	2	-1	-1
Contrast 3	0	0	1	-1

RAF Probability. Running an ANOVA on RAF with utility, IS and order as factors, a main effect of order was found, $F(1,691)=4.41$, $p=.04$, $\eta^2=.01$. Participants provided greater estimates when answering CBP before ($M=48.96$, $SD=9.44$) compared to afterwards ($M=47.56$, $SD=9.52$). Furthermore, there was a significant main effect of IS, whereby those high in IS provided greater estimates ($M=48.98$, $SD=9.65$) than those low in IS ($M=47.57$, $SD=9.33$), $F(1,691)=3.96$, $p=.05$, $\eta^2=.01$. The utility*IS interaction did not reach significance, $F(1,691)=2.55$, $p=.11$, $\eta^2=.004$. No other main effects or interactions were found significant (all F s >1).

A set of planned orthogonal contrasts (see Table 2.5) showed a non-significant trend that those in the high IS negative condition differed from all other conditions combined, $t(695)=1.6$, $p=.11$ (see Figure 2.15). Unexpectedly, those low in IS in the negative condition differed significantly from the neutral conditions combined, $t(691)=2.19$, $p=.03$, however, individuals low in IS answering a negative scenario provided *lower* estimates than the two neutral conditions combined, see Figure 2.14.

Last, contrast 3 was not significant, $t(695)=.15, p=.88$. In sum, the RAF scenario does not support our hypothesis that only individuals high in IS show a negativity bias. As such, the results from this scenario contradict the findings of the same scenario in Experiment 1.

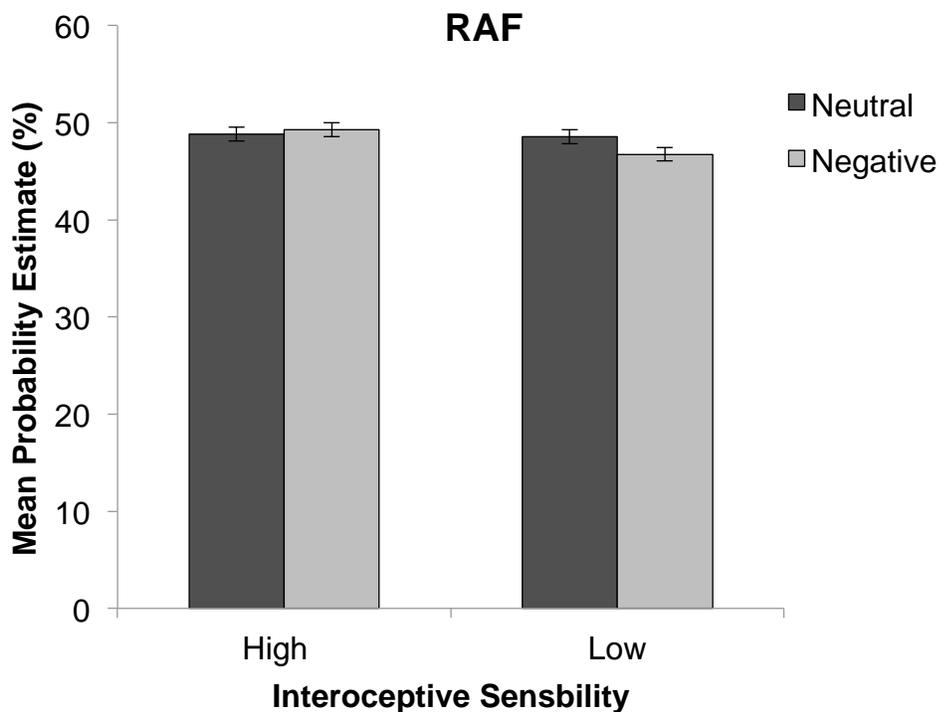


Figure 2.14. Probability estimate means for the RAF scenario for participants high and low in IS for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

Dice Probability. Running an ANOVA on the dice scenario with utility, IS, and order as factors revealed a significant main effect of utility, $F(1,700)=9.35, p<.01, \eta^2=.01$ ($M_{\text{negative}}=20.63$ (13.26) vs. $M_{\text{neutral}}=17.77$ (11.93)), as well as a significant three-way interaction between utility, order and IS, $F(1,700)=7.96, p<.01, \eta^2=.01$.

Next, two ANOVAs were run on each level of order. When IS was assessed after the scenarios (unlike in Experiment 1), there were no main effects of utility,

$F(1,346)=2.63, p=.11, \eta^2=.01$, or IS, $F<1$. Simple effect analyses revealed that participants low in IS showed a negativity bias, $F(1,346)=6.61, p=.01, \eta^2=.02$, whilst there was no difference between the negative and neutral condition for participants high in IS when IS was assessed after the scenarios, $F<1$ (see Figure 2.15). Orthogonal contrasts showed that only contrast 2 was significant, $t(346)=2.32, p=.02$, demonstrating that those low in IS in the negative condition differed significantly from the two neutral conditions. No other contrasts were significant, all $ps>.23$ ¹¹.

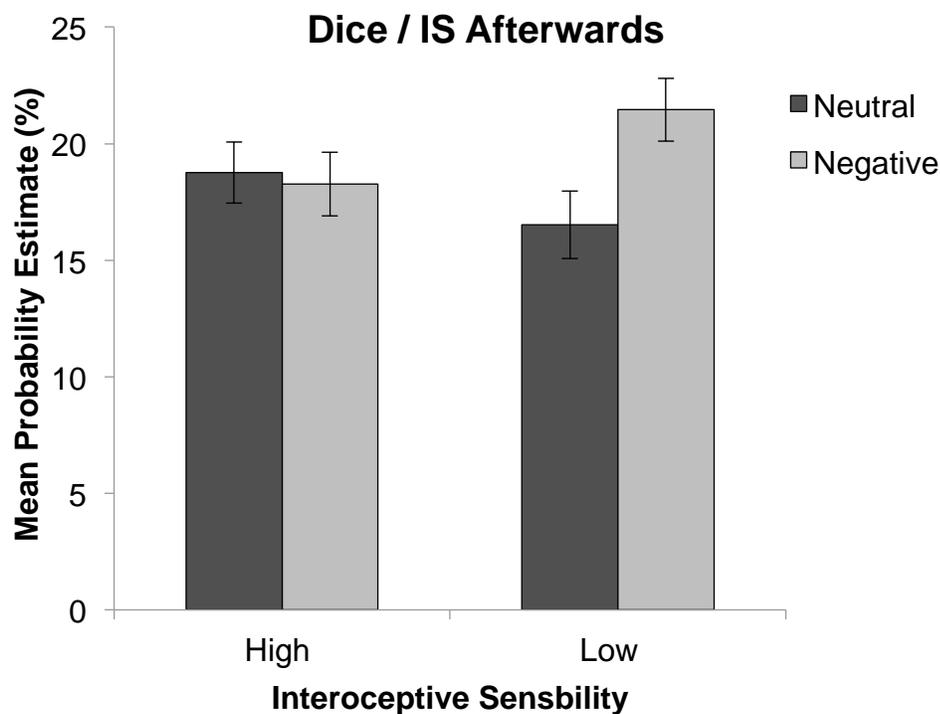


Figure 2.15. Probability estimate means for the dice scenario for participants high and low in IS (assessed after the scenarios) for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

When IS was assessed before the scenario there was a significant main effect of utility, $F(1,354)=7.43, p<.01, \eta^2=.02$. Participants answering negative scenarios

¹¹ Before removing outliers, contrast 2 was non-significant, $t(368)=1.52, p=.12$.

provided greater probability estimates ($M=21.21$, $SD=13.35$) than participants answering neutral scenarios ($M=17.85$, $SD=11.36$). Furthermore, there was a significant IS*utility interaction, $F(1,354)=4.01$, $p<.05$, $\eta^2=.01$. Simple effect analyses revealed that those high in IS provided greater estimates for negative compared to neutral events, $F(1,354)=10.51$, $p<.001$, $\eta^2=.03$, whilst this was not the case for those low in IS, $F<1$ (see Figure 2.16). Orthogonal contrasts confirmed these results, showing that only the critical contrast 1 was significant, $t(354)=3.39$, $p<.001$, all other contrasts were non-significant, all $ps>.35$, and there was no rest-variance unaccounted for, $F<1$. Thus, this result constitutes a direct replication of the results for the dice scenario of Experiment 1. However, when the CBP scale was administered after probability estimates, the results did not support our hypotheses, with individuals low in IS demonstrating a greater negativity bias than individuals high in IS.

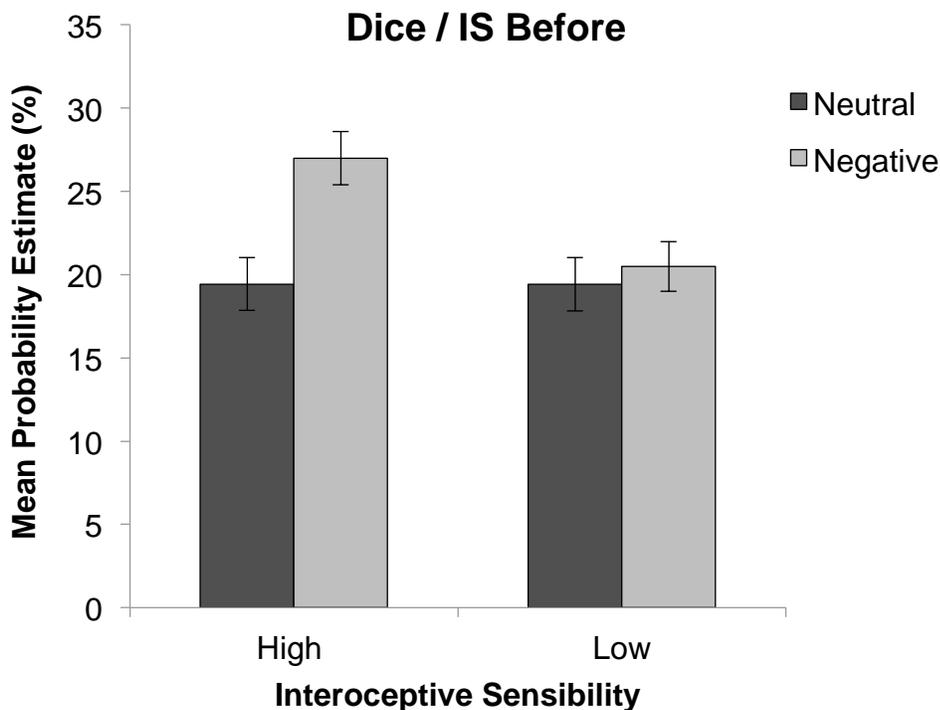


Figure 2.16. Probability estimates for the dice scenario for participants high and low in IS (assessed before the scenarios) for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

Container Probability. Running an ANOVA on the container scenario including utility, IS and order as factors revealed a significant main effect of utility, with those in the negative condition giving higher probability estimates ($M=34.28$, $SD=18.92$) compared to those in the neutral condition ($M=29.76$, $SD=18.53$), $F(1,712)=9.26$, $p<.01$, $\eta^2=.01$. A significant main effect of order was present, in that participants gave higher ratings of probability when answering the IS scale before ($M=33.68$, $SD=19.37$) as compared to when answering the IS scale the after the scenarios ($M=32.3$, $SD=20.57$),

$F(1,712)=4.87, p=.03, \eta^2=.01$ ¹². No other main effects or interactions were found significant, all $ps>.12$.

Running a set of orthogonal contrasts on the container scenario revealed that there was an unpredicted difference between those low in IS answering a negative scenario compared to those high or low answering a neutral scenario combined, $t(712)=3.4, p<.001$, no other contrasts were significant, all $ps>.42$, no rest-variance was unaccounted for, $F<1$ (see Figure 2.17). Thus, these results are not in line with our hypothesis, where we had predicted that only those high in IS would show a negativity bias.

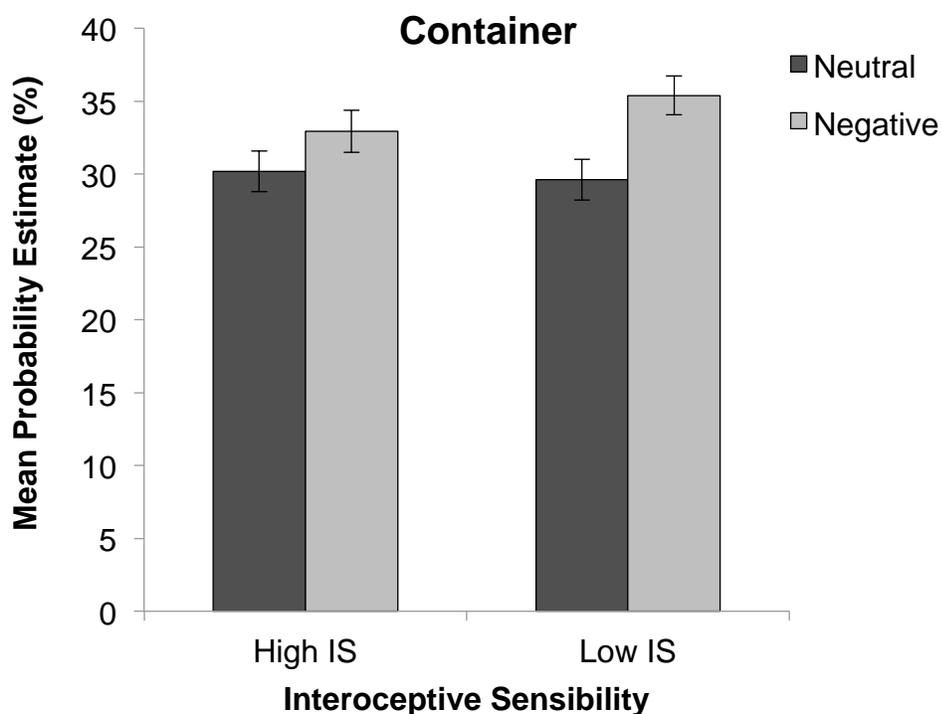


Figure 2.17. Probability estimate means for the container scenario for participants high and low in IS (averaged across order) for neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

¹² Before removing outliers, the main effect of order is marginally significant, $F(1,736)=3.15, p=.08, \eta^2=.004$

Discussion

The results from Experiment 2 were less conclusive than the results from Experiment 1. Whilst in the aggregated analyses only the predicted contrast was significant, there was also a marginally significant trend for individuals low in IS to provide greater estimates than the neutral conditions combined (see contrast 2, Table 2.5). Furthermore, we did not replicate the results in the RAF scenario, and an impact of the order of scale administration was difficult to interpret. The dice scenario replicated the findings of Experiment 1. However, examining order effects, it became evident that the replication only held when IS was assessed before, but not after the scenarios – where a negativity effect was, surprisingly, only found for participants *low* in IS. For the container scenario, a novel scenario with a new method of displaying probabilities, participants overall demonstrated a negativity bias, but no moderation by IS was observed.

There are several noteworthy results: First, if IS were to assess the accuracy with which participants experience interoceptive processes, the order with which the scale is administered should not be of relevance. However, the moderation of IS on the impact of negative utility on probability estimates was strongest when IS was assessed beforehand. One possibility is that completing the scales leads to those high in IS to attend more to their internal states, and those low in IS to attend less to their internal states.

Moreover, whilst we do find a reliable moderation by IS for the dice scenario, the moderation for the RAF scenario was trending, but not significant, and there was no moderation in the container scenario. One variable that could potentially account for

these differences is self-relevance. Arguably, the dice scenario is the most self-relevant scenario, describing a first-person scenario. The RAF scenario has some personal relevance, as participants are being directly asked by the RAF to provide their estimate. Least self-relevant is the container scenario. However, in an experiment not reported in the current thesis, de Molière and Harris (2014) manipulated self- and other-relevance and whilst in one out of two scenarios the moderation by IS occurs only for self- but not for other-relevant scenarios, the overall evidence did not support this hypothesis very strongly.

In sum, both Experiment 1, and Experiment 2 proved to be noisy. Whilst in some cases we demonstrated the hypothesized moderation after removing outliers, in others we demonstrated it only when not removing outliers. We also lacked clarity as to the specific scenarios in which we appear to observe the moderation, as these were not consistent between experiments. Part of this noise might be due to self-reported interoceptive ability. Silvia and Gendolla (2001) heavily criticized the self-report of interoceptive ability, and in their literature review on the evidence for the “perceptual accuracy” hypothesis (that is, individuals that *report* that they feel internal bodily sensations more accurately will also feel them more accurately objectively), they conclude that the evidence for this hypothesis is thin. Furthermore, Garfinkel et al. (2014) notice that interoceptive sensibility might in fact demonstrate “individual differences in subjective thresholds” (Garfinkel et al. 2014, p. 5). Therefore, Experiment 3 aims to objectively measure interoceptive ability (“interoceptive awareness”), as well as to objectively measure arousal, in order to provide a more direct assessment of the

SLH and the impact of interoception on the relationship between negative utility and probability estimates.

Experiment 3

Background

Experiment 3 aimed to test the full moderated mediation described at the start of this chapter (see Figure 2.1), by directly measuring the mediation by arousal, and assessing the hypothesized moderation of this mediation by interoceptive accuracy.

In order to directly assess arousal, we measured electrodermal activity (EDA). EDA is the “change in electrical properties of the skin in response to the secretion of sweat by the eccrine sweat glands” (Lajante, Droulers, Amarantini & Dondaine, 2012) and is closely related to autonomic nervous system activities. The most commonly used method in judgment and decision-making research to assess EDA is to assess skin conductance, changes in eccrine sweating (Figner & Murphy, 2011). Changes in EDA are closely related to arousal and the level of activation following emotional stimuli (Andreassi, 2007; Boucsein, 2012; Dawson, Schell, & Courtney, 2011; Dawson, Schell, & Filion, 2007; Sequeira et al. 2009).

Skin conductance can be separated into the tonic (slow) and phasic (fast) response, and should be analysed in these separate components (Boucsein et al. 2012). The tonic conductance is the absolute level of resistance, which is independent of stimulus presentation, ranges between 2 and 20 microSiemens (μS) (Dawson, Schell, & Filion, 2007), and is a measure of general arousal. The phasic conductance is defined as a steeper change in SCR 1-3s after the onset of a stimulus (Dawson et al. 2007) (see Figure 2.18).

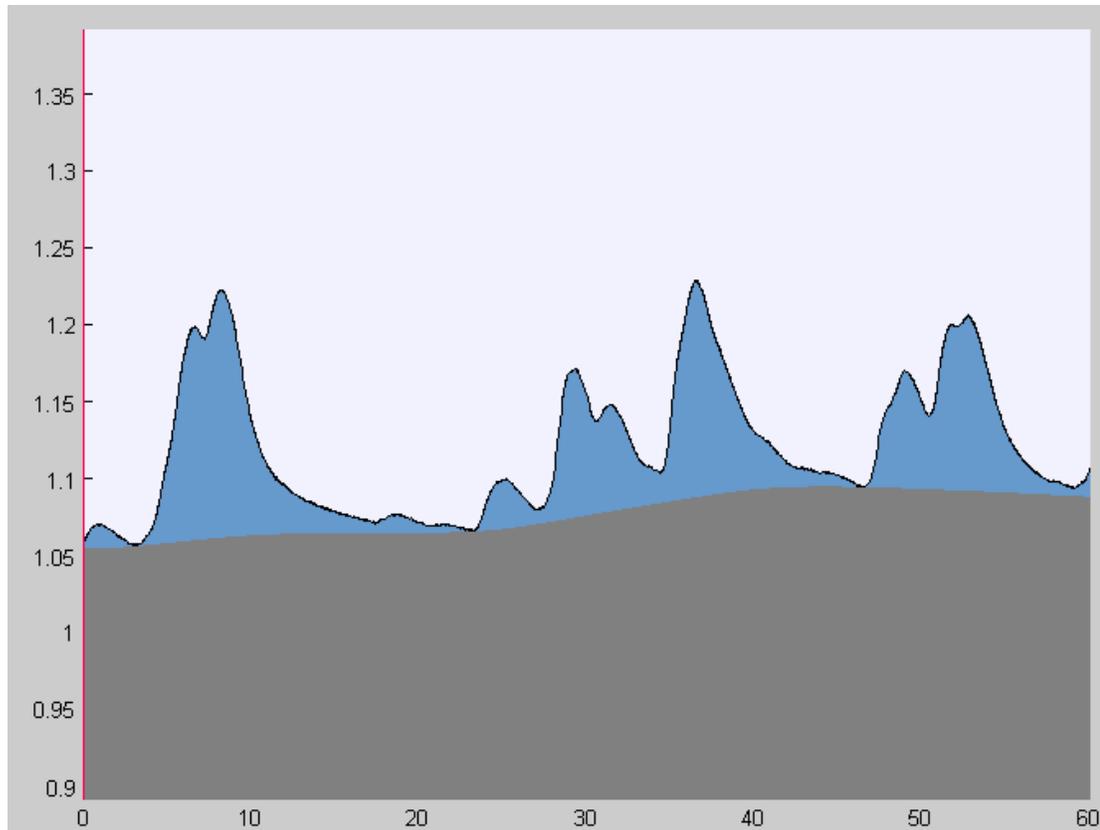


Figure 2.18. Illustration of phasic (blue) and the tonic (grey) response based on stimulus onset (red line).

As in the current experiment participants read passages of text, rather than being presented with one emotional stimulus, and since reading speed varies, it was not possible to match phasic responses to the emotional stimuli (here: words). The main variable of interest was therefore the tonic driver, which should increase being subjected to emotional material (e.g. Sundar & Kalyanaraman, 2004). Hence, the current experiment assesses tonic arousal levels during the presentation of negative or neutral scenarios. Following Vosgerau (2010), negative scenarios should lead to an increase in arousal (and as such in the tonic driver), which in turn should mediate the relationship between utility and probability estimates.

Furthermore, we assessed interoceptive accuracy, one dimension of interoception that does not rely on self-reports but is assessed objectively. Two versions of this task are commonly used: A heartbeat counting (Schandry et al. 1981), and a heartbeat tracking (Whitehead, Drescher, Heiman & Blackwell, 1977) version. Heartbeat tracking involves participants hearing a sequence of tones (or flashes of light), which are either synchronous or dyssynchronous to their own heartbeat. Participants are asked to indicate whether the tone is in or out of sync with their own heartbeat. However, some studies found that the heartbeat tracking task does not resolve in sufficient variance to be considered sensitive enough for an individual difference measure (see Dunn et al. 2010, supplementary material). A floor-effect is often present, where participants that are classified as good perceivers score rather low (40-50% correct) (e.g. Eichler & Katkin, 1994).

In the heartbeat counting task, participants are asked to simply count their own heartbeats over a series of time intervals, which vary in length. The difference between objectively measured heartbeats and counted heartbeats serves as an accuracy score. This version of the task is commonly used and well validated (Ainley & Tsakiris, 2013), and was employed in Experiment 3.

In Experiment 3 participants were presented with three scenarios: the dice and container scenarios already employed in Experiments 1 and 2, and a new scenario. This new scenario was of exploratory nature and designed to be low in decision-control. We were interested in including a scenario low in decision-control as the SLH and an ALF account are not mutually exclusive: Arousal misattribution, and subsequently a moderation by interoceptive accuracy could occur only for scenarios high, but not low,

in decision-control. That is, scenarios low in decision-control could evoke less arousal – similar to the experiments by Sigrist and Sütterlin (2014), where man-made disasters, arguably associated with greater control, were rated as more negatively than nature-made disasters, arguably associated with lower control. If scenarios low in decision-control are rated not as negatively as scenarios high in decision-control, it might be possible that they also evoke less arousal, and are less likely to be moderated by interoceptive accuracy. As the decision to include this scenario was made after the start of the experiment, this scenario was only included from participant number 36 onwards.

Furthermore, as a first exploration into the role of social power, we included a questionnaire at the end of the experiment measuring one's personal sense of power (Anderson, John, & Keltner, 2012), previously hypothesized to moderate the impact of negative utility due to the greater interoceptive awareness of the powerful (see Figure 1.7).

In the current experiment, participants first completed the measure of interoceptive accuracy. As interoceptive accuracy has been associated with affective states (Ehlers & Breuer, 1992), we measured interoceptive accuracy before the completion of the scenarios, to ensure orthogonality of the independent variables utility and interoceptive accuracy.

Method

Participants and Design. 127 Participants were recruited (101 female, median age=19 years¹³). Participants either received a payment of £3 or 0.5 course credits as a reimbursement for participating in this laboratory experiment. Participants were

¹³ One participant typed “2” as the answer to how old they are. Since all participants were above 18 years old, we assumed that this was a typing mistake and did not include them when calculating the median of age.

randomly assigned to a 2 (utility: neutral/negative) x 3 (scenario: dice/container/roulette) mixed design, the last factor manipulated within participants.

Materials and Procedure.

Physiological Sensors. Upon arrival in the lab, participants were informed that we are interested in decision-making and physiology. Participants were first seated in front of the computer, and then all physiological equipment was attached. For the ECG (to measure heart beats objectively), two disposable Ag-AgCl electrodes (11mm diameter, 7% chloride salt wet liquid gel) were attached to the dorsal forearm of the non-dominant hand and the opposite ankle. Moreover, an electrodermal activity (EDA) transducer¹⁴ was attached to the index and second finger of the non-dominant hand (see Figure 2.19). The transducer consists of two Ag-AgCl electrodes, each with a 6mm diameter, and was attached with a Velcro strap after isotonic electrode paste¹⁵ (0.5% saline in neutral base) was filled into their 1.6mm deep cavities and applied onto the finger tips. The electrodes were connected to a Biopac Student Lab (MP36) system and all cables were subsequently secured to the table with tape to minimize movement and resulting artefacts. In order for the electrodes to establish contact with the skin, participants were instructed to relax for 5 minutes. Afterwards a baseline measure of EDA was obtained for a time period of 1 minute.

¹⁴ EDA Finger Transducer, BSL-SS3LA, Biopac Systems, Inc.

¹⁵ GEL101, Biopac Systems, Inc.



Figure 2.19. Display of the finger transducer.

Likelihood Judgments. Participants completed the dice and container scenarios from previous Experiments 1 and 2, presented with Eprime (version 2.0). The order with which the two scenarios were presented was counterbalanced. Furthermore, between each of the scenarios was a 1-minute break during which participants were asked to relax. This was done as to allow the EDA responses to return to baseline.

Besides the dice and container scenarios, a third scenario was added (starting from participant number 36 onwards). This scenario, since it was not pre-tested, was of exploratory nature and was always presented after the first two scenarios. The scenario was designed to be of low control and from a third person perspective. In this third scenario (“Roulette”) participants read the following text in the negative condition (see Figure 2.20 for the graphics provided with the scenario):

David has been forced to play a special version of Russian Roulette. This special version of Russian Roulette consists of two steps: A box contains 13 bullets which look and feel identical. However, 8 of these bullets are real, fatal bullets, and 5 bullets are fake, and will cause no harm whatsoever. David must now pick 2 bullets from this box, to put into 2 chambers of a 5-chamber revolver. He then spins the cylinder, holds the gun to his head and pulls the trigger. If a real bullet is now in the active chamber, David will shoot himself in the head. What is the likelihood that David shoots himself in the head?

In the neutral condition participants read instead:

David is playing a special, completely harmless version of Russian Roulette. This harmless version of Russian Roulette consists of two steps: A box contains 13 toy bullets which look and feel identical. However, 8 of these are operational bullets that release a small puff of air, and 5 bullets are non-operational, and will not release any air puff. David must now pick 2 bullets from this box, to put into 2 chambers of a 5 chamber play revolver. He then spins the cylinder, and pulls the trigger. If an operational bullet is now in the active chamber, the gun will release a small puff of air. What is the likelihood that the gun releases a small puff of air?

This convoluted set-up of describing the probabilities was necessary as simply asking participants to estimate the chance of being shot with 2 out of 5 bullets in the chamber would have been too easy to calculate. However, it also would not have been possible to ask participants to estimate the chance of David shooting himself at least once out of three times he pulls the trigger (which would have resembled the “dice” scenario), as it could have been possible that David is dead the first time he pulls the trigger.

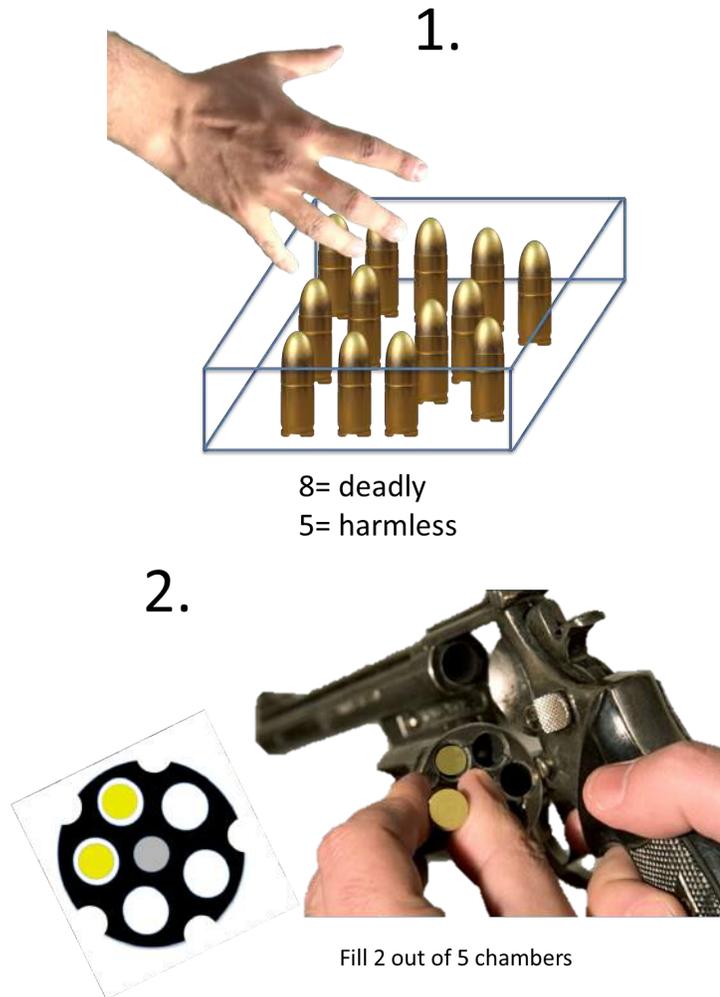


Figure 2.20. Visual display of the negative condition used in the “Roulette” scenario.

Heartbeat Perception Task. Afterwards, interoceptive accuracy was measured with the heartbeat tracking task by Schandry (1981). Participants were asked to count their heartbeats, without measuring their pulse, over the time of varying time intervals. Participants were prompted by an auditory cue to start and stop the counting task (presented via Eprime 2.0). However, since this task has been shown to be confounded with the ability to count time and approximate the number of heartbeats (Brener, Knapp & Ring, 1995), participants additionally engaged in one block of time estimation

(between two blocks of heart beat counting). Thus, after an initial practice trial, participants first completed three heartbeat trials (1x 25s, 1x35s, 1x45s), then completed three time estimation trials (1x 23s, 1x56s, 1x 40s), and again one set of heartbeat trials (1x 25s, 1x35s, 1x45s) (Ehlers & Breuer, 1992). In all blocks the order of time intervals was randomized. As the EDA finger transducer could have potentially allowed participants to feel their pulse on their fingertips, the transducer was removed before the heartbeat counting task started.

Questionnaire. After the heartbeat counting task, participants were detached from the physiological equipment and completed an additional questionnaire (presented via the software Qualtrics). Participants completed the CBP scale (Miller et al. 1981, see Appendix A), and the Sense of Power scale (Anderson et al. 2012, see Appendix E). Next, participants reported their current affective state with the 20 item positive and negative affect scale (PANAS, Watson et al. 1988, see Appendix F), the short (6 item long) form of the state and trait anxiety scale (Marteau & Bekker, 1992, see Appendix G), and how many days a week they engage in physical activity, as these variables have been shown to be associated with performance in the heart beat counting task (Ehlers & Breuer, 1992). Following Dunn et al. (2010), we also asked participants whether they used time to approximate how many heartbeats they had.

Afterwards participants were again presented with the three scenarios and completed three manipulation checks for utility (“how bad would it be if the container would overlap with a water vein”, “how bad would it be if at least 2 sixes are rolled?”, “how bad would it be if an operational bullet is now in the active chamber”) on 9-point scales (1, not bad at all, to 9, very bad). Last, participants were asked to indicate whether

there were any parts of the study they did not understand, and asked to guess the aim of the experiment.

Results

Of the 129 participants 8 participants were excluded from any analysis. Either their English proficiency was too poor to understand the instructions, or they misunderstood the instructions (for example did not move the slider for the likelihood estimates, or did not count their heartbeats).

Data Preparation.

Interoceptive Awareness. In order to obtain an accuracy score for the heartbeat counting task, the following formula was applied (see Schandry et al. 1981):

$$\text{Heartbeat Perception Score} = \frac{1}{6} \sum_{i=1}^6 \left(1 - \frac{|na_i - nr_i|}{na_i} \right)$$

i = number of time interval,
 na = number of actual heartbeats,
 nr =number of reported heartbeats.

The score reflects the accuracy as a percentage, averaged across the 6 trials. Thus, the scores range from 0 (0% accuracy) to 1 (100% accuracy). The mean score was .53 (SD=.23) (see Figure 2.21 for the frequency distribution of accuracy scores), meaning that on average, participants were 53% accurate in estimating their heartbeat. Participants that indicated that they used time to approximate how many heart-beats they had did not perform differently from participants who indicated that they had not used time, $F < 1$, and subsequently all participants were included in the analysis.

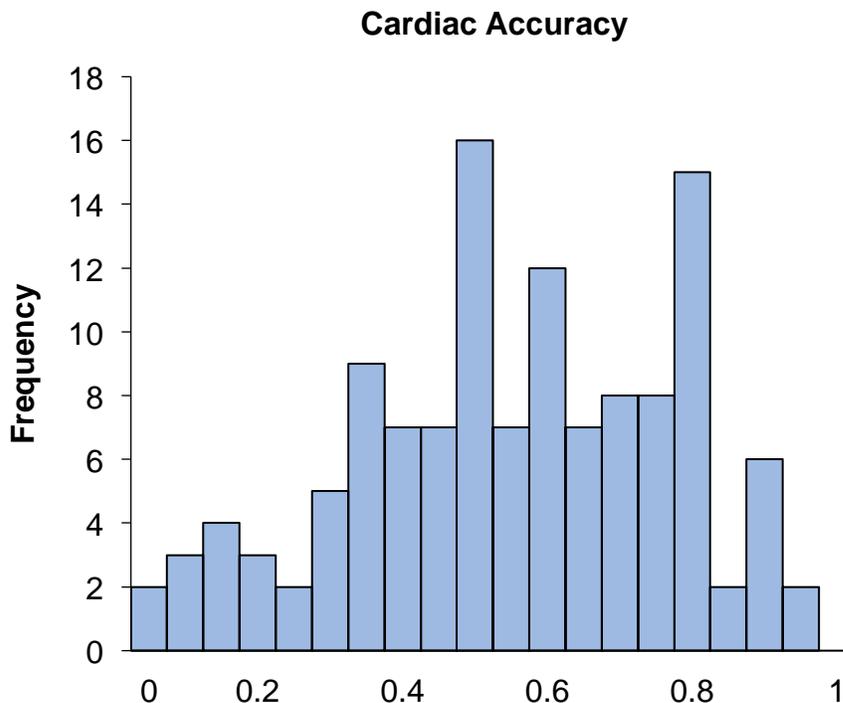


Figure 2.21. Frequency distribution of accuracy scores in the heartbeat counting task. The higher the score, the more accurate the counting, with 1 representing perfect accuracy.

Next, a median split was performed on the accuracy score, in order to obtain two groups, one high and one low in interoceptive accuracy (see Dunn et al. 2010).

Electrodermal Activity. EDA features were extracted using the MATLAB toolbox Ledalab (v.3.4.4) by means of continuous decomposition analysis (Benedek & Kaernach, 2010). This analysis allows separation between the tonic and phasic activity, whilst avoiding superimposing SCRs (common in the standard through-to-peak analysis).

EDA responses were analysed during the completion of the dice and container scenario¹⁶. Specifically, we extracted the mean tonic activity. Individuals differ in their baseline of tonic conductivity, and in order to infer the *relative change* of tonic activity following the stimulus presentation, we created a difference score by subtracting the tonic baseline arousal measures from the tonic arousal experienced during the completion of the dice and container scenario, creating one score for each of the two scenarios.

Furthermore, the data of 8 participants were removed from any EDA analyses. Either the participants were “non-responders” (did not show any response to the request to cough) or had too many artefacts in their data to allow for meaningful interpretation.

Data Analysis.

Manipulation Checks. For the dice scenario the manipulation of utility was successful, with participants in the negative condition rating the outcome as worse ($M=4.84$, $SD=2.28$) than participants in the neutral condition ($M=1.64$, $SD=1.31$), $t(100.26)=9.57$, $p<.001$ (df adjusted due to inequality of variances). The manipulation of utility was equally successful for the container scenario, where participants indicated the negative outcome to be worse ($M=8.13$, $SD=1.55$) than the neutral outcome ($M=4.54$, $SD=2.49$), $t(94.17)=9.48$, $p<.001$ (df adjusted). Furthermore, there was a significant effect of utility on the manipulation check of the roulette scenario, with participants again indicating that the negative outcome would be worse ($M=8.12$, $SD=1.36$) than the neutral outcome ($M=4.12$, $SD=2.84$), $t(51.58)=8.08$, $p<.001$ (df adjusted). Thus, the manipulation of utility was successful across all scenarios.

¹⁶ Since there was not sufficient physiological data available for the roulette scenario, it was decided to omit these responses and focus on the main dependent variables instead.

ANOVAS. First, we assessed interoceptive awareness as a moderator of the impact of negative utility on probability estimates for the three scenarios. Some people did not move the slider for some of the scenarios, and were therefore not included in these analyses. Furthermore, we also report the results after entering BMI, positive and negative affect (as measured by the PANAS), anxiety (as measured by the STAI), level of physical activity, gender, as well as the error score of time estimation task as covariates. All these variables have been shown to confound interoceptive accuracy as assessed with the heartbeat counting task.

Container. Conducting a 2 (utility: negative/neutral) x 2 (Interoceptive Accuracy: high/low) between subjects ANOVA on the container scenario, no main effect of utility, $F(1,109)=1.69$, $p=.20$, $\eta^2=.02$, or of interoceptive accuracy was found, $F<1$, and no interaction between the two was present $F<1$. These results remained the same when including the above mentioned covariates.

Dice. There were no main effects of utility, $F(1,103)=2.04$, $p=.16$, or of interoceptive accuracy, $F(1,103)=1.02$, $p=.31$. However, a significant utility*interoceptive accuracy interaction was found, $F(1,103)=5.48$, $p=.02$, $\eta^2=.05$. Individuals low in interoceptive accuracy provided greater estimates for negative (M=24.78, SD=10.67) compared to neutral events (M=16.04, SD=9.3), $F(1,107)=7.17$, $p<.01$, $\eta^2=.07$, whilst this was not the case for individuals high in interoceptive accuracy, $F<1$ ¹⁷. These results remain the same when including the aforementioned covariates.

¹⁷ Before removing outliers, the main effect of utility was marginally significant for the dice scenario, $F(1,109)=3.45$, $p=.07$, $\eta^2=.03$. Participants in the negative condition gave higher estimates (M=26.83, SD=17.46) compared to participants in the neutral condition (M=21.27, SD=14.42).

Roulette. Neither a main effect of utility, $F < 1$, nor of interoceptive accuracy, $F(1,79) = 1.57$, $p = .22$, $\eta^2 = .02$, nor an interaction between the two was present, $F < 1$. These effects remained non-significant when including the covariates.

Testing the Moderated Mediation Model. Next, we examined the potential mediation of arousal, as assessed by the change of tonic activity, on the relationship between utility and probability estimates, as well as the moderation of the mediation as displayed in Figure 2.21. Note that we only conducted these analyses for the container and dice scenario, as the roulette scenario lacked a proficient number of participants allowing any meaningful interpretation of the results.

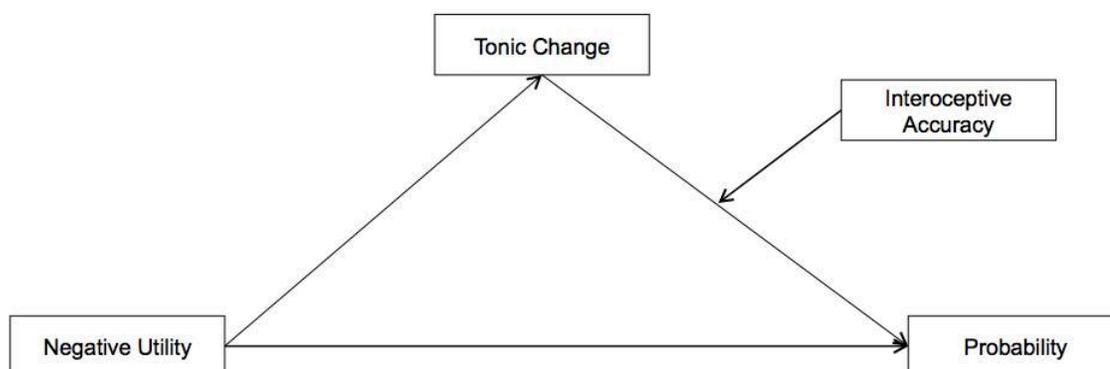


Figure 2.22. Moderated mediation model assessed in the current experiment.

Container. In the first instance, we tested for a mediation of arousal on the impact of negative utility and probability estimates. 109 Participants had physiological data available for the Container scenario and were therefore considered for analysis.

The relationship between utility and probability estimates was significant, $t(103) = 2.05$, $p = .04$, $\beta = 3.79$ ¹⁸. Next, we tested the impact of utility on the *change* in tonic arousal depending on utility. The relationship between utility and the change in the

¹⁸ Before the removal of outliers, this relationship was not significant, $t(105) = 1.49$, $p = .14$. Note also that this relationship was not significant when conducting the ANOVA including interoceptive awareness above. This is likely due to the fact that the ANOVA included individuals that did not have physiological data available.

tonic driver was significant in the predicted direction, in that individuals in the negative condition experienced a greater change in tonic arousal levels, $t(97)=2.78$, $p<.01$, $\beta=.21$ ¹⁹.

Next, we determined the indirect effects of utility on probability estimates through arousal by means of a bootstrapping method (2000 samples, model 4, PROCESS, see Hayes, 2013). This mediation was non-significant, as 0 was included in the range of the confidence intervals for the indirect effect of tonic change [-.37; 1.37].

Given the non-significant mediation, and the non-significant moderation by interoceptive accuracy, we therefore did not assess the full mediated moderation model.

Dice Scenario. 104 participants had intact physiological data and were considered for the analyses. Again, we first tested the mediation by the change in tonic activity. The impact of utility on probability estimates was not significant, $t(94)=1.59$, $p=.12$, $\beta=1.96$ ²⁰. Given the non-significance of the impact of utility on probability estimates, we did not conduct any further analyses on the dice scenario.

Further Analyses.

Interoceptive Sensibility. We aimed to replicate the analyses conducted in previous experiments, assessing the role of interoceptive sensibility, as measured with the CBP scale at the end of the experiment. Since a 2 (scenario: dice/container) x 2 (utility: negative/neutral) x 2 (IS: high/low) repeated measures ANOVA did not show any interactions including scenario, we z-scored and averaged across the two scenarios.

¹⁹ Before outliers were removed, the impact of utility on tonic change was not significant, $t(105)=1.8$, $p=.07$, $\beta=.17$

²⁰ Before removing outliers, the impact of utility on probability was marginally significant for the dice scenario, $t(102)=1.8$, $p=.07$, $\beta=2.97$.

A 2 (utility: negative/neutral) x 2 (IS: high/low) between subjects ANOVA conducted on the average of the two scenarios showed neither a significant main effect of utility, $F(1,103)=2.35, p=.13, \eta^2=.02$, nor a main effect of IS or a significant utility*IS interaction, all $F_s < 1$.

We then conducted the same set of orthogonal contrasts as in Experiments 1 and 2. The critical contrast comparing individuals high in IS to all other conditions was not significant, $t(103)=1.5, p=.14$ (see Figure 2.23), and neither were contrast 2, $t(103)=.82, p=.42$, or contrast 3, $t(103)=.22, p=.83$ ²¹.

Next, we examined the dice and container scenario separately from one another. Submitting the container scenario to a 2 (utility: negative/neutral) x 2 (IS: high/low) ANOVA, there was a marginally significant main effect of utility, $F(1,113)=3.09, p=.08$, but no main effect of IS, and no interaction between IS and utility, all $F_s < 1$. No orthogonal contrasts were significant, all $p_s > .19$.

Equally, submitting the dice scenario to the same ANOVA there were no main effects or interactions, all $p_s > .24$, and no orthogonal contrasts were significant, all $p_s > .21$.

²¹ CBP was not a significant predictor of the heart-beat counting task, $F(1,116)=1.98, p=.16, \eta^2=.02$.

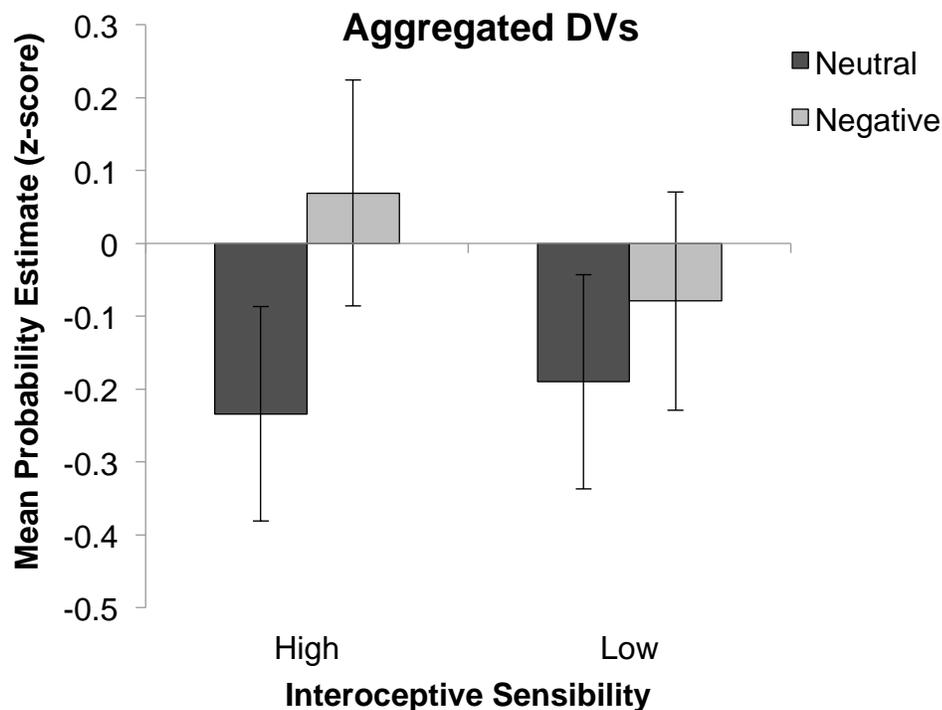


Figure 2.23. Probability estimate means for participants high and low in IS for neutral and negative conditions, averaged across scenarios. Error bars represent +/- 1 standard error of the mean.

Sense of Power. First, in line with our previous analyses, we performed a median split on the sense of power scale. As the power scale was administered at the very end of the experiment, we first observed whether utility impacted sense of power (given the proposed links between power and affect, see Anderson & Berdahl, 2002). A t-test confirmed that utility did not impact sense of power, $t(111)=.03, p=.98$.²²

Next, we conducted a 2 (utility: negative/neutral) x 2 (SOP: high/low) x 2 (scenario: dice/container) mixed ANOVA, with repeated measures on the last factor. Again, since no interactions involving the within subjects variable were found, all $ps > .29$, we z-scored and averaged the two scenarios. A 2 (utility: negative/neutral) x 2

²² Sense of power was neither related to CBP, nor to the performance in the heart-beat counting task, all $ps > .12$.

(SOP: high/low) between subjects ANOVA revealed no significant main effect of utility, $F(1,105)=2.42$, $p=.13$, $\eta^2=.02$ ²³, no main effect of sense of power, and no significant utility*sense of power interaction, all $F_s<1$ (see Figure 2.24). Next, we conducted orthogonal contrasts on the average of the two scenarios²⁴. Concretely, we tested for the hypothesis that only participants high in social power will show a negativity bias, see Table 2.6.

Table 2.6

Orthogonal contrasts conducted in Experiment 3.

	Powerful	Powerless	Powerful	Powerless
	Negative	Negative	Neutral	Neutral
Contrast 1	3	-1	-1	-1
Contrast 2	0	2	-1	-1
Contrast 3	0	0	1	-1

No orthogonal contrasts were significant, all $ps>.16$ ²⁵.

²³ Before the removal of outliers, the main effect of utility is marginally significant, $F(1,109)=3.79$, $p=.054$, $\eta^2=.03$

²⁴ We also examined whether sense of power impacted the utility manipulation checks of the two scenarios. For the dice scenario, the powerless perceived the negative outcome as worse ($M=5.5$, $SD=2.31$) than the powerful ($M=4.07$, $SD=2.02$), $F(1,117)=4.011$, $p=.048$, $\eta^2=.03$. No differences were found for the container scenario.

²⁵ Before the removal of outliers, contrast 1, coding for a difference between the powerful and all other conditions was significant, $t(109)=2.5$, $p=.02$. No other contrasts were found significant, all $ps>.54$, and no rest-variance was unaccounted for, $F<1$.

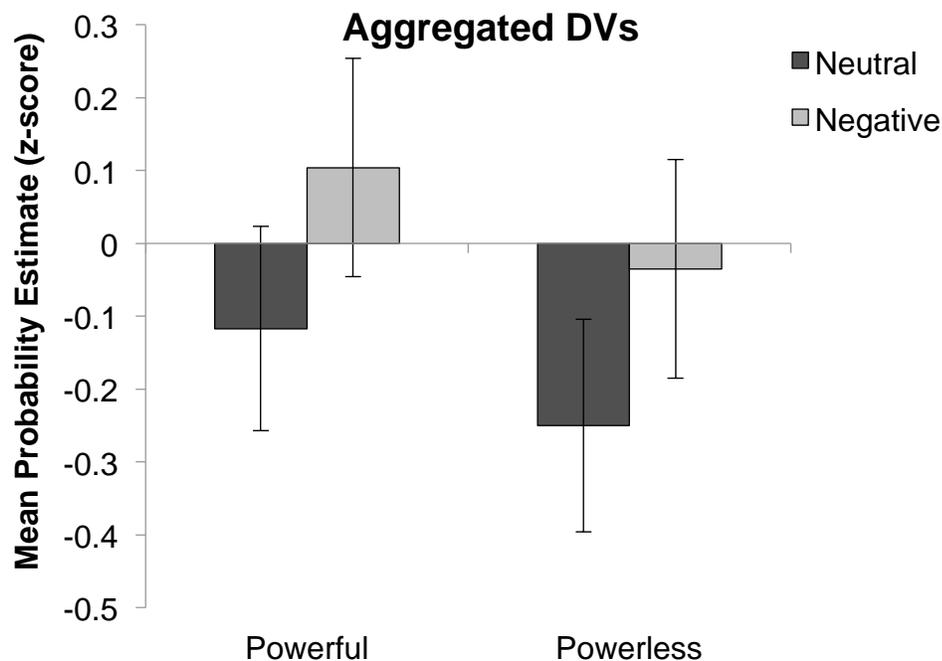


Figure 2.24. Probability estimate means for participants high and low in social power, for neutral and negative conditions, averaged across scenarios. Error bars represent +/- 1 standard error of the mean.

Next, as in Experiments 1 and 2 we examined the two scenarios separately. A 2 (utility: negative/neutral) x 2 (SOP: high/low) between subjects ANOVA on the dice scenario revealed no significant main effects or interactions, all $F_s < 1$, and no orthogonal contrasts were significant, all $p_s > .24$.

Next, we submitted the container scenario to the same 2 (utility: negative/neutral) x 2 (SOP: high/low) between subjects ANOVA. There was a significant main effect of utility, $F(1,112)=4.69$, $p=.03$, $\eta^2=.04$. Participants answering negative scenarios provided greater estimates ($M=36.64$, $SD=20.26$) than participants answering neutral scenarios ($M=29.59$, $SD=17.83$). Furthermore, there was a significant main effect of sense of power, $F(1,112)=4.37$, $p=.04$, $\eta^2=.04$. Individuals high in sense

of power provided greater probability estimates ($M=36.71$, $SD=19.35$) compared to individuals low in sense of power ($M=29.90$, $SD=19.02$). The sense of power*utility interaction was not significant, $F<1$. Conducting orthogonal contrast analyses, the critical contrast 1 was significant, $t(112)=2.65$, $p<.001$, whilst no other contrasts were significant, all $ps>.23$ (see Figure 2.25), and no rest-variance was unaccounted for, $F(2,113)=1.07$, $p=.35$.

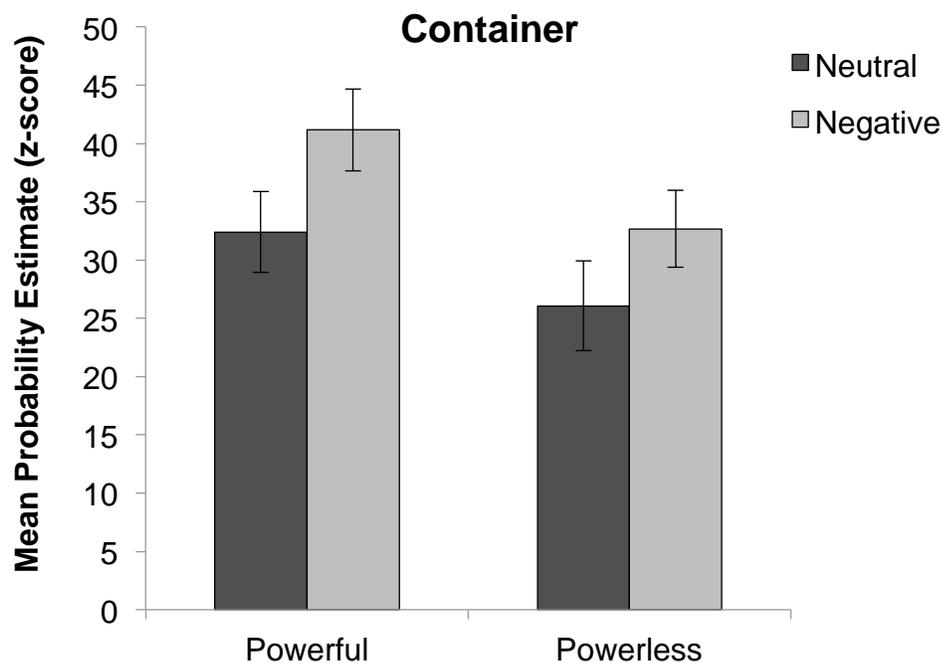


Figure 2.25. Probability estimate means for participants high and low in social power, for neutral and negative conditions for the container scenario. Error bars represent +/- 1 standard error of the mean.

Therefore, overall the hypothesis that individuals high in social power exhibit a greater negativity bias than individuals low in social power was weakly supported, receiving some support for the hypothesis in the container, but none in the dice scenario or when aggregating across scenarios.

Discussion

Experiment 3 failed to demonstrate evidence for a mediation of utility on probability estimates by arousal. Whilst the impact of utility on the change in tonic arousal was significant in the container scenario, this change did not mediate the relationship between utility and probability estimates. Furthermore, Experiment 3 did not provide evidence for the hypothesis that the extent to which individuals experience arousal impacts the relationship between arousal and probability estimates.

Interoceptively accurate individuals, who have been shown to experience arousal more intensively (e.g. Feldman-Barrett, Quigley, Bliss-Moreau, & Aronson, 2004; Werner, Jung, Duschek, & Schandry, 2009; Wiens, Mezzacappa, & Katkin, 2000) did not provide greater probability estimates for negative events compared to individuals low in interoceptive accuracy.

However, one limitation of the current set-up was that connecting participants to the physiological equipment might have made the source of arousal more salient. If this were the case, participants should not misattribute arousal, as they are aware of the source (see Zillman et al. 1974). However, according to the SLH participants should then also not demonstrate a negativity bias. Whilst the negativity bias in the current study was indeed weaker than in previous experiments in the container scenario, the negativity bias for the dice scenario was significant. Therefore, a factor other than arousal misattribution might have been the source of this bias.

Additionally, we provided some evidence that social power might potentially moderate the impact of negative utility on probability estimates: Only individuals high in sense of power showed a negativity bias in one out of two presented scenarios.

However, the effect in the container scenario was not very strong, and analysing the aggregated scores the hypothesis was not supported. The relationship between social power and the interdependence of utility and probability estimates will be examined further in Chapter 4.

Furthermore, whilst the moderation by IS was not significant, the means were congruent with the previous experiments. The discrepancy of finding some support for a moderation for the sensibility across the experiments presented in this chapter, but not for the accuracy dimension of interoception will be discussed in this chapter's General Discussion.

Chapter Discussion

Across 3 experiments, we explored the impact of interoception on the relationship between negative utility and probability estimates. Experiments 1 and 2 provided some preliminary evidence for the notion that interoceptive sensibility (IS) moderates the negativity bias. Here, individuals high in IS tended to exhibit a greater negativity bias than individuals low in IS. However, these results were not consistent across different scenarios, and no moderation occurred when IS was assessed after rather than before participants provided probability estimates.

Experiment 3 failed to demonstrate the mediation by measured arousal, and interoceptive accuracy did not impact the relationship between negative utility and probability estimates, questioning the state of the SLH. On the other hand, we provided some first evidence for the moderating role of social power in one of two scenarios, demonstrating that the powerful may, under some circumstances, be more sensitive to negative stimuli. Note further that we replicated Harris et al.'s (2009) finding that

individuals provide significantly greater probability estimates for negative events across almost all scenarios, and in all of them directionally.

Whereas Experiments 1 and 2 are somewhat supportive of the hypothesis that interoception, through an increase of the experience of arousal, leads to stronger arousal misattribution and therefore to a greater negativity bias, Experiment 3 suggests otherwise. The failure to show a mediation by arousal and the moderation by interoceptive accuracy might indicate that for one, interoceptive sensibility and accuracy are less related than we assumed and measure different constructs. Furthermore, they also hint that mechanisms other than arousal misattribution are responsible for the impact of interoceptive sensibility on the relationship between negative utility and probability estimates.

Whilst some studies do demonstrate the same effects of interoceptive sensibility and accuracy on dependent variables or even similar neurological correlates, (e.g. Critchley et al. 2004; Häfner, 2013) sensibility and accuracy measures are not necessarily correlated (Critchley et al. 2004; McFarland, 1975; Whitehead et al. 1977) – and neither were they in the current research. Exemplarily, Khalsa et al. (2008) showed that experienced meditators report greater levels of interoceptive sensibility, but do not perform better in the heartbeat counting task than non-meditators (despite also reporting to find the task easier than non-meditators).

The heartbeat counting task and its relationship to arousal intensity is, however, considerably less subject to criticism (Garfinkel & Critchley, 2013) than interoceptive sensibility as a valid self-report measure of one's ability to perceive internal states. For example, Silvia and Gendolla (2001) reviewed the literature with regard to the

“perceptual accuracy hypothesis”, but conclude in their review that the evidence for this perceptual accuracy hypothesis with regard to interoceptive sensibility is thin, and studies demonstrating this link are almost always more equivocally explained by other mechanisms. Nevertheless, sensibility and accuracy are often used interchangeably (Garfinkel et al. 2014; Ginzburg et al., 2013), and the lack of establishment of the differences between the two also affected the current research.

If, however, arousal misattribution did not underlie the impact of interoceptive sensibility, what else might have accounted for the effects? To the best of our knowledge, IS does not impact imaginability, construal levels, or approach-avoidance motivation, which could potentially explain its impact in light of an existing theory. However, IS could lead to greater bracing for loss. Individuals high in IS report to be more impacted by emotional states (Mor & Winquist, 2002), and might therefore engage in greater bracing to protect themselves from the impact of negative affect. Note that this does not necessarily mean that participants high in IS *actually* experience a greater impact of negative affect due to greater physiological responses, but more that individuals high in IS believe that they experience affect more intensively (see Roth, Breivik, Jørgensen and Hoffman, 1998). The bracing for loss hypothesis also leads to loss function asymmetries (where underestimating the probability leads to a greater negative affective impact than overestimating it), which we will turn to in Chapter 4.

In sum, the current findings question the state of the SLH. Given the lack of support for moderators (interoceptive sensibility and awareness) as well as the lack of support for arousal as a mediator, Chapter 3 aimed to take stock of the SLH by directly replicating Vosgerau’s (2010) experiments.

Chapter 3 Replicating the Stake-Likelihood Hypothesis

Chapter Overview

Chapter 2 failed to provide convincing support for arousal misattribution as a mechanism underlying the impact of negative utility on probability estimates, testing predictions derived from the SLH. In particular, we were unable to demonstrate that individuals who feel the arousal more strongly show a greater negativity bias. In addition, we failed to demonstrate that measured arousal mediates the negativity bias. Given the lack of indirect support for the SLH in the previous chapter, the current chapter aims to assess the SLH more directly by replicating some of the original findings that provided support for arousal misattribution as a mechanism. We present four experiments that aim to replicate two experiments by Vosgerau (2010). Despite our best efforts to create experimental conditions under which we would be able to demonstrate the stake-likelihood effect, we were unable to replicate previous results.

Introduction

Confirmation comes from repetition. Any attempt to avoid this statement leads to failure and more probably to destruction.

-John Tukey (1969, p.84)

The experiments reported in the last chapter question the status of the SLH. We predicted that only participants who “feel” the arousal evoked by negative utility (individuals high in interoceptive ability) should subsequently misattribute it to

probability estimates. However, whilst we provided some evidence that interoceptive sensibility, a self-report measure of interoception, moderates the impact of negative utility on probability estimates, we failed to demonstrate the same for objectively measured interoceptive ability. Furthermore, we were unable to find support for one of the most direct predictions of the SLH, namely that negative utility (having a stake in the outcome) increases arousal, and that it is this arousal that mediates the impact of negative utility on probability estimates.

In sum, the previous chapter did not find support for arousal misattribution as a mechanism when testing predictions that derive from the SLH, which brings doubt about the existence, or at least about the robustness, of arousal as a mechanism. On the other hand, it remains possible that whilst arousal misattribution indeed underlies the relationship between utility and probability, our predictions with regard to the above mentioned effects do not hold, or that our methods were inappropriate. For instance, the arousal evoked by utility in our scenario-based Experiment 3 could have been too weak to be detected by our equipment.

Before going any further with this line of enquiry, we decided it would be prudent to assess the reliability of the initial support of the SLH. Therefore, in the current chapter, rather than testing predictions that derive from the SLH, we aim to assess the status of the SLH by replicating the effects originally reported by Vosgerau (2010).

The majority of researchers agree that replicability is one of, if not *the* most important determinant as to whether or not an effect truly exists, guarding against the impact of false positives (Cohen, 1994, Fisher, 1935/1956; Roediger, 2012). Guarding

against false positives is of importance: whilst it appears that with the adoption of an alpha level of 5% less than 5% of research constitutes false positives, the actual number is likely to be much greater (Pashler & Harris, 2012) - for instance due to low statistical power and small effect sizes. As a result, Ioannidis (2005) even concludes that “most published research findings are false”. However, few replications are published in psychology – with one meta-analysis demonstrating that only 1.07% of psychological research are replications (Makel, Plucker & Hegarty, 2012). One reason for the small number of direct replications is the assumption that erroneous research is “self-correcting”, that in the process of assessing predictions of theories and working with effects, those effects that do not hold will be published as null results. However, research might be less self-correcting than assumed. In the present research for example, upon our inability to provide evidence for our hypotheses that are based on the SLH, we could have simply inferred that our results reflect a wrong prediction, rather than to examine the replicability of the original results. In practise, this leads to a slow self-correction of erroneous findings, as researchers more readily doubt their own experimental designs than the original effect (Pashler & Harris, 2012).

Recently, however, the importance of directly replicating psychological research has gained attention in particular in the field of (but not limited to) social psychology (e.g. Pashler & Wagenmakers, 2012; Wagenmakers, Wetzels, Borsboom, 2011) following some extraordinary research results by Bem (2011), fraudulent data (see Stroebe, Postmes, & Spears, 2012), and failed replications of high profile research (e.g. Doyen, Klein, Pichon, & Cleeremans, 2012; Pashler, Coburn & Harris, 2013; Shanks et al. 2013). As a positive result, researchers have started to also report successful and non-

successful replications on websites such as psychfiledrawer.org (see also openscienceframework.org), and following the call for changes in the publishing incentive system (Koole & Lakens, 2012) some of the most highly regarded journals of psychology are now accepting pre-registered replication attempts (see Brandt, 2013, Chabris et al. 2012; Pashler et al. 2013). The current research follows the recent call of researchers in psychology to assess the replicability of existing effects.

The Current Research

In this chapter, we will present three experiments that aim to replicate Vosgerau's (2010) Experiment 2, and one experiment that aims to replicate Vosgerau's (2010) Experiment 1. Such replication attempts can be direct or conceptual (Schmidt, 2009), constituting replications that employ the original materials or extend the theory to another set of stimuli, respectively. Whilst Experiments 6 and 9 constitute conceptual replications of Vosgerau's (2010) Experiments 2 and 1, respectively, Experiments 7 and 8 are direct replications of Experiment 2. Whilst Experiment 6 was conducted online, Experiments 7, 8 and 9 were run in the laboratory.

Out of the four experiments presented by Vosgerau (2010), we considered his Experiment 2 to be the strongest demonstration of arousal misattribution as a mechanism. In Experiment 2, Vosgerau (2010) demonstrated that probability estimates decrease once the source of arousal is salient, supporting the notion that arousal is misattributed under conditions where the source of the arousal is not salient. To preempt the results from our attempt at replicating Experiment 2, despite conducting three experiments, one online, and two direct replications in the laboratory, we consistently find no evidence for the existence of arousal misattribution in the formation of

probability estimates. Therefore, we also replicated Vosgerau's (2010) Experiment 1, testing the basic premise that increased arousal results in greater probability estimates. Again, we were unable to replicate these results.

Experiment 4

Background

In Vosgerau's (2010) Experiment 2, participants were asked either before or after they estimated the probability of a positive event to indicate their levels of excitement. Vosgerau proposes that this question draws attention to the source of their arousal (from having a stake in the outcome), and as a result, arousal is not misattributed. Indeed, in Vosgerau's (2010) Experiment 2, participants provided lower probability estimates when they were asked to rate their excitement before. The current experiment aims to examine this effect more closely by means of a conceptual replication.

In his research, Vosgerau (2010) only tested the manipulation of the arousal misattribution process for positive, but not for negative events. Moreover, Vosgerau tested this effect for both focal and non-focal events. Whilst the effect appeared very strong for the focal event ("what is the probability that at least one 3 will show up"), this was reduced for the non-focal event ("what is the probability that no 3 will show up"). It therefore remains a possibility that arousal misattribution occurs for positive, but not for negative events. For instance, the arousal evoked by negative events could be too strong for misattribution processes to occur (Gorn et al. 2011), explaining our failure to support predictions derived from the SLH in Chapter 2. Therefore, the current study aimed to

replicate Vosgerau's Experiment 3, as well as to extend it to negative and neutral events. The latter were included in order to rule out the possibility that answering questions probing for arousal before providing probability estimates decreases estimates for all events irrespective of utility. If this were to be the case, a mechanism other than the prohibiting of the misattribution process would underlie this effect, as neutral events should not evoke arousal.

In addition, the current experiment changed the probability levels to the original levels by Vosgerau, in order to resemble his experiment more closely ("at least one 3 within four rolls"). Moreover, we did not include scenarios other than the dice scenario. In the original experiment participants indicated to be more aroused if arousal was rated before providing probability estimates ($M=5.06$, $SD=1.15$) compared to when the arousal question was asked afterwards ($M=4.35$, $SD=1.30$). Vosgerau interprets this shift as the arousal being attributed to probability judgments, which serves as evidence for a manipulation of the arousal misattribution process. If we were to include more than one scenario, a meaningful interpretation of arousal ratings would not be possible since arousal questions of the later scenarios would have been preceded by likelihood estimates. Therefore, in order to achieve non-confounded estimates of arousal, the current experiment only includes the dice scenario, the scenario that is conceptually the most similar to the original task by Vosgerau (2010).

Method

Participants and Design. 303 participants were recruited via Amazon Mechanical Turk and were paid \$0.2 to participate in this online experiment. Following previous exclusion criteria, the final sample consisted of 285 participants (100 female,

median age = 27 years). Note that we more than tripled the cell size from Vosgerau's (2010) 13 participants per cell to approximately 47 participants per cell. Participants were randomly assigned to a 3 (utility: positive/neutral/negative) x 2 (order: arousal rating first/probability estimate first) between participants design.

Materials and Procedure. Participants were informed that they would participate in a study on risk perception. Participants read the dice scenario, where the negative and neutral scenarios were identical to the dice roll scenario used in Experiments 1 and 2. In the positive utility condition participants read:

Imagine you are walking down the street with another person. Seeing a dice for sale in the window of a nearby thrift shop, they propose the following: You will role a six-sided dice 4 times. If a 3 comes up on at least 1 of these throws, the other person will give you \$100. Otherwise they will keep it.

In all conditions, participants estimated the probability that at least one 3 shows up on four throws and provided their answer on a sliding scale from 0 to 100. Furthermore, depending on the order condition, participants were either asked before or after the provision of the probability estimate to answer the questions: "how much would you like playing this game?" and "how exciting would it be to play this game"?, both on 8-point Likert-scales from 0 (not at all) to 7 (very much).

Results and Discussion

Outlier Exclusion. Whilst the previous and the following empirical chapters removed outliers by plotting Cook's Distance against Studentized Deleted Residuals, the experiments in the current chapters only remove outliers when Vosgerau (2010) explicitly stated that he did so, too. In that case (see Experiment 7), we followed his instructions for the way we removed outliers. Thus, unless explicitly noted, the analyses in this chapter are based on data where no outliers have been removed.

Probability Estimates. Following an arousal misattribution account, it would be predicted that probability estimates for the negative and positive conditions are both higher than estimates for the neutral condition, when the arousal question is asked afterwards. This would be expected, as the positive and negative conditions lead to arousal that is misattributed to probability estimates (unless the source of the arousal is made salient). Furthermore, it would be expected that probability estimates decrease when the arousal question is asked before compared to after the provision of estimates. In particular, the order of the arousal question was expected to impact probability estimates for negative and positive, but not for neutral outcomes (where no arousal should be elicited).

A 3 (utility: negative/neutral/positive) x 2 (order: arousal rating first/probability estimate first) ANOVA yielded no significant main effect of utility, $F(2,270)=1.8$, $p=.17$, $\eta^2=.01$. However, directionally, participants in the negative condition gave the highest probability estimates ($M=47.9$, $SD=23.84$), followed by participants in the neutral ($M=43.37$, $SD=21.47$) and positive condition ($M=42.04$, $SD=21.33$). The direction of these means are not in line with the SLH, which would have predicted greater estimates for positive and negative compared to neutral events. Furthermore, contrary to the predictions by the SLH, the effect of order, $F(1,279)=2.15$, $p=.14$, $\eta^2=.01$, and the utility*order interaction, $F<1$, were non-significant (see Figure 3.1).

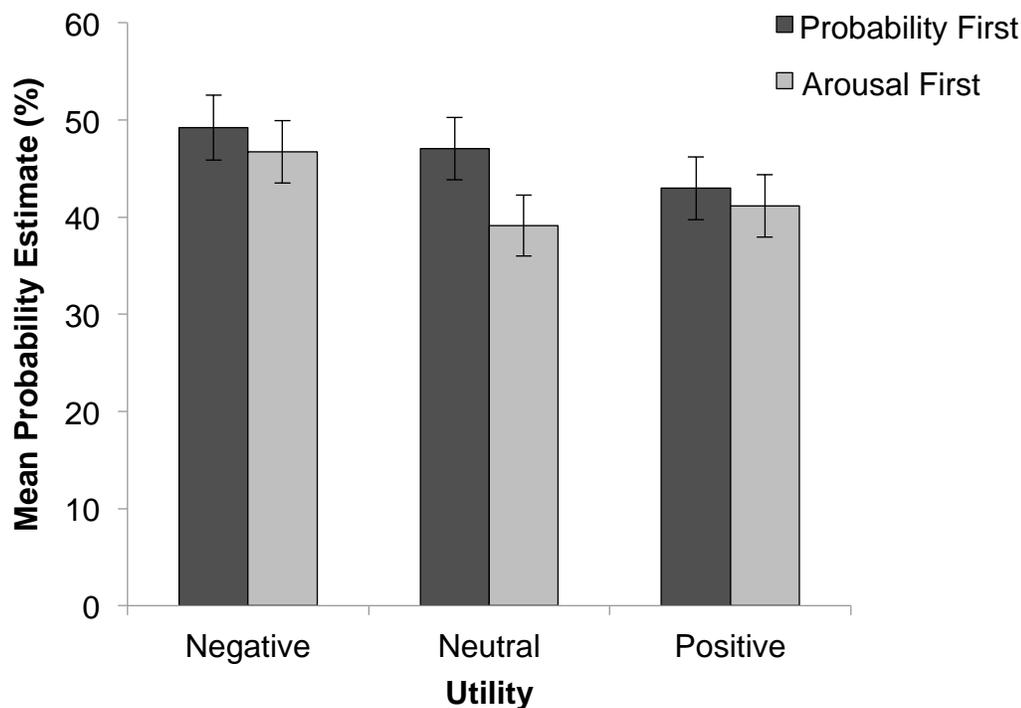


Figure 3.1. Probability estimate means for the utility and order conditions. Error bars represent +/- 1 standard error of the mean.

Arousal Ratings. As the two arousal questions were highly correlated, $r(285)=.81, p<.001$, they were averaged into one arousal score and submitted to a 3 (utility: negative/neutral/positive) x 2 (order: arousal rating first/probability estimate first) between participants ANOVA. A main effect of utility was present, $F(2,279)=94.49, p<.001, \eta^2=.41$. Participants in the positive condition gave the highest arousal ratings ($M=5.6, SD=1.46$), followed by participants in the neutral ($M=3.84, SD=1.62$) and negative conditions ($M=2.52, SD=1.52$). Whilst a main effect of utility would be expected, in that both individuals in the negative and positive conditions report to be more aroused than participants in the neutral condition, it was not expected that individuals in the negative condition reported lower arousal than individuals in the neutral condition. This result might hint towards an inappropriate arousal question for the

negative condition. We employed the original questions from Vosgerau (2010) that asked participants how “excited” they were, and how much they “liked” playing the game. It is likely that these questions capture arousal in the positive domain, but that more appropriate questions in the negative domain would be to ask participants how “anxious” they would be playing the game, and how much they “disliked” playing the game.

Furthermore, following the SLH, it was expected that participants’ ratings of arousal decrease if they rate their arousal after the provision of probability estimates (as arousal is then misattributed). In particular, this effect was expected to occur for the negative and positive, but not for the arousal-free neutral condition. Contrary to these predictions, however, arousal ratings were not affected by the main effect of order, $F(2,279)=1.75, p=.19, \eta^2=.01$, and the order*utility interaction was not significant, $F<1$ (see Figure 3.2).

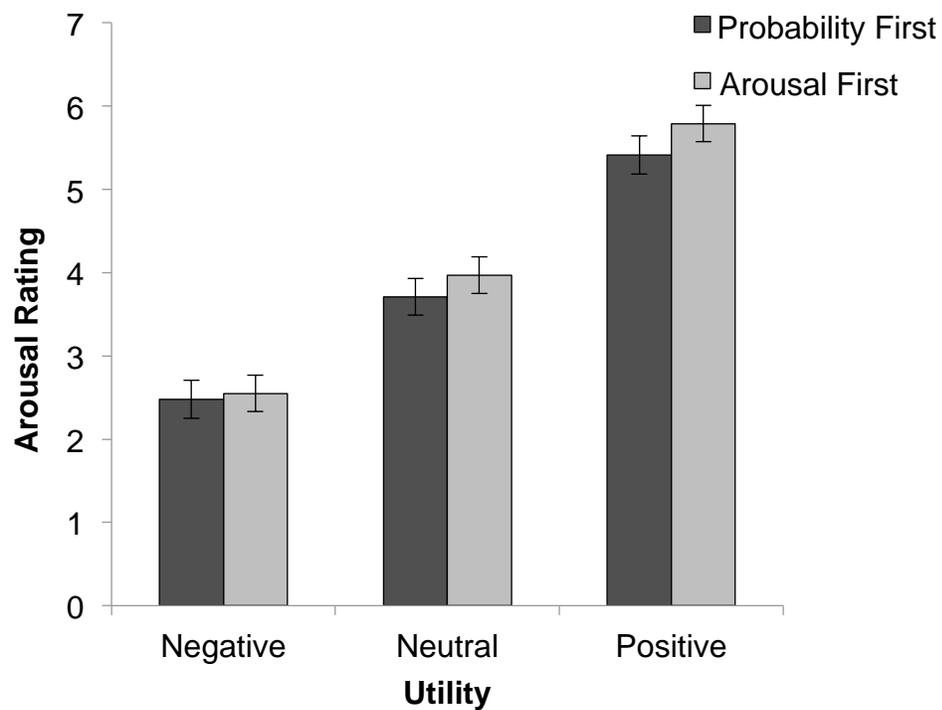


Figure 3.2. Arousal ratings for the utility and order conditions. Error bars represent +/- 1 standard error of the mean.

Bayesian Analysis. It is recognized that conventional hypothesis testing cannot provide support in favour of a null hypothesis: it can only make the alternative hypothesis less likely. Also as a result of the greater attention paid to replications recently, some researchers have argued for the application of Bayesian statistical tests that can provide support for the null hypothesis over and above an alternative hypothesis (e.g. Gallistel, 2009; Rouder et al. 2009; Wagenmakers, Wetzels, Borsboom, & Van Der Maas, 2011). Therefore, in order to assess the evidence for the null compared to the Stake-Likelihood hypothesis on the influence of order (arousal rating first / probability estimate first) on estimates of probability, we conducted 3 Bayesian t-tests on each of

the 3 utility levels, using the R package BayesFactor²⁶. In order to assess the concrete prediction that probability estimates will be higher when the arousal rating is provided after compared to before the probability estimate, an interval null extension was applied (Morey & Rouder, 2011), which allows for directional hypothesis testing. Since this interval null extension only holds for comparisons between two independent groups, we did not conduct a Bayesian ANOVA across the three utility levels, but instead analysed them individually.

The data from the negative and positive probability levels were found to be 3.03 and 3.25 times more likely respectively under the null hypothesis compared to the SLH (where estimates for the “probability first” condition are predicted to be greater than estimates for the “arousal first” condition). These results contribute ‘some’ evidence for the null hypothesis (Rouder et al., 2009, p. 228).

Likewise, we conducted three directional Bayesian t-tests on the arousal scores on each of the three utility levels, testing the hypothesis that the liking scores are greater in the “arousal first” compared to the “probability first” condition. The data were found to be 3.81 and 1.23 times more likely under the null hypothesis than under the SLH for the negative and positive probability levels, respectively. Whilst the evidence for the null hypothesis constitutes “some” evidence for the negative condition (Rouder et al. 2009, p.228), the evidence in favour of the null hypothesis is “not worth more than a bare mention” (Kass & Raftery, 1995, p. 777) for the positive conditions.

²⁶ The data was analysed using the software R (version 3.0.2) (<http://www.r-project.org/>). Bayesian analyses were conducted using the R package BayesFactor (version 0.9.6), with the package’s default scale of $\sqrt{2} / 2$ for the Cauchy prior of δ .

In sum, although the means of the present experiment (both of arousal ratings and probability estimates) were in the direction predicted by the SLH, the results provided greater evidence in favour of the null hypothesis in all instances.

Experiment 5

Background

Experiment 4 failed to replicate Vosgerau's (2010) Experiment 2 in the positive domain and was also unable to extend it to the negative domain. Participants' probability estimates did not decrease when they were directed to the source of the arousal, and arousal levels did not decrease when they rated it after the provision of probability estimates.

However, Experiment 4 differed from the original experiment. Above all, this experiment was conducted online with no real monetary outcomes, whereas in Vosgerau's (2010) Experiment 2 participants had the chance to earn \$5. One could argue that the arousal created from the scenario was not strong enough to be misattributed in this online setting due to the lack of real incentives. On the other hand, it could also be possible that the arousal evoked was too strong, due to the mentioning of a significantly greater amount of money that could be won or lost (\$100 rather than \$5). Vosgerau (2010) highlights that if arousal is too intense, the source of arousal becomes salient and no misattribution should occur. Furthermore, in the original experiment, participants could earn an extra \$1 when the estimate was within +/- 5% of the objective probability. Unlike Vosgerau (2010), we did not incentivise accurate probability estimates in Experiment 4.

Following the failure to replicate arousal misattribution conceptually online in Experiment 4, Experiment 5 sought to directly replicate the original findings of Vosgerau's (2010) Experiment 2, using the same manipulations, the same computer program to display instructions, and the same rewards in the laboratory. Given the difficulties in asking a functionally equivalent arousal question in the negative domain, and the unexpected trends in the neutral condition, we only conducted this experiment with positive outcomes, as in the original experiment. Furthermore, Vosgerau (2010) also manipulated the focus, in that participants either reported the likelihood of the focal outcome (what is the likelihood of getting a 3 at least once in four dice rolls), or the non-focal outcome (what is the likelihood of getting no 3s in four dice rolls). We omitted the non-focal outcome condition and only asked participants to report the focal outcome, as it was this latter condition where the order of arousal rating had the strongest effect on probability estimates in the original experiment.

Method

Participants and Design. 30 Participants took part in this experiment (15 females, median age =22 years) conducted in the laboratory. With 15 participants per cell, we therefore employed a similar cell size as the original experiment (cell size = 13) by Vosgerau (2010). Participants received a chocolate bar in return for their participation, in addition to any money won during the experiment, and were randomly assigned to one of the two order conditions (probability first/arousal first).

Procedure and Materials. Participants were approached opportunistically and were asked to participate in a brief study in which they would get a chocolate bar and could win some money. Participants were informed that they would take part in a brief

study on risk perception. The experimenter then started the computer program that was the original program from Vosgerau's (2010) Experiment 2, generously provided to us by the author. Participants read the following instructions on the screen: "You will toss a die four times. If you toss a "3" (one or several 3s) in the four tosses you win £3".

Afterwards, participants were given the following instructions: "In order to make sure that you understand the game, please describe it in your own words below. Please ask the experimenter if anything is not clear". After participants described the task in their own words, they were asked the questions: "How much do you like playing this game" and "How exciting is it to play this game", either before or after they were asked to estimate the probability of a 3 occurring. The instructions for the probability judgment were: "How likely do you think it is that you will toss one or more 3s in the four rolls? If your probability estimate is within 5% of the true probability, you will receive an extra £1." Afterwards, participants completed demographic information, rolled the dice, and were paid with any money won during the task as well as with a chocolate bar.

Results

Probability Estimates. Contrary to the predictions of the SLH, there was no main effect of order, $F < 1$. Furthermore, in conflict with previous findings by Vosgerau (2010), when estimating the probability *after* the arousal ratings, participants gave *higher* probability estimates ($M=42.67$, $SD=25.49$) compared to when participants were asked to estimate the probability *before* the arousal ratings ($M=37.33$, $SD=25.2$).

Arousal Ratings. Since the correlations between the two questions were high, $r(30)=.80$, $p < .001$, we averaged the two questions. The average was then submitted to an ANOVA. The main effect of order was not significant, $F < 1$. Directionally, when

answering the arousal questions before the probability judgment, participants rated their arousal to be lower ($M=4.03$, $SD=1.23$) compared to when they answered the arousal questions afterwards ($M=4.43$, $SD=1.04$). Again, this difference was in the opposite direction to the findings by Vosgerau (2010), where arousal ratings were greatest when participants answered the arousal question first.

Bayesian Statistics. As in the previous experiment, we conducted the Bayesian equivalent of a t-test with an interval null extension both on the probability estimates as well as on the arousal ratings. The data were 4.12 and 4.97 times more likely under the null hypothesis than the SLH for probability estimates and arousal ratings, respectively, providing “some” (Rouder et al. 2009, p.228) evidence in favour of the null hypothesis.

Experiment 6

Background

Experiment 5 failed to replicate the original Experiment 2 by Vosgerau (2010), despite using identical materials. However, we offered participants a mars-bar for participating, irrespective of the outcome (to make sure that participants would be compensated for taking part in the experiment in case they did not win the money). As offering a sugary, fatty reward could potentially lead to a “hot” visceral state that could interfere with probability estimates (see Risen & Critcher, 2011), we repeated the identical experimental set-up, but did not mention the chocolate-bar at the start of the experiment. Therefore, Experiment 6 removes the only difference in the experimental set-up of our Experiment 5 and Vosgerau’s (2010) Experiment 2, and therefore constitutes a direct replication attempt.

Method

Participants and Design. 40 Participants were recruited to take part in this experiment conducted in the laboratory (24 female, median age =23.5 years), slightly increasing the cell size to 20 participants per cell (compared to Vosgerau's, 2010, 13 participants per cell). Participants were reimbursed with money won during the task (if participants did not win anything, they were reimbursed with a chocolate bar, but this was unknown to participants when starting the task). Participants were again randomly assigned to one of the two order conditions (probability first/arousal first).

Materials and Procedure. Materials and procedure were identical to Experiment 6, except when approaching participants they were told that there is the chance to win money, and no chocolate bar was mentioned.

Results

Probability Estimates. Order (probability first/arousal first) did not affect probability estimates, $F < 1$. As in Experiment 6, and contrary to Vosgerau's (2010) findings, when estimating the probability *after* the arousal ratings, participants gave *higher* probability estimates ($M=49.55$, $SD=24.82$) compared to when participants were asked to estimate the probability *before* the arousal ratings ($M=47.77$, $SD=27.18$).

Arousal Ratings. The correlations between the two questions were high, $r(40)=.74$, $p < .001$, and the average of the two questions was submitted to a one-way ANOVA with order (arousal rating first/probability estimate first) as a factor. There was a significant impact of order on arousal, $F(1,38)=6.92$, $p=.01$, $\eta^2=.15$. As predicted by the SLH, when answering the arousal questions before the probability judgment,

participants rated their arousal to be higher ($M=5.55$, $SD=1.22$) compared to when they answered the arousal questions afterwards ($M=4.39$, $SD=1.55$).

Bayesian Statistics. Again, we observed the evidence with Bayesian Statistics applying an interval null extension to test for the directional hypothesis as proposed by SLH. For the probability estimates the data were 3.73 times more likely under the null hypothesis compared to the SLH, constituting “some” evidence in favour of the null hypothesis (Rouder et al. 2009, p. 228). On the other hand, the arousal rating was 8.4 times more likely under the SLH compared to the null hypothesis, approaching “strong” evidence (Rouder et al. 2009, p. 228). However, despite the decrease in arousal following the probability estimates in this experiment, probability estimates themselves remained unaffected.

Bayesian Meta-Analysis across Experiments 4, 5 and 6

The evidence in favour of the null hypothesis over the SLH for Vosgerau’s (2010) Experiment 2 is further strengthened by a Bayesian meta-analysis conducted across the three experiments (for instructions see Rouder & Morey, 2011). This meta-analysis assumes that the true effect size is constant across experiments. However, it does not assume equal variances, which makes it applicable to the current set of experiments where different variances could have been produced, for example by different means of recruiting participants (e.g. online in Experiment 4, in the laboratory in Experiments 5 and 6).

We applied this meta-analytic Bayes factor to the probability estimates of the three experiments we conducted (we only considered the data from the positive condition of Experiment 4). Here, we only observe the Bayes factor for probability

estimates, as the decrease in arousal levels is only of interest as a potential mediating mechanism should there really be a decrease in probability.

For each, we first extracted the t -value of order (arousal rating first/probability estimate first) (see Table 3.1 for the t -values associated with the three experiments). We then analysed the data using the R package Bayes Factor, together with a script for a Bayesian meta-analysis, written by Rouder & Morey (2011), again including an interval null extension allowing us to assess the SLH specifically. This analysis revealed a value of 5.8 to 1 in favour of the null hypothesis over the SLH. As such, the current results are constituting “some” evidence in favour of a null over the SLH (Rouder et al. 2009, p.228).

Table 3.1

The number of participants, the t -value associated with the main effect of order, and whether the results were in the direction predicted by the SLH for the Experiments 4, 5 and 6.

	N	t	<i>Direction Predicted by SLH</i>
Experiment 6	95	.42	Yes
Experiment 7	30	-.58	No
Experiment 8	40	-.21	No
Data pooled	165	-.37	1/3

Experiment 7

Background²⁷

Given our repeated failure to replicate Vosgerau's (2010) Experiment 2 across three experiments and the evidence for the null hypothesis over the SLH in a Bayesian meta-analysis, we now aimed to replicate Vosgerau's (2010) Experiment 1, in order to further assess the replicability of effects reported in this article. In his Experiment 1, Vosgerau (2010) asked participants to estimate the probabilities of 9 future outcomes, such as "How likely do you think Barack Obama is to become the Democratic presidential candidate?", or "How likely do you think the Steelers are to win the Super Bowl in 2008?". Vosgerau (2010) included two sets of questions (manipulated between participants), one of which was the complement of the other set (e.g. "How likely do you think Barack Obama is *not* to become the Democratic presidential candidate?").

Importantly, these questions were either presented on bright pink paper, which was assumed to increase arousal, or on grey paper, which was assumed to not evoke arousal. Irrespectively of the set of stimuli used, Vosgerau (2010) demonstrated that when participants read the statements on pink paper, they reported overall greater estimates than when they read the statements on grey paper, suggesting that arousal is misattributed to probability estimates.

As the original experiment was conducted before 2008 in the US, most of the original items were not appropriate for a 2014 UK participant pool. Therefore, we kept 2 of the original items, which did not depend on time or location, and complemented them with 8 items that were more suitable.

²⁷ I would like to thank Emma Breeze, Jasmine Taylor, Sarita Aujla and So Pei Quan for their help in collecting data and creating the stimuli used in this experiment.

Method

Participants and Design. 177 participants (97 female, median age = 20 years) took part in the study and were recruited opportunistically on the campus of University College London (a similar recruitment method was employed by Vosgerau, 2010, who recruited passers-by close to the campus of the University of Pittsburgh). Participants were randomly assigned to a 2 (colour: pink/grey) x 2 (focus: likelihood of event occurring/not occurring) x 10 (probability questions) mixed design, with repeated measures on the last factor.

Materials and Procedure. Materials and procedure closely followed Vosgerau's (2010) set up. Half the participants were asked 10 likelihood questions, whereas the other half answered 10 questions with complementary likelihoods (see Table 3.2). Following Vosgerau (2010), arousal was manipulated by printing questionnaires either on bright pink paper (arousal) or on light grey paper (no arousal). As in Vosgerau (2010), likelihood judgments were made on a 21-point scale ranging from 0% to 100%, in 5% increments.

This set-up closely follows the original set-up by Vosgerau (2010). However, one of the differences between the original and the present experiment are the questions used, as we had to adapt the questions to fit the current participant sample. Furthermore, Vosgerau (2010) presented participants with 9 questions, whereas we used 10. Additionally, as there were no records of the brand of the paper used, we most likely used a different brand of paper (but chose the most fluorescent pink possible, see Figure 3.3).

Table 3.2

Probability estimation questions asked in Experiment 7. Questions 1 and 5 are original items from Vosgerau (2010).

Focus	
Version 1 Ten outcome questions	Version 2 Ten complementary outcome questions
1. How likely do you think it is that of 80 passengers on an airplane, none of the passengers will have been born on the same day (i.e. same day but not necessarily same year)?	1. How likely do you think it is that of 80 passengers on an airplane, at least two will have been born on the same day (i.e. same day but not necessarily same year)?
2. How likely do you think it is that Roger Federer will win Wimbledon 2014?	2. How likely do you think it is that Roger Federer won't win Wimbledon 2014?
3. How likely do you think it is that Great Britain will not break their existing record for gold medals in the next summer Olympics?	3. How likely do you think it is that Great Britain will break their existing record for gold medals in the next summer Olympics?
4. How likely do you think it is that Prince William and his wife Catherine won't have another baby by the end of 2016?	4. How likely do you think it is that Prince William and his wife Catherine will have another baby by the end of 2016?
5. If a six-sided die is tossed four times, how likely do you think a 6 is to appear exactly twice in the four tosses?	5. If a six-sided die is tossed four times, how likely do you think a 6 is to appear fewer or more than two times in the four tosses?
6. How likely do you think it is that the average temperature in South West England for June this year won't be higher than previous records of the last 10 years?	6. How likely do you think it is that the average temperature in South West England for June this year will be higher than previous records of the last 10 years?
7. How likely do you think it is that it will snow in London in January 2015?	7. How likely do you think it is that it won't snow in London in January 2015?
8. How likely do you think it is that Scotland will not become independent following the upcoming referendum?	8. How likely do you think it is that Scotland will become independent following the upcoming referendum?
9. How likely do you think it is that we will find a cure to all cancers in the next 20 years?	9. How likely do you think it is that we won't find a cure to all cancers in the next 20 years?
10. How likely do you think it is that organ donation will become opt out rather than opt in?	10. How likely do you think it is that organ donation will remain opt in rather than become opt out?



Figure 3.3. Images of the paper used in Experiment 7, with (un-arousing) grey paper on the left, and (arousing) pink paper on the right hand side.

Results

Out of the 177 responses all but one participant answered all probability questions. Following the procedure of Vosgerau (2010), outliers further than 2.5 standard deviations away from the mean were removed from further analyses ($n=3$). Thus, the final sample consisted of 173 participants, an increase of 22 participants compared to the original experiment by Vosgerau (2010).

After averaging across the ten probability questions, a 2 (colour: pink/grey) x 2 (focus: version 1/version 2) between subjects ANOVA was conducted. Expectedly, this ANOVA yielded a significant main effect of focus, $F(1,169)=13.67$, $p<.001$, $\eta^2=.08$. Participants who answered version 2 of the questionnaire (see Table 3.2) gave higher probability estimates ($M=51.67$, $SD=9.35$) compared to participants who answered version 1 of the questionnaire ($M=46.52$, $SD=9.03$). However, the colour of the questionnaire did not significantly impact probability estimates, $F(1,169)=1.05$, $p=.31$, $\eta^2=.01$, and there was also no colour*focus interaction, $F<1$ (see Figure 3.4). Although directionally, in line with Vosgerau's (2010) findings, participants who answered on

pink paper gave higher probability estimates ($M=49.8$, $SD=9.52$) compared to participants who answered on grey paper ($M=48.46$, $SD=9.54$).

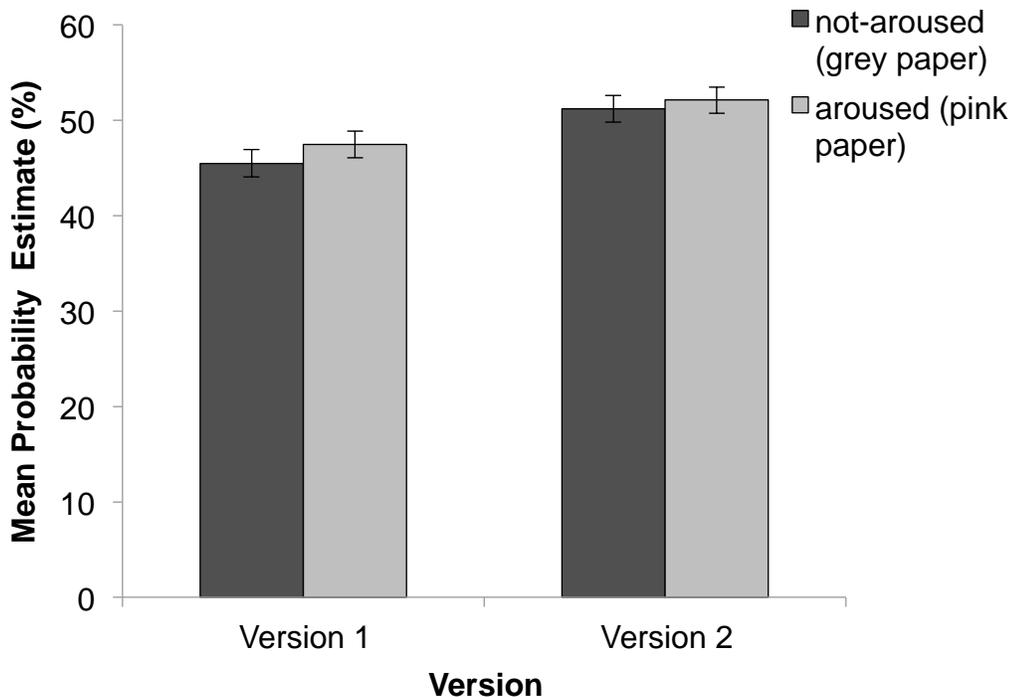


Figure 3.4. Probability estimate means for participants answering version 1 or version 2 on either pink, or grey paper. Error bars represent +/- 1 standard error of the mean.

Bayesian Analysis

Last, we conducted a directional Bayesian t-test (Morey & Rouder, 2011), assessing the prediction of the SLH that participants should give higher probability estimates on pink than on grey paper. The data were 2.56 times more likely under the null hypothesis than under the SLH, thus approaching “some” evidence for the null hypothesis over the SLH (Rouder et al. 2009, p.228).

Chapter Discussion

We presented four experiments that aimed to replicate Experiments 1 and 2 of Vosgerau's (2010) research demonstrating the stake-likelihood effect. However, despite our best efforts to closely match the original experimental materials and procedures, we failed to provide any evidence for the existence of such an effect. In an online experiment, Experiment 4 failed to replicate Vosgerau's (2010) finding that when the misattribution process is interrupted by directing participants to the source of their arousal, likelihood estimates decrease. Experiments 5 and 6 failed to replicate the same finding in a set-up that matched Vosgerau's (2010) more closely, using his materials in the laboratory. Quantifying these results, a Bayesian meta-analysis across Experiments 4, 5 and 6 suggested that the data of the three experiments were 5.8 times more likely under the null hypothesis than under the SLH, constituting "some" evidence (Rouder et al., 2009, p.228). Last, we also did not find evidence for an increase in probability estimates when participants provided their estimates on pink (arousing) paper, compared to when they estimated probabilities on grey (unarousing) paper. Importantly, for all four presented experiments the evidence was in favour of the null hypothesis over the SLH.

Our inability to reproduce the original findings of the SLH across four experiments challenges the replicability of the results reported by Vosgerau (2010). More fundamentally, since, to the best of our knowledge, no other research has been published exploring the role of arousal in probability estimates, the present failed replications potentially also question the role of arousal misattribution as a mechanism for the impact of negative utility on probability estimates.

For a misattribution process to occur, two independent mechanisms need to take place: First, arousal needs to be created, and next, this arousal needs to inform the decision-maker. In our experiments, (at least) one of the two processes did not occur. In Experiment 3 of the previous chapter arousal was increased for one of the two scenarios, however, this increased arousal did not mediate the impact of utility on probability estimates, suggesting that arousal did not inform probability estimates.

Similarly, in Experiment 6 we observed a decrease in arousal ratings when participants rated their arousal after the probability estimates – the only time the evidence appeared to be in favour of the SLH rather than the null hypothesis. One reason why we observed a drop in arousal ratings might be that participants did not actually enjoy providing probability estimates, and therefore the decrease might more reflect how much they enjoyed the game (rather than actual arousal levels). Participants were only informed after they provided the arousal rating that their estimates were incentivised, which might have increased the pressure to get the probability right, and decreased the joy with which participants took part in this game. However, even if this decrease in arousal ratings reflected an actual reduction in experienced arousal after individuals provided probability estimates, what we can conclude is that the greater experienced arousal did not have any impact on probability estimates. Only in conjunction with a decrease in probability estimates would this observation be of importance. In addition, what remains unexplained is why we observed this decrease in arousal in Experiment 6, but not in Experiment 5 (where an additional mars-bar was promised to participants at the start of the experiment).

Furthermore, some of the results in Experiments 1 and 2 provided indirect evidence for a potential misattribution of arousal, whereby individuals high in interoceptive sensibility tended to provide a greater negativity bias than individuals low in interoceptive sensibility (although this effect was not consistent across different samples and scenarios). The preliminary conclusions about the existence of this moderation now need to be viewed in a new light, as Experiments 3-7 do not find evidence for this process in more direct examinations of the SLH: Interoceptive accuracy did not moderate and measured arousal did not mediate the effect, increasing the awareness of the source did not decrease, and manipulating arousal independently of utility did not increase probability estimates. Therefore, the initial findings of Experiments 1 and 2 might be due to a mechanism other than arousal misattribution (such as greater bracing for loss, see the General Discussion of the previous chapter).

However, critics of replications as a means of assessing the existence of effects argue that replications have no meaningful scientific value. According to these critics, even if one follows the concrete instructions of the original methodology closely, there will be unreported experimental events (such as experimenter behaviours) that are not reported in the method section (e.g. Mitchell, 2014). This claim appears to be supported by the finding that replications are less likely to be successful if the original author is not involved in the research (Makel et al., 2012) (although it is possible that the original authors are less likely to publish conflicting results). Whilst we followed the original protocol as closely as possible, and Experiments 7 and 8 were even administered by a computer program that was identical to the program used in the original experiments (provided to us by the author), we cannot eliminate the possibility that also in the current

experiments previously unidentified moderators of the effect led to our inability to replicate the original findings.

In addition, it has been noted that effect sizes in the literature are likely to be greater than the true effect size (Greenwald, 1975), and when replicating, one should increase the number of participants to ensure sufficient statistical power (Brandt et al. 2012). Whilst in all our experiments we ensured that the number of participants was greater than in Vosgerau's (2010) original experiments, one might argue that the current experiments could be underpowered and therefore fail to demonstrate the stake-likelihood effect. However, in the two experiments (Experiments 5 and 6) that used the original materials by Vosgerau (2010), the means were in the *opposite* direction to what was predicted by the SLH, highlighting that we did not simply lack statistical power. Additionally, Bayes-Factors are more independent from sample size than conventional hypothesis testing (see Rouder et al. 2009, for a discussion), allowing us to confidently draw conclusions in favour of the null hypothesis.

To conclude, the experiments we presented across Chapters 2 and 3 provide little evidence for arousal misattribution as a mechanism for the impact of negative utility on probability estimates. At this stage, champions of the SLH must provide more evidence for the existence of such an effect, both in the form of conceptual and direct replications.

On the other hand, whilst we reliably observed a negativity bias in Chapter 2, we were unable to demonstrate a concrete mechanism. Furthermore, we provided some preliminary evidence for a potential moderation by interoceptive sensibility, and across one of two scenarios, showed that social power moderated the negativity bias – but none of these moderations appeared to be due to greater interoceptive *ability*, as hypothesized

by the SLH. Thus, at this stage, we will turn to another mechanism which could potentially explain these moderations: loss function asymmetries.

Interoceptive sensibility could lead to a greater anticipation of feeling negative affect elicited by negative outcomes and could impact loss function asymmetries due to bracing for loss. Social power, due to increased control could impact the perception of decision-control and therefore loss function asymmetries. The next chapter will explore this moderation by social power and its relationship to loss function asymmetries in more detail.

Chapter 4 Social Power and Loss Function Asymmetries

Chapter Overview

Chapter 4 provides an investigation into the impact of social power on the relationship between negative utility and probability estimates. As outlined in Chapter 1, based on an account of loss function asymmetries, there is reason to believe that social power moderates the impact of negative utility on probability estimates due to the powerful's greater personal sense of control. Loss function asymmetries should only exist in situations where the individual has decision-control, and social power increases the amount of control individuals feel they have over themselves and their environment. As a result, in situations where decision-control is absent or ambiguous, individuals high in social power should perceive that they have more decision-control and loss function asymmetries should be present. However, the powerless should perceive that they have no decision-control, and loss function asymmetries should be absent. In the present chapter, we report 4 experiments examining the potential moderation by social power of the relationship between negative utility and probability estimates. Experiment 8 examined this moderation assessing power via a self-report measure, Experiment 9 manipulated power in the laboratory, and measured arousal and interoceptive accuracy. Experiment 10 assessed the moderation with real outcomes and examined the role of self-relevance. Finally, Experiment 11 gave participants real decision-control and incentivised accuracy.

Introduction

In order to provide the theoretical setting for the empirical work introduced in the current chapter, first we will show which assumptions of the ALF account have already been tested, and discuss how to further test this theory. As a potential moderator, we present the concept of a personal sense of control and demonstrate how control can impact the relationship between negative utility and loss function asymmetries. Next, we discuss the potential moderation of social power as a result of the greater control of the powerful, and present previous research on the impact of social power on risk-taking and optimism.

Testing Loss Function Asymmetries

As discussed in Chapter 1, an account of ALFs predicts that negative utility leads to loss function asymmetries, which in turn bias probability estimates (see Figure 4.1).

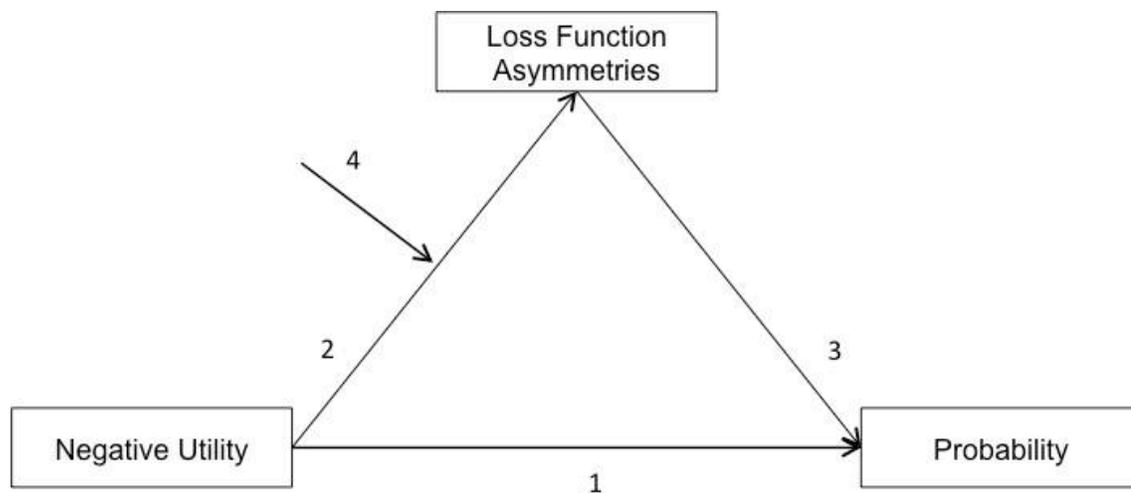


Figure 4.1. Mediation model of the impact of negative utility on probability estimates.

However, previous research testing the mediation by ALFs did not assess ALFs directly. Whilst it is questionable in how far self-reports would be suitable to measure ALFs as a mediator, one possibility could be to ask individuals “how bad would it be if

you overestimated the likelihood” as well as “how bad would it be if you underestimated the likelihood”, to assess path 2 more directly and to establish its causal role. Although potentially trivial, path 3, the impact of loss function asymmetry on probability estimates has not been assessed directly either (a simple demonstration would be to manipulate the pay-off of probability estimates depending on the type of error, e.g. to decrease payment more when participants over vs. underestimate the probability of a neutral event).

On the other hand, evidence in favour of this model came from the inclusion of a moderator (path 4). According to an account of ALFs, individuals should only overestimate the likelihood of negative events should there be the possibility to act based on the estimate, that is, if there is decision-control (moderator, path 4). This means that there should be no loss function asymmetry when there is no decision-control, when one cannot act based on one’s estimates. Based on this reasoning, it is possible to derive further moderators that impact the relationship between utility and loss function asymmetry. In particular, we propose that decision-control does not need to be integrated into the decision context, but can be manipulated as an individual difference variable, too. That is, the amount of control an individual feels over their ability to make an impact based on their estimate should equally moderate the relationship between utility and loss function asymmetry – whether or not decision-control is truly present. In order to illustrate this point in more detail, below we discuss different conceptualisations of personal and interpersonal control and how they could moderate the impact of negative utility on probability estimates.

Dimensions of Control

There are several types of control that could potentially moderate the relationship between negative utility and probability estimates (see Table 4.1). One potential moderator of the effect is the locus of control (Rotter, 1954), a dimension of core self-evaluations (the other three dimensions being neuroticism, self-efficacy and self-esteem, Judge, Locke, & Durham, 1997). According to this theory, individuals can be classified as having either an internal or an external locus of control. Individuals with an internal locus of control believe that they have control over their own life, whereas individuals with an external locus of control believe that forces outside their own control determine their life's outcomes. Following an account of ALFs, individuals with an internal locus of control should demonstrate a greater negativity bias than individuals with an external locus of control, since they should feel an increased capability to actually act on their estimates.

In addition, the literature on control also distinguishes between primary and secondary control when it comes to control over future events (see the two-process *primary-secondary control model*, Rothbaum, Weisz, & Snyder, 1982; Weisz, Rothbaum, & Blackburn, 1984). Primary control constitutes the ability to control external factors in order to change outcomes in a way that they increase reward or decrease a negative outcome. Secondary control is defined as the ability to control and adapt oneself (e.g. one's hopes, attributions etc.) to the outcomes. Both types of control could lead to an inflation of probability estimates in the negative domain. Primary control could lead to a negativity bias to motivate the process of changing the circumstances in order to prevent the outcome from happening. Secondary control could

lead to a negativity bias through a greater bracing for loss (see Shepperd et al. 2000) – the latter having the potential to moderate also objectively uncontrollable events.

Table 4.1

Different types of control and their definitions.

Type of Control	Definition	Reference
<i>Internal Locus of Control</i>	Belief that one's own ability, effort, or actions determine what happens	Rotter (1954)
<i>External Locus of Control</i>	Belief that fate, luck, or outside forces are responsible for what happens	Rotter (1954)
<i>Primary Control</i>	Attempts to change the world so that it fits with the self's needs	Rothbaum, Weisz, & Snyder, 1982, p.8
<i>Secondary Control</i>	Attempts to fit in the world and to "flow with the current"	Rothbaum, Weisz, & Snyder, 1982, p.8
<i>Illusion of Control</i>	People's belief that they have influence over the outcome of uncontrollable events	Langer, 1975

An important potential moderator of the effect is the illusion of control (IOC, Langer, 1975). The IOC is defined as "people's belief that they have influence over the outcome of uncontrollable events" (Montier, 2007). In situations with few to no elements of decision-control, an illusion of control should lead to individuals providing higher probability estimates for negative events, since they feel that the event is more controllable than it actually is. Thus, for this type of control to potentially moderate the negativity bias, decision-control might not need to be a necessary characteristic of the event in question. Instead, an IOC should impact the *perception* of how much decision-

control an individual has in a given situation, and therefore moderate the relationship between negative utility and probability estimates.

Furthermore, control can also be *over others* rather than *over oneself*. Individuals that feel like they have control over others might be more convinced that others actually act based on their estimate. Thus, in situations where a probability estimate is communicated, and where another individual has the chance to act based on one's estimate (e.g. Harris et al., 2009), individuals that have control over others might similarly provide higher estimates for negative events. The reasoning here is that when one feels control over another individual's actions, one provides higher estimates in order to make the individual act in a way that could prevent the negative outcome from happening.

Summing up, an ALF account predicts several ways through which different types of control can moderate the impact of negative utility on probability estimates. An internal compared to an external locus of control, an illusion of control, primary and secondary control, as well as control over others when the probability estimate is communicated, could all potentially lead to a moderation of the mediation by loss function asymmetries. One variable that is strongly related to different dimensions of control and therefore can be hypothesized to moderate the impact of negative utility on probability estimates via a mediation of control is social power, as shall be discussed below.

Social Power and Control. Control is a central social construct, and the amount of control that individuals feel they have is heavily determined by their social roles. For instance, a high socio-economic status (Lachman & Weaver, 1998) and being part of a

majority group (Guinote, Brown, & Fiske, 2006) have been associated with reporting greater control over future outcomes. By definition, social power is associated with an asymmetry in the control over valued resources, in that powerful individuals can determine whether powerless individuals will gain access to these resources (e.g. Emerson, 1962; Fiske, 1993; Keltner et al. 2003; Magee & Galinsky, 2008). Thus, individuals high in power are by definition high in interpersonal control. Furthermore, from the definition of social power, it can be inferred that individuals in powerful positions are more likely to believe that others will act based on their advice. We mentioned previously that Harris et.al. (2009) tested the ALF account in scenarios that included an element of persuasion. Since powerful individuals are often in a position to advise the powerless about desired courses of action (imagine a manager giving orders to a subordinate), it is reasonable to assume that the powerful will feel more decision-control than the powerless in situations where risk information is communicated.

Furthermore, power can also lead to intrapersonal control. For instance, power and powerlessness have been associated with an internal versus external locus of control, respectively. Seeman (1963) describes the locus of control as a range of powerlessness (from external – high powerlessness to internal – low powerlessness), and Riskind (1984) showed that an upright versus slumped body posture (associated with power and powerlessness, respectively, e.g. Carnery, Cuddy & Yap, 2010) led to changes in locus of control. Moreover, power has been shown to increase the IOC. For instance, individuals high in power are more likely to take another card in a game of blackjack, a behaviour associated with an IOC (Galinsky, Gruenfeld, & Magee, 2003). In addition, Fast et al. (2009) demonstrated that powerful individuals indicated that they

were more likely to vote compared to powerless individuals and that this relationship was mediated by the heightened sense of an IOC in the powerful, assessed here as the perceived control over hard-to-control outcomes (“to what extent are you able to have some control over what happens in the economy?”). Whilst these studies highlight the IOC in the intrapersonal domain, it could also be that a greater IOC could be expressed in the interpersonal domain, in that powerful individuals feel they can control and influence individuals even if this is not the case.

Summing up, there are unambiguous links between social power and a personal sense of control. Powerful individuals are both more likely to feel a greater sense of control over their own actions, as well as over other individuals’ actions. As a result of this greater control, the powerful should demonstrate a greater negativity bias than the powerless both for estimates that concern themselves, as well as for estimates where the outcome concerns someone else’s action.

Seemingly inconsistent with these predictions is, however, a body of research that has demonstrated greater risk-taking and optimism with increased social power. If the powerful would indeed be more risk-taking and optimistic, the opposite finding to the results anticipated from an ALF account would be expected: Powerful individuals should then show a decreased negativity bias compared to powerless individuals. Next, we will review and discuss the evidence linking social power to risk-taking and optimism, highlighting also the significance of the current research for mechanisms that underlie the impact of social power on cognition, behaviour and affect.

Social Power, Risk, and Optimism

Social power has been famously associated with greater approach motivation, whilst powerlessness has been associated with avoidance motivation (Keltner, Gruenfeld & Anderson, 2003). It has been argued that it is this greater approach motivation that leads to greater optimism and risk taking behaviour of the powerful (e.g. Anderson & Galinsky, 2006). In particular, it has been proposed that the asymmetric focus on rewards over punishments of powerful individuals increases risk-taking behaviour, whereas the stronger focus on punishments over rewards might decrease risk-taking behaviour in powerless individuals.

The most extensive empirical investigation into the impact of social power on risk-taking and perception of risks came from Anderson and Galinsky (2006). In their Experiment 1, Anderson and Galinsky (2006) showed that powerful individuals reported greater optimism for positive and negative future life events, as measured with the comparative optimism method (Weinstein, 1980). However, as discussed in Chapter 1, the comparative optimism method is subject to statistical artefacts, and Harris and Hahn (2009) were able to demonstrate optimism with perfectly rational agents. Whilst it is therefore possible to say that the powerful expect a better future than the powerless as they score generally lower in this task, in light of Harris and Hahn's (2009) results is not possible to say whether the powerful are actually more optimistic than the powerless. Indeed, one can imagine that the powerful have *actually* better futures (i.e. greater access to resources), in which case the pattern would demonstrate realism rather than an optimistic bias.

Furthermore, in Experiment 2 by the same authors, powerful individuals reported lower perceived risk, as measured by Johnson and Tversky's (1983) measure of risk perception. In this paradigm, participants are told that "about 50 000 in the United States die in motor vehicle accidents per year", and are asked afterwards to estimate the number of annual fatalities due to another 17 causes of death. Powerful participants provided lower estimates than powerless participants. However, unless evidence is provided that there are no differences between powerful and powerless individuals for estimates of neutral events, the results can also be attributed to other factors such as a greater propensity for powerless individuals to be influenced by anchors (such as the base-rate of 50 000 deaths by motor vehicle accidents, which was presented first to participants).

Next, Anderson and Galinsky (2006) seemingly demonstrated that individuals high in social power perceive the probability of a gain as greater, and of a loss as lower, than individuals low in social power. Participants were presented with an adaptation of Tversky and Kahneman's (1981) Asian Disease Problem. Participants read "imagine that you work for a large car manufacturer that has recently been hit with a number of economic difficulties. It appears as if three plants need to be closed and 6000 employees laid off. As vice president of production, you have been exploring alternative ways to avoid this crisis." Further, participants in the gain frame read:

Plan A will save one of the three plants and 2000 jobs.

Plan B has a 1/3 probability of saving all three plants and all 6000 jobs, but has a 2/3 probability of saving no plants and no jobs.

Participants in the loss frame read instead:

Plan A will result in the loss of two of the three plants and 4000 jobs.

Plan B has a 1/3 probability of losing no plants and no jobs, but has a 2/3 probability of resulting in the loss of all three plants and 6000 jobs.

Participants answered on a scale from 1 (very much prefer program A) to 6 (very much prefer program B). Powerful participants were more likely to prefer the risky option across both frames than powerless or power-neutral control participants (see Figure 4.2). Anderson and Galinsky (2006) interpret these results as an indication that powerful participants estimate the probability of a gain as more likely, and the probability of a loss as less likely compared to powerless and control participants.

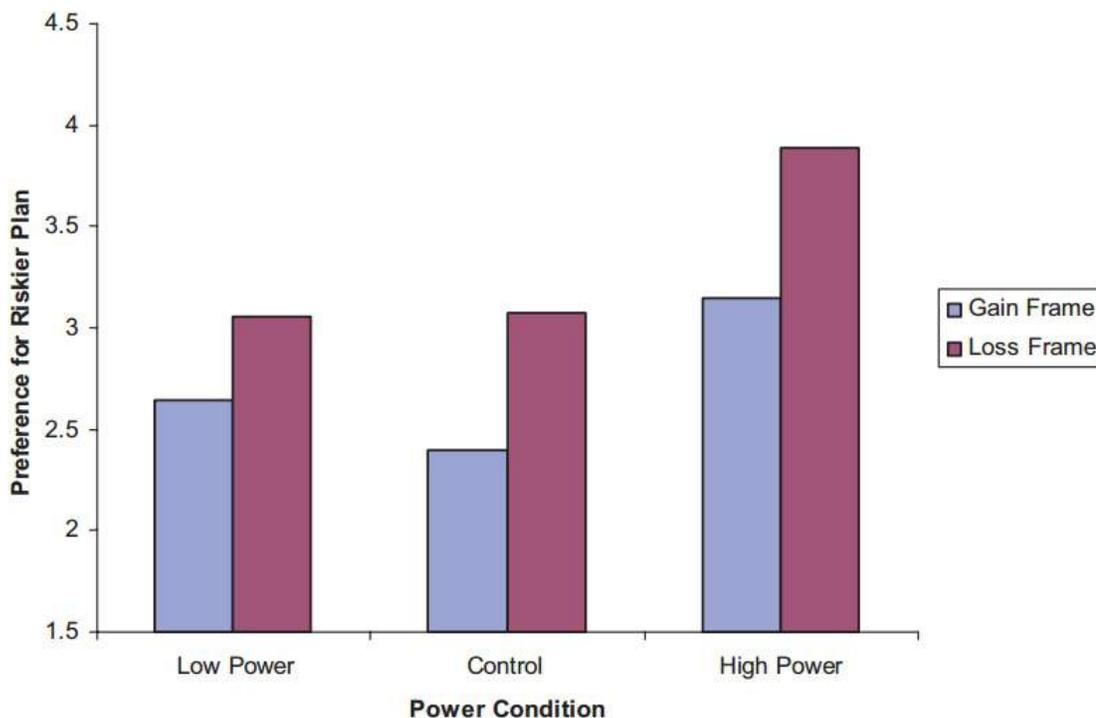


Figure 4.2. Preference for a riskier plan by power and frame condition in Anderson and Galinsky's (2006) Experiment 3. Figure reproduced with permission from the author.

On the other hand, the particular scenario (in which participants had to imagine that they were a vice president) might have influenced the results. In particular, the power prime might have led to a greater fit between the scenario and participants for those primed with greater power. In this particular scenario, powerful participants might have thought more about the fact that losing some plants is the same as losing all plants for their reputation as a vice president. This interpretation is in line with research by Jordan, Sivanathan and Galinsky (2011), who showed that with view on the stability of their social hierarchy, unstable powerful individuals are more risk-taking than stable powerful individuals.

Alternatively, however, and potentially in line with the original conceptualisation of such framing effects, powerful participants could have actually demonstrated greater loss avoidance than powerless and control participants. In the classic administration of the framing paradigm participants have to choose one option over the other (e.g. Tversky & Kahneman, 1981), rather than to indicate their preferences on a Likert-scale. Indeed, participants high in power showed only a preference for the riskier option in the loss frame (where the group mean is above the mid-point of 3.5 of the scale). This leaves room for the possibility that the responses of high power participants are in fact more coherent with the original findings of Tversky and Kahneman (1981) than the responses of low power and control participants: in the original findings, participants showed a preference for program A in the gain, and for program B in the loss frame.

In their Experiment 4, Anderson and Galinsky (2006) tested participants' proposed willingness to take risks by providing participants with a vignette describing a sexual encounter that could end in having unprotected sex:

“Imagine that you are single and that you run into a very attractive acquaintance while ordering a drink at the bar. The two of you begin to talk, and both of you find the conversation very enjoyable. She (or he) has a good sense of humor, and seems genuinely interested in what you are saying. It is clear that there is definite chemistry between you and that you are interested in this person. You continue to spend time together throughout the night. When the bar closes, she (he) offers to walk you home. When you get home, she (he) kisses you goodnight at the door. You decide to go inside and talk for a while. After talking, you and she (he) begin to make out on the couch. Things progress and you realize that you are both very interested in having sex with each other. She is (you are) on the pill, but neither of you have a condom. You discuss the possibility of going to a store, but there is not one nearby. You awkwardly discuss your sexual history, and she (he) tells you that she (he) does not sleep around.”

Afterwards, participants were asked to indicate the likelihood that they were to engage in unprotected sex, and rated the perceived risks of having sex without a condom. Anderson and Galinsky (2006) showed that powerful participants were more likely to indicate that they would engage in unprotected sex, and this result was mediated by their decreased perceptions of the likelihood of a negative outcome associated with unprotected sex. Thus, this experiment would constitute direct evidence for the hypothesis that powerful individuals provide lower estimates of negative outcomes than powerless individuals.

Potentially in conflict with this conclusion is the notion that the vignette used in this experiment most likely led participants to adopt a “hot visceral state” of sexual arousal. Moeini-Jazani, Guinote and Warlop (2014) demonstrated that powerful

participants were more reward oriented than powerless participants *only* when in a *hot cognitive state*. Therefore, it is possible that this particular vignette played a role in showing that powerful individuals are more risk seeking and have lower perceptions of risk, as they are more reward oriented when under a hot visceral state compared to powerless and control participants. The generalizability of this finding is therefore questionable.

Another study by Fast et al. (2009) demonstrated that an IOC underlies powerful individuals' greater propensity to provide greater estimates for positive events. In their experiment, after manipulating social power in a role-play scenario where participants anticipated playing the role of a manager (high power) or a subordinate (low power), participants were asked to imagine working for a marketing agency, and were given some background about the company. Next, optimism was assessed by asking participants to estimate their organisation's performance ("how likely is your agency to increase its profitability in the next two years?"). Individuals primed with a managerial mind-set indicated a greater likelihood of the agency to increase profitability, and this relationship was mediated by the greater control of powerful participants (assessed by asking participants to indicate the likelihood that they can control the agency's future outcomes). These findings might, however, also be explained by a motivational account: For example, the more workers are involved in decisions in companies (and arguably, the more control they possess), the more engaged they are with their organisation, and the greater the identification (e.g. Konrad, 2006). Thus, managers might identify more with their organisation. Research in team reasoning has demonstrated effects whereby greater likelihoods are reported for the success of one's in-group (Blake & Mouton,

1961; Sherif & Sherif, 1956), for example to protect the group's positive self-image (Jourdan & Heath, 1996). Therefore, motivational factors other than optimism could account for the effects demonstrated in Fast et al. (2009).

Carney et al. (2010) manipulated power by asking participants to engage either in an expansive (high power) or a contractive (low power) posture. Following these manipulations, participants were given \$2 and asked if they wanted to either keep the money, or to roll a dice and risk losing the money for a payoff of \$4. The odds constituted a 50/50 gamble. However, 86.36% of individuals that engaged in expansive postures went for the risky option, compared to only 60% of the individuals that engaged in contractive postures. Whilst this experiment does not address the question of likelihoods of the occurrence of negative or positive events, it shows that powerful individuals, under some circumstances, are more willing to accept a risky option for a higher payoff.

Summing up, research in the domain of power and risk-taking behaviour has provided some support for the notion that powerful individuals are willing to take risks when they have the possibility of gaining a reward. However, the discussed studies have mainly concerned choice or action-taking paradigms rather than probability estimates. The results in these paradigms could be explained for example by outcome salience, in that powerful individuals live in reward-rich environments and consequently the positive outcome is more salient than for the powerless. As discussed in Chapter 1, outcome salience impacts the evidence accumulation stage of the process of forming a probability estimate. Therefore, it is not necessarily a cognitive bias at the stage of the internal

estimate, but rather differences in evidence selection and accumulation, which might have lead to the discussed results.

The present research aims to systematically investigate the impact of social power on the impact of negative utility on probability estimates, with objective representations of probabilities in order to keep the evidence accumulation stage constant. An account of ALFs (Weber, 1994) predicts that, due to increased control, powerful individuals should show a greater negativity bias than powerless individuals. On the other hand, as discussed in Chapter 1, powerful individuals have been shown to be more approach, and powerless individuals more avoidance oriented (Keltner et al. 2003). Following this most dominant framework of social power, the powerless should be more impacted by negative utility compared to the powerful. Therefore, not only does this research provide an important investigation into the predictions following an account of ALFs, it can also provide a potential boundary condition to the approach-avoidance theory of power.

Experiment 8

Background

The purpose of Experiment 8 was to assess the moderation by social power on the impact of negative utility on probability estimates. Experiment 3 provided some evidence for the notion that powerful individuals show a greater negativity bias than powerless individuals in one out of the two presented scenarios, but this evidence was overall not very strong. However, before power was assessed in Experiment 3, participants completed a task measuring interoceptive accuracy, and completed a battery

of questionnaires. Thus, we sought to replicate the results from Experiment 3 in a simpler experimental set-up with fewer potential confounds.

Experiments examining effects of social power typically either assess social power as an individual difference variable, or manipulate social power. In this instance, social power was assessed by means of a self-report questionnaire measure (Anderson, John, & Keltner, 2012). Whilst social power is a variable defined through social relationships (Emerson, 1962) and conceptualised through greater control over resources in social contexts, social power can also be regarded as a “psychological property of the individual” (Anderson & Galinsky, 2006, p. 514, see also Anderson, John, & Keltner, 2012; Bargh, Raymond, Pryor, & Strack, 1995; Galinsky et al. 2003). The so-called “personal sense of power” is defined as “the perception of one’s ability to influence another person or other people” (Anderson, John, & Keltner, 2012, p. 316), which is often, but not always, related to one’s actual ability to influence others (Anderson et al. 2012).

The sense of power scale is well validated (Anderson et al. 2012) and, importantly, many experiments report identical effects of manipulated and measured social power (see Anderson & Galinsky, 2006; Chen, Langner & Mendoza-Denton, 2009; Galinsky, Magee & Inesi, 2006; Karremans & Smith, 2010; van Kleef et al. 2008; van Kleef et al. 2006; Lammers & Stapel, 2010).

In the current experiment, participants first completed the sense of power questionnaire (Anderson et al. 2012), and afterwards completed the dice and RAF scenarios from Experiments 1 and 2.

Method

Participants and Design. After excluding 3 participants following previous criteria, the final sample consisted of 197 participants (128 female, median age =23), who completed the experiment online for the chance of winning £25. Participants were randomly assigned to a between participants design with one factor (utility: negative/neutral).

Materials and Procedure. Participants were informed that they would take part in three brief experiments. For the first study, participants were told that we are interested in their personality in interpersonal relationships. Subsequently, participants completed the generalised sense of power scale (SOP, Anderson, John & Keltner, 2012). This scale consists of eight items assessing participants' beliefs about the power they have in relationships with others (e.g. "In my relationships with others, I think I have a great deal of power"), which participants rate on a scale from 1 ("strongly disagree") to 7 ("strongly agree"). Following, participants completed the dice and RAF scenarios as in Experiments 1 and 2, introduced as an experiment interested in risk taking behaviour. Afterwards, participants also completed the CBP scales, but the analysis for this scale will not be reported here²⁸. Finally, participants were thanked and debriefed.

Results

Manipulation Checks.

RAF Utility. The answer to the question "Thinking back about the RAF risk estimation task you just did: how bad would it be if debris would be dropped?"

(answered on a scale from 1, not at all bad, to 7, very bad) was submitted to a between

²⁸ Note that the moderation by CBP supported our previous findings, where only individuals high in CBP (interoceptive sensibility) showed a negativity bias.

participants t-test. Utility was manipulated successfully, with participants in the negative condition indicating the outcome to be worse ($M=5.92$, $SD=1.35$) compared to participants in the neutral condition ($M=3.06$, $SD=1.59$), $t(196)=13.59$, $p<.001$. To assess the moderating role of sense of power (SOP), a median split was performed on the SOP scale²⁹, and SOP was entered together with utility in a 2 (SOP: high/low) x 2 (utility: high/low) between participants ANOVA. SOP did not impact the RAF utility manipulation check, and no interaction between SOP and utility was found, all $F_s<1$.

Dice Utility. The manipulation of severity for the dice example was successful, with those in the negative condition indicating a worse outcome ($M=3.94$, $SD=1.85$) compared to those in the neutral condition ($M=1.3$, $SD=.77$), $t(129)=13.15$, $p<.001$ (df adjusted due to inequality of variances). Including SOP as a factor did not show a significant main effect, $F(1,194)=1.69$, $p=.2$, $\eta^2=.01$, or interaction, $F<1$.

Thus, there was no indication that powerful individuals perceive the negative outcome as worse than powerless individuals in any of the two scenarios. This is important for the same reasons as mentioned in Chapter 2: If the powerful perceive negative outcomes as worse than the powerless, this could explain the hypothesized moderation, as probability estimates have been shown to increase with more negative utility (e.g. Harris et al. 2009).

Aggregated Probability Estimates. A 2 (utility: negative/neutral) x 2 (SOP: high/low) x 2 (scenario: RAF/dice) mixed ANOVA with repeated measures on the last factor showed no interactions which included the within-subjects variable, all $p_s>.35$.

We therefore collapsed across the standardized means of the two scenarios. A 2 (utility:

²⁹ This median split was performed in order to have comparable analyses as in Chapter 2. Furthermore, a median split enabled us to analyse orthogonal contrasts and therefore specific hypotheses.

negative/neutral) x 2 (SOP: high/low) ANOVA conducted on the aggregated scores showed a marginally significant effect of utility, $F(1,186)=3.8$, $p=.05$, $\eta^2=.02$ ³⁰. Participants in the negative condition provided greater estimates ($M=.11$, $SD=.58$, z-scored) compared to participants in the neutral condition ($M=-.05$, $SD=.53$, z-scored). However, there was neither a main effect of power, $F(1,186)=1.07$, $p=.30$, $\eta^2=.01$, nor a power*utility interaction, $F<1$.

As in Chapter 2, we next conducted orthogonal contrasts testing specific hypotheses. Following an ALF account, we predict that participants high in social power estimating negative scenarios should provide greater probability estimates than the other three conditions combined (contrast 1, see Table 4.2). We do not expect powerless individuals estimating negative events to be significantly different to powerful and powerless participants estimating the probability of neutral events (contrast 2). Last, there is no reason to expect any differences between powerful and powerless individuals when estimating neutral events (contrast 3).

Table 4.2

Orthogonal contrasts conducted in Experiment 8.

	Powerful Negative	Powerless Negative	Powerful Neutral	Powerless Neutral
Contrast 1	3	-1	-1	-1
Contrast 2	0	2	-1	-1
Contrast 3	0	0	1	-1

³⁰ Before the removal of outliers, the main effect of utility was not significant, $F(1,196)=1.78$, $p=.18$, $\eta^2=.01$.

Conducting these orthogonal contrasts on the aggregated DV showed that participants high in power answering a negative scenario gave significantly higher probability estimates compared to the other three conditions combined, $t(186)=2.06$, $p=.04$. No other contrasts were found significant, all $ps>.35$ (see Figure 4.3). An omnibus test assessing whether there was any rest-variance left unaccounted for was non-significant, $F<1$.

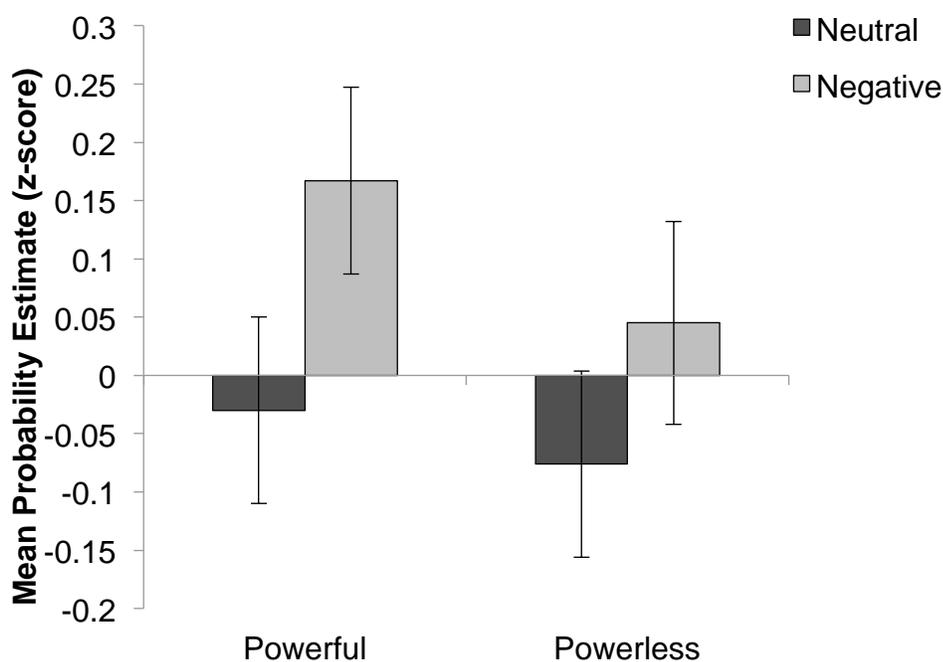


Figure 4.3. Probability estimate means (z -scored) for powerful and powerless participants answering neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

However, upon closer examination of the separate scenarios, it became evident that the results do not support our hypotheses throughout both scenarios. Thus, in order to be open about our results, we decided to also report the results for the two scenarios separately.

Dice Probability Estimates. A 2 (utility: negative/neutral) x 2 (SOP: high/low) ANOVA conducted on the dice scenario showed no main effects of utility, $F(1,188)=1.39, p=.24, \eta^2=.01$, power, $F<1$, and no power*utility interaction, $F<1$. Furthermore, conducting the same set of orthogonal contrasts as above on the dice scenario, none of the three contrasts were significant (all $ps>.13$). Therefore, the results of the dice scenario did not support our hypotheses, and shows directionally even a trend that the powerless provide a greater negativity bias than the powerful (see Figure 4.4), however, these differences did not reach significance.

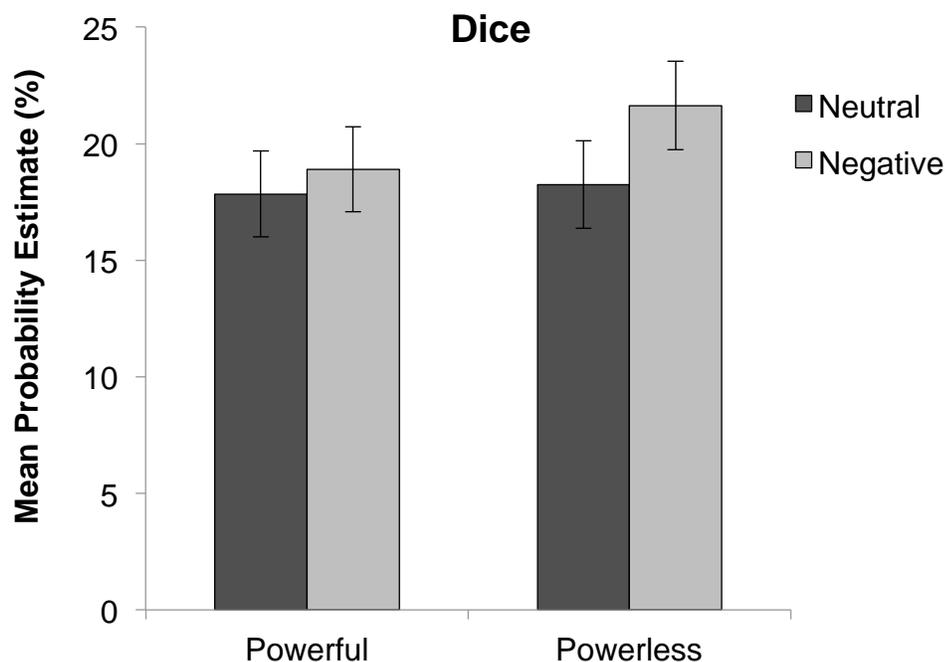


Figure 4.4. Probability estimate means for the dice scenario for powerful and powerless participants answering neutral and negative conditions. Error bars represent +/- 1 standard error of the mean.

RAF Probability Estimates. A 2 (utility: negative/neutral) x 2 (SOP: high/low) ANOVA showed no significant main effect of power, $F(1,190)=1.61, p=.21, \eta^2=.01$ ³¹, or of utility, $F<1$. However, there was a significant power*utility interaction, $F(1,190)=4.15, p=.04, \eta^2=.02$. Simple effects showed that powerful participants provided greater estimates than powerless participants in the negative condition, $F(1,190)=5.44, p=.02, \eta^2=.03$, whilst there was no difference in the neutral condition, $F<1$. Moreover, the powerful provided marginally greater estimates in the negative compared to the neutral condition, $F(1,190)=3.65, p=.06, \eta^2=.02$, whilst this was not the case for the powerless $F<1$ (see Figure 4.5).

Furthermore, conducting orthogonal contrasts revealed that the predicted contrast 1 was significant, $t(190)=2.34, p=.02$ (see Figure 4.5). Neither contrast 2, $t(190)=.84, p=.40$, nor contrast 3, $t(190)=.64, p=.52$, were significant, and no rest-variance was unaccounted for, $F<1$. Thus, only the RAF, but not the dice scenario appeared to be in line with the current hypotheses.

³¹ Before the removal of outliers, the main effect of power was significant, $F(1,196)=4.08, p=.045, \eta^2=.02$. Powerful participants provided overall greater probability estimates ($M=48.92, SD=10.54$) than powerless participants ($M=45.87$).

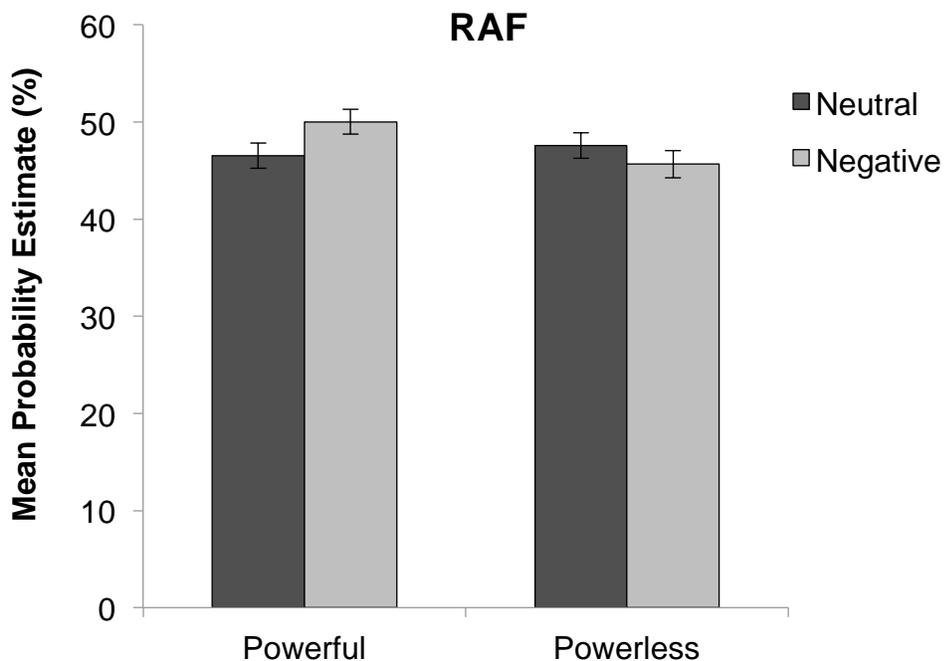


Figure 4.5. Probability estimate means for the RAF scenario for powerful and powerless participants answering neutral and negative conditions. Error bars represent ± 1 standard error of the mean.

Discussion

Experiment 8 provided preliminary support for the notion that individuals high in social power show a greater negativity bias than the powerless when averaging across the two scenarios. However, whilst this hypothesis received support in the RAF scenario, the dice scenario failed to show any differences between the groups. In the dice scenario, directionally, individuals low in social power showed a negativity bias, which could be due to the fact that decision-control in a scenario involving a friend is arguably relatively high. If it is the perception of decision-control that is underlying the impact of social power on the relationship between negative utility and probability estimates, providing decision-control to the powerless should increase their negativity bias to the level of the powerful. Only in situations where decision-control is more

ambiguous would we expect to see a negativity bias. Experiment 9 will employ a modified version of the dice scenario, where the ambiguity of decision-control is increased.

Experiment 9

Background³²

Experiments 3 and 8 provided some first evidence that sense of power, as measured with a questionnaire, moderates the impact of negative utility on probability estimates. Next, we sought to examine whether manipulated, rather than measured, social power moderates the relationship, to establish a causal rather than a correlational link. Furthermore, Experiment 9 included a control group to determine the direction of potential effects.

In the current experiment, power was manipulated by means of a role-play scenario (Guinote, Judd, & Brauer, 2002; Guinote, 2008). Here, participants expect to take part in a creative task together with another participant. Ostensibly, the results of a personality questionnaire that participants completed as a pre-test several weeks before the start of the laboratory experiment assigned participants either to the role of the “manager” (high power) the “subordinate” (low power) or the “co-worker” (control) (the assignment to the conditions was, in fact, random). In this task, participants are told that managers evaluate the subordinates’ solutions in the creativity task, and determine

³² The design of this experiment was developed in collaboration with Dr. Mehrad Moeini-Jazani, Dr. Klemens Knöpferle, Dr. Elia Gatti and Prof. Luk Warlop, who I wish to sincerely thank for their invaluable help and input. The experiment was supported by a grant from the Department of Marketing, Business Institute Norway. The analysis of the main effect of social power on interoceptive accuracy will be submitted as “Moeini-Jazani, Knöpferle, de Molière, Gatti, & Warlop” and is reported here with permission of the first author.

how much money the subordinates receive based on their performance, whilst managers receive a fixed amount. No power-difference is mentioned to co-workers.

Furthermore, conducting a laboratory experiment offered once more the chance to examine the roles of interoceptive accuracy and arousal as assessed by means of skin conductance responses. Whilst Experiment 3 did not provide evidence for the role of either, we sought to examine these variables to provide stronger evidence for the suspected null results. However, this time we assessed arousal at the time when participants actually physically provided their estimate, as opposed to Experiment 3, where arousal was measured whilst participants read the whole scenario. We reasoned that this shorter time window could potentially give a more precise estimate of the arousal induced by negative utility.

Moreover, the nature of the manipulation of social power, whereby participants were sent a questionnaire battery a few weeks before the completion of the laboratory experiment allowed us to once more assess the moderation by measured, as well as by manipulated social power. In this questionnaire battery we included the sense of power scale (Anderson et al. 2012), aiming to replicate our previous findings.

In addition to the sense of power scale, we included a questionnaire measuring approach and avoidance motivation. As discussed in Chapter 1, following the “compatibility-incompatibility” account (Rose, 2009), individuals high in approach motivation should provide greater probability estimates for the occurrence of positive events, whereas individuals high in avoidance motivation should provide greater estimates for the occurrence of negative events. However, powerful individuals have been proposed to be more approach -, and powerless individuals to be more avoidance-

oriented (Keltner et al. 2003). Therefore, our previous findings would not support the notion that differences in approach and avoidance motivation of powerful and powerless individuals underlie the observed effects.

However, Lench (2009) proposes that when individuals are high in avoidance motivation, they should *avoid* negative outcomes, and therefore provide *lower* estimates. Thus far, we cannot make any claims about the direction of the impact of power on probability estimates for negative events – it is unclear, whether the powerful provide greater estimates than a power neutral control group, or whether the powerless provide lower estimates. Therefore, it remains possible that the greater avoidance orientation of powerless individuals could underlie the findings from Experiments 3 and 8. Speculatively, according to an ALF account, approach motivation could also lead to greater probability estimates for negative events, as individuals high in approach orientation are action oriented (Elliot, 2006), which might lead to greater beliefs about the possibility to act based on their estimates.

In sum, the current experiment tested for a potential moderation by manipulated social power, and assess whether this moderation is due to greater interoceptive awareness of the powerful, whilst also measuring the mediating influence of arousal. Secondly, this experiment allows us to yet again assess a moderation by trait social power and by trait approach-avoidance motivation, both assessed before participants entered the laboratory.

Method

Participants and Design. 142 participants from the Business Institute Norway were recruited for course credits or a payment of 180 NOK (approximately £18). Four

participants were excluded for not completing the experiment or for not understanding the tasks due to insufficient knowledge of English. The final sample consisted thus of 138 participants (89 female, median age=23)³³. Participants were randomly assigned to a 2 (utility: negative/neutral) x 3 (power: powerful/control/powerless) between participants design.

Materials and Procedure. Participants were informed that the experiment concerned the impact of creativity styles on collaborative problem solving. Before participants came to the laboratory, they completed a questionnaire battery that included the sense of power scale and the BIS-BAS scale to measure approach and avoidance motivation (Carver & White, 1994, see Appendix H) (as well as further scales which shall not be reported here, as they are not relevant to the current research questions). Upon arrival in the laboratory, participants were told that their teamwork partner was still engaged in another experiment, and asked whether it would be possible for them to help with the calibration of physiological sensors until their partner arrived. All participants agreed to help, and were connected to physiological sensors on their non-dominant hand. In particular, skin conductance was measured by an EDA sensor (attached to second and fourth finger of the non-dominant hand), and heart rate was measured by a photoplethysmograph³⁴ placed on the third finger³⁵. A two-minute baseline measure of skin conductance was obtained. Next, power was manipulated by informing participants that the completion of their pre-test assigned them to a particular role in the teamwork task, and that we would already present them with their task so that

³³ Only 126 participants had pre-test data available, thus accounting for differences in degrees of freedom depending on the analysis.

³⁴ Sensors were purchased from Thought Technology Ltd., Montreal, CA

³⁵ Respiration rate and skin temperature was measured, but are not analysed for the purpose of this experiment.

we can start the teamwork exercise quickly as soon as their partner was ready. This power manipulation was adapted from Guinote, Judd and Brauer (2002, see also Guinote, 2008). Participants were told that they would either have the role of the manager (high power), the subordinate (low power), or the colleague (power-neutral control condition). Participants in the manager condition were told that their role would be to manage and supervise the subordinate during the creativity task, as well as to determine the subordinate's reward upon completion of the task, whilst they would be paid a fixed amount. Subordinates were informed that they would be responsible for coming up with solutions during the creativity task, that their performance was evaluated by their manager, and that their reward depended upon the manager's evaluation. Participants in the control condition were simply told that they would work together with another colleague on the task, and that they would both be paid equally for completing the exercise.

Afterwards, participants completed the heartbeat counting task (Schandry, 1981). The set-up was identical to the heartbeat counting task in Experiment 3. Participants first counted their heartbeats over three randomly varying time intervals, then counted seconds for three intervals, and again heartbeats for three intervals (25s, 35s, and 45s, randomised).

Following the heartbeat counting task, participants completed a dice scenario that was slightly adapted from previous versions in order to increase decision-control ambiguity by changing the person proposing the game from a friend to an unlikeable stranger.

In the negative condition participants read:

*Imagine you walk down the street and find \$100. You picked it up, so the \$100 is in your pocket. However, a very rich, arrogant and rude person was bending down just as you picked up the \$100. Seeing a dice for sale in the window of a nearby thrift shop, they propose the following:
You will roll the regular, six-sided dice 4 times. If a 6 comes up on at least 2 of these throws, the rich person will get the \$100 and you will lose the \$100. Otherwise you can keep it.
What do you think is the chance that a 6 would turn up on at least 2 out of 4 throws, so you would lose the money to the very rich, arrogant and rude person?*

In the neutral scenario, another person proposes the same game. There is no reference to money or any outcome involved, and the dice roll has no consequences.

Next, participants completed the power and utility manipulation checks.

Furthermore, as in Experiment 3, participants completed the PANAS (Watson et al. 1988), the STAI (Marteau & Bekker, 1992), and answered how many days a week they engage in physical activity, as these variables have been shown to be associated with performance in the heart beat counting task (Ehlers & Breuer, 1992).

Results

Manipulation Checks.

Power. The difference of two questions “how much control do you have over the other participant” and “how much control does the other participant over you” was submitted to a one-way ANOVA. Power was manipulated successfully, $F(2,135)=23.64$, $p<.001$, $\eta^2=.26$, with powerful participants reporting to have more power over the other participant than the other participant over them ($M=1.35$, $SD=2.08$), followed by power-neutral control participants ($M=.78$, $SD=1.37$) and powerless participants who reported that they have less power over the other participant than the other participant over them

($M=-1.42$, $SD=2.54$). Utility did not affect the power manipulation, $F<1$. This rules out that participants felt more powerless after the negative scenario.

Utility. The answer to the question “how bad would it be if at least 2 times a 6 would be rolled” was submitted to an independent samples t-test. Utility was successfully manipulated, with participants in the negative condition indicating that the outcome would be worse ($M=4.08$, $SD=2.06$) than participants in the neutral condition ($M=2.82$, $SD=1.75$), $F(1,136)=14.92$, $p<.001$, $\eta^2=.10$. Power did not affect the utility manipulation check, $F<1$, demonstrating that participants did not perceive the scenario as worse depending on their power condition.

Dice Probability Estimates.

Manipulated Power. Submitting the dice-scenario to a 2 (utility: negative/neutral) x 3 (manipulated power: powerful/control/powerless) between subjects ANOVA, there was a marginally significant main effect of utility, $F(1,126)=3.53$, $p=.06$, $\eta^2=.03$. Participants in the negative condition provided overall greater probability estimates ($M=27.89$, $SD=18.99$) than participants in the neutral condition ($M=22.72$, $SD=15.42$). The main effect of power was non-significant, $F<1$. However, there was a marginally significant utility*power interaction, $F(2,126)=2.59$, $p=.08$, $\eta^2=.04$ (see Figure 4.6). Simple effects showed that powerful participants provided greater estimates for the negative compared to the neutral condition, $F(1,126)=8.1$, $p<.01$, $\eta^2=.06$, whilst this was not the case for powerless or control participants, all $F_s<1$.³⁶

³⁶ Before removing outliers, no main effects or interactions were found, all $p_s>.23$.

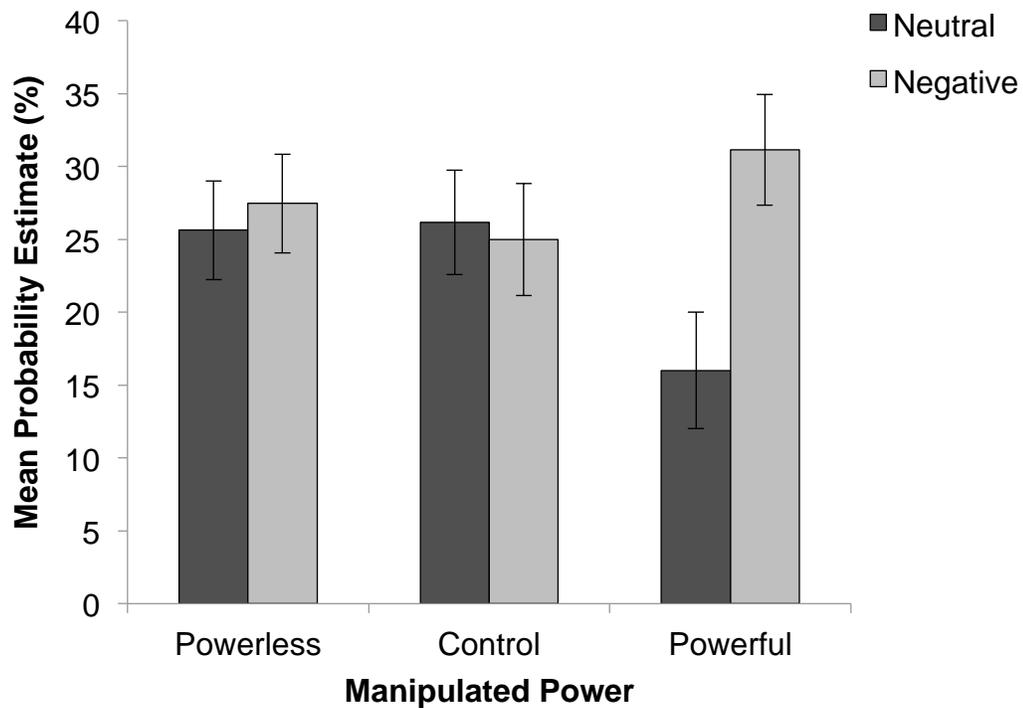


Figure 4.6. Probability estimate means for the dice scenario for powerless, control, and powerful participants (manipulated power), in the neutral, and negative conditions. Error bars represent +/- 1 standard error of the mean.

Physiological Data.

Mediation by Arousal. In order to examine whether arousal, measured when participants provide their probability estimates, mediates the relationship between negative utility and probability, we first computed the difference between tonic arousal measure when participants started to move the slider, and the tonic arousal during the baseline. First, we extracted the mean tonic activity for both time points by means of continuous decomposition analysis using the MATLAB toolbox Ledalab (v.3.4.4) (Benedek & Kaernach, 2010). Secondly, as in Experiment 3, we computed the difference between tonic arousal and baseline arousal. However, utility did not impact

the difference in arousal, $t(131)=.58$, $p=.56$, which ruled out the possibility for a mediating impact of arousal.

Mediated Moderation: Power and Interoceptive Accuracy. As reported in Moeini-Jazani et al. (2014, results reproduced with permission from the first author), social power significantly predicted interoceptive accuracy: Powerful individuals were most accurate ($M=64.3$, $SD=14.82$), followed by control ($M=44.34$, $SD=29.46$) and powerless participants ($M=38.65$, $SD=26.24$), $F(1,129)=13.19$, $p<.001$, $\eta^2=.17$ ³⁷. It could therefore be possible, that the powerful's greater interoceptive accuracy led to an increase in the negativity bias. However, in a 2 (utility: negative/neutral) x 2 (interoceptive accuracy: high/low) between subjects ANOVA there was neither a main effect of interoceptive accuracy, nor an interaction between interoceptive accuracy and utility on probability estimates in the dice scenario, all $F_s<1$. These results did not change upon including the covariates mood, anxiety, sports activity, and the error in the time counting task. There was therefore no evidence that the powerful's greater interoceptive accuracy had an impact on the relationship between social power and the interdependence of negative utility on probability estimates.

Further Analyses

As we also collected measures of sense of power and approach-avoidance motivation, we now report these analyses on the dice scenario, too. Note that due to some participants entering a wrong ID when completing the pre-test, we were unable to match all laboratory responses to the questionnaire data, accounting for the difference in degrees of freedom.

³⁷ The means and standard deviations reported here differ from the ones reported in Moeini-Jazani et al. (2014). This is due to the fact that, in line with previous experiments, we removed outliers. The results do not change when removing or not removing outliers.

Measured Power. Submitting the dice-scenario to a 2 (utility: negative/neutral) x 2 (measured power: high SOP/low SOP) ANOVA, there was neither a main effect of SOP, $F(1,116)=2.14$, $p=.15$, $\eta^2=.02$, nor of utility, $F<1$. However, there was a marginally significant interaction between utility and sense of power, $F(1,116)=2.99$, $p=.09$, $\eta^2=.03$. Simple effects analyses revealed that the powerful showed a negativity bias, $F(1,116)=5.23$, $p=.02$, $\eta^2=.04$, whilst this was not the case for the powerless, $F<1$ (see Figure 4.7). Furthermore, conducting the same orthogonal contrasts as in previous experiments, the contrast comparing those high in SOP answering a negative scenario to all other conditions was significant, $t(116)=2.1$, $p=.04$, whilst no other contrasts were significant, all $ps>.31$, and no rest-variance was unaccounted for, $F<1$.

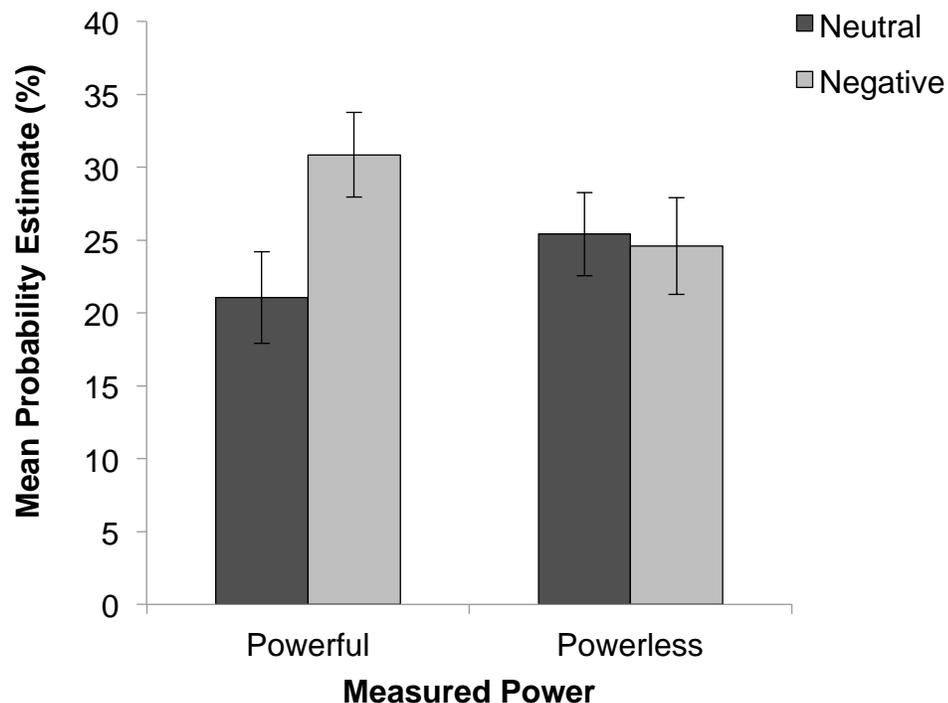


Figure 4.7. Probability estimate means for the dice scenario for individuals high and low in sense of power (measured power), in the neutral, and negative conditions. Error bars represent +/- 1 standard error of the mean.

Approach-Avoidance. After summing the items to form an avoidance and an approach subscale of the BIS-BAS questionnaire, we performed a median split and submitted both to two 2 (utility: negative/neutral) x 2 (high avoidance (approach) /low avoidance (approach)) between participants ANOVAS.

Neither did trait avoidance motivation impact the interdependence of negative utility and probability estimates, all $ps > .24$, nor did approach motivation have any impact, all $ps > .26$.

Discussion

Experiment 9 demonstrated that (after the removal of disproportionately influential data points) only individuals that anticipated a powerful role in a future interaction demonstrated a negativity bias. This experiment therefore demonstrated a causal link of power on the impact of negative utility on probability estimates. The inclusion of a control condition showed that powerful individuals' negativity bias was greater than that of powerless and control participants.

Last, neither arousal, nor interoceptive accuracy had an impact on the relationship between negative utility and probability estimates. Therefore, a mechanism other than arousal misattribution is likely to underlie the impact of power on the interdependence of negative utility on probability estimates.

More, this experiment does not support the notion that approach-avoidance motivation underlies the impact of social power on the relationship between negative utility and probability estimates, contrary to the predictions derived from experiments by Lench (2009). Neither trait avoidance nor trait approach motivation moderated the impact of negative utility on probability estimates.

Experiment 10

Background³⁸

Experiments 3, 8, and 9 provided some first evidence that social power moderates the impact of negative utility on probability estimates. However, it seems unlikely that interoceptive accuracy or approach-avoidance motivation were mechanisms underlying this finding.

Therefore, Experiment 10 aims to further explore the role of control, as suggested by an ALF account. As mentioned earlier, feelings of control can be both *over another individual*, as well as *over oneself*. Both types of control would be predicted by an account of ALF to moderate the impact of negative utility on probability. If one has control over other individuals, one might be more inclined to believe that they act based on one's estimate. As described in Chapter 1, Harris et al.'s (2009) examination of ALFs incorporated only events where the negative outcomes would impact someone else, but not oneself.

On the other hand, social power is also associated with control over events happening to oneself (e.g. illusion of control, internal locus of control, primary control). Based on an ALF account, the belief that one can act based on one's estimate, and therefore impact whether or not the negative outcome occurs, should moderate the impact of negative utility on loss function asymmetries.

The current experiment examines whether social power moderates the impact of negative utility on probability estimates for self-, for other-relevant scenarios, or for both. In contrast to the previous experiments in this chapter, Experiment 10 was not scenario based and included real incentives. Participants either played a game in the

³⁸ I thank Melinda Soh for help with collecting the data for Experiment 10.

'self' condition, or the game was described to them as being played by another player in the 'other' condition. Participants in the negative self condition were given £5 (whilst participants in the negative other condition were told that the player of this game would be given £5), and were told that they would draw four counters from a bag, out of which one counter was red. Participants were informed that if they would draw the red counter at least once in the four draws they (or the player of the game) would lose the £5, and were asked to estimate the likelihood that they (or the player of the game) would draw the red counter. No reference to the £5 was made to participants in the neutral condition.

Method

Participants and Design. 200 participants (115 female, median age=21) were recruited opportunistically and completed the experiment on the campus of University College London. Participants were randomly assigned to a 2 (utility: negative/neutral) x 2 (self-relevance: self/other) between participants design.

Materials and Procedure. Participants were informed that the study is concerned with participants' perception of games. In the 'self' condition, participants were told that they were going to play a game now, whereas in the 'other' condition, participants were explicitly told another player would play the game. Participants were then given verbal instructions about the game.

In the 'negative self' condition, participants were handed £5 and the experimenter put six differently coloured counters (one of which was red) in a bag. It was explained that the participant would draw four counters from the bag, whilst always placing the counter back inside the bag after each draw. Moreover, participants were informed that if they were to pull out the red counter on at least one of the four draws,

they would be asked to give the £5 back, otherwise they could keep it. In the ‘negative other’ condition, participants were shown the £5 but were not handed the money. The instructions were identical but rather than directly addressing participants, the game was described from the perspective of “a player of this game”. In the neutral conditions there was no reference to the £5, and therefore no outcome was attached to drawing the red counter.

Following these verbal instructions, participants read the descriptions of the game again on an iPad (presented via the software Qualtrics) and were told that we would like them to answer some questions about the game. In all conditions, participants were then asked to estimate the chance that they (or the player of this game) would draw the red counter at least once. Answers were given on sliding scales from 0-100 on the iPad. Following, as a manipulation check, participants answered the two questions “how bad would it be if at least one red counter was drawn” (from ‘not bad at all’ to ‘very bad’) and “how much would you personally be affected if at least one red counter was drawn” (from ‘not at all’ to ‘very much’), both on 7-point scales.

Afterwards participants completed the Sense of Power scale (Anderson & Berdahl, 2012), were thanked and debriefed.

Results

Manipulation Checks.

Utility. The answer to the question how bad it was if at least one red counter was drawn was submitted to an independent samples t-test. The manipulation of utility was successful, with participants in the negative condition indicating that the outcome was worse ($M=3.86$, $SD=2.11$) compared to participants in the neutral condition ($M=1.79$,

SD=1.34), $t(167)=8.28$, $p<.001$ (df adjusted due to inequality of variances). Next, in order to examine potential influences of self-relevance and power on the perception of utility, we submitted the manipulation check to a 2 (utility: negative/neutral) x 2 (power: powerful/powerless) x 2 (self-relevance: self/other) between participants ANOVA. Unexpectedly, a significant main effect of self-relevance was found, with participants in the ‘other’ condition reporting that they would find the outcome (across severity conditions) worse ($M=3.24$, $SD=2.11$) than participants in the ‘self’ condition ($M=2.41$, $SD=1.9$), $F(1,192)=10.33$, $p<.001$, $\eta^2=.05$. There was no main effect of power, and no further significant interactions, all $F_s<1$. The lack of an interaction between utility and self-relevance demonstrates that the impact of self-relevance on utility was not specific to negative events. For aforementioned reasons, this would have been problematic, as it could explain potential moderation effects.

Self-Relevance. The answer to the question how much participants would be personally affected if at least one red counter were drawn was submitted to an independent samples t-test. The manipulation check did not indicate that we successfully manipulated self-relevance, with participants in the other-relevant condition reporting that they would be more personally affected by the outcome ($M=2.45$, $SD=1.8$) compared to participants in the self-relevant condition ($M=2.08$, $SD=1.64$). This difference was not significant, $t(198)=1.52$, $p=.13$ ³⁹.

Next, we submitted the manipulation check to a 2 (utility: negative/neutral) x 2 (powerful/powerless) x 2 (self-relevance: self/other) between subjects ANOVA. There

³⁹ This failure might be attributable to an inappropriate manipulation check question. For instance, instead of asking participants “how much would you personally be affected if at least one red counter was drawn” (from ‘not at all’ to ‘very much’) maybe a more appropriate question would have been “who would be affected if at least one red counter was drawn” (‘me’ or ‘someone else’).

was a marginally significant main effect of severity, in that participants in the negative condition gave higher ratings ($M=2.49$, $SD=1.71$) than participants in the neutral condition, ($M=2.04$, $SD=1.72$), $F(1,192)=3.22$, $p=.07$, $\eta^2=.03$. No further main effects or interactions were significant, all $ps>.15$.

Similar to the manipulation check of utility, this marginally significant main effect was unexpected, but no reason for concern regarding further analyses since no interactions were present.

Probability Estimates.

Aggregated Probability Estimates Across Self-Relevance. As the manipulation of ‘self’ vs. ‘other’ failed, we collapsed across the target conditions. A 2 (utility: negative/neutral) x 2 (power: powerful/powerless) ANOVA showed only a main effect of utility, $F(1,196)=6.05$, $p=.02$, $\eta^2=.03$. Neither the main effect of power, nor the power*utility interaction was significant, all $Fs<1$.

Next, we conducted the identical set of orthogonal contrasts as in previous experiments (see Table 4.2). Only the contrast comparing powerful individuals providing estimates for a negative scenario to all other conditions was significant, $t(196)=2.13$, $p=.04$. The contrast comparing the powerless negative condition to the neutral conditions did not reach significance, $t(196)=1.5$, $p=.13$, and neither did the difference between powerful and powerless individuals in the neutral condition, $t(196)=.46$, $p=.65$ (see Figure 4.8).

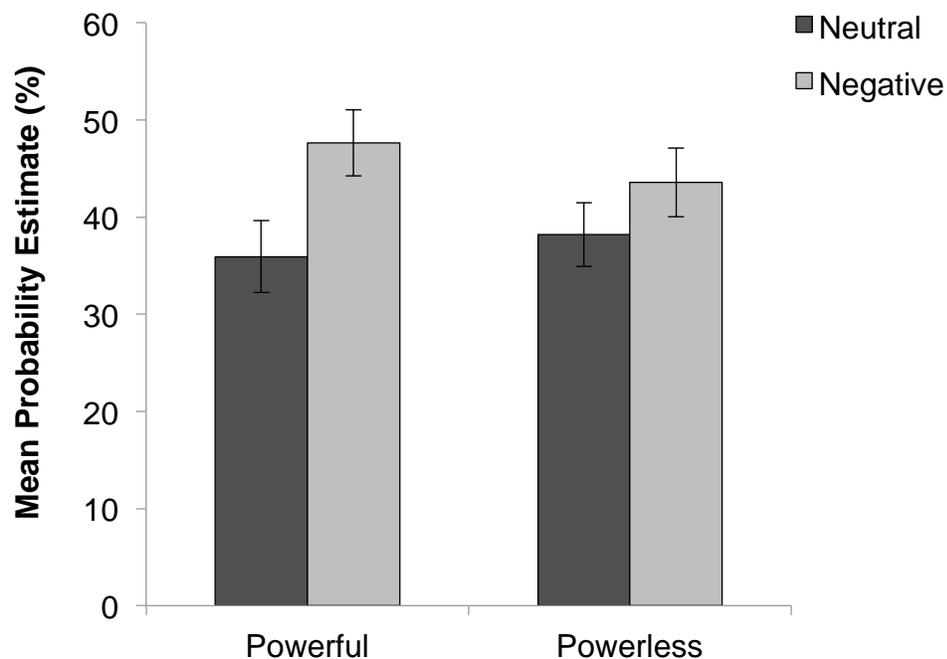


Figure 4.8. Probability estimate means for powerful and powerless participants in the neutral or negative condition. Error bars represent +/- 1 standard error of the mean.

However, whilst there were no effects of self- vs. other-relevance on the manipulation check question, probability estimates appeared to be impacted by the manipulation (constituting further evidence that this was not a genuine failed manipulation, but instead an insensitive manipulation check question). Therefore, we decided to conduct the full set of orthogonal contrasts, assessing the impacts of power and utility on levels of self-relevance.

Examining Self-Relevance. A 2(utility: negative/neutral) x 2 (power: powerless/powerful) x 2 (self-relevance: self/other) ANOVA revealed only a main effect of utility, $F(1,192)=5.71, p=.02, \eta^2=.03$, all other $F_s < 1$.

Next, the same set of orthogonal contrasts as in previous experiments from this chapter was conducted on self- and other-relevant scenarios separately, examining

whether the hypothesis that powerful individuals answering negative scenarios would provide greater estimates than the other 3 conditions combined, would hold both across self-, and other-relevant scenarios.

For self-relevant scenarios, contrast 1 was marginally significant, $t(96)=1.8$, $p=.07$. Contrast 2, $t(96)=1.53$, $p=.13$, and contrast 3, $t(96)=.22$, $p=.83$, were not significant. No rest-variance was unaccounted for, $F(2,96)=1.18$, $p=.31$.

For other-relevant scenarios, neither contrast 1, $t(96)=1.24$, $p=.22$, contrast 2, $t(96)=.48$, $p=.63$, nor contrast 3, $t(96)=.87$, $p=.39$, were significant (see Figure 4.9).

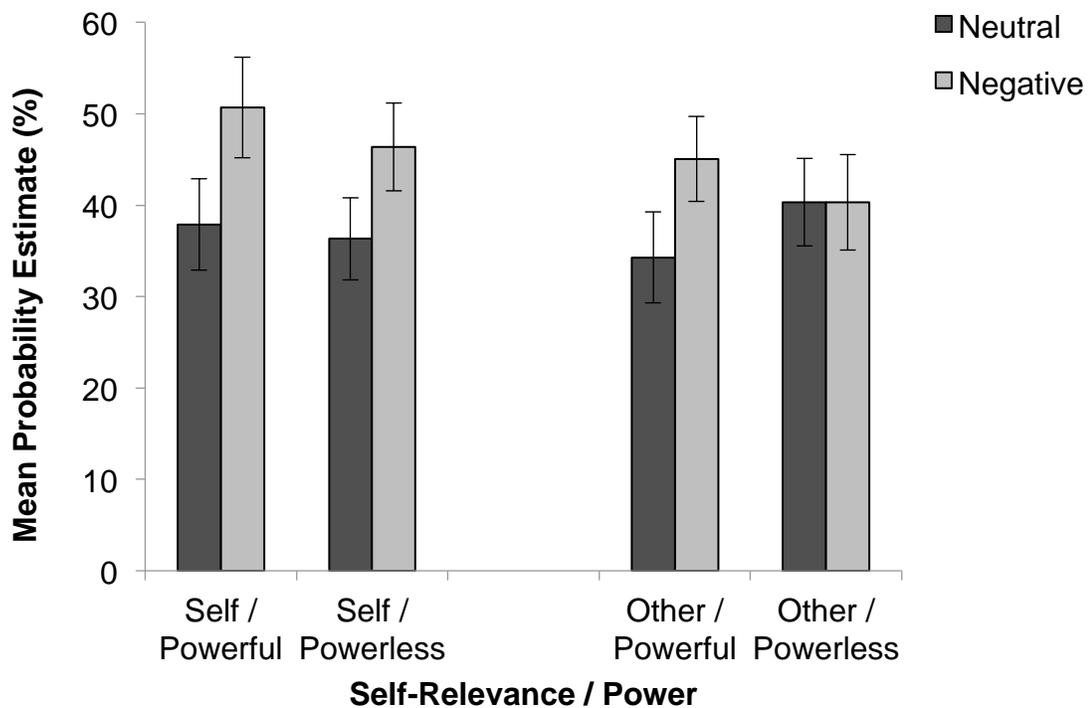


Figure 4.9. Probability estimate means for powerful and powerless participants answering self- or other-relevant scenarios in the neutral or negative condition. Error bars represent +/- 1 standard error of the mean.

Discussion

Averaging across conditions of self-relevance, Experiment 10 replicated our previous findings whereby individuals high in social power demonstrate a greater negativity bias than those low in social power.

However, Experiment 10 failed to demonstrate that self-relevance was successfully manipulated. Nevertheless, an impact of the self-relevance manipulation was observed. Running a set of orthogonal contrasts, powerful participants estimating the likelihood of negative events only provided (marginally) greater estimates compared to the other conditions for self-, but not other-relevant scenarios. However, examining the graphs it became evident that whilst powerless individuals exhibited a directional negativity bias for self- but not for other-relevant scenarios, the powerful appeared to show a negativity bias both for the self- and other-relevant scenarios. Given these descriptive statistics, there should not be too much emphasis placed on the conclusion that we arrive at from analysing orthogonal contrasts, whereby our hypotheses only held for self, but not for other-relevant scenarios.

Experiment 11

Background⁴⁰

In previous experiments, we hypothesized, in line with an ALF account, that it is their heightened sense of control which leads powerful individuals to provide greater estimates for negative compared to neutral events. Experiment 11 aims to explore the

⁴⁰ The design of Experiment 11 was developed in collaboration with Kai Barron, Department of Economics, University College London.

role of decision-control in a novel design. If a heightened sense of control is underlying the findings of previous experiments, then the difference in magnitude of the negativity bias between the powerful and powerless should be particularly pronounced in scenarios where decision-control is ambiguous or non-existent. In situations where decision-control is clearly present, the difference between powerful and powerless individuals should be reduced, as then both the powerful and the powerless experience decision-control.

In the current experiment, for the first time, participants are given real decision-control with real outcomes and incentives for accuracy. Furthermore, this experiment will test yet unassessed predictions of the ALF account, which shall be explained in more detail below.

If individuals exercise decision-control, they aim to be in an end-state with the lowest probability of the negative event occurring. For example, imagine a person estimating the chance of getting lung cancer in 20 years when they continue to smoke. When the person continues to smoke, arguably the chances of getting lung cancer are relatively high, and this constitutes an undesirable end-state. On the other hand, imagine the same person estimating the chances of getting lung cancer in 20 years when they stop smoking now. The chances of getting lung cancer should be decreased compared to when they continue to smoke and, relatively, this should be the more desirable end-state (despite still carrying some risk). In previous experiments, we have only asked participants to estimate the chance for an end-state without asking them to estimate the chance for a more desirable end-state. However, an account of ALFs has different

predictions about how individuals are biased by the utility of the event for the more and the less desirable end-state.

In line with our previous experiments, an ALF account predicts a negativity bias for end-states that are more undesirable, as it is more costly to underestimate the risk of these states where it is more likely for the negative outcome to occur than to overestimate it. Regarding the smoker example above, it would be more costly for the smoker to underestimate the risk of getting lung cancer than to overestimate it.

However, loss function asymmetries should *reverse* for end-states that, despite carrying some risk for the negative outcome to occur, are relatively more desirable due to lower risk. Here, it should be more costly to *overestimate* the risk rather than to underestimate it. Therefore, according to an ALF account, participants should provide lower estimates in the negative compared to the neutral condition. Considering the smoking example, if one overestimates the risk of getting lung cancer when one stops smoking, there might be a lack of an incentive to do so. A lack of an incentive is here conceptualised as the difference in probability estimates for the less compared to the more desirable end-state. The assumption here is that the greater this difference, the more people are motivated to work towards the better and away from the worse end-state. Importantly, however, both loss functions (and their asymmetries) should only exist when individuals have decision-control.

In line with the reasoning above, in the present experiment there are two different states that participants can be in at the end of the experiment: In one state the probability of the negative outcome is high, and in the other state the probability of the negative outcome is lower, which constitutes the more desirable state. Importantly,

participants in a condition without decision-control are informed that the chance of being in one or the other state is 50%. However, participants with decision-control can influence the likelihood of being in the more desirable-end state by completing a task that depends only on effort. Thus, this experiment allows us to observe the estimates for the two end-states with and without decision-control, and observe the magnitude of the negativity bias depending on both decision-control and power.

Furthermore, this experiment also employs real incentives. Concretely, participants in the negative utility condition are told that they have been entered into a (real) lottery to win £50, but that they might be removed from the list of ticket-holders by the end of the experiment (the negative outcome), whereas participants in the neutral condition were equally entered into the draw, but no reference to the possibility of losing their ticket was made. Participants in the negative condition without decision-control (the “luck” condition) were informed that keeping their lottery ticket depends on luck. In this condition, participants were told that in order to determine whether they would keep their lottery ticket, a coin would be dropped onto one of two maps A or B. Both maps have blue lines crossing them in a grid-like pattern, and if the coin overlaps with any of the two lines they would lose their ticket. Importantly, Map A had fewer lines than Map B and constituted therefore the more desirable map. Participants were further informed that which one of the two maps would be chosen for the coin drop would depend on luck only, with a 50% chance that either map would be chosen.

Participants in the negative decision-control condition were instead told that they would be able to influence the likelihood with which Map A, the desirable map, would be chosen, by completing an effortful task (where the amount of effort was then

proportional to the likelihood of choosing Map A). Thus, their estimate could provide information with regard to how much effort they should exert for this task, constituting real decision-control. Participants in the neutral condition were simply told that they were entered into the lottery, and no connection was made between keeping their lottery ticket and estimating the likelihood of the coin crossing a line for the two maps. All participants then estimated the likelihood of the coin crossing a line for the two maps, and were incentivised for accurate estimates.

This design allowed to test three hypotheses: one regarding the probability for the less desirable Map B, one regarding probability estimates for the more desirable Map A, and one hypothesis regarding the difference of the two maps.

Our main hypothesis was regarding probability estimates of powerful and powerless individuals regarding Map B, the less desirable map. The estimates for this map most closely resemble the set-up of our previous experiments, where participants simply estimate the chance of a negative outcome to happen. If a greater sense of control is underlying the previous findings, we would expect a negativity bias for powerful individuals both for situations where there is no decision-control, and for situations where decision-control is clearly stated. However, powerless participants should exhibit a negativity bias only in the condition of decision-control.

A more exploratory hypothesis was made for probability estimates for Map A. In the contrast to previous experiments, this map constituted a desirable end-state, one that participants would want to work towards to in order to have a decreased chance of losing their ticket. As reasoned above, in the presence of decision-control it would be more costly to *overestimate* rather than to *underestimate* the probability of this map. As

a result, individuals should provide lower estimates for negative compared to neutral outcomes in the presence of decision-control. As we expect the powerful to perceive control also under conditions where objectively there is little control, we hypothesize that the powerful provide lower estimates across both the negative luck and the negative decision-control conditions compared to the neutral condition. Powerless individuals, however, should provide only a lower estimate for the negative decision-control compared to the neutral condition. For the powerless, there should be no difference between the neutral and negative-luck condition.

As a result of the hypotheses above, and as a more general assessment of an ALF account beyond the effects of power, we would expect that the difference between the two maps is the greatest for the decision-control condition. If participants provide estimates to inform action, they should provide greater estimates for the undesirable Map B, and the smallest estimates for the desirable Map A, in order to motivate the exertion of effort in the decision-control, but not in the luck or the neutral condition.

Method

Participants and Design. 642 Participants were recruited via UCL's email distribution system (349 female, median age=23) and completed the experiment online. Participants were randomly assigned to one of the 3 between participants conditions (neutral luck, negative luck, negative decision-control), which constituted an incomplete 2 (utility: negative/neutral) x 2 (decision-control/luck) x 2 (Map: A/B) design, with the last factor manipulated within participants.

Materials and Procedure. First, participants completed the sense of power scale (Anderson et al. 2012). Next, participants were informed that they were to take part in

an experiment on the perception of games. All participants were asked to first enter their email address in order to be entered into the lottery. Participants in the negative conditions were told that some participants will lose their ticket, and that keeping the lottery ticket depends on luck (in the negative luck condition) and on luck and effort (in the negative decision-control condition).

Next, all participants were informed that in this game there are two maps (A and B) and that they will be asked to make some predictions about the two maps. In order to incentivise accuracy, participants were informed that if they are accurate in their predictions they can increase the amount they will win if they win the lottery by £15 for each of the two predictions, should they be within 10% of the true value. Following, it was explained that a coin will be dropped onto one of the two maps, either Map A, or Map B. Importantly, participants were informed that each of the two maps has blue lines crossing it – but that Map B has more blue lines crossing through it than does Map A. Participants in the negative conditions were told that if the coin overlaps with a blue line, they would lose their lottery ticket.

Participants in the luck conditions were informed that there is an equal chance that Map A or Map B will be chosen for the game, and that this will be determined randomly.

However, participants in the decision-control condition were told that they would have the chance to influence which of the two maps the coin will be dropped onto by completing a 30 seconds long task. Participants were informed that this task requires them to click on as many boxes in a large display as possible, and that the better they would perform, the greater their chances of having the coin dropped onto Map A.

Afterwards, participants saw the two Maps A (true probability $=\frac{1}{3}$) and B (true probability $=\frac{2}{3}$) (see Figure 4.10).

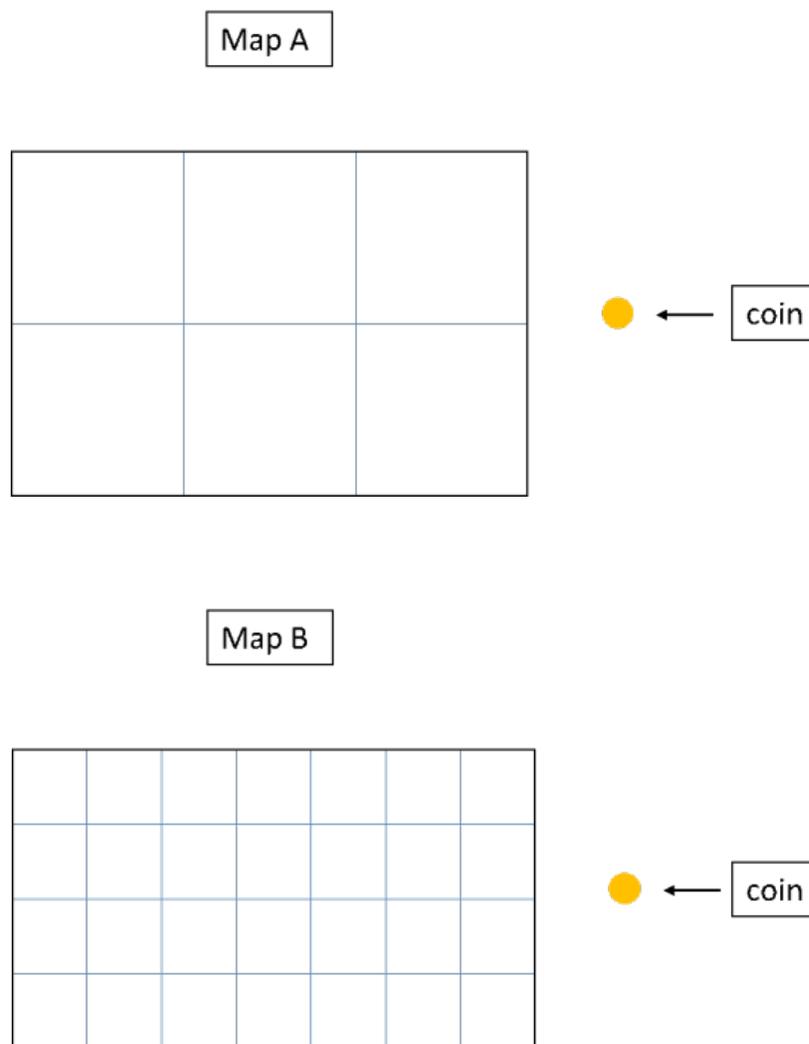


Figure 4.10. Visual probability displays used in Experiment 11.

The probabilities of the two maps were calculated by applying the formula:

$$Probability = \frac{(l - d)^2}{l^2}$$

l =length of a single square

d =diameter of the coin

Underneath each Map was a sliding scale from 0 to 100 as in previous experiments, and participants were asked for each of the maps to estimate the chance that the coin would overlap with one or more of the blue lines (in the negative condition an “and you will therefore lose your lottery ticket” was added here to make the outcome salient) if the coin would be randomly dropped anywhere on the maps. Whilst both maps were presented on the same page, the order with which the maps appeared on the screen was counterbalanced.

Following, participants in the decision-control condition were given 30 seconds to click on as many boxes as possible in a large multiple-choice matrix. Next, all participants were informed that their chance of keeping their lottery ticket was equal to the probability of the coin dropping onto Map A, were thanked and dismissed.

Results

Participants. Participants who provided greater estimates for Map A, which had a true probability of 33%, than for Map B, which had a true probability of 66%, were excluded from the analysis, as they were considered to have misunderstood the task. Furthermore, participants who reported that they did not understand or read the instructions were also excluded, resulting in 605 participants in the final sample (349 female, median age=23).

Manipulation Checks.

Utility. The answer to the question “how bad would it be if the coin would overlap with one or more of the blue lines” was submitted to an independent samples t-test. The manipulation of utility was successful, with participants in the negative condition indicating that the outcome would be worse ($M=4.29$, $SD=2.08$) than participants in the neutral condition ($M=1.87$, $SD=1.55$), $t(603)=15.15$, $p<.001$.

Unexpectedly, participants in the negative decision-control condition gave slightly more negative utility ratings ($M=4.57$, $SD=1.99$) than participants in the negative luck condition ($M=4.03$, $SD=2.14$), $F(1,601)=6.67$, $p=.01$. These results support the findings by Siegrist and Sütterlin (2014) who demonstrated that man-made (arguably more controllable) hazards are perceived as worse than nature-made (arguably less controllable) hazards.

Importantly, there was no main effect of power, $F(1,601)=1.1$, $p=.3$, and no interaction between utility and power, $F<1$, when running a 2 (utility: negative/neutral) x 2 (decision-control/luck) x 2 (power: powerful/powerless) between participants ANOVA.

Decision-Control. The answer to the question “how much can you influence which one of the two maps will be chosen for the coin drop” was equally submitted to an independent samples t-test. Participants in the decision-control condition reported more control ($M=4.02$, $SD=1.72$) than participants in the luck condition ($M=1.79$, $SD=1.57$), $t(603)=15.6$, $p<.001$.

Unexpectedly, participants in the neutral (luck) condition reported slightly more control over the outcome ($M=1.93$, $SD=1.63$) than participants in the negative (luck) condition ($M=1.63$, $SD=1.48$), $t(603)=1.99$, $p=.05$.

Importantly, neither the main effect of power, nor the interaction were found significant, all $F_s < 1$.

Map B Probability Estimates. As Map B was the less desirable of the two maps, we hypothesized in line with our previous findings that powerful participants would provide greater estimates for both the negative decision-control and the negative luck compared to the neutral condition. Powerless participants, on the other hand, should only provide greater estimates for the negative decision-control compared to the neutral condition.

A 2 (power: powerful/powerless) x 2 (utility: negative/neutral) x 2 (control: decision-control/luck) between subjects ANOVA showed a main effect of utility, $F(1,586)=7.72$, $p < .01$, $\eta^2 = .01$. Individuals in the negative conditions estimated the probability to be greater ($M=55.58$, $SD=20.77$) than participants in the neutral condition ($M=48.82$, $SD=21.68$). No other effects were significant, all $p_s > .31$.

As for Map A, we conducted planned comparisons. As predicted, powerful participants showed a negativity bias both for the comparisons of the neutral to the negative luck, $F(2,586)=4.74$, $p < .01$, and for the neutral to the negative decision-control condition $F(2,586)=5.0$, $p < .01$. The difference between the negative luck and the negative decision-control condition was not significant, $F < 1$.

Importantly, however, powerless participants showed a negativity bias only for the negative effort condition $F(2,586)=4.3$, $p = .01$, but not for the negative luck

condition $F(2,586)=1.6, p=.20$ (see Figure 4.11). The difference between the negative luck and effort conditions did not reach significance, $F(2,586)=1.6, p=.26$. In sum, these results support the notion that a difference in the perception of decision-control underlies the moderation of the impact of negative utility on probability estimates by power.

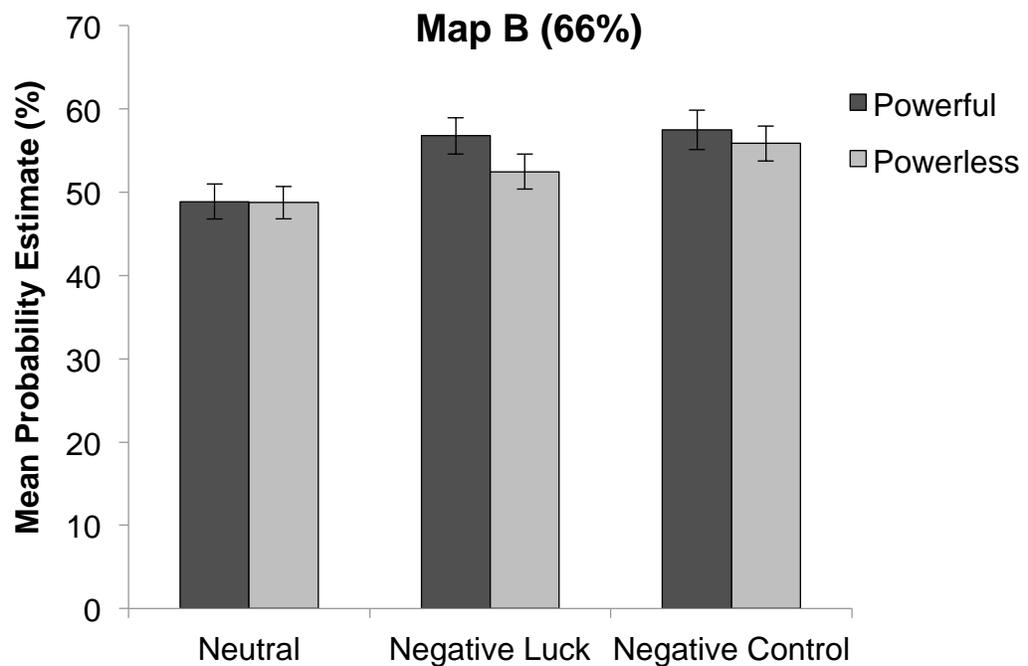


Figure 4.11. Probability estimate means for Map B for powerful and powerless participants, in the neutral, negative luck and negative decision-control conditions. Error bars represent +/- 1 standard error of the mean.

Map A Probability Estimates. As Map A was the more desirable of the two maps, we hypothesized powerful participants to provide lower estimates for the negative decision-control and the negative luck condition compared to the neutral condition. Powerless individuals were instead expected to only provide lower estimates for the

negative decision-control compared to the neutral condition, but not difference between the neutral and the negative luck condition was predicted.

A 2 (power: powerful/powerless) x 2 (utility: negative/neutral) x 2 (control: decision-control/luck) between subjects ANOVA yielded a significant main effect of utility, $F(1,586)=10.18, p=.03, \eta^2=.01$. Individuals in the neutral condition provided the lower estimates ($M=18.66, SD=12.75$) than participants in the negative condition ($M=21.11, SD=12.89$). Furthermore, there was a significant main effect of power, $F(1,586)=4.19, p=.03, \eta^2=.01$ ⁴¹. Participants high in power provided greater estimates ($M=22.26, SD=13.14$) than participants low in power ($M=19.51, SD=12.61$). No other main effect or interaction was significant, all $ps>.26$. However, due to our inability to analyse this incomplete design factorially, we next examined planned pairwise comparisons. The powerful provided greater estimates for the negative luck compared to the neutral condition, $F(2,581)=3.23, p=.04$ (see Figure 4.12). However, for the powerful there was no difference between the neutral and the decision-control conditions, $F<1$, and no difference between the negative luck and the decision-control condition, $F(2,581)=2.12, p=.12$. Thus, powerful participants showed a tendency to display a negativity bias only for the luck condition, but not for the effort condition. Thus, they did not, as was hypothesized, provide lower estimates for the negative decision-control and negative luck compared to the neutral condition.

For the powerless, on the other hand, the difference between the neutral and the negative luck condition, $F(2,581)=4.3, p=.01$, as well as the difference between the neutral and the decision-control condition, $F(2,581)=4.46, p=.01$, was significant. There was no difference between the negative decision-control and the negative luck condition

⁴¹ Before the removal of outliers, the main effect of power was not significant, $F(1,599)=1.85, p=.18$.

for the powerless, $F < 1$. Thus, unexpectedly, the powerless demonstrated a negativity bias both for the negative luck and effort conditions. However, for the effort condition, the difference was not in the predicted direction, as we would have expected probability estimates to decrease below the neutral condition in the negative decision-control condition, as loss function asymmetries reverse for this more desirable map⁴².

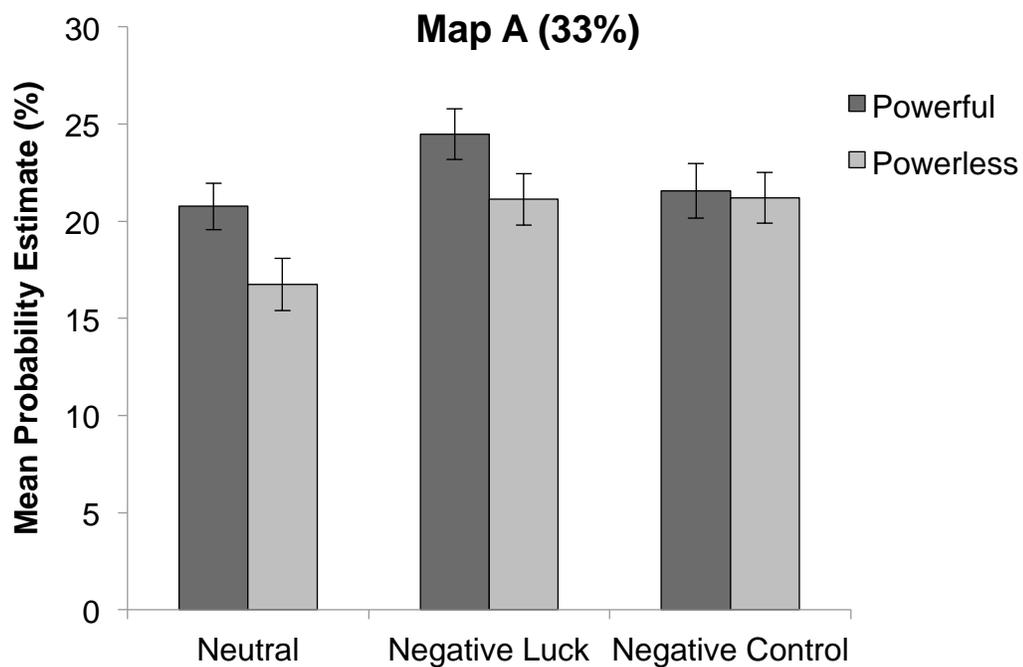


Figure 4.12. Probability estimate means for Map A for powerful and powerless participants, in the neutral, negative luck and negative decision-control conditions. Error bars represent +/- 1 standard error of the mean.

Difference-Score. In the current experimental set-up, simply providing greater estimates for Map B by itself is not the strongest demonstration of loss function

⁴² Before removing outliers, for the powerless there was only a marginally significant difference between the neutral and the negative decision-control condition, $F(2,599)=2.34$, $p=.10$. No other comparisons were significant, all $p_s > .14$.

asymmetries as a mechanism. It might be possible, that individuals simply assigned greater estimates for both maps. If the difference between the two maps is identical across conditions, then there is little reason to assume that participants would act differently based on their estimates. What is important here, is that the difference between the two maps should, according to an ALF account, increase with greater decision-control: in order to motivate oneself to exert effort to avoid the negative outcome, one should assign a greater probability for the undesirable Map B, and a smaller probability to the desirable Map A.

Therefore, we analysed the difference score by subtracting the probability estimates for Map A from the estimates for Map B. For individuals high in social power we hypothesize that due to their greater sense of control, the difference score should be greater for both negative conditions compared to the neutral condition, motivating them to exert effort. This should not be the case for individuals low in social power, who we hypothesized to only show a difference in the difference score between the neutral and the negative decision-control condition.

A 2 (power: powerful/powerless) x 2 (utility: negative/neutral) x 2 (control: decision-control/luck) between subjects ANOVA showed a significant main effect of effort, $F(1,588)=4.4$, $p=.04$, $\eta^2=.01$. Those in the decision-control condition had a greater difference between the two maps ($M=33.3$, $SD=14.63$) than those in the luck conditions ($M=28.51$, $SD=14.95$). Furthermore, there was a main effect of utility, $F(1, 588)=5.07$, $p=.03$, $\eta^2=.01$ ⁴³. Individuals in the negative conditions had greater difference scores ($M=31.71$, $SD=15.25$) than individuals in the neutral condition ($M=26.99$, $SD=14.11$). No other main effects or interactions were significant, all $ps < .13$.

⁴³ Before removing outliers, the main effect of utility was not significant, $F < 1$.

Simple effect analyses showed that, as hypothesized, for the powerful there was a significant difference between the neutral and negative luck condition, $F(2,588)=4.74$, $p<.01$, as well as a significant difference between the neutral and negative decision-control condition, $F(2,588)=6.99$, $p<.001$. The difference between the negative luck and decision-control condition was marginally significant, $F(2,588)=2.44$, $p=.09$ (see Figure 4.13).

For the powerless, there was only a marginally significant difference between the neutral and the decision-control condition, $F(2,588)=2.83$, $p=.06$, all other comparisons were non-significant, all $ps>.21$.⁴⁴

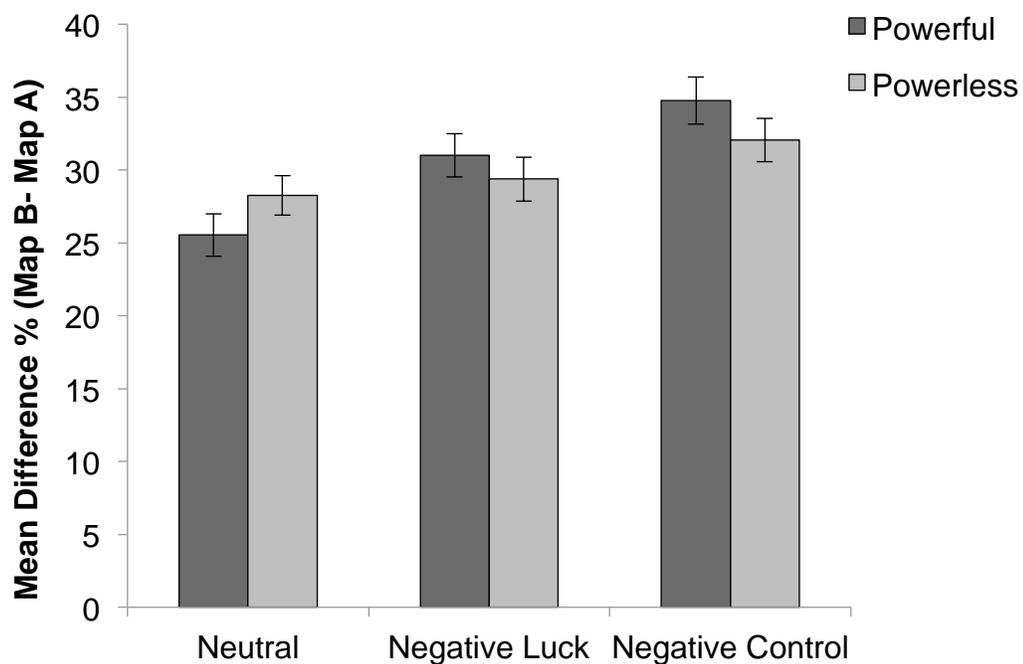


Figure 4.13. Probability estimate means for the difference score between Map B and Map A for powerful and powerless participants, in the neutral, negative luck and negative decision-control conditions. Error bars represent +/- 1 standard error of the mean.

⁴⁴ Before removing outliers, there was only a difference for the powerful neutral to the decision-control condition, $F(2,599)=4.64$, $p<.01$, and a marginally significant difference between the powerful negative luck and neutral condition, $F(2,599)=3.37$, $p=.08$. No other differences were significant, all $ps>.11$.

Discussion

Experiment 11 is the first experiment to provide evidence for the role of decision-control as a mechanism underlying the moderation by social power on the impact of negative utility on probability estimates, in a situation with *real* rather than fictional decision-control.

Assessing the predictions that loss function asymmetries can reverse for the more desirable option, we predicted that probability estimates for Map A should decrease compared to the neutral condition for both the negative decision-control condition (for both powerful and powerless participants), and for the negative luck condition (for the powerful only, due to their heightened sense of control). However, the predictions for Map A were not supported. The powerful increased in their estimates from the neutral to the negative luck condition. However, for the first time did the powerful show *no* negativity bias in the negative decision-control condition, which did not differ from the neutral condition. This finding is, directionally, in line with the hypotheses brought forward by an ALF account.

Unexpectedly, however, the powerless provided greater probability estimates for the negative luck and decision-control conditions compared to the neutral condition for Map A. This finding by itself suggests that the powerless appear to demonstrate a negativity bias unconditionally of decision-control for the more desirable outcome. However, in conjunction with the finding that the difference between the more desirable and less desirable maps increased for the powerless, this might also reflect a tendency for the powerless to not work *towards* the better, but *away* from the less desired end-state, when decision-control is explicit.

Whilst our hypotheses for Map A were not supported, the predictions following an ALF account were supported for Map B. Whilst powerful individuals showed a negativity bias both for the luck and the decision-control conditions, powerless individuals only showed a negativity bias in the decision-control condition. This indicates that powerful individuals feel decision-control also when there is objectively little control, whilst powerless individuals only act when actual decision-control is present. This finding strongly supports the notion that powerful individuals show a greater negativity bias across conditions with differing decision-control due to their greater sense of control. In previous experiments where decision-control was more ambiguous, the powerful could have had different strategies to avoid the negative outcome more readily accessible than the powerless. However, the present experiment highlighted very explicitly that participants in the negative luck condition did not have any decision-control.

More generally, the finding that both the powerful and powerless individuals seemed to act more in line with loss function asymmetries for the less desirable Map B than for the more desirable Map A might reflect a positive-negative asymmetry (see Baumeister et al. 2001, for a review), whereby negative events are given a greater weight than positive events. The motivation to work *away from* the less desirable state might have dominated and overshadowed the motivation to work *towards* the more desirable state.

Furthermore, our hypotheses were supported for the difference score between Maps B and A: For the powerful there was a difference between the neutral and negative luck, as well as the neutral and negative decision-control conditions, indicating greater

sense of control of the powerful. Despite their heightened sense of control, however, the powerful still had a greater difference in the negative decision-control compared to the negative luck condition – meaning that their hypothesized IOC in the negative luck condition is not as great as having *actual, explicit* control. The powerless showed a marginally significant difference only between the neutral and the negative decision-control conditions, further substantiating the notion that it is indeed the lack of decision-control which leads the powerless to not provide greater estimates for negative situations, unless explicitly provided with greater control. Therefore, the differences between the powerful and powerless in their impact on the relationship between negative utility and probability estimates will likely be greatest in situations with little or ambiguous decision-control.

Furthermore, besides providing evidence for decision-control as a mechanism for the impact of social power on the relationship between negative utility and probability estimates, the current experiment also provided more general evidence for the ALF account. Previously, no experiment had been conducted giving participants *actual* decision-control, as thus far, experiments providing evidence for decision-control as a moderator were scenario based. The current study, demonstrating that the difference in probability estimates between Maps A and B increased when participants had decision-control, provides evidence for loss function asymmetries as a mechanism underlying the interdependence of negative utility and probability estimates.

Chapter Discussion

The present research sought to explore the impact of social power on probability estimates for negative or neutral future events across 4 experiments. Experiment 8 showed that individuals with a heightened sense of power provided greater estimates for negative compared to neutral events, whilst this was not the case for individuals with a decreased sense of power. Experiment 9 replicated the results of Experiments 3 and 8, both across manipulated and measured power, providing evidence also for a causal impact of power on the relationship between negative utility and probability (after the removal of disproportionately influential data points). By including a control group, we established that it is the powerful that drive the effect, showing an increase in the negativity bias (rather than the powerless showing a decrease). Experiment 10 tried to tease apart the role of control over oneself and over others. Orthogonal contrasts reflected that the powerful only show a negativity bias for self-relevant events. However, examination of descriptive statistics weaken this claim, as it appeared that the powerful showed a negativity bias both for self- and for other-relevant scenarios, whereas the powerless demonstrated a negativity bias for the self- but not for other-relevant scenarios. Finally, Experiment 11 employed a novel design providing participants with real decision-control, and incentivising accuracy. Whilst powerful individuals showed a negativity bias across scenarios high and low in decision-control, powerless individuals only showed a negativity bias when given decision-control.

As such, the present research has informed the literature on mechanisms underlying the impact of negative utility on probability estimates, the literature on mechanisms underlying the impact of social power on behaviour and cognition, as well

as identifying an important individual differences variable, informing the literature on individual decision-making.

The present experiments support the role of loss function asymmetries in the impact of negative utility and probability estimates. Previous research demonstrated that the negativity bias only occurs for events that incorporate an element of decision-control (Harris et al. 2009). Experiment 11 provided, for the first time, evidence in favour of an account of ALFs in a task involving real outcomes, rather than fictitious scenarios. Furthermore, this experiment also incorporated a novel design, whereby participants could influence whether they would end up in a more or less desirable end-state, therefore providing real decision-control. Whilst the hypotheses were not supported for the predictions made by an ALF account on estimates for the more desirable end-state (where loss function asymmetries should reverse), both the estimates for the less desirable end-state, and for the difference between the probability estimates for the two states were supportive of decision-control as a moderator of the effect. When given real decision-control, participants' estimates for the two end-states were the furthest apart from one another, hinting that this greater perceived difference might motivate action as to prevent the negative event from occurring.

Furthermore, social power, associated with a greater sense of control over oneself and others (Fast et al. 2009), led to greater probability estimates for negative events, highlighting the role of control and providing further evidence for the notion of a mediating effect by ALFs. Concretely, when powerless participants were given decision-control in an unambiguous scenario, they exhibited a negativity bias as well,

providing evidence for the notion that differences in control are underlying the different impact of utility on probability estimates for powerful and powerless individuals.

However, the current experiments did not tease apart the type of control that mediates these findings: As mentioned above, Experiment 10 aimed to distinguish between control over oneself and other's outcomes, however, the results were ambiguous. Whilst we did explore the concrete mechanism of decision-control in Experiment 11, this only concerned self-relevant scenarios. Whether powerful individuals' probability estimates for other-relevant scenarios are also more biased by utility, and whether control would underlie this process is to be assessed in future work.

Moreover, the current research did not support the alternatively proposed model derived from the SLH (Vosgerau, 2010). It was hypothesised that if the powerful pay more attention to internal states, they should perceive arousal more intensely and misattribute arousal to probability estimates more readily. However, there was little indication that negative utility increased arousal in Experiments 3 and 9. Only for one of the two scenarios presented in Experiment 3 did we measure an increase in arousal when the outcome was negative, and no such increase was observed in Experiment 9. In sum, arousal did not mediate the relationship between negative utility on probability estimates. And whilst Moeini-Jazani et al. (2014) demonstrated that social power increases interoceptive accuracy, interoceptive accuracy did not moderate the relationship of negative utility on probability estimates, replicating the null effects demonstrated in Chapter 2.

In addition, no evidence was found for the role of approach-avoidance motivation in estimating the probability of future negative events. Lench (2009)

hypothesized that greater avoidance-motivation would lead to an “avoidance” of the future event and subsequently to a decrease in probability estimates, whilst the opposite is the case for approach-motivation. Thus, this account would be in line with the current finding, where individuals low in social power, associated with a decrease in avoidance motivation (Keltner et al. 2003; but see Smith & Bargh, 2008) provided lower estimates for negative events than individuals high in social power. Experiment 9 could not provide evidence for the mediating role of trait approach-avoidance motivation.

However, the conclusions we can draw from this findings are arguably not very strong, as trait avoidance was assessed a couple of weeks before the laboratory session. It could therefore be possible that state approach-avoidance motivation still impacts the present results. On the other hand, in an experiment not reported in this thesis, de Molière and Harris (2014) demonstrated that neither state approach, nor state avoidance had a moderating impact on the relationship between negative utility and probability, supporting our present findings.

In sum, the theory most supported by the present results is an account of ALFs.

However, alternative explanations not addressed in the current research should also be considered. For instance, it has been hypothesized that the increased imagination for negative events leads to a negativity bias (e.g. Bilgin, 2012; Risen & Gilovich, 2007). Thus, it remains possible that powerful individuals imagine future events more vividly compared to powerless individuals, which was not assessed in the current experiments. Whilst the impact of power on imagination has not been demonstrated yet, powerlessness is associated with an impairment of executive functioning (Smith et al. 2008), which can lead to a decrease in the capacity to simulate future events (Dalglish

et al. 2007; Hertel, 2000). However, this account would propose that the powerless would drive the direction of the current effects, which was not supported in Experiment 9, where the powerful provided a greater negativity bias than powerless and control participants, who in turn did not differ from one another.

Importantly, however, the present research also informs the literature on mechanisms underlying the impact of social power on behaviour and cognition. The most dominant theory in this field, the approach-avoidance theory of power, proposes that power increases approach-, and powerlessness avoidance-motivation. Based on this theory, a vast amount of research has concluded that powerful individuals are more risk-taking than their powerless counterparts (Anderson & Galinsky, 2006; Carney et al. 2010). What is more, Anderson and Galinsky (2006) concluded “they [the powerful] should act in a more risky manner because they will be assigning higher probabilities to positive outcomes and lower probabilities to the possibility of negative outcomes” (Anderson & Galinsky, 2006, p. 520). Therefore, the current results serve as a clear boundary condition to the approach-avoidance theory of power. For the first time, we demonstrated that powerful individuals can also be *more* sensitive to negative information than powerless individuals, potentially through a strong motivation to *avoid* the negative outcome. On the other hand, the Situated Focus Theory of Power (SFTP, Guinote, 2008) predicts that the powerful should provide greater probability estimates for negative events, due to greater attention paid to internal states and acting on accessible information. However, whilst powerful individuals were indeed more internally focussed, interoceptive awareness did not moderate the impact of negative

utility on probability estimates. Thus, whilst the predictions of the SFTP are in line with the current data, the proposed mechanisms were not.

Furthermore, the results in the current experiments cannot be explained by the fact that power would impact the perceived magnitude of losses. In fact, the manipulation checks provided no indication that the powerful thought that the negative outcome was worse than the powerless. Indeed, less surprising would be if powerless participants would indicate these losses to be worse – in particular, as the outcomes in the dice scenario and in the games in Experiments 10 and 11 constituted a financial loss. As social power correlates negatively with socio-economic status (Domhoff, 1988), losing money should be asymmetrically worse for the powerless. Thus, a different perception in valence cannot account for the current findings.

The current data also oppose the notion that the powerful are generally not impacted by potentially negative outcomes that are experienced by other individuals. The container scenario and the RAF scenario both concern other individuals, and the powerful equally showed a negativity bias. Van Kleef et al. (2008) proposed that the powerful are not affected by the suffering of others. The current data show that their compassion might express itself through other measures, as in the current experiments their estimates could have helped prevent the negative outcome from happening in the first place. Certainly, in the present research, the powerful did not “turn a blind eye to the suffering of others” as suggested by van Kleef et al. (2008, p. 1315), at least not to potential suffering in the future. Importantly, individuals high in social power are more likely to make impactful decisions, which may sometimes also be based on their estimates of the occurrence of negative events. The fact that we demonstrate that these

individuals do not, as previously suggested, engage in optimistic decision-making and ignore the negative consequences, is a reassuring finding.

In conclusion, the present research provides support for the notion that loss function asymmetries are underlying the interdependence of negative utility and probability estimates. We demonstrated for the first time that real decision-control in an experiment with real outcomes and incentives is a mechanism of the impact of negative utility on probability estimates. Furthermore, the present research showed that powerful individuals might, in some circumstances, be more sensitive to negative information than the powerless, providing an important boundary condition to the approach-avoidance theory of power. It is this boundary condition of the greater approach motivation of the powerful that is also the content of the next chapter, where we examine the impact of negative affective states on powerful and powerless individuals.

Chapter 5 Joint Impact of Power and Negative Affect on Approach – Avoidance Motivation

Chapter Overview

According to the approach-avoidance theory of power, powerholders are approach motivated, and their powerless counterparts are avoidance motivated (Keltner, Gruenfeld & Anderson, 2003). As a result, the powerful should be *less*, and the powerless should be *more* responsive to negative affective information. Chapter 4 provided some evidence that the powerful show a greater negativity bias when estimating probabilities than the powerless, thus providing a boundary condition to the approach-avoidance theory of power. The current chapter will examine another potential boundary condition of this theory, and examine whether the powerful will also react more to affordances of negative affective states. Recent research has shown that the powerful enact internal states in line with situated experiential information (SFTP, Guinote, 2007). As negative affective states typically lead to avoidance motivation, they can help set apart the accounts of approach-avoidance motivation and of the greater enactment of experiential information. The approach-avoidance theory predicts that due to their greater responsiveness to negative information, the powerless should be more avoidance motivated than the powerful under negative affect. On the other hand, the SFTP predicts that the powerful, acting on information accessible in the situation, become more avoidance oriented than the powerless under negative affect. In the present chapter, we describe three experiments examining the impact of negative

affective states on approach-avoidance motivation in powerful and powerless individuals.

Introduction⁴⁵

“Nature has placed mankind under the governance of two sovereign masters, pain and pleasure. It is for them alone to point out what we ought to do, as well as to determine what we shall do”

(Bentham, 1779/1879 p.1)

Previous research has associated social power with wealth (Domhoff, 1998), health (Carney, Cuddy, & Yap, 2010) and happiness (Keltner et al., 2003). Power facilitates acting on rewards and opportunities, leading to greater action orientation (Galinsky, Gruenfeld, & Magee, 2003). For example, the powerful negotiate more (Magee, Galinsky, & Gruenfeld, 2007) and are more risk taking and optimistic (Anderson & Galinsky, 2006, but see Chapter 4 of the present thesis). These behavioural and affective tendencies of the powerful have been attributed by the most predominant framework of power to the powerful’s greater approach motivation (Keltner et al., 2003). According to this theory, the powerful should be less, and the powerless more responsive to negative information, due to increased approach and avoidance motivation linked to a greater responsiveness to reward and threat, respectively.

Experiments in Chapters 2 and 4 have demonstrated that there is evidence to believe that individuals high in social power can, under certain circumstances, also be more responsive to negative information. In the previous chapters we demonstrated that when estimating the likelihood of negative events, powerful individuals were *more*

⁴⁵ The work in this chapter was conducted under the supervision of Ana Guinote, who I wish to sincerely thank for her ideas and input that led to the research presented here.

sensitive to negative information than powerless individuals, assigning greater likelihood estimates to negative compared to neutral events, whilst this was not the case for the powerless. This could hint towards a greater avoidance of the powerful of the negative outcome than the powerless.

The present research aims to further examine the relationship between social power and negative stimuli and to distinguish between different theories of social power in the domain of negative affect. Chapter 4 examined how social power impacts judgments about negative information. In the current experiments, we aim to scrutinize how powerful and powerless individuals act in line with the affordances of being in a negative affective state: in everyday life, individuals in a powerful social position will not just have to make decisions that have potentially negative consequences and try to avoid these consequences, they will also actually experience negative affect. Important to the current aims, negative affective states are usually associated with avoidance motivation (Gray, 1990). Thus, the powerful, who have been shown to be more approach motivated than the powerless (Keltner, Gruenfeld & Anderson, 2003, Smith & Bargh, 2008), are in everyday life also confronted with negative affective states that signal avoidance motivation.

The focus on negative affective states derives from its theoretical interest: on one hand, individuals high in approach are less reactive to negative affect (Larsen & Ketelaar, 1989) and therefore power could serve as a buffer against negative affect, due to the powerful's greater approach motivation. According to the approach-avoidance theory of power, the powerless should be more avoidance motivated under negative affective states than the powerful. On the other hand, according to the SFTP (Guinote,

2007), power increases attentional focus on internal states (Guinote, 2007a, 2010, see also Moeini-Jazani et al. 2014), and as a result, the powerful enact internal states and automatic responses in line with situated experiential information (Guinote, 2007). Therefore, power could also lead to greater enactment of negative affective states and to greater avoidance motivation compared to the powerless. Thus, negative affective states could potentially serve as an important boundary condition to the powerful's greater approach motivation.

The next sections will illustrate the relationship between emotions and approach-avoidance motivation and review the literature on power and affect in light of the approach-avoidance theory of power (Keltner et al. 2003) and the SFTP (Guinote, 2007).

Affect and Approach-Avoidance Motivation

Affect is defined as subjective states that can be either positive (happy, cheerful, calm) or negative (sad, disgusted, angry) (Gray & Watson, 2001)⁴⁶. Affective states have been argued to have the strongest influence on the appraisal of a target stimulus (Clore, 1992). For example, when asked how much an individual is liked, one most likely bases one's answer on evaluations elicited by affect (Schwarz & Clore, 1983). Affect therefore serves the adaptive purpose to inform quickly about the value of a target stimulus (mostly rewarding or threatening in nature). Importantly, affect precedes motivational states (Tomkins, 1970; Zeelenberg, Nelissen, Breugelmans & Pieters, 2008). Signalling positive or negative appraisal, affect facilitates a fast approach or

⁴⁶The term "affect" in this article includes emotions as well as moods. Emotions focus on evaluative properties of specific objects (Ortony, Clore & Collins, 1988) whereas moods have no specific onset or object to refer to (Clore, 1992).

avoidance response (Frijda, 1986). For example, upon seeing a snake, anxiety triggers the action orientation to quickly move away.

Thus, approach and avoidance motivation can be triggered by affective states. These motivations are facilitated by two neurological systems: the behavioural approach (BAS) and inhibition systems (BIS). Based on research led by Gray (1972, 1981), the BAS has been associated with reward seeking behaviour, approaching novel situations, and impulsivity (Gray, 1990). The BIS is associated with vigilance, anxiety and inhibition (Gray, 1990). Evidence supports the notion that the approach-avoidance systems are activated by affective states. For instance, upon viewing disgust or anxiety evoking TV news images, participants were faster to respond with an avoidance motor movement, compared to when participants saw an anger evoking image (associated with approach motivation), where instead they were more likely to respond with an approach movement (Newhagen, 1998, see also Plutchik, 1998). In sum, different affective states signal different motivational orientations.

Power and Affect

The following will review the current literature on power and affect and discuss the predictions about the joint effects of power and affect in light of the approach-avoidance theory of power (Keltner et al. 2003) and the SFTP (Guinote, 2007).

Approach Avoidance Theory of Power. Proposing the approach-avoidance theory of power, Keltner et al. (2003) argue that enhanced power is associated with living in reward rich environments, freedom to act without consequences and striving autonomously for goals (Galinsky et al. 2008). Powerless individuals on the other hand have decreased access to resources and live in an environment dominated by punishment

and threat. According to Keltner et al. (2003), these asymmetries in exposure to reward and threat lead to the greater relative activation of the BAS for the powerful, and the greater relative activation of the BIS for the powerless. This leads to the powerful being more approach, and the powerless being more avoidance oriented.

Following the close links of affect and approach-avoidance motivation, the approach-avoidance theory of power predicts that power should increase the frequency and intensity of the experience of positive affect, associated with BAS (Keltner et al. 2003). On the other hand, powerlessness should facilitate the experience and intensity of negative affect, associated with BIS. These predictions have received some moderate support. For example, Langner and Keltner (2003) showed that powerful individuals were more likely to report positive and the powerless negative affect (see also Berdahl & Martorana, 2006; Gonzaga, Keltner, & Ward, 2008; but see Fast et al., 2009; Smith & Trope, 2006).

Heightened BAS has been shown to buffer against negative mood induction. For example, Larsen & Ketelaar (1989) induced positive and negative affect in individuals high and low in extraversion and neuroticism, associated with greater approach and avoidance motivation, respectively. The change in negative (positive) mood was only significant for individuals high in neuroticism (extraversion), but not for those high in extraversion (neuroticism). Thus, following the approach-avoidance theory, the powerless should act on these affective states more than the powerful. It is noteworthy that the approach-avoidance theory itself also reports boundary conditions where the powerful can become avoidance oriented. As one boundary condition the theory names threat to social hierarchies and accountability –negative affect or conditions deriving

from the threat to the position of power itself. Nonetheless, negative affective states elicited independently of a threat to one's power position should be more relied on by the powerless.

On the other hand, current evidence supports the notion that the greater BAS activation of the powerful is merely a *bias* in information processing (Guinote, 2007a). For example, the approach-avoidance theory predicts greater heuristic processing strategies of the powerful in line with positive affect, and greater systematic processing of the powerless in line with negative affect (Keltner et al., 2003). However, Guinote, Weick, & Cai (2012) demonstrated that the powerful also engage in systematic processing, and change processing styles depending on the affordances of the situation. This flexibility in processing styles points towards the possibility that the approach tendencies of the powerful can be overridden by situational cues, should the situation afford this. Moreover, new developments in the area of approach-avoidance motivation, as for example the revised Reinforcement Sensitivity Theory (rRST, Gray & McNaughton, 2000) predicting that BIS is only active when an approach-avoidance conflict is detected, have been disregarded by the approach-avoidance theory of power. These highlighted shortcomings of the predictions of the approach-avoidance theory weaken the predictions about power and negative affect by Keltner et al. (2003), and will be discussed in more detail below.

Situated Focus Theory of Power. The SFTP (Guinote, 2007a) adds the components of flexibility and functionality during goal pursuit with regards to strategies used by the powerful. A greater sense of control and freedom of constraints enables the powerful to use information in their environment in order to pursue goals. Thus, the

powerful rely on constructs that are in line with affordances and goals of the situation, and engage in “moment-to-moment” cognition. This results in a greater flexibility of means to pursue goals (Guinote, 2007b), predicting greater situational influences and behavioural variation (Guinote, Judd, & Brauer, 2002) in the powerful compared to the powerless. The powerless on other hand experience constraints, and a loss of control (Fiske, 1993). The powerless distribute their attention between multiple sources of information in their environment (Guinote, 2007a) and in order to regain control act on a greater variety of cues than the powerful (Guinote, 2010). The influence of situational factors is therefore decreased compared to the powerful. The SFTP therefore predicts that powerful individuals rely more on experiential information. Experiential information can be divided into bodily feelings, thought processes, and affective states (Schwarz & Clore, 1996), and whilst earlier research supports the hypothesis that the powerful enact experiential information more than the powerless in the domain of bodily feelings (hunger, Guinote, 2010) and thought processes (ease of retrieval, Weick & Guinote, 2008), the predictions of the SFTP in the domain of affective states have yet to be established.

The SFTP predicts that the powerful compared to the powerless will act more in line with negative affective states: power increases the focus on accessible constructs (Guinote, 2007), increasing the saliency of affect. Since negative affect activates the BIS, according to this perspective powerholders in negative emotional states should be more avoidance oriented compared to control and powerless participants. The SFTP reasons that this happens in spite of the approach orientation (a bias in information processing) of the powerful under neutral mood, as affordances of the situation can

override these biases in the powerful. On the other hand, the SFTP predicts that the powerless, who are typically more avoidance motivated than the powerful, should be *less* avoidance motivated under negative affect compared to the powerful. It is reasoned that the powerless show decreased reliance on internal states due to greater attention paid to threatening information and seeking to regain control. This leads to decreased prioritisations of internal goals and states.

Empirical evidence supports the SFTP with regards to situational influences of bodily sensations and affect. Hsee, Hatfield, Carlson, & Chemtob (1990) asked either a powerful participant to watch a video of a powerless participant displaying positive or negative emotions, or a powerless participant to watch a powerful participant. The powerful were more likely to display the powerless' feelings than the other way around⁴⁷. In addition, the notion that the powerful's approach orientation can be overwritten by negative states was supported by a study by Maner, Gailliot, Menzel, & Kunstman (2012). Maner et al. (2012) showed that individuals high in trait anxiety were not approach orientated when primed with power. Moreover, the powerless' decreased acting on their internal states in the domain of affect was demonstrated by Hecht & LaFrance's (1998) research showing that powerful individuals' smiling correlates with their true affective state, whilst this was not the case for the powerless.

Summing up, the approach-avoidance theory of power and the SFTP differ in their predictions about the joint effects of power and negative affect on approach-avoidance motivation. The approach-avoidance theory predicts greater avoidance motivation of the powerless under negative affect whilst the SFTP predicts that the

⁴⁷ Whilst the difference was statistically significant for positive emotions, a trend was observed for sadness in both subjective ratings as well as judges' ratings of facial expressiveness, too. Unfortunately, the authors do not report p-values for these effects.

powerful should become more avoidance motivated compared to powerless and control participants, acting on accessible experiential information.

The Current Research

The present research sought to examine the joint effects of power and negative affective states (sad mood, disgust, fear) on BIS and BAS activation in powerful versus powerless individuals. In this initial set of experiments, we chose to examine basic emotions (Ekman, 1992). These innate, “primary” emotions (happiness, fear, anger, sadness, surprise and disgust) show more distinct and concrete features than the more complex, “secondary” emotions⁴⁸ (such as pride and embarrassment) (Rainville, Bechara, Naqvi, & Damasio, 2006). Three basic emotions are associated with avoidance motivation: sadness, disgust, and fear (Gray, 1990), and the impact of these emotions on BIS and BAS activation in powerful and powerless participants will be assessed in Experiments 12, 13 and 14, respectively.

Experiment 12

This experiment examined the reliance on sad mood with regard to attention allocation to threatening and rewarding words. Sad affective states have been linked to an increase in avoidance motivation (Gray, 1981, 1990). An active BIS in turn is associated with attention allocation toward threatening stimuli (MacLeod, Mathews & Tata, 1986; Mogg et al., 2004) whereas an active BAS is associated with attentional biases toward rewarding stimuli (e.g. Tamir & Robinson, 2007). If greater approach motivation buffers the powerful against the impact of negative affective states, it would

⁴⁸ Primary emotions are innate and universal, whereas complex emotions are socially determined and arise from higher cognitive processes (e.g. Damasio, 1994)

be expected that the powerless but not the powerful show a threat related bias under sad affective states (whilst showing a reward bias under neutral affect). If, however, the powerful act more on accessible experiential information, they should become more avoidance oriented than the powerless when under the influence of sad affect, displaying a greater attentional bias toward threat.

Conversely, sad affective states have also been associated with a decrease in reward orientation (Carver, 2004). Therefore, we included reward related words to explore a third possibility: if the powerful are more approach motivated (Keltner et al. 2003) they should display a bias toward rewarding information under neutral mood. However, if sadness leads to a decrease in approach motivation, it would be possible to observe a decrease in attentional bias toward rewarding information under sad mood for the powerful but not the powerless.

In the present research, power and mood were manipulated simultaneously (see Lammers, Galinsky, Gordijn, & Otten, 2008) and attentional biases toward threatening and rewarding words were assessed in a dot-probe paradigm (MacLeod, Mathews, & Tata, 1986). In the dot-probe task participants see two words simultaneously on the screen, one word on the left, and one word on the right hand side. One of the two words is always neutral (e.g. “table”), whereas the other one is either associated with reward (e.g. “money”), or with threat (e.g. “murder”). After the words disappear, a dot replaces either the word on the left or on the right hand side, and participants are instructed to indicate its location as quickly as possible. Here, participants can either be slower to react when the dot replaces the neutral word, meaning that they shift their attention away from the rewarding or threatening word, or can indeed be quicker, which means

that they attend away from the valenced word before. As such, the dot-probe task assesses both attention *toward* as well as *away from* valenced (here: threatening or rewarding) words. Importantly, it is the former, attending towards threat that has been associated with avoidance motivation.

Method

Participants and Design. 100 Participants were recruited for a payment of £3 (71 female, median age =22). Participants were randomly assigned to a 3 (power: powerful/powerless/control) x 2 (mood: sad/neutral) between participants design.

Materials and Procedure. Upon arrival participants were informed that they would participate in two separate studies. First, power and mood were manipulated by means of a modified past recall prime (Galinsky et al., 2003), adapted from Lammers et al. (2008). Depending on their condition, participants were asked to recall a situation in which they felt powerful or powerless *and* were in either a sad or a neutral mood. In the control condition, participants were asked about a time where they were at home and either in a sad or a neutral mood. A manipulation check ensured that the priming procedure was effective, asking how much in control they were in the situation that they described (1= not at all to 9= very much), and how they felt in the situation (1 = very sad to 9=very happy). Participants then completed a dot probe task (adapted from Tamir & Robinson, 2007). Word pairs were presented on the screen, one to the right and one to the left side of the centre (10 threat, 10 reward and 20 neutral words, matched for length and frequency (selected from Tamir & Robinson, 2007). The words were presented in font size 18 and were spaced approximately 10 cm from the centre of the screen on the horizontal axis. Afterwards, a “probe” appeared randomly on the left or right hand side

of the screen. Participants were instructed to press “1” if the dot was on the left, and “3” if the dot was on the right. In order to increase attention, participants were told that there might be a memory recall of the words afterwards, but that their main goal was to respond as quickly as possible to the probe (Mogg & Bradley, 1998). One trial consisted of a fixation cross presented for 1000ms, before the word pairs were presented for 500ms, followed by the dot probe which remained on the screen until a response was detected or until 1000ms had elapsed. A practice trial consisting of 10 word pairs preceded 240 experimental trials. Last, participants were thanked and debriefed.

Results

One participant did not complete the power manipulation and was henceforth excluded from any further analyses.

Manipulation Checks.

Power Manipulation Check. Three independent samples t-tests were run on the amount of control reported (see Table 5.1 for means and standard deviations). There was a significant difference between power and powerlessness, $t(52.13)=9.49, p<.001$ (degrees of freedom adjusted due to inequality of variances), powerful participants indicated being more in charge than powerless participants. Powerful people also felt more in charge than control participants, $t(46.21)=4.12, p<.001$ (degrees of freedom adjusted due to inequality of variances), who in turn felt more control than the powerless, $t(65)=3.6, p<.001$. To check for potential effects of mood on the amount of power reported, a 3 (power: powerful/powerless/control) x 2 (mood: sad/neutral) ANOVA was run separately. A main effect of mood, $F(1,93)=10.15, p<.01, \eta^2=.1$, and a marginally significant power*mood interaction, $F(2,93)=2.44 p=.09, \eta^2=.05$, were

followed up by simple effect analyses . Individuals in the control group reported a loss of control when under sad mood, $F(2, 93)=13.15, p<.001, \eta^2=.12$, whilst this was not the case for individuals high or low in power, all $ps>.25$ (see Table 5.2). Importantly, there was no loss of control for powerful individuals under sad mood.

Mood Manipulation Check. A between samples t-test was run to ensure a successful manipulation of mood (see Table 5.2). Participants describing a sad situation perceived this as more sad ($M=2.78, SD=1.36$) than participants in a neutral condition ($M=5.54, SD=1.79$), $t(97)=8.65, p<.001$. A 3 (powerful/powerless/control) x 2 (sad/neutral mood) between subjects ANOVA was conducted. A main effect of power on the mood manipulation check was found, $F(2, 93)=8.86, p<.001, \eta^2=.16$, however, no power*mood interaction was present $F(2, 93)=1.23, p=.3$. Thus, the sad mood manipulation had a similar impact on individuals in different power conditions.

Table 5.1

Amount of control reported for powerful, powerless and control participants under neutral or sad mood.

	Sad	Neutral	Overall
Powerful	7.19 (1.28)	7.69 (1.08)	7.44 (1.19)
Control	4.13 (2.10)	6.56 (2.28)	5.46 (2.49)
Powerless	3.06 (1.83)	3.81 (2.46)	3.41(2.14)

Table 5.2

Mood reported for powerful, powerless and control participants under neutral or sad mood.

	Powerful	Control	Powerless
Neutral	6.56 (1.59)	5.56 (1.62)	4.5 (1.63)
Sad	3.38 (1.15)	2.6 (1.68)	2.39 (1.09)
Overall	4.97 (2.11)	4.21 (2.2)	3.38 (1.72)

Dot Probe Task.

Scoring. The accuracy rate was high with 97% averaged across participants (97.16% powerful, 97.03% powerless, 96.25% control). Three participants were removed from the analysis for having error rates above 25% of the overall trials.

Following the procedure of Ratcliff (1993) inaccurate trials and trials with reaction time (RT) points above or below 2 standard deviations of the individual mean were not considered for analysis⁴⁹. Additionally, the data was log transformed to obtain a more natural distribution (see Tamir & Robinson, 2007). Next, we obtained scores for the reward and threat bias. To obtain the score for the threat bias, first the reaction time when the probe replaced the neutral word and then when the probe replaced the threat related word was computed. Subsequently, the latter was subtracted from the former (Macleod et al. 1986). The same calculations were done for reward related words; leading to both a reward bias and a threat bias.

Note that the dot probe task distinguishes between two types of biases: vigilance *toward* a stimulus (a slower response when the probe is on the side opposing the

⁴⁹ Note that we did not remove outliers based on Studentized Deleted Residuals and Cook's Distance as in previous experiments. Instead, we followed the exact procedure described by Tamir and Robinson (2007).

location of the valenced stimulus) and avoidance of the stimulus (a faster response when the probe is on the side opposing the valenced stimulus). Importantly, it is vigilance toward a threatening stimulus that is associated with avoidance motivation, but not the avoidance of the stimulus.

Repeated Measures ANOVA. To test the full model, a 2 (reward bias/threat bias) x 3 (powerful/powerless/control) x 2 (sad/neutral mood) repeated measures ANOVA was run. The three-way interaction between valence, power and mood was marginally significant, $F(2,90)=2.34$, $p=.10$, $\eta^2=.05$.

Attention to Threat. A 3 (power: powerful/powerless/control) x 2 (mood: sad/neutral) ANOVA was run on the threat bias score. A significant power*mood interaction was found, $F(2,90)=3.8$, $p=.03$, $\eta^2=.08$, (see Figure 5.1). Neither the main effect of power, $F(2,90)=.18$, $p=.16$, $\eta^2=.04$, nor the main effect of mood were significant, $F<1$.

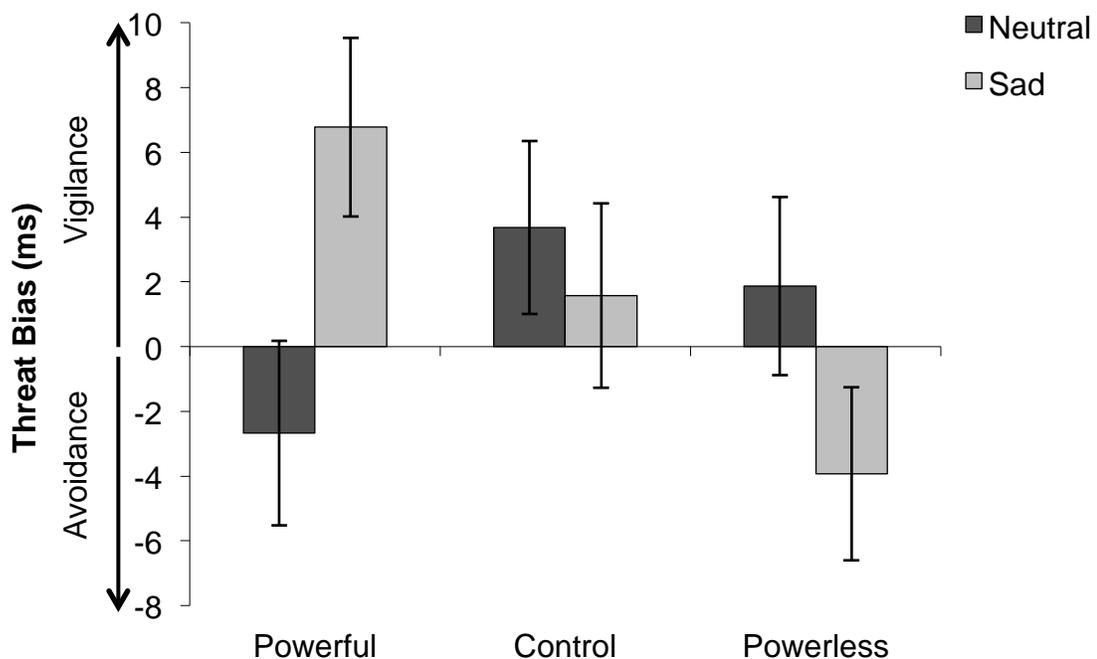


Figure 5.1. Threat bias scores for participants in a powerful, control or powerless condition under neutral or sad mood. Positive bias scores indicate attention *toward* the threat related word, whereas negative bias scores indicate attention *away* from the threat related word. Error bars represent +/- 1 standard error of the mean.

As expected, the differences in threat bias as a function of mood was significant for powerful participants, $F(1,90)=5.29, p=.02, \eta^2=.06$, but not for powerless, $F(1,90)=1.37, p=.26$, or control participants, $F(1,90)=1.02, p=.32$. This finding supported our hypothesis that powerful individuals are more likely to act on affective information, here displayed through greater flexibility in attentional biases as a function of mood.

Importantly, for powerful sad participants the bias was significantly different from zero, $t(15)=2.5, p=.03$, showing vigilance towards threatening stimuli. For no other conditions were the deviations from zero significant, all $ps > .10$.

Attention to Reward. To examine the impact of power and mood on attention to rewarding information, a 3 (powerful x powerless x control) x 2 (sad x neutral mood) ANOVA was carried out on the reward bias scores. The overall ANOVA was not significant, $F < 1$ and neither were any main effects or interactions. This finding is in line with recent research by Moeini-Jazani et al. (2014), showing that powerful individuals only pay attention to rewards once in a hot cognitive state. Under neutral conditions, Moeini-Jazani et al. (2014) failed to find differences in reward sensitivity between individuals high and low in power.

Experiment 13

Experiment 12 provided some initial evidence for the notion that powerful individuals become more avoidance oriented under negative affect. Experiment 13 aimed to extend this finding to a different negative affective state (disgust) and to a different measure of approach-avoidance motivation (BIS-BAS scales, Carver & White, 1994).

Disgust is defined as an aversive reaction towards potential contagious sources (Rozin, Haidt, & McCauley, 2000) and leads to an increase in avoidance motivation (Reuter et al., 2004). In line with the findings of Experiment 12, we hypothesized that powerful participants would become more avoidance motivated after having been exposed to disgusting stimuli compared to powerless participants. Power was manipulated by means of a past event recall (Galinsky et al. 2003). Disgust was manipulated by displaying disgust eliciting pictures. Afterwards, participants completed a version of the BIS BAS scales (Carver & White, 1994), modified to assess *state* rather

than *trait* approach-avoidance motivation to capture the impact of manipulated power and mood on this variable.

Method

Participants and Design. 211 participants whose first language was English⁵⁰ where recruited by Amazon Mechanical Turk for a monetary incentive of \$0.6 and completed this experiment online. Six participants were excluded for not completing the power manipulation. The final sample consisted of 205 participants (119 female, median age=23). Participants were randomly assigned to a 2 (power: powerful/powerless) x 2 (emotion: disgust/neutral) between participants design.

Procedure and Materials. Participants first completed the power manipulation (Galinsky et al. 2003), which asks participants to describe either an event in which they had control over another person (powerful) or in which someone else had control over them (powerless). Afterwards disgust was manipulated by showing participants 6 pictures from the International Affective Picture System (IAPS; Lang, Bradley, & Cuthbert, 2005) displaying disgusting or neutral events (Borg, De Jong, Renken, & Georgiadis, 2012). This was ostensibly introduced as a pre-test for a future study. As a manipulation check, participants answered 4 questions from Tolin, Lohr, Sawchuk and Lee (1997), designed to measure a subjective disgust feeling, avoidance reaction, physiological responses and contamination appraisal, respectively. Participants were asked to indicate their level of disgust from 0 (not at all true) to 10 (very true) answering the statements “This picture makes me feel disgusted”, “This picture makes me feel like

⁵⁰ Only participants whose first language was English were eligible for participation. Thus, anyone signing up for the experiment whose first language was not English, was not considered.

pushing the picture away from me”, “This picture makes me feel sick to my stomach” and “This picture makes me feel like I might be contaminated or infected”.

Afterwards, in order to measure approach-avoidance orientation, the BIS-BAS questionnaire (Carver & White, 1994) was employed. This questionnaire consists of 24 questions out of which 7 questions measure avoidance, and 13 items approach motivation (4 filler items are included). Participants read statements such as “I crave excitement and new sensations” and rate how much they agree with the statements on a scale from 1 (very true for me) to 4 (very false for me). However, we adapted the original questionnaire to measure participants’ current state rather than trait approach and avoidance motivation by adding words such as “at the moment” “currently” or “now” to the original items (see Lammers et al., 2008) (see Appendix I). For example, the item “When I want something I usually go all-out to get it” was changed to “If I wanted something at the moment, I would go all-out to get it”. This was done as social power and disgust were manipulated, and should therefore impact state approach-avoidance motivation.

Results

Manipulation Checks. Participants indicated how much in charge they felt in the situation they described and how much influence they had. These items were averaged (Cronbach’s Alpha = .98) and submitted to an independent samples t-test. The manipulation of power was successful, $t(203)=28.91, p<.001$.

In a similar fashion, the 4 items serving as the emotion manipulation check were averaged (Cronbach’s Alpha =.97) and entered to an independent samples t-test. The manipulation of disgust was successful, $t(203)=26.49, p<.001$.

The 24 items were averaged to form a BIS and BAS subscale. As in line with the experiments of Chapters 2 and 4 of this thesis, we present the data with the removal of disproportionately influential data points, following the same criteria (plotting Cook's Distance versus Studentized Deleted Residuals)⁵¹.

BIS BAS Scales.

BIS. There were no main effects of power or emotion (all $F_s < 1$), but a significant power*emotion interaction for the BIS subscale was found as predicted, $F(1,189)=5.89, p=.012, \eta^2=.03$ (see Figure 5.2)⁵². Expectedly, a significant difference between powerful participants in the neutral ($M=2.69, SD=0.66$) compared to in the disgust condition ($M = 2.98, SD=0.54$), $F(1,189)=4.9, p=.03, \eta^2=.03$, was found. The difference between the powerless' neutral ($M = 2.82, SD=0.57$) and disgust condition ($M = 2.68, SD=0.58$), did not reach significance, $F(1,189)=1.4, p=.24$. The difference between the powerful and the powerless was significant in the disgust $F(1,189)=5.04, p=.03, \eta^2=.03$, but not in the neutral condition, $F(1,189)=1.25, p=.26$. These findings support the hypothesis that the powerful increase avoidance motivation under negative affect compared to the powerless, as the powerful act in line with internal states (Guinote, 2007a).

⁵¹ Unfortunately, a power manipulation is noisy and outliers are often removed (e.g. Carney, Cuddy & Yap, 2010; Lee & Schnall, 2014; Schmid-Mast, Jonas & Hall, 2009; Scholl & Sassenberg, 2014; Sligte, de Dreu & Nijstad, 2011). We decided to follow Lammers, Gordijn and Otten (2008) and to remove outliers by plotting studentized deleted residuals against Cook's distance (see Cohen, Cohen, West & Aiken, 2003).

⁵² Before the removal of outliers, the power*disgust interaction was not significant, $F(201)=1.91, p=.17$, and neither were the main effects of power or mood, all $F_s < 1$.

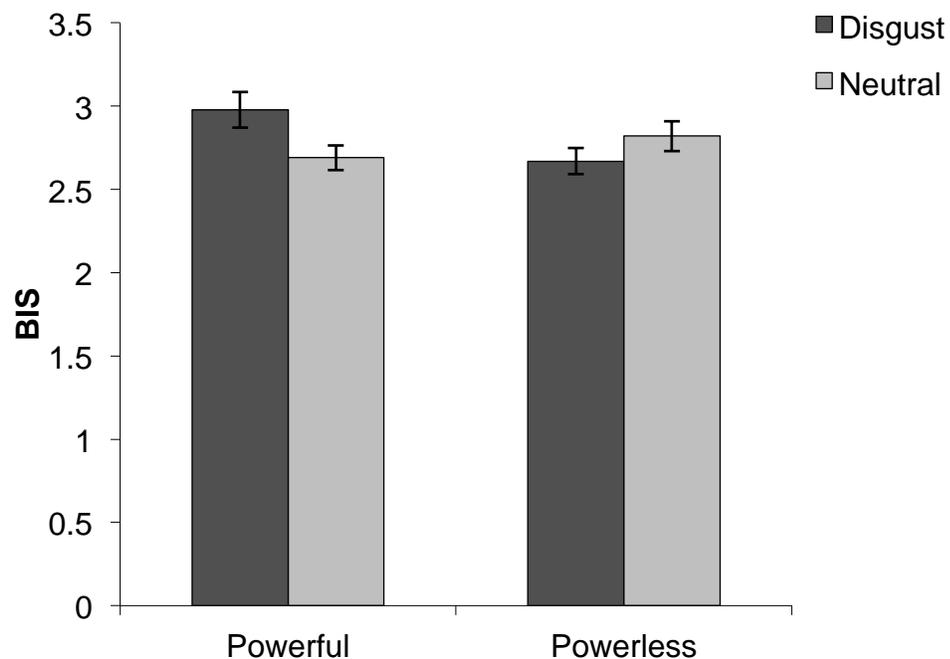


Figure 5.2. Means for the BIS scale for powerful and powerless individuals under a disgusted or neutral affective state. Error bars represent +/- 1 standard error of the mean.

BAS. Surprisingly, participants who saw disgusting stimuli increased in BAS ($M=3.13$, $SD=.42$) compared to participants who saw neutral stimuli ($M=2.92$, $SD=.41$), $F(1,194)=11.49$, $p<.01$, $\eta^2=.06$. Neither the power*emotion interaction, $F(1,194)=1.54$, $p=.22$, nor the main effect of power were significant, $F<1$.

Experiment 14

Experiment 14 extended the findings from Experiments 13 and 12 to the negative affective state of fear and a different power manipulation. Fear was manipulated by asking participants to watch a scary or neutral video clip, and power was manipulated by means of the “day of a manager” manipulation (Guinote, 2007).

Method

Participants and Design. Two hundred and seventy nine participants (103 female, median age=33) were recruited for this online experiment on Amazon Mechanical Turk for a payment of \$0.5. However, participants who did not complete the power manipulation or did not watch the video until the end, as well as participants who took less than 10 seconds to answer the 24 BIS BAS questions were also not considered for analysis⁵³. The remaining 250 participants were randomly assigned to a 2 (power: powerful/powerless) x 2 (emotion: fear/neutral) between participants design.

Materials.

Power Manipulation. Power was manipulated by asking participants to imagine themselves in the role of a manager or a subordinate, and write down what a typical day would look like (Guinote, 2007). The manipulation was slightly adapted to increase its strength. In the high power conditions participants read the following:

This study focuses on social roles. Your task is to read some information about the role of a person in a given social context, and imagine yourself in that role. You will then be asked to describe what a typical day in your life would be if you would be in that particular role.

Role: The managing director in this marketing organization has 20 employees working under him. The organization promotes various products to the public, and the role of the director is to distribute the work that subordinates must complete, set priorities for the team, approve project proposals, and accept or decline new clients. The managing director knows the work well and makes all decisions within the company. He sets priorities and determines the salary and the workload of all employees.

Below, you see the office of the manager.

⁵³ Note that this exclusion criterion is slightly less conservative than in Experiment 13. This was due to there being around 1.5 years in between the collection of these two studies. Following our experience within this time, we included a timer to allow for filtering out participants that were too quick. No such timer was included in Experiment 13.



Figure 5.3. Graphical display of the office in the manager condition.

Please imagine yourself occupying this role of the managing director and describe what a typical day in your life would be, what you would do, how you would feel, and what you would think. You can start from when you went to work to when you left to go back home. You will need to spend 5 minutes on this task, and write at least 500 characters. There are no right or wrong answers, we are simply interested in people's roles in everyday life.

In the low power condition participants read instead:

The employee in this marketing organization works in a team of 20 people. The organization promotes various products to the public, and the role of the employee is to complete any task that the manager assigns to him/her, and to follow instructions regarding priorities in this marketing organization. The employee must also keep records and prepare paperwork for projects and new clients that were approved by the manager. The employee knows the work well and strictly follows the procedures and priorities set by the manager. His or her salary and workload are determined by the manager. Below you see the office of the employee.



Figure 5.4. Graphical display of the office in the subordinate condition.

Emotion Manipulation. Fear was manipulated by showing participants approximately 2 minute long video clips, pre-tested by Hewig, Hagemann, Seifert, Gollwitzer, Naumann and Bartussek (2005). Participants in the fear condition were shown a video clip either from *silence of the lambs* or from *Halloween*, whilst individuals in the neutral condition watched a clip from *Hannah and her Sisters* or *All the Presidents' men*.

Emotion Manipulation Check. In order to verify that the emotion manipulation was successful, participants answered the questions “how did this movie make you feel” (1 very sad – 7 very happy / 1 very anxious – very relaxed) alongside several distractor items (“how familiar were you with the video clip?”, “Would you watch this type of movie at home?”).

Procedure. Participants were informed that they would participate in 3 studies: the first study was introduced as a study concerned about roles in organisational settings, and depending on their condition, participants either completed the high or low power manipulations. Following, participants were asked to imagine that they would come home from their work day, and would now want to watch a video in order to unwind. Participants were then shown the fear inducing or neutral video clips, and afterwards asked to rate the clip on the dimensions described above. Next, in order to increase the manipulation, participants wrote about the morning in the office the next day, still following their power role. Participants then completed the BIS BAS scales, introduced as a study interested in participants' personality. Afterwards, participants completed the power manipulation check, gave demographic information, were thanked and debriefed.

Results

Manipulation Checks.

Power Manipulation Checks. In order to guarantee a successful manipulation of power, participants' answers to the questions "how much influence did you have in the situation you described in study 1" and "How much were you in charge in the workday you described in study 1" were collapsed (Cronbach's Alpha = .93) and submitted to a 2(powerful/powerless) x 2 (fear/neutral) between subjects ANOVA. A main effect of power showed that power was manipulated successfully, $F(1,244)= 166.83, p<.001$. Individuals high in power indicated having more influence and being more in charge ($M=7.8, SD=1.38$) compared to individuals low in power ($M=4.64, SD=2.32$). Moreover, a non-significant main effect of emotion, $F<1$, and a non-significant

interaction between power and emotion, $F < 1$, demonstrate that there was no perceived loss of control when under a negative affective state.

Mood Manipulation Checks. The answer to the question “how did this clip make you feel” (1 very anxious – 7 very relaxed) was equally submitted to a 2 (power: powerful/powerless) x 2 (emotion: fear/neutral) between subjects ANOVA. A main effect of emotion showed that fear was successfully manipulated, with participants in the fear condition indicating that they were more anxious ($M = 2.53$, $SD = 1.34$) compared to participants in the neutral condition ($M = 4.00$, $SD = 1.07$), $F(1,244) = 90.62$, $p < .001$. No main effects of power, $F(1,244) = 1.45$, $p = .23$, and no interaction were present, $F(1,244) = 1.01$, $p = .32$, indicating that powerful and powerless participants did not perceive the video clip differently.

BIS BAS Scales.

BIS. The BIS scores were entered into a 2 (power: powerful/powerless) x 2 (emotion: fear/neutral) between subjects ANOVA. A marginally significant main effect of power was present, $F(1,231) = 2.82$, $p = .09$, $\eta^2 = .01$, with powerful participants scoring higher on the BIS scale ($M = 2.74$, $SD = .58$) than powerless participants ($M = 2.64$, $SD = .52$)⁵⁴. In addition, a marginally significant main effect of emotion was found, $F(1,231) = 3.4$, $p = .07$, $\eta^2 = .02$, with participants in the fear condition scoring higher on the BIS scale ($M = 2.75$, $SD = .55$) than control participants ($M = 2.62$, $SD = .54$). Importantly, a significant power*emotion interaction was found, $F(1,231) = 5.14$, $p = .02$, $\eta^2 = .02$. To further scrutinize this interaction, simple effect analyses were conducted. Only the powerful showed an increase in BIS when under fear ($M_{\text{Fear}} = 2.90$, $SD = .57$ vs.

⁵⁴ Before the removal of outliers, the power*fear interaction was not significant, $F(244) = 1.74$, $p = .19$, $\eta^2 = .01$, and neither were the main effects of power, $F(244) = 1.07$, $p = .30$, $\eta^2 = .004$, or fear, $F(244) = 1.45$, $p = .23$, $\eta^2 = .01$

$M_{\text{Neutral}}=2.60$, $SD=.56$), $F(1,231)=8.18$, $p<.01$, $\eta^2=.03$. No such difference was found for powerless participants, $F<1$. Moreover, when under fear, the powerful were *higher* in BIS than the powerless, $F(1,231)=7.37$, $p<.01$, $\eta^2=.03$. No significant difference was found for the neutral condition, $F<1$.

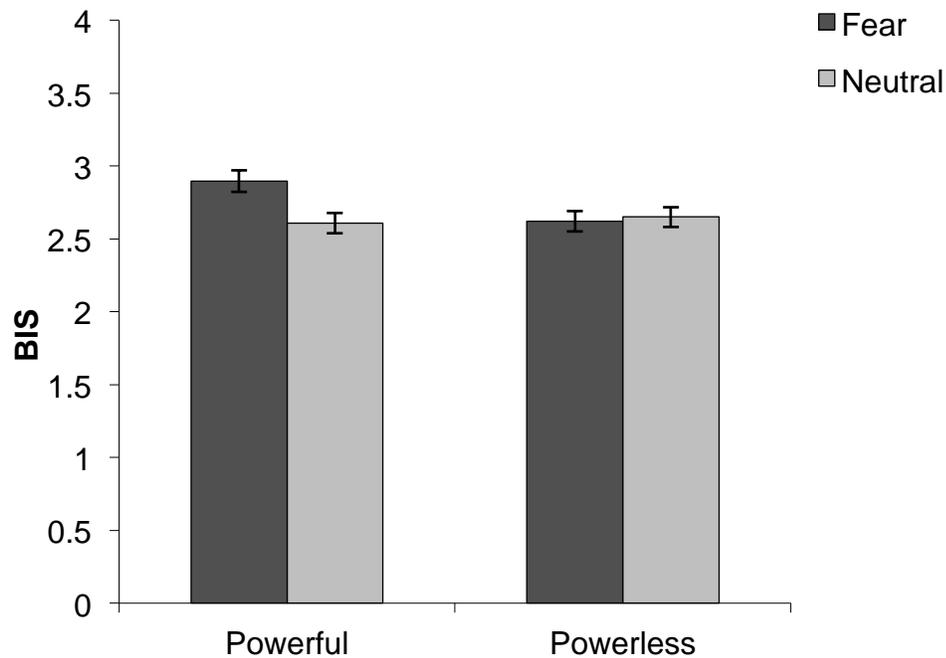


Figure 5.5. Means of BIS for powerful and powerless participants under fear or neutral affective states. Error bars represent +/- 1 standard error of the mean.

BAS. No main effect of power was present, $F<1$. However, there was a marginally significant main effect of emotion, with those in the fear condition reporting lower values of BAS ($M=2.90$, $SD=.46$) than those in the neutral condition ($M=3.01$, $SD=.49$), $F(1,239)=2.94$, $p=.09$, $\eta^2=.01$. In addition, a marginally significant interaction was found, $F(1,239)=3.85$, $p=.05$, $\eta^2=.02$. Simple effects revealed that the powerless experienced a significant drop in BAS under fear ($M=2.84$, $SD=.47$) than when in the

neutral condition ($M = 3.06$, $SD = .47$), $F(1,239) = 6.88$, $p < .01$, $\eta^2 = .03$, whilst there was no difference for the powerful, $F < 1$. However, there were no differences in BAS between individuals high and low in power for neither the neutral or fear conditions for powerful and powerless individuals, all $ps > .12$ ⁵⁵.

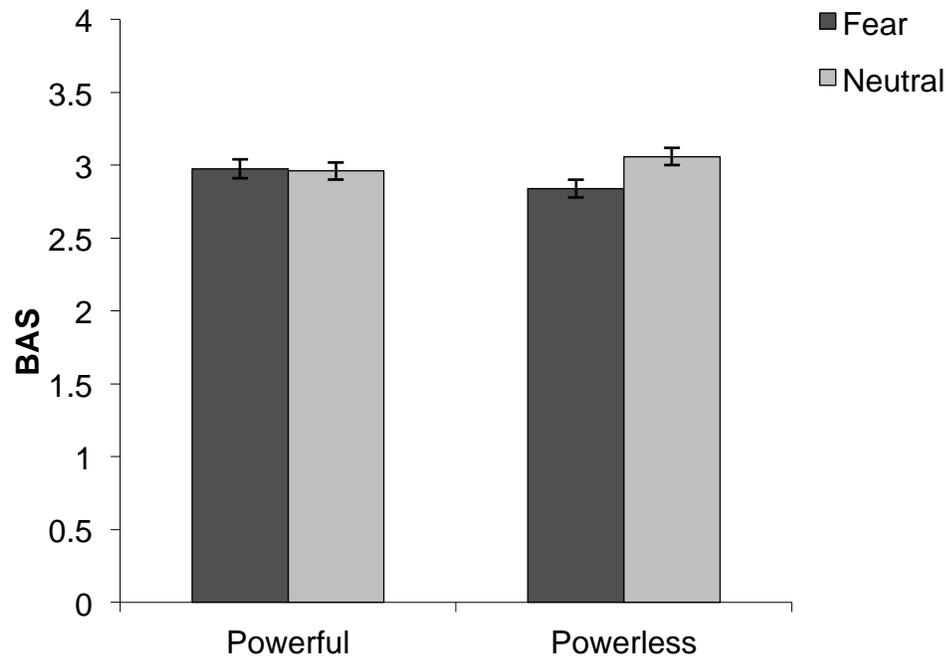


Figure 5.6. BAS means for powerful and powerless participants under fear or neutral affective states. Error bars represent +/- 1 standard error of the mean.

Chapter Discussion

Emotions guide our everyday decision making and goal-pursuit (Ortony, Clore, & Collins, 1988), helping us to navigate through the world with limited cognitive resources. However, the extent to which individuals attend to their internal emotional states has been shown to differ on an individual level (Gohm & Clore, 2000). We

⁵⁵ Before the removal of outliers, there was neither a main effect of power, nor of emotion, all $F_s < 1$, and the power*emotion interaction was marginally significant, $F(1,244) = 2.64$, $p = .11$, $\eta^2 = .01$.

hypothesized that social power moderates the extent to which individuals act on affective information. Powerful individuals have been shown to be under less constraint and evaluation, as well as to display a more narrow attentional focus (Guinote, 2007a). Thus, our contention was that compared to the powerless, saliency of affect will be increased and the powerful will be freed to act on their internal states and default processes (Guinote, 2007a). We hypothesized that the powerful would become more avoidance oriented under negative affect compared to powerless and control participants.

Three experiments supported these claims. In Experiment 12, primed with sadness, powerful participants showed a pronounced vigilance bias toward threatening information in a dot-probe task. In Experiments 13 and 14, after having been exposed to disgust eliciting pictures or scary video clips, respectively, an interaction between power and emotion showed that the powerful are more likely to increase in avoidance motivation after being exposed to disgusting or fearful stimuli (after the removal of disproportionately influential data points), whilst this was not the case for the powerless.

Thus far, it was shown that the powerful show greater reliance on bodily feelings (Guinote, 2010) and cognitive experiences (Weick & Guinote, 2008). However, this is the first time that reliance on internal states was shown in a third and influential domain of experiential information- affective states (Schwarz & Clore, 2007). We reason that the differences in affective influence are due to the powerful's greater enactment of situational information and greater accessibility of affective cues. Compared to powerless and control participants, the powerful show a more narrow attentional focus, increasing the saliency of affective information. Moreover, once a construct is activated,

the powerful then focus on this accessible cue, engaging in “moment-to-moment” cognition. Thus, once affect became accessible, the powerful became avoidant oriented, supporting our hypotheses. Theoretically important is the fact that the powerful also act on accessible cues when this cue is opposing the approach orientation normally observed. This is in line with previous research showing that the powerful usually rely more on dispositional constructs than the powerless, however, once counterdispositional constructs are activated, these differences are overridden (Guinote et al. 2012).

Additionally, it is possible that the powerful also make use of affective information differently, compared to the powerless. Due to greater freedom of constraints, the powerful might be able to be disinhibited and rely on their true affective state (Hecht & LaFrance, 1998), without having to consider social norms. In line with this argument, this research is also supportive of the notion that powerful individuals act in behaviourally more variable ways (Guinote et al. 2002), here displayed through a greater flexibility of utilizing approach and avoidance strategies.

Besides, this research has important theoretical implications by providing initial support for a boundary condition for the approach-avoidance theory of power in the domain of affective states. Whilst the greater approach orientation of the powerful under neutral affective states is well established (Anderson & Berdahl, 2002), we believe that this approach orientation is merely a *bias* in information processing that can be overridden by situational cues, here, affective states. Indeed, the powerful became *more* avoidance oriented under negative affect, and were more avoidance oriented than the powerless or control groups. The approach-avoidance theory predicts avoidance orientation for the powerful only in case of a threatened power position, for example

through instability of the social hierarchy. We were able to demonstrate that merely experiencing sad mood (Experiment 12) did not lead to a loss of control for the powerful, which, according to this account should be the only prerequisite for the powerful to become more inhibited.

However, an alternative explanation for the effects found for the powerful should be considered, taking into account the revised reinforcement sensitivity theory (rRST) that remains un-integrated into the approach-avoidance theory. The rRST (Gray & McNaughton, 2000) extends the approach-avoidance systems by the fight-flight-freeze system (FFFS). The FFFS is responsible for the avoidance and withdrawal of aversive stimuli. The BIS, however, is related to conflict detection between BAS and the FFFS, leading to inhibition of on-going behaviour and behaviour in line with FFFS. Assuming that power does lead to an increase in BAS activity, this should lead to a greater conflict between BAS and negative affective states, compared to the powerless who are already more avoidance oriented. Thus, the rRST predicts that those highest in BAS should have the greatest BIS responses towards aversive stimuli, since the response conflict between BAS and the FFFS is the greatest (Berkman, Lieberman, & Gable, 2009). This is in line with findings showing that participants high in BAS but not BIS incorporate negative affect in their decision making process (Kramer, 2007).

Whilst the emphasis of this work was on the powerful, there were (non-significant) trends for the powerless to show reactions unpredicted by both the SFTP and the approach-avoidance theory. The powerless attended away from threatening information (directionally, Experiment 12), decreased in avoidance motivation (directionally, Experiment 13) as well as to significantly decrease in approach

motivation (Experiment 14)⁵⁶. The SFTP predicts that the powerless are more likely to consider peripheral information and are subject to influence of multiple factors in their environment, thus showing decreased reactivity to single situational influences.

However, this could not be supported, since the powerless showed also some tendencies to be reactive to negative affect. Instead, they suppressed the response as far as to even (directionally) decrease in avoidance orientation.

Neither could the predictions of the approach-avoidance theory be supported. Keltner et al.'s (2003) predictions state that due to their increased avoidance motivation the powerless are vigilant to threat and express more negative affect compared to the powerful. However, the powerless were neither more avoidance oriented under neutral, nor under negative affect. Instead, these tendencies of the powerless to pay less attention to negative information and to decrease in avoidance are congruent with a "repressive coping style" (Weinberger & Davidson, 1994). This form of coping has been shown to lead to avoidance of negative information, aiming to decrease the experience of negative affect in the face of threat. This reactivity of the powerless to negative affect is different from reactions to other experiential information (Guinote, 2010), arguing for a special case of affect that remains to be explored. It is possible that the powerless aim to escape further strain that is already caused by powerlessness. However, repressing negative affect and diverting attention away from threat comes at a cognitive and physiological cost (Weinberger, 1990), decreasing resources for focal tasks.

⁵⁶ We are unable to explain why the powerless decreased in avoidance in Experiment 16 (directionally), but significantly decreased in approach in Experiment 17. As both fear and disgust are only related to the avoidance dimension of approach-avoidance motivation, the decrease in approach is unexpected. However, there has been no research observing the impacts of a repressive coping style on approach-avoidance motivation, which could potentially identify some still unknown moderators.

Implications

This research highlights the advantages that the powerful have compared to the powerless when it comes to emotion regulation. The powerful are seemingly advantaged by their use of affect, helping to speed up action plans. Integrating affect into cognition is highly adaptive, helping us to coordinate in a world with limited resources. However, the described effects of the powerful occur in an area with consequences that can potentially threaten one's power position: Tiedens (2001) showed that expressing avoidance related affect can decrease one's status in a social hierarchy. Being more avoidance oriented can harm the powerful's legitimacy. What remains to be examined is in how far powerful individuals suppress expression of negative affect when interacting with low-power individuals.

If low power individuals really were to engage in repressive coping, they would be disadvantaged by their style of emotion regulation. Not expressing affect has detrimental health effects (Pennebaker & Beal, 1986), for instance, repressive coping heightens mortality as a result of cancer (McKenna, Zevon, Corn, & Rounds, 1999).

Further Research

Further studies could aim to extend the current findings to different types of negative affective states, for instance to secondary (e.g. shame, guilt) rather than primary emotions as in the current research. In order to guarantee the universality of the current findings, approach related affect such as anger should be explored further, where it would be expected that the powerful would become more approach oriented.

In addition, we did not provide a concrete mechanism underlying the current effects. Self-awareness (Scheier & Carver, 1977) and interoception (Schnall et al. 2008)

have been demonstrated to increase behaviour in line with affective states, and as mentioned in Chapter 1, there is evidence pointing towards a relationship between power and these two constructs (Moeini-Jazani et al. 2014). Thus, both self- as well as interoceptive awareness could be explored as mechanisms in future research.

Moreover, it will be impactful to explore in how far powerful versus powerless rely on mood as information (the tendency to ask oneself how one *feels* about a decision, Schwarz & Clore, 1983) in decision making. We demonstrated in the Chapters 2 and 4 that the powerful are, under some circumstances, more sensitive to negative information and incorporate the utility of an event more into their judgment. However, according to the current findings, the powerful should be more prone to use affect as information when making decisions, and this remains subject to further research.

What is more, it could also be explored whether powerful individuals' goals change depending on their affective state, as affective states not only lead to motivational orientations but also to certain "emotivational goals", that are concrete action plans in the face of each emotional state. In preliminary research, powerful participants, after having been exposed to disgusting stimuli, had the emotivational goal of disgust, cleanliness, more accessible than the powerless (de Molière & Guinote, 2012). Note in passing that this finding could possibly also be in line with our results of Chapter 4. Potentially, the powerful feel a greater control over the negative outcome (here: contamination) and therefore have action plans more accessible than the powerless.

Conclusion

Summing up, this research provided initial support for the notion that powerful individuals act on negative affective states by becoming more avoidance oriented, setting apart the approach-avoidance theory and SFTP in the domain of negative affect. Ironically, when being under the influence of negative affect, this leads to the powerful being more avoidance oriented compared to the powerless, reversing the tendencies under neutral affect. Thus, this research provides a boundary condition to both the approach and action orientation of the powerful, as well as to the greater inhibition of the powerless, when under the influence of negative affect.

Chapter 6 General Discussion

Summary

The main aim of the present thesis was to distinguish mechanisms that underlie the relationship between negative utility and probability estimates by means of social power. In particular, we examined whether social power would moderate the relationship due to the hypothesized mechanisms arousal misattribution and decision-control, which were deduced from the Stake Likelihood Hypothesis (Vosgerau, 2010) and an account of loss function asymmetries (e.g. Harris et al. 2009; Weber, 1994), respectively.

In Chapter 2 we assessed the hypothesis that if arousal misattribution is underlying the impact of negative utility on probability estimates, then interoceptive ability should moderate the effect. That is, only individuals that “feel” the arousal should then misattribute it and demonstrate a negativity bias. In our first experiments we assessed interoceptive ability as a self-report measure and showed that across different scenarios and different interoception scales, only individuals high in interoceptive sensibility increased probability estimates for negative outcomes. However, as the moderation by interoceptive sensibility proved to be less consistent than we had hoped, we next assessed interoception objectively as compared to a self-report measure, and also directly measured arousal. Neither did interoceptive awareness moderate the relationship as assumed, nor did arousal mediate the relationship, as would be suggested by the Stake-Likelihood Hypothesis. Thus, the full model of the moderation by power via the mechanism of interoceptive ability was not assessed in this chapter due to a lack

of evidence for the impact of interoception.

In Chapter 3, following the mixed results of Chapter 2, we took stock of the Stake Likelihood Hypothesis across four replication attempts. We were unable to replicate original findings by Vosgerau (2010), whilst providing quantified evidence for the null hypothesis over the Stake Likelihood Hypothesis. The lack of evidence for arousal misattribution as a mechanism underlying the impact of negative utility on probability estimates motivated the next chapter, where we continued to examine the moderator proposed by an account of loss function asymmetries: decision-control.

Chapter 4 explored the moderating role of social power on the impact of negative utility on probability estimates, where it was hypothesized that powerholders would show a greater negativity bias than their powerless counterparts due to the enhanced perception of decision-control of the former. Both across measured and manipulated social power, we provided some first evidence for the moderation by social power, where powerful individuals showed a stronger negativity bias than the powerless across different scenarios. Importantly, when we directly examined the mechanism of decision-control as a moderator in Experiment 11, the powerful but not the powerless showed a negativity bias for situations without decision-control, whilst these differences disappeared for situations with decision-control. This experiment provides preliminary evidence for our hypothesis that it is the lack of control that prevents, and the greater (illusion of) control that leads to the existence of loss function asymmetries and the resulting interdependence of utility and probability estimates for powerless and powerful individuals, respectively. By demonstrating that participants' perception of the difference of probabilities for a more and a less desired end-state increases with

decision-control, presumably as a means to motivate oneself to act on one's estimates, this experiment also provided more general evidence for loss function asymmetries beyond social power.

The evidence presented in Chapter 4 questioned the well accepted claim from the power literature that *powerless* individuals are more responsive toward negative information due to greater avoidance motivation, and that the powerful are generally more optimistic and risk-taking due to being more approach motivated than their powerless counterparts. Following the findings of Chapter 4, Chapter 5 examined whether powerful individuals would not only incorporate negative information more in their judgments, but would also act more on the affordances of negative affect in terms of the motivational orientations accompanying affective states. Across three experiments we demonstrated that the powerful, when under a negative affective state, become *more* avoidance oriented compared to the powerless, providing an important boundary condition to the greater approach and action orientation of the powerful.

In sum, this research made important contributions across two domains: First, systematically assessing the SLH and ALFs, we found no evidence for arousal misattribution, but found evidence for loss function asymmetries as a mechanism underlying the interdependence of negative utility and probability estimates. Secondly, the current research informs the power literature. We demonstrated that the powerful both show a greater negativity bias than the powerless, and become more avoidance oriented under negative affect. These findings question the generally accepted claim that the powerful are optimistic and responsive to reward, and the powerless pessimistic and responsive to threat. In the following, we will discuss these main contributions in the

light of the current respective literatures as well as to highlight implications and suggest further work.

Mechanisms Underlying the Interdependence of Utility and Probability

Taking a very broad perspective on the current research, we demonstrated that the negativity bias is truly a robust effect. In almost all experiments, we observed a significant impact of negative utility on probability estimates, and in none of the experiments did participants provide lower estimates for negative compared to neutral events, not even directionally. Thus, the current results strongly support the findings by Harris et al. (2009) and Vosgerau (2010), and add to a body of evidence showing that when objective probabilities and therefore the evidence accumulation stage of the probability estimation process is anchored, individuals display a negativity bias.

Furthermore, the empirical evidence presented in this thesis supports a cognitive mechanism, loss function asymmetries, more strongly than an affective mechanism, arousal misattribution. These results therefore support the notion brought forward by Harris et al. (2009), in that probability estimates are not *inherently* biased by utility, as would be more suggested by an affective account, but are biased in situations where a judgment can inform an action. In the present work, the moderation by sense of power demonstrated that it is not *necessarily* decision-control in a given situation, but instead its perception that leads to biased or unbiased probability estimates. Powerful individuals perceived that their estimate might inform their action, also in circumstances where it did not in reality. In contrast, powerless individuals might underestimate the ability to act based on their estimate. That we demonstrate that there are individual

differences that determine whether an individual perceives that their estimate informs their action is an important advancement of the current thesis.

Furthermore, this thesis also identified another possible individual difference variable, interoceptive sensibility, where individuals who believed that they can “feel” their bodily processes, demonstrated a greater negativity bias than individuals that were less interoceptively sensible. Whilst we did not provide a mechanism through which this potential moderation occurs, it is possible that individuals high in interoceptive sensibility engage in greater “bracing for loss” (Shepperd et al. 1996, 2000), a strategic overestimation of the likelihood of negative events to protect oneself from the unexpected negative affect, as those individuals might feel that they are more impacted by emotional states. Whilst we do not have any evidence for this hypothesis, this explanation would also lead to loss function asymmetry, but through a different mechanism (a weakened emotional impact rather than acting to prevent the negative outcome). Although there is no empirical evidence that bracing for loss would be dependent on decision-control, the research by Siegrist and Sütterlin (2014) demonstrating that individuals perceive human-made (arguably less controllable) hazards as worse than nature-caused (less controllable) hazards could hint towards a greater bracing for loss for controllable rather than uncontrollable events.

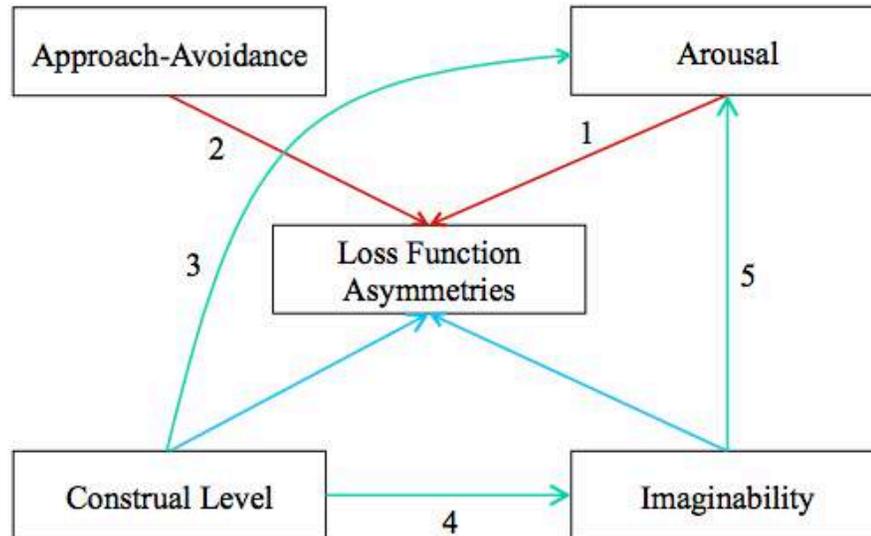
Last, it is worth mentioning that the moderation by IS was noisy. However, when IS was assessed before participants provided probability estimates, the means were almost always in a direction congruent with the hypothesis that individuals high in IS provide a greater negativity bias than those low in IS. Nevertheless, the noisiness of

the moderation by IS suggests more research necessary to potentially uncover some hitherto unthought-of moderator.

Furthermore, whilst we provided support for ALFs in the present thesis, other previously proposed mechanisms that were not assessed in the current research could have indirectly contributed to the current findings. It is noteworthy that the potential mechanisms introduced in the introduction (imaginability, construal level, approach-avoidance motivation, SLH and ALFs) are not mutually exclusive. For instance, Vosgerau (2010) proposed that as arousal and therefore emotional reactions increase, loss function asymmetries should become more pronounced. However, we were unable to provide any evidence for the role of arousal and therefore the relationship between ALFs and the SLH remains unsupported. Furthermore, one could also speculate that approach-avoidance motivation could impact loss function asymmetries. Potentially, individuals high in avoidance motivation also try to avoid the negative outcome more, and therefore have steeper ALFs. On the other hand, it remains equally possible that the greater action orientation of those that are approach motivated results in those individuals to show a greater negativity bias. However, Experiment 9 provided no evidence for a moderation by trait approach-avoidance motivation in either direction.

On the other hand, the relationship between ALFs and imaginability, for instance, remains yet to be explored. It could be possible that loss function asymmetries increase as the outcome, or indeed the consequences of under- or overestimating probability, are imagined more vividly. Likewise, events that are construed on a lower level could increase loss function asymmetries. Previous research has shown that the closer in time an individual gets to the occurrence of the event, the more the probability

estimate for the negative outcome increases (Gilovich, Kerr, & Medvec, 1993). Previously, this was interpreted as a shift from evaluating prospects in terms of desirability to feasibility (Sagristano, Trope & Liberman, 2002). An explanation in light of ALFs could be that the closer in time a person is to an event, the greater the estimate of the negative outcome in order to facilitate action would have to be (as less time would be left to correct for an underestimation, and potentially greater action would have to be taken to prevent the negative outcome). Potential, supported, as well as hypothesized relationships between different theories are depicted in Figure 6.1.



1 Hypothesized by Vosgerau (2010, p. 36), no evidence found in Chapters 2, 3 and 4

2 No evidence found in Experiment 9

3 Ekman & Lundberg (1971)

4 Böhm & Bawerk, (1889), Williams et al. (1996)

5 Acosta & Vila (1990)

Figure 6.1. Potential (blue), supported (green) and unsupported (red) relationships between different theories on the mechanism underlying the interdependence of utility and probability estimates.

The Role of Social Power

Furthermore, the results from the current research also advance the literature on social power, having provided evidence that powerful but not powerless individuals show a negativity bias. Powerful individuals are often in a position where judgments are truly impactful, and it is imperative that they get them right. Imagine a doctor misestimating the chance of a patient dying during an operation, a judge the chance that a defendant is guilty, or a politician the likelihood of an economic crisis. The current results suggest that the powerful do not as previously assumed optimistically underestimate the occurrence of negative events (Anderson & Galinsky, 2006, see also Keltner et al. 2003). In fact, the powerful appear to be even more attuned to negative utility than the powerless, potentially both for judgments that concern themselves as well as others.

Having identified this ubiquitous individual difference, it is possible to develop interventions that guard against biases and improve judgment and decision-making. On the other hand, it is also possible to motivate individuals to take into account negative utilities when it is of importance, as a negativity bias can potentially serve as a protective mechanism. In Experiment 11 we demonstrated that the powerless are demonstrating a negativity bias when decision-control is made explicit. Therefore, highlighting an individual's ability to act based on their estimates is of importance. For example, health campaigns advocating the risk of life-style choices should stress the possibility to take different actions. Being a patient in a doctor's office for instance leads to power asymmetries, with the doctor being in a more powerful position than the

patient. To empower the patient in moments of evaluating risks, placing an emphasis on different action plans for instance would be one possibility for such an intervention.

Worth mentioning, however, the present thesis does not exclude the possibility that whilst the powerful provided greater estimates for negative events in the laboratory, the powerless might perceive negative events as more likely in practice. For instance, some research has demonstrated that the powerless are in a more negative affective state compared to the powerful (e.g. Anderson & Berdahl, 2002). Negative affect can lead to mood dependent retrieval (see associative network model, Gilligan & Bower, 1984) and by making negative outcomes more salient, impact probability estimates (see Bar-Hillel et al. 2008; Wright & Bower, 1992). Else, when no objective bases for probabilities are provided, base-rate estimates for negative events might differ due to the fact that the powerless might *actually* experience more negative events compared to the powerful, who live in more reward-rich environments (Keltner et al. 2003). However, the present research demonstrated that this would most likely not be due to a negativity bias as a general feature of the powerless' cognition.

Furthermore, while we proposed that it is the powerful's greater sense of control underlying power's moderating role, other explanations should be taken into account. If control were to underlie the present results, then this would be due to the notion that loss function asymmetries should only exist when one can *act* based on one's estimate. The powerful, due to a heightened sense of control might perceive their ability to act based on these estimates as greater than the powerless. And indeed, throughout situations with and without explicit decision-control did the powerful display a negativity bias, implying that they might have perceived control across situations.

However, it remains possible that it is not greater decision-control, but increased bracing for loss (Shepperd et al. 1996, 2000) that might lead to the powerful displaying an interdependence of negative utility and probability estimates across situations independently of decision-control. Bracing for loss, a way of protecting oneself from the negative emotional impact should the outcome occur, also leads to loss function asymmetries, but through a different mechanism. As Chapter 5 showed that the powerful appear to be more impacted by negative affective states, they might provide greater estimates for negative outcomes to weaken the intensity of the experience of potential negative affect. A more repressive coping style of the powerless, as we hypothesized, would also explain why the powerless appear to not be impacted by negative utilities. On the contrary, unless one proposes different mechanisms for powerful and powerless individuals, this viewpoint would not explain why the powerless displayed a negativity bias for events with decision-control in Experiment 11, but not for events where decision-control was more ambiguous. One way of reconciling these apparent contradictions would be to hypothesize that by providing powerless participants with decision-control, they increased in power and therefore also engaged in bracing for loss.

Further Research

The present findings open up several new avenues of further work relating to the moderating role of interoceptive sensibility, loss function asymmetries in the positive domain, as well as to the role of social power. These different avenues shall be discussed below in more detail.

The current research motivated further examinations of the role of interoceptive sensibility. As mentioned earlier, we did not provide a mechanism through which we

observed this moderation effect, where individuals high in interoceptive sensibility provided a negativity bias but those that are low did not. Providing evidence that it is not as previously assumed arousal misattribution that underlies this moderation, the mechanism that is suggesting itself is that those high in interoceptive sensibility expect to feel the negative emotion more intensively and therefore engage in strategic pessimism, potentially resulting in loss function asymmetries. Further research should assess the intensity with which individuals high and low in interoceptive sensibility expect to feel the negative outcome, in order to assess this potential mechanism.

Another persisting question is also whether loss-function asymmetries exist in the positive domain as suggested by Weber (1994). Weber (1994) notes that forgoing a reward can lead to regret, and therefore underestimating the likelihood of a positive outcome might be equally more consequential than overestimating it. However, de Molière et al. (2014) considered the role of decision-control in the positive domain across three experiments, and found no evidence for the existence of such an effect in scenarios equivalent to those where previous research (Harris et al. 2009) demonstrated a moderation by decision-control in the negative domain. On the other hand, multiple studies have demonstrated negative-positive asymmetries (for a review see Baumeister et al. 2001), and asymmetries, should they exist, are likely to be weaker in the positive domain given the urgency of negative outcomes over positive outcomes. One question that arises is whether loss function asymmetries might indeed exist for the non-focal rather than the focal outcome in the positive domain. That is, rather than to assess the probability estimate for a positive event to occur, one could assess the estimate for a positive event *not* to occur and examine whether loss function asymmetries are more

pronounced for the absence of a reward, which would be more aligned with the original hypothesis by Weber (1994). Likewise, it is possible that loss-function asymmetries are weaker for the non-focal event in the negative domain.

Furthermore, the finding that individuals high in social power demonstrate a greater negativity bias motivates a broad range of further research opportunities. Whilst we did observe said moderation both for correlational as well as causal designs, the role of decision-control was examined only in a correlational experiment. Therefore, one experiment could manipulate power as well as to manipulate decision-control. Furthermore, Experiment 11 included another layer of randomness by including a lottery draw, where participants avoided losing a ticket rather than avoided the losing of a certain reward. To exclude the possibility that powerful and powerless participants have different estimates of having a winning ticket in the first place, this experiment should employ a task in the laboratory where participants directly have or do not have control over the potential outcome with certain negative or neutral outcomes. In addition, such an experiment should also assess power after participants are given or not-given decision-control in order to eliminate the possibility that decision-control leads to elevated power, which accounts for the present results. Furthermore, it is important to note that the current research, whilst providing some real-world application in Experiments 10 and 11, related mainly to fictional scenarios without real consequences. In particular, it remains to be seen whether individuals high in social power would actually *act* based on their estimates, or make more risk averse decisions for events with potentially negative outcomes than the powerless, too.

Moreover, the current research did not distinguish between types of control, and did not measure control separately as a mediator. Whilst we attempted to tease apart the role of control over oneself and others in Experiment 10, the results were not conclusive. Further research could manipulate not only the estimator's but also the recipient's power: If the powerful feel like they have more control over other individuals, they should only provide estimates that are greater in the negative compared to the neutral domain for *powerless*, but not for the less controllable *powerful* individuals.

Examining the impact of the receiver could potentially also give rise to the location of the effect. As we held the evidence accumulation and selection stage constant (see Figure 6.2) by providing participants with objective representations of probabilities, the bias could both be on the stages of the internal estimate as well as on the stage of the report.



Figure 6.2. The process of making and reporting a probability estimate (as in Harris et al. 2009).

Whilst Harris et al. (2009, see also Harris, 2009) assumed that the bias occurs on the stage of the internal estimate, it remains also possible that individuals have unbiased internal estimates but a biased report. If the internal estimate is biased but there is a separate bias on the report stage, then characteristics of the *recipients* (such as high or low power, as mentioned above) of this information might be expected to be of relevance to the final provided estimate. In addition, reaction time experiments, or experiments providing a cognitive load or time pressure could try to tease apart internal estimates and biases that occur at the level of reporting likelihoods. However, what

speaks against the existence of such a bias on the “report” stage is our finding that estimates were still biased by utility in Experiment 11, where we incentivised accuracy.

Conclusion

In sum, the here presented work has added to a growing body of evidence that shows that we are prone to making pessimistic judgments about negative future events. The reason why we are providing greater estimates for negative compared to neutral future events appears to be a result of loss function asymmetries, in that we are sensitive towards the differences in cost for an underestimation or an overestimation of the likelihood of a negative event. Furthermore, having identified individual difference variables as moderators of the effect also indicates that there appears to be a lot of variation within and between individuals as to if and under which circumstances such a bias is exhibited.

Thus, returning to the example of an office worker pondering about the chance of rain presented at the very beginning of this thesis, we can conclude that there will indeed be circumstances where the probability estimate will differ depending on whether the worker is a manager or a subordinate. Our research has shown that if it is ambiguous as to whether or not the office worker has time to act on her estimate, or indeed when she has no time to get back inside to get an umbrella, a manager will likely estimate the chance of rain to be greater than a subordinate.

As such, this thesis also highlighted the versatility of social power: whilst lay people and researchers assume that powerful individuals are the happiest, optimistic, risk-taking individuals, this thesis paints a much more complex, affectively richer picture. Potentially, social hierarchies have an impact on mechanisms that protect an

individual from experiencing negative life events in the first place, which would be disadvantageous for those in the already less desirable social position. This thesis provided some theoretical background and empirical findings on whose basis one can design interventions that help counteract these outcomes.

References

- Acosta, A. U., & Villa, J. (1990). Emotional imagery: Effect of autonomic response information on physiological arousal. *Cognition & Emotion, 4*, 145-160.
- Ainley, V., & Tsakiris, M. (2013). Body conscious? Interoceptive awareness, measured by heartbeat perception, is negatively correlated with self-objectification. *PloS one, 8*(2), e55568.
- Anderson, C., & Berdahl, J. L. (2002). The experience of power: examining the effects of power on approach and inhibition tendencies. *Journal of personality and social psychology, 83*(6), 1362.
- Anderson, C., & Galinsky, A. (2006). Power, optimism, and risk-taking. *European Journal of Social Psychology, 36*(4), 511–536.
- Anderson, C., John, O. P., & Keltner, D. (2012). The personal sense of power. *Journal of personality, 80*(2), 313-344.
- Andreassi, J.L. (2007). *Psychophysiology: Human Behavior and Physiological Response*. Lawrence Erlbaum Associates, Mahwah, NJ.
- Babad, E. (1995). Can accurate knowledge reduce wishful thinking in voters' predictions of election outcomes? *The Journal of Psychology, 129*, 285-300.
- Babad, E., & Katz, Y. (1991). Wishful thinking—against all odds. *Journal of Applied Social Psychology, 21*, 1921-1938.
- Bargh, J. A., Raymond, P., Pryor, J. B., & Strack, F. (1995). Attractiveness of the underling: An automatic power→ sex association and its consequences for sexual harassment and aggression. *Journal of personality and social psychology, 68*(5), 768.

- Bar-Hillel, M., & Budescu, D. (1995). The elusive wishful thinking effect. *Thinking and Reasoning, 1*, 71-103.
- Bar-Hillel, M., Budescu, D. V., & Amar, M. (2008). Predicting World Cup results: Do goals seem more likely when they pay off? *Psychonomic Bulletin & Review, 15*(2), 278–283.
- Barrett, L. F., Quigley, K. S., Bliss-Moreau, E., & Aronson, K. R. (2004). Interoceptive sensitivity and self-reports of emotional experience. *Journal of personality and social psychology, 87*(5), 684.
- Barry, R. J., Rushby, J. A., Wallace, M. J., Clarke, A. R., Johnstone, S. J., & Zlojutro, I. (2005). Caffeine effects on resting-state arousal. *Clinical Neurophysiology, 116*(11), 2693-2700.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology, 5*(4), 323–370.
- Bechara, A., Damasio, H., & Damasio, A. R. (2000). Emotion, decision making and the orbitofrontal cortex. *Cerebral cortex, 10*(3), 295-307.
- Bem, D. J. (2011). Feeling the future: experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of personality and social psychology, 100*(3), 407.
- Benedek, M. & Kaernbach, C. (2010). A continuous measure of phasic electrodermal activity. *Journal of Neuroscience Methods, 190*, 80-91.
- Bentham, J. (1779/1879). *Introduction to the principles of morals and legislation*. Oxford: Clarendon Press.

- Berdahl, J. L., & Martorana, P. (2006). Effects of power on emotion and expression during a controversial group discussion. *European Journal of Social Psychology*, 36(4), 497–509. doi:10.1002/ejsp.354
- Berkman, E. T., Lieberman, M. D., & Gable, S. L. (2009). BIS, BAS, and response conflict: Testing predictions of the revised reinforcement sensitivity theory. *Personality and Individual Differences*, 46(5-6), 586–591.
- Berscheid, E. (1983). Emotion. In H. H. Kelley, E. Berscheid, A. Christensen, J. H. Harvey, T. L. Huston, G. Levinger, E. McClintock, L. A. Peplau, & D. R. Peterson (Eds.), *Close Relationships* (pp. 110-168). San Francisco: Freeman.
- Bilgin, B. (2012). Losses loom more likely than gains: Propensity to imagine losses increases their subjective probability. *Organizational Behavior and Human Decision Processes*, 118(2), 203-215.
- Blake, R. R., & Mouton, J. S. (1961). Reactions to intergroup competition under win-lose conditions. *Management Science*, 7(4), 420-435.
- Boden, J. M., & Baumeister, R. F. (1997). Repressive coping: distraction using pleasant thoughts and memories. *Journal of Personality and Social Psychology*, 73(1), 45–62.
- Böhm-Bawerk, E. V. (1889). Kapital und Kapitalzins. *Zweite Abteilung: Positive Theorie des Kapitals*, 3.
- Bollen, K. A., & Jackman, R. W. (1990). Regression diagnostics: An expository treatment of outliers and influential cases. *Modern methods of data analysis*, 257-291.

- Borg, C., De Jong, P. J., Renken, R. J., & Georgiadis, J. R. (2012). Disgust trait modulates frontal-posterior coupling as a function of disgust domain. *Social Cognitive and Affective Neuroscience*.
- Boucsein, W. (2012). *Electrodermal activity*. New York: Springer.
- Bradburn, N.M. (1969). The structure of psychological well-being. Norman M. Oxford, England: Aldine.
- Brandt, M. J. (2013). Do the disadvantaged legitimize the social system? A large-scale test of the status–legitimacy hypothesis. *Journal of personality and social psychology*, 104(5), 765.
- Brener, J., Knapp, K., & Ring, C. (1995, August). The effects of manipulating beliefs about heart-rate on the accuracy of heartbeat counting in the Schandry task. *Psychophysiology* (32), 22.
- Cacioppo, J. T., Priester, J. R., & Berntson, G. G. (1993). Rudimentary determinants of attitudes: II. Arm flexion and extension have differential effects on attitudes. *Journal of personality and social psychology*, 65(1), 5.
- Campbell, J., Greenauer, N., Macaluso, K., & End, C. (2007). Unrealistic optimism in internet events. *Computers in human behavior*, 23(3), 1273-1284.
- Cantor, J. R., Zillmann, D., & Bryant, J. (1975). Enhancement of experienced sexual arousal in response to erotic stimuli through misattribution of unrelated residual excitation. *Journal of personality and social psychology*, 32(1), 69.
- Carney, D. R., Cuddy, A. J. C., & Yap, A. J. (2010). Power posing: brief nonverbal displays affect neuroendocrine levels and risk tolerance. *Psychological Science*, 21(10), 1363–8. doi:10.1177/0956797610383437

- Carroll, J. S. (1978). The effect of imagining an event on expectations for the event: An interpretation in terms of the availability heuristic. *Journal of experimental social psychology, 14*(1), 88-96.
- Carver, C. S. (2004). Negative Affects Deriving From the Behavioral Approach System. *Emotion, 4*(1), 3–22. doi:10.1037/1528-3542.4.1.3
- Carver, C. S., & White, T. L. (1994). Behavioral inhibition, behavioral activation, and affective responses to impending reward and punishment: The BIS/BAS Scales. *Journal of Personality and Social Psychology, 67*(2), 319–333.
- Chabris, C. F., Hebert, B. M., Benjamin, D. J., Beauchamp, J., Cesarini, D., van der Loos, M., ... & Laibson, D. (2012). Most reported genetic associations with general intelligence are probably false positives. *Psychological science, 0956797611435528*.
- Chambers, J. R., Windschitl, P. D., & Suls, J. (2003). Egocentrism, event frequency, and comparative optimism: When what happens frequently is “more likely to happen to me”. *Personality and Social Psychology Bulletin, 29*(11), 1343-1356.
- Chen, S., Langner, C. A., & Mendoza-Denton, R. (2009). When dispositional and role power fit: Implications for self-expression and self–other congruence. *Journal of personality and social psychology, 96*(3), 710.
- Clore, G. L. (1992). Cognitive phenomenology: Feelings and the construction of judgment. The construction of social judgments. In *The construction of social judgments* (pp. 133–163). Lawrence Erlbaum Associates, Inc.
- Cohen, J. (1994). The earth is round ($p < .05$). *American psychologist, 49*(12), 997.

- Cohen, J., Cohen, P., West, S., & Aiken, L. (2003). *Applied Multiple Regression / Correlation Analysis for the Behavioral Sciences*. 2nd ed Hillsdale NJ Lawrence Erlbaum Associates (Vol. Third Edit, p. 703 S.).
- Costa, P. T., & McCrae, R. R. (1980). Influence of extraversion and neuroticism on subjective well-being: happy and unhappy people. *Journal of personality and social psychology*, 38(4), 668.
- Cox, C. R., Goldenberg, J. L., Pyszczynski, T., & Weise, D. (2007). Disgust, creatureliness and the accessibility of death-related thoughts. *European Journal of Social Psychology*, 37(3), 494–507.
- Craig, A. D. (2002). How do you feel? Interoception: the sense of the physiological condition of the body. *Nature Reviews Neuroscience*, 3(8), 655-666.
- Critchley, H. D., Wiens, S., Rotshtein, P., Öhman, A., & Dolan, R. J. (2004). Neural systems supporting interoceptive awareness. *Nature neuroscience*, 7(2), 189-195.
- Dai, X., Wertenbroch, K., & Brendl, C. M. (2008). The value heuristic in judgments of relative frequency. *Psychological Science*, 19(1), 18-19.
- Dalgleish, T., Williams, J. M. G., Golden, A. M. J., Perkins, N., Barrett, L. F., Barnard, P. J., ... & Watkins, E. (2007). Reduced specificity of autobiographical memory and depression: the role of executive control. *Journal of Experimental Psychology: General*, 136(1), 23.
- Damasio, A. R. (1994). *Descartes' error: emotion, research and the human brain*. New York: Avon.

- Damasio, A. R. (1996). The somatic marker hypothesis and the possible functions of the prefrontal cortex. *Philosophical Transactions of the Royal Society of London (Biology)*, *351*, 1413-1420.
- Dawson, M. E., Schell, A. M., & Courtney, C. G. (2011). The skin conductance response, anticipation, and decision-making. *Journal of Neuroscience, Psychology, and Economics*, *4*(2), 111.
- Dawson, M. E., Schell, A. M., & Fillion, D. L. (2007). *The electrodermal system*. In J.T. Cacioppo, L.G. Tassinari, & G.G. Berntson (Eds.). *Handbook of psychophysiology* (3rd ed., pp. 159-181). New York: Cambridge University Press.
- Diener, E., Emmons, R. A., Larsen, R. J., & Griffin, S. (1985). The Satisfaction with Life Scale. *Journal of Personality Assessment*, *49*(1), 71-75.
- Dijksterhuis, A., & Aarts, H. (2003). On Wildebeests and Humans The Preferential Detection of Negative Stimuli. *Psychological Science*, *14*(1), 14-18.
- Domhoff, G. W. (1998). *Who rules America?: power and politics in the year 2000*. Mountain View, CA: Mayfield Publishing Company.
- Doyen, S., Klein, O., Pichon, C. L., & Cleeremans, A. (2012). Behavioral priming: it's all in the mind, but whose mind?. *PloS one*, *7*(1), e29081.
- Dunn, B. D., Galton, H. C., Morgan, R., Evans, D., Oliver, C., Meyer, M., ... Dalgleish, T. (2010). Listening to your heart. How interoception shapes emotion experience and intuitive decision making. *Psychological Science*, *21*(12), 1835-44.
- Dutton, D. G., & Aron, A. P. (1974). Some evidence for heightened sexual attraction under conditions of high anxiety. *Journal of personality and social psychology*, *30*(4), 510.

- Edwards, W. (1962). Utility, subjective probability, their interaction, and variance preferences. *Journal of Conflict Resolution*, 42-51.
- Ehlers, A., & Breuer, P. (1992). Increased cardiac awareness in panic disorder. *Journal of abnormal psychology*, 101(3), 371.
- Eichler, S., & Katkin, E. S. (1994). The relationship between cardiovascular reactivity and heartbeat detection. *Psychophysiology*, 31(3), 229-234.
- Ekman, G., & Lundberg, U. (1971). Emotional reaction to past and future events as a function of temporal distance. *Acta Psychologica*, 35(6), 430-441.
- Ekman, P. (1992). An argument for basic emotions. *Cognition & Emotion*, 6(3), 169–200.
- Elliot, A. J. (Ed.). (2008). *Handbook of approach and avoidance motivation*. Taylor & Francis.
- Elliot, A. J., & Church, M. a. (2003). A motivational analysis of defensive pessimism and self-handicapping. *Journal of Personality*, 71(3), 369–96.
- Elliot, A.J., & Thrash, T.M. (2002). Approach-avoidance motivation in personality: Approach and avoidance temperaments and goals. *Journal of Personality and Social Psychology*, 82, 804-818.
- Elliott, R. (1964). Physiological activity and performance: A comparison of kindergarten children with young adults. *Psychological Monographs: General and Applied*, 78(10), 1.
- Emerson, R. M. (1962). Power-dependence relations. *American sociological review*, 31-41.

- Erev, I., & Cohen, B. L. (1990). Verbal versus numerical probabilities: Efficiency, biases, and the preference paradox. *Organizational Behavior and Human Decision Processes*, *45*, 1-18.
- Fast, N. J., Gruenfeld, D. H., Sivanathan, N., & Galinsky, A. D. (2009). Illusory control: a generative force behind power's far-reaching effects. *Psychological Science*, *20*(4), 502-8..x
- Fast, N. J., Sivanathan, N., Mayer, N. D., & Galinsky, A. D. (2012). Power and overconfident decision-making. *Organizational Behavior and Human Decision Processes*, *117*(2), 249-260.
- Ferguson, M. L., & Katkin, E. S. (1996). Visceral perception, anhedonia, and emotion. *Biological Psychology*, *42*(1), 131-145.
- Fiedler, K., Jung, J., Wänke, M., & Alexopoulos, T. (2012). On the relations between distinct aspects of psychological distance: An ecological basis of construal-level theory. *Journal of Experimental Social Psychology*, *48*(5), 1014-1021.
- Fiedler, K., Jung, J., Wänke, M., Alexopoulos, T., & de Molière, L. (2014). Toward a Depper Understanding of The Ecological Origins of Distance Construal. Manuscript Under Review.
- Figner, B., & Murphy, R. O. (2011). Using skin conductance in judgment and decision making research. *A handbook of process tracing methods for decision research*, 163-184.
- Fisher, R. A. (1935). The logic of inductive interference (with discussion). *Journal of the Royal Statistical Society*, *98*(1), 39-82.
- Fisher, R. A. (1956). *Statistical Methods and Scientific Inference*. New York: Hafner.

- Fishburn, P. C. (1981). Subjective expected utility: A review of normative theories. *Theory and Decision, 13*, 139–199.
- Fiske, S. T. (1993). Controlling other people. The impact of power on stereotyping. *American Psychologist, 48*(6), 621–628.
- Fiske, S. T., & Dépret, E. (1996). Control, interdependence and power: Understanding social cognition in its social context. *European review of social psychology, 7*(1), 31-61.
- Foa, U. G. (1971). Interpersonal and economic resources. *Science, 171*(3969), 345-351.
- Fox, J. (Ed.). (1991). *Regression diagnostics: An introduction* (Vol. 79). Sage.
- French, J. R. P., & Raven, B. (1959). The bases of social power. In P. Bob & J. Randall (Eds.), *The Negotiation Sourcebook* (Vol. 6, pp. 340–346). University of Michigan.
- Frijda, N. H. (1986). *The emotions*. Cambridge University Press.
- Frijda, N. H. (2009). Emotion experience and its varieties. *Emotion Review, 1*(3), 264-271.
- Galinsky, A. D., Gruenfeld, D. H., & Magee, J. C. (2003). From power to action. *Journal of Personality and Social Psychology, 85*(3), 453–66.
- Galinsky, A. D., Magee, J. C., Gruenfeld, D. H., Whitson, J. A., & Liljenquist, K.A. (2008). Power reduces the press of the situation: implications for creativity, conformity, and dissonance. *Journal of Personality and Social Psychology, 95*(6), 1450–66.
- Galinsky, A. D., Magee, J. C., Inesi, M. E., & Gruenfeld, D. H. (2006). Power and perspectives not taken. *Psychological science, 17*(12), 1068-1074.

- Gallistel, C. R. (2009). The importance of proving the null. *Psychological review*, *116*(2), 439.
- Garfinkel, S. N., Seth, A. K., Barrett, A. B., Suzuki, K., & Critchley, H.D. (in press). Knowing your heart: Distinguishing interoceptive accuracy from interoceptive awareness. *Biological Psychology*.
- Garfinkel, S. N., & Critchley, H. D. (2013). Interoception, emotion and brain: new insights link internal physiology to social behaviour. Commentary on: “Anterior insular cortex mediates bodily sensibility and social anxiety” by Terasawa et al. (2012). *Social Cognitive and Affective Neuroscience*, *8*(3), 231–234.
- Gerits, P. (2000). Life events, coping and breast cancer: state of the art. *Biomedicine Pharmacotherapy Biomedecine Pharmacotherapie*, *54*(5), 229–233.
- Gibbons, F. X., Carver, C. S., Scheier, M. F., & Hormuth, S. E. (1979). Self-focused attention and the placebo effect: Fooling some of the people some of the time. *Journal of Experimental Social Psychology*, *15*(3), 263-274.
- Gibbons, F. X., & Gaeddert, W. P. (1984). Focus of attention and placebo utility. *Journal of Experimental Social Psychology*, *20*(2), 159-176.
- Gigerenzer, G. (2003). *Reckoning with risk: learning to live with uncertainty*. Penguin UK.
- Gilbert, D.T., & Wilson, T.D. (2008). Propection: Experiencing the future. *Science*, *317*, 1351-1354.
- Gilliam, G. E. (1991). The effects of Baker-Miller pink on physiological and cognitive behavior of emotionally disturbed and regular education students. *Behavioral Disorders*, *17*, 47–55.

- Gilligan, S.G., & Bower, G. H. (1984). Cognitive consequences of emotional arousal. In C. E. Izard, J. Kagan & R. Zajonc (Eds.), *Emotions, cognitions, and behavior*. New York: Cambridge Press.
- Gilovich, T., Kerr, M., & Medvec, V. H. (1993). Effect of temporal perspective on subjective confidence. *Journal of personality and social psychology*, *64*(4), 552.
- Ginzburg, K., Tsur, N., Barak-Nahum, A., & Defrin, R. (2013). Body awareness: differentiating between sensitivity to and monitoring of bodily signals. *Journal of Behavioral Medicine*.
- Gohm, C. L., & Clore, G. L. (2000). Individual Differences in Emotional Experience: Mapping Available Scales to Processes. *Personality and Social Psychology Bulletin*, *26*(6), 679–697. doi:10.1177/0146167200268004
- Gonzaga, G. C., Keltner, D., & Ward, D. (2008). Power in mixed-sex stranger interactions. *Cognition & Emotion*, *22*(8), 1555–1568.
doi:10.1080/02699930801921008
- Gordon, R., Franklin, N., & Beck, J. (2005). Wishful thinking and source monitoring. *Memory & Cognition*, *33*(3), 418–29.
- Gorn, G., Tuan Pham, M., & Yatming Sin, L. (2001). When arousal influences ad evaluation and valence does not (and vice versa). *Journal of consumer Psychology*, *11*(1), 43-55.
- Granberg, D., & Brent, E. (1983). When prophecy bends: the preference-expectation link in U.S. presidential elections, 1952-1980. *Journal of Personality and Social Psychology*, *45*, 477- 491.

- Gray, E., & Watson, D. (2001). Emotion, mood, and temperament: Similarities, differenced, and a synthesis. In R. L. Payne & C. L. Cooper (Eds.). *Emotion at work*. (pp. 21–45). Chichester, U. K.: Wiley.
- Gray, J. A. (1972). The psychophysiological basis of introversion-extraversion: A modification of Eysenck's theory. *The biological bases of individual behaviour*, 182-205.
- Gray, J.A. (1987). *The psychology of fear and stress* (2nd ed.). New York: Cambridge University Press.
- Gray, J.A. (1990). Brain systems that mediate both emotion and cognition. *Cognition and Emotion*, 4, 269-288.
- Gray, J.A. (1994). Three fundamental emotion systems. In P. Ekman & R.J. Davidson (Eds.), *The nature of emotion* (pp. 243-247). New York: Oxford University Press.
- Gray, J. A., & McNaughton, N. (2000). *The neuropsychology of anxiety: an enquiry into the functions of the septo-hippocampal system (second edition)*. Oxford: Oxford University Press.
- Green, D. P., Goldman, S. L., & Salovey, P. (1993). Measurement error masks bipolarity in affect ratings. *Journal of Personality and Social Psychology*, 64(6), 1029.
- Greenwald, A. G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin*, 82, 1-20.
- Gregory, W. L., Cialdini, R. B., & Carpenter, K. M. (1982). Self-relevant scenarios as mediators of likelihood estimates and compliance: Does imagining make it so?. *Journal of Personality and Social Psychology*, 43(1), 89.

- Greifeneder, R., Bless, H., & Pham, M. T. (2011). When do people rely on affective and cognitive feelings in judgment? A review. *Personality and Social Psychology Review*, 15, 107–141.
- Gonzaga, G. C., Keltner, D., & Ward, D. (2008). Power in mixed-sex stranger interactions. *Cognition and Emotion*, 22(8), 1555-1568.
- Guinote, A. (2007a). Behaviour variability and the situated focus theory of power. In W. Stroebe & M. Hewstone (Eds.), *European Review of Social Psychology* (Vol. 18, pp. 256–295).
- Guinote, A. (2007b). Power and goal pursuit. *Personality & Social Psychology Bulletin*, 33(8), 1076–87.
- Guinote, A. (2007c). Power affects basic cognition: Increased attentional inhibition and flexibility. *Journal of Experimental Social Psychology*, 43(5), 685–697.
- Guinote, A. (2010). In Touch With Your Feelings: Power Increases Reliance on Bodily Information. *Social Cognition*, 28(1), 110–121.
- Guinote, A., Brown, M., & Fiske, S. T. (2006). Minority status decreases sense of control and increases interpretive processing. *Social Cognition*, 24(2), 169-186.
- Guinote, A., Judd, C. M., & Brauer, M. (2002). Effects of power on perceived and objective group variability: Evidence that more powerful groups are more variable. *Journal of Personality and Social Psychology*, 82(5), 708–721.
- Guinote, A., Weick, M., & Cai, A. (2012). Does power magnify the expression of dispositions? *Psychological Science*, 23(5), 475–82.
- Häfner, M. (2013). When body and mind are talking: Interoception moderates embodied cognition. *Experimental Psychology*, 1–18.

- Harris, A. J. L. (2009). *Investigating the influence of outcome utility on estimates of probability* (Unpublished doctoral dissertation). Cardiff University, Cardiff, Wales.
- Harris, A. J. L., Corner, A., & Hahn, U. (2009). Estimating the probability of negative events. *Cognition*, *110*(1), 51–64.
- Harris, A. J. L., & Corner, A. (2011). Communicating environmental risks: Clarifying the severity effect in interpretations of verbal probability expressions. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(6), 1571.
- Harris, A. J. L., & Hahn, U. (2011). Unrealistic optimism about future life events: A cautionary note. *Psychological Review*, *118*, 135-154.
- Harris, A. J. L., Shah, P., Catmur, C., Bird, G., & Hahn, U. (2013). Autism, optimism and positive events: Evidence against a general optimistic bias. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 555-560). Austin, TX: Cognitive Science Society.
- Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process: A regression based approach*. New York: The Guilford Press.
- Hecht, M. A., & LaFrance, M. (1998). License or obligation to smile: The effect of power and sex on amount and type of smiling. *Personality and Social Psychology Bulletin*, *24*, 1332-1342.
- Herbert, B. M., Pollatos, O., & Schandry, R. (2007). Interoceptive sensitivity and emotion processing: an EEG study. *International Journal of Psychophysiology*, *65*(3), 214-227.

- Hertel, P. T. (2000). The cognitive-initiative account of depression-related impairments in memory. In D. Medin (Ed.), *The psychology of learning and motivation* (Vol. 39, pp. 47–71). New York: Academic Press.
- Hewig, J., Hagemann, D., Seifert, J., Gollwitzer, M., Naumann, E., & Bartussek, D. (2005). Brief report. *Cognition & Emotion*, *19*(7), 1095-1109.
- Higgins, E. T. (1998). Promotion and prevention: Regulatory focus as a motivational principle. *Advances in experimental social psychology*, *30*, 1-46.
- Hilbig, B. E. (2012). Good things don't come easy (to mind): Explaining framing effects in judgments of truth. *Experimental psychology*, *59*(1), 38.
- Hinkle, S., & Schopler, J. (1986). Bias in the evaluation of in-group and out-group performance. *Psychology of intergroup relations*, *2*, 196-212.
- Hirshleifer, J., Glazer, A., & Hirshleifer, D. (2006). *Price Theory and Applications: Decisions, Markets, and Information*. New York: Cambridge University Press.
- Hölzl, R., Erasmus, L. P., & Möltner, A. (1996). Detection, discrimination and sensation of visceral stimuli. *Biological psychology*, *42*(1), 199-214.
- Hsee, C. K., Hatfield, E., Carlson, J. G., & Chemtob, C. (1990). The effect of power on susceptibility to emotional contagion. *Cognition & Emotion*, *4*(4), 327–340.
- Inesi, M. E. (2010). Power and loss aversion. *Organizational Behavior and Human Decision Processes*, *112*(1), 58-69.
- Ingram, F., & Lieberman, L. R. (1985). Effects of expectations on the performance of hand grip after viewing selected hues. *Perceptual and motor skills*, *61*(2), 370.
- Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS medicine*, *2*(8), e124.

- Irwin, F. W. (1953). Stated expectations as functions of probability and desirability of outcomes. *Journal of Personality, 21*, 329-335.
- Isen, A. M. (1984). Toward understanding the role of affect in cognition. In R. S. W. jr. & T. K. Skrull (Eds.), *Handbook of Social Cognition* (Vol. 3, pp. 179-236). Hillsdale, NJ: Erlbaum.
- James, W. (1894). The physical basis of emotion. *Psychological Review, 1*, 516-529.
- James, W. (1890). *The principles of psychology* (Vols. 1 & 2). New York: Holt.
- Jenkins, S. (2013). Explaining the severity effect in interpretations of verbal probability expressions- the role of imagination. Unpublished MSc thesis. University College London, London, United Kingdom.
- Johnson, E. J., & Tversky, A. (1983). Affect, generalization, and the perception of risk. *Journal of personality and social psychology, 45*(1), 20.
- Jordan, J., Sivanathan, N., & Galinsky, A. D. (2011). Something to Lose and Nothing to Gain The Role of Stress in the Interactive Effect of Power and Stability on Risk Taking. *Administrative Science Quarterly, 56*(4), 530-558.
- Jourden, F. J., & Heath, C. (1996). The evaluation gap in performance perceptions: Illusory perceptions of groups and individuals. *Journal of Applied Psychology, 81*, 369- 379.
- Judge, T. A., Locke, E. A., & Durham, C. C. (1997). The dispositional causes of job satisfaction: A core evaluations approach. *Research in Organizational Behavior, 19*, 151-188.
- 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). On the study of statistical intuitions. *Cognition, 11*(2), 123-141.
- Karremans, J. C., & Smith, P. K. (2010). Having the power to forgive: When the experience of power increases interpersonal forgiveness. *Personality and Social Psychology Bulletin, 36*(8), 1010-1023.
- Kass, R. E., & Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association, 90*, 773-781.
- Katkin, E. S. (1985). Blood, Sweat, and Tears: Individual Differences in Autonomic Self-Perception. *Psychophysiology, 22*(2), 125-137.
- Keltner, D., Gruenfeld, D. H., & Anderson, C. (2003). Power, approach, and inhibition. *Psychological Review, 110*(2), 265–284.
- Khalsa, S. S., Rudrauf, D., Damasio, A. R., Davidson, R. J., Lutz, A., & Tranel, D. (2008). Interoceptive awareness in experienced meditators. *Psychophysiology, 45*(4), 671–677.
- Khalsa, S. S., Rudrauf, D., Sandesara, C., Olshansky, B., & Tranel, D. (2009). Bolus isoproterenol infusions provide a reliable method for assessing interoceptive awareness. *International Journal of Psychophysiology, 72*(1), 34-45.
- Kifer, Y., Heller, D., Perunovic, W. Q. E., & Galinsky, A. D. (2013). The good life of the powerful: The experience of power and authenticity enhances subjective well-being. *Psychological Science, 24*(3), 280-288.
- Van Kleef, G. A., De Dreu, C. K., Pietroni, D., & Manstead, A. S. (2006). Power and emotion in negotiation: Power moderates the interpersonal effects of anger and

- happiness on concession making. *European Journal of Social Psychology*, 36(4), 557-581.
- Van Kleef, G. A., Oveis, C., Van Der Löwe, I., LuoKogan, A., Goetz, J., & Keltner, D. (2008). Power, distress, and compassion turning a blind eye to the suffering of others. *Psychological Science*, 19(12), 1315-1322.
- Konrad, A. M., & Pringle, J. (Eds.). (2006). *Handbook of workplace diversity*. London: Sage.
- Konstantinidis, E., & Shanks, D. R. (2014). Don't bet on it! Wagering as a measure of awareness in decision making under uncertainty. *Journal of Experimental Psychology: General*. In press.
- Koole, S. L., & Lakens, D. (2012). Rewarding Replications A Sure and Simple Way to Improve Psychological Science. *Perspectives on Psychological Science*, 7(6), 608-614.
- Kramer, T. (2007). Approach-Avoidance Motivation and the Use of Affect as Information. *New York*, 17(2), 128-138.
- Kraus, M. W., Chen, S., & Keltner, D. (2011). The power to be me: Power elevates self-concept consistency and authenticity. *Journal of Experimental Social Psychology*, 47(5), 974-980.
- Krizan, Z., Miller, J. C., & Johar, O. (2010). Wishful thinking in the 2008 US presidential election. *Psychological Science*, 21(1), 140-146.
- Krizan, Z., & Windschitl, P. D. (2007). The influence of outcome desirability on optimism. *Psychological Bulletin*, 133, 95-121.

- Krizan, Z., & Windschitl, P. D. (2009). Wishful thinking about the future: Does desire impact optimism?. *Social and Personality Psychology Compass*, 3(3), 227-243.
- Kruger, J., & Burrus, J. (2004). Egocentrism and focalism in unrealistic optimism (and pessimism). *Journal of Experimental Social Psychology*, 40(3), 332-340.
- Lachman, M. E., & Weaver, S. L. (1998). The sense of control as a moderator of social class differences in health and well-being. *Journal of personality and social psychology*, 74(3), 763.
- Lajante, M., Droulers, O., Dondaine, T., & Amarantini, D. (2012). Opening the “black box” of electrodermal activity in consumer neuroscience research. *Journal of Neuroscience, Psychology, and Economics*, 5(4), 238-249.
- Lammers, J., Galinsky, A. D., Gordijn, E. H., & Otten, S. (2008). Illegitimacy moderates the effects of power on approach. *Psychological Science*, 19(6), 558–64.
- Lang, P., Bradley, M., & Cuthbert, BN. (2005). International affective picture system (IAPS): affective ratings of pictures and instruction manual. Technical Report A-6. University of Florida, Gainesville, FL.
- Lange, C. (1922). *The emotions* (I. A. Haupt, Trans.). Baltimore: Williams & Wilkins. (Original work published 1885).
- Langer, E. J. (1975). The illusion of control. *Journal of personality and social psychology*, 32(2), 311.
- Langner, C. A., & Keltner, D. (2003). Social power and emotional experience : Actor and partner effects within dyadic interactions.
- Larsen, J. T., Berntson, G. G., Poehlmann, K. M., Ito, T. A., & Cacioppo, J. T. (2008). The psychophysiology of emotion. In R. Lewis, J. M. Haviland-Jones, & L. F.

- Barrett (Eds.), *The handbook of emotions* (3rd ed., pp. 180-195). New York: Guilford.
- Larsen, R. J., & Ketelaar, T. (1989). Extraversion, neuroticism and susceptibility to positive and negative mood induction procedures. *Personality and Individual Differences, 10*(12), 1221-1228.
- Lee, E. H., & Schnall, S. (2014). The influence of social power on weight perception. *Journal of Experimental Psychology: General*. In press.
- Lench, H. C. (2009). Automatic optimism: the affective basis of judgments about the likelihood of future events. *Journal of Experimental Psychology: General, 138*(2), 187–200.
- Lench, H. C., & Bench, S. W. (2012). Automatic Optimism: Why People Assume Their Futures will be Bright. *Social and Personality Psychology Compass, 6*(4), 347–360.
- Lench, H. C., & Darbor, K. E. (2014). Negative affective reactions reduce perceived likelihood of risk. *Motivation and Emotion, 38*(4), 569–577.
- Lench, H. C., & Ditto, P. H. (2008). Automatic optimism: Biased use of base rate information for positive and negative events. *Journal of Experimental Social Psychology, 44*(3), 631-639.
- Levita, L., Hoskin, R., & Champi, S. (2012). Avoidance of harm and anxiety : A role for the nucleus accumbens. *NeuroImage, 62*(1), 189–198.
- Liberman, N., & Trope, Y. (1998). The role of feasibility and desirability considerations in near and distant future decisions: A test of temporal construal theory. *Journal of personality and social psychology, 75*(1), 5.

- MacLeod, C., Mathews, A., & Tata, P. (1986). Attentional bias in emotional disorders. *Journal of Abnormal Psychology, 95*(1), 15–20.
- Magee, J. C., & Galinsky, A. D. (2008). 8 Social Hierarchy: The Self-Reinforcing Nature of Power and Status. *The Academy of Management Annals, 2*(1), 351-398.
- Magee, J. C., Galinsky, A. D., & Gruenfeld, D. H. (2007). Power, propensity to negotiate, and moving first in competitive interactions. *Personality & Social Psychology Bulletin, 33*(2), 200–12.
- Maia, T. V., & McClelland, J. L. (2004). A reexamination of the evidence for the somatic marker hypothesis: What participants really know in the Iowa gambling task. *Proceedings of the National Academy of Sciences of the United States of America, 101*(45), 16075-16080.
- Makel, M. C., Plucker, J. A., & Hegarty, B. (2012). Replications in Psychology Research How Often Do They Really Occur?. *Perspectives on Psychological Science, 7*(6), 537-542.
- Mandel, D. R. (2008). Violations of coherence in subjective probability: A representational and assessment processes account. *Cognition, 106*(1), 130-156.
- Maner, J. K., Gailliot, M. T., Menzel, A. J., & Kunstman, J. W. (2012). Dispositional anxiety blocks the psychological effects of power. *Personality & Social Psychology Bulletin, 38*(11), 1383–95.
- Marks, R. W. (1951). The effect of probability, desirability and ‘privilege’ on the stated expectations of children. *Journal of Personality, 19*, 332-351.

- Marteau, T. M., & Bekker, H. (1992). The development of a six-item short-form of the state scale of the Spielberger State—Trait Anxiety Inventory (STAI). *British Journal of Clinical Psychology, 31*(3), 301-306.
- Massey, C., Simmons, J. P., & Armor, D. A. (2011). Hope Over Experience Desirability and the Persistence of Optimism. *Psychological Science, 22*(2), 274-281.
- Mcfarland, R.A. (1975). Heart-rate perception and heart-rate control. *Psychophysiology, 12* (4), 402-405.
- McKenna, M. C., Zevon, M. A., Corn, B., & Rounds, J. (1999). Psychosocial factors and the development of breast cancer: a meta-analysis. *Health Psychology Official Journal of the Division of Health Psychology American Psychological Association, 18*(5), 520–531.
- Mehling, W. E., Price, C., Daubenmier, J. J., Acree, M., Bartmess, E., & Stewart, A. (2012). The multidimensional assessment of interoceptive awareness (MAIA). *PloS one, 7*(11), e48230.
- Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychological Science, 423*-429.
- Mevissen, F. E., Meertens, R. M., Ruiter, R. A., & Schaalma, H. P. (2012). Bedtime stories: The effects of self-constructed risk scenarios on imaginability and perceived susceptibility to sexually transmitted infections. *Psychology & health, 27*(9), 1036-1047.
- Miller, L. C., Murphy, R., & Buss, A. H. (1981). Consciousness of body: Private and public. *Journal of Personality and Social Psychology, 41*(2), 397–406.

- Mitchell, J. (2014). On the emptiness of failed replications. [Retrieved 30.09.2014 from http://wjh.harvard.edu/~jmitchel/writing/failed_science.htm]
- Moeini-Jazani, M., Guinote, A., & Warlop, L. (2014). *When Power has It's Pants Down: Social Power Increases Reliance on Visceral States*. Manuscript in Preparation.
- Moeini-Jazani, M., Knöpferle, K., de Molière, L., Gatti, E., & Warlop, L. (2014). *Social Power Increases Interoceptive Awareness*. Manuscript in Preparation.
- Mogg, K., & Bradley, B. P. (1998). A cognitive-motivational analysis of anxiety. *Behaviour Research and Therapy*, 36(9), 809–848.
- Mogg, K., Bradley, B., Miles, F., & Dixon, R. (2004). Time course of attentional bias for threat scenes: Testing the vigilance-avoidance hypothesis. *Cognition and emotion*, 18(5), 689-700.
- de Molière, L., & Guinote, A. (2012). Impact of power and disgust on the accessibility of cleanliness related words. Unpublished raw data.
- de Molière, L., & Harris, A. J. L. (2014). Impact of interoceptive sensibility and self-relevance on the relationship between negative utility and probability. Unpublished raw data.
- de Molière, L., & Harris, A. J. L. (2014). Impact of power and approach-avoidance motivation on the relationship between negative utility and probability. Unpublished raw data.
- de Molière, L., Harris, A., Quantmeyer, V. & Hahn, U. (2014). *On the Hunt for Wishful Thinking*. Manuscript submitted for publication.

- Montier, J. (2007). *Behavioural Investing: A Practitioner's Guide to Applying Behavioural Finance*. London: Wiley.
- Mor, N., & Winquist, J. (2002). Self-focused attention and negative affect: a meta-analysis. *Psychological bulletin*, 128(4), 638.
- Morey, R. D., & Rouder, J. N. (2011). Bayes factor approaches for testing interval null hypotheses. *Psychological methods*, 16(4), 406.
- Nelson, L. D., Galak, J., & Vosgerau, J. (2008). The unexpected enjoyment of expected events: The suboptimal consumption of televised sports. In A. Y. Lee & D. Soman (Eds.), *Advances in consumer research* (Vol. 35, pp. 185-188). Duluth, MN: Association for Consumer Research.
- Newhagen, J. E. (1998). TV news images that induce anger, fear, and disgust: Effects on approach-avoidance and memory. *Journal of Broadcasting & Electronic Media*, 42(2), 265-276.
- Norem, J. K. (2008). Defensive Pessimism, Anxiety, and the Complexity of Evaluating Self-Regulation. *Social and Personality Psychology Compass*, 2(1), 121-134.
- Ortony, A., Clore, G. L., & Collins, A. (1988). *The Cognitive Structure of Emotions*. *Contemporary Sociology* (Vol. 18, p. 957). Cambridge University Press.
- Pashler, H., Coburn, N., & Harris, C. R. (2012). Priming of social distance? Failure to replicate effects on social and food judgments. *PloS one*, 7(8), e42510.
- Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives on Psychological Science*, 7(6), 531-536.

- Pashler, H., & Wagenmakers, E. J. (2012). Editors' Introduction to the Special Section on Replicability in Psychological Science A Crisis of Confidence?. *Perspectives on Psychological Science*, 7(6), 528-530.
- Pellegrini, R. J., & Schauss, A. G. (1980). Muscle strength as a function of exposure to hue difference in visual stimuli: an experimental test of the kinesoid hypothesis. *Journal of Orthomolecular Psychiatry*, 9, 144-147.
- Pellegrini, R. J., Schauss, A. G., & Birk, T. (1980). Leg strength as a function of exposure to visual stimuli of different hues. *Bulletin of the Psychonomic Society*, 17, 27-28.
- Pennebaker, J., & Beal, S. (1986). Confronting a traumatic event: Toward an understanding of inhibition and disease. *Journal of Abnormal Psychology*, 95, 274-281.
- Pettigrew, R. (2012). Accuracy, Chance, and the Principal Principle, *121*(2), 241-276.
- Pleskac, T. J., & Hertwig, R. (in press). Ecologically rational choice and the structure of the environment. *Journal of Experimental Psychology: General*.
- Plutchik, R. (1998). *Emotions, diagnoses and ego defenses: A psychoevolutionary Perspective*, in W.F. Flack & J.D. Laird (Eds.). *Emotions in Psychopathology: Theory and research* (pp. 367-379). New York: Oxford University Press.
- Podolny, J. M. (1993). A status-based model of market competition. *American Journal of Sociology*, 829-872.
- Pollatos, O., Gramann, K., & Schandry, R. (2007). Neural systems connecting interoceptive awareness and feelings. *Human Brain Mapping*, 28(1), 9-18.

- Pollatos, O., Schandry, R., Auer, D. P., & Kaufmann, C. (2007). Brain structures mediating cardiovascular arousal and interoceptive awareness. *Brain research, 1141*, 178-187.
- Price, P. C. (2000). Wishful thinking in the prediction of competitive outcomes. *Thinking and Reasoning, 6*, 161-172.
- Price, P. C., Pentecost, H. C., & Voth, R. D. (2002). Perceived event frequency and the optimistic bias: Evidence for a two-process model of personal risk judgments. *Journal of Experimental Social Psychology, 38*(3), 242-252.
- Pruitt, D. G., & Hoge, R. D. (1965). Strength of the relationship between the value of an event and its subjective probability as a function of the method of measurement. *Journal of Experimental Psychology, 69*, 483-489.
- Quantmeyer, V. (2014). Do losses loom more likely than gains? The role of imagination. Unpublished MSc thesis. University College London, London, United Kingdom.
- Rainville, P., Bechara, A., Naqvi, N., & Damasio, A. R. (2006). Basic emotions are associated with distinct patterns of cardiorespiratory activity. *International Journal of Psychophysiology, 61*(1), 5-18.
- Ratcliff, R. (1993). Methods for dealing with reaction time outliers. *Psychological Bulletin, 114*(3), 510-532
- Reisenzein, R. (1995). On Oatley and Johnson-Laird's theory of emotion and hierarchical structures in the affective lexicon. *Cognition & Emotion, 9*(4), 383-416.

- Reisenzein, R., & Gattinger, E. (1982). Salience of arousal as a mediator of misattribution of transferred excitation. *Motivation and Emotion*, 6(4), 315-328.
- Reuter, M., Stark, R., Hennig, J., Walter, B., Kirsch, P., Schienle, A., & Vaitl, D. (2004). Personality and emotion: test of Gray's personality theory by means of an fMRI study. *Behavioral Neuroscience*, 118(3), 462-469.
- Ridgeway, C. L., & Walker, H. A. (1995). Status structures. *Sociological perspectives on social psychology*, 281-310.
- Risen, J. L., & Critcher, C. R. (2011). Visceral fit: While in a visceral state, associated states of the world seem more likely. *Journal of personality and social psychology*, 100(5), 777.
- Risen, J. L., & Gilovich, T. (2007). Another look at why people are reluctant to exchange lottery tickets. *Journal of Personality and Social Psychology*, 93(1), 12.
- Riskind, J. H. (1984). They stoop to conquer: Guiding and self-regulatory functions of physical posture after success and failure. *Journal of Personality and Social Psychology*, 47(3), 479.
- Roediger, H. L. (2012). Psychology's woes and a partial cure: The value of replication. *APS Observer*, 25(2), 9.
- Rose, J. P. (2009). *Approach-avoidance and optimism*. Unpublished doctoral dissertation, University of Iowa [retrieved July 10th 2014, <http://ir.uiowa.edu/etd/317>].
- Roth, W. T., Breivik, G., Jørgensen, P. E., & Hofmann, S. (1996). Activation in novice and expert parachutists while jumping. *Psychophysiology*, 33(1), 63-72.

- Rothbaum, F., Weisz, J. R., & Snyder, S. S. (1982). Changing the world and changing the self: A two-process model of perceived control. *Journal of personality and social psychology*, 42(1), 5.
- Rottenstreich, Y., & Hsee, C. K. (2001). Money, kisses, and electric shocks: on the affective psychology of risk. *Psychological Science*, 12, 185-190.
- Rotter, J. B. (1954). *Social learning and clinical psychology*. Englewood Cliffs, NJ: Prentice-Hall.
- Rouder, J. N., & Morey, R. D. (2011). A Bayes factor meta-analysis of Bem's ESP claim. *Psychonomic Bulletin & Review*, 18(4), 682-689.
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic bulletin & review*, 16(2), 225- 237.
- Rozin, P., Haidt, J., & McCauley, C. R. (2000). Disgust. In M. Lewis & J. Haviland (Eds.), *Handbook of emotions, 2nd Edition* (pp. 637–653). New York: Guilford Press.
- Russell, B. (1938). *Power, a new social analysis*. London: Unwin Books.
- Russell, J. A. (1980). A circumplex model of affect. *Journal of personality and social psychology*, 39(6), 1161-1178.
- Sagristano, M. D., Trope, Y., & Liberman, N. (2002). Time-dependent gambling: odds now, money later. *Journal of Experimental Psychology: General*, 131(3), 364.
- Savage, L.J. 1954. *The foundation of statistics*. Wiley, New York.
- Schachter, S., & Singer, J. (1962). Cognitive, social, and physiological determinants of emotional state. *Psychological review*, 69(5), 379.

- Schandry, R. (1981). Heart beat perception and emotional experience. *Psychophysiology*, 18(4), 483–8.
- Scheepers, D., de Wit, F., Ellemers, N., & Sassenberg, K. (2012). Social power makes the heart work more efficiently: Evidence from cardiovascular markers of challenge and threat. *Journal of Experimental Social Psychology*, 48(1), 371-374.
- Scheier, M. F., & Carver, C. S. (1977). Self-focused attention and the experience of emotion: attraction, repulsion, elation, and depression. *Journal of Personality and Social Psychology*, 35(9), 625–36.
- Scheier, M. F., & Carver, C. S. (1985). Optimism, coping, and health: assessment and implications of generalized outcome expectancies. *Health psychology*, 4(3), 219.
- Schmid-Mast, M., Jonas, K., & Hall, J. A. (2009). Give a person power and he or she will show interpersonal sensitivity: the phenomenon and its why and when. *Journal of personality and social psychology*, 97(5), 835.
- Schnall, S., Haidt, J., Clore, G. L., & Jordan, A. H. (2008b). Disgust as embodied moral judgment. *Personality and Social Psychology Bulletin*, 34, 1096–1109.
- Schwarz, N., & Clore, G. L. (1983). Mood , Misattribution , and Judgments of Well-Being : Informative and Directive Functions of Affective States. *Journal of Personality and Social Psychology*, 45(3), 513–523.
- Seeman, M. (1963). Alienation and social learning in a reformatory. *American Journal of Sociology*, 270-284.
- Shaffer, V. A., & Arkes, H. R. (2009). Preference reversals in evaluations of cash versus non-cash incentives. *Journal of Economic Psychology*, 30(6), 859-872.

- Shanks, D. R., Newell, B. R., Lee, E. H., Balakrishnan, D., Ekelund, L., Cenac, Z., ... & Moore, C. (2013). Priming intelligent behavior: An elusive phenomenon. *PloS one*, 8(4), e56515.
- Sharot, T., Korn, C. W., & Dolan, R. J. (2011). How unrealistic optimism is maintained in the face of reality. *Nature neuroscience*, 14(11), 1475-1479.
- Sharot, T., Guitart-Masip, M., Korn, C. W., Chowdhury, R., & Dolan, R. J. (2012a). How dopamine enhances an optimism bias in humans. *Current Biology*, 22(16), 1477-1481.
- Sharot, T., Guitart-Masip, M., Korn, C. W., Chowdhury, R., & Dolan, R. J. (2012b). How dopamine enhances an optimism bias in humans. *Current Biology*, 22(16), 1477-1481.
- Shepperd, J. A., Findley-Klein, C., Kwavnick, K. D., Walker, D., & Perez, S. (2000). Bracing for loss. *Journal of personality and social psychology*, 78(4), 620.
- Shepperd, J. A., Klein, W. M., Waters, E. A., & Weinstein, N. D. (2013). Taking stock of unrealistic optimism. *Perspectives on Psychological Science*, 8(4), 395-411.
- Shepperd, J. A., Ouellette, J. A., & Fernandez, J. K. (1996). Abandoning unrealistic optimism: Performance estimates and the temporal proximity of self-relevant feedback. *Journal of Personality and Social Psychology*, 70(4), 844.
- Sherif, M., & Sherif, C. W. (1956). *An Outline of Social Psychology (revised edition)*. New York: Harper & Brothers.
- Sherman, S. J., Cialdini, R. B., Schwartzman, D. F., & Reynolds, K. D. (1985). Imagining can heighten or lower the perceived likelihood of contracting a disease

- the mediating effect of ease of imagery. *Personality and Social Psychology Bulletin*, 11(1), 118-127.
- Shields, S. A., Mallory, M. E., & Simon, A. (1989). The body awareness questionnaire: Reliability and validity. *Journal of Personality Assessment*, 53(4), 802-815.
- Siegrist, M., & Sütterlin, B. (2014). Human and Nature-Caused Hazards: The Affect Heuristic Causes Biased Decisions. *Risk Analysis*, 34(8), 1482-1494.
- Simmons, J. P., & Massey, C. (2012). Is optimism real?. *Journal of Experimental Psychology: General*, 141(4), 630.
- Slabu, L., & Guinote, A. (2010). Getting what you want: Power increases the accessibility of active goals. *Journal of Experimental Social Psychology*, 46(2), 344-349.
- Sligte, D. J., de Dreu, C. K., & Nijstad, B. A. (2011). Power, stability of power, and creativity. *Journal of Experimental Social Psychology*, 47(5), 891-897.
- Slovic, P. (1966). Value as a determiner of subjective probability. *IEEE Transactions on Human Factors in Electronics*, HFE-7, 22-28.
- Slovic, P., Fischhoff, B., & Lichtenstein, S. (1982). Facts versus fears. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment Under Uncertainty: Heuristics and Biases* (pp. 463-489). Cambridge, UK: Cambridge University Press.
- Smith, P. K., & Bargh, J. A. (2008). Nonconscious Effects of Power on Basic Approach and Avoidance Tendencies. *Social Cognition*, 26(1), 1-24.

- Smith, N. K., Cacioppo, J. T., Larsen, J. T., & Chartrand, T. L. (2003). May I have your attention, please: Electrocortical responses to positive and negative stimuli. *Neuropsychologia*, *41*(2), 171-183.
- Smith, P. K., Jostmann, N. B., Galinsky, A. D., & van Dijk, W. W. (2008). Lacking power impairs executive functions. *Psychological science*, *19*(5), 441-447.
- Smith, P. K., & Trope, Y. (2006). You focus on the forest when you're in charge of the trees: power priming and abstract information processing. *Journal of Personality and Social Psychology*, *90*(4), 578-96.
- Sokol-Hessner, P., Hartley, C. A., Hamilton, J. R., & Phelps, E. A. (2014). Interoceptive ability predicts aversion to losses. *Cognition and Emotion*, (ahead-of-print), 1-7.
- Sokol-Hessner, P., Hsu, M., Curley, N. G., Delgado, M. R., Camerer, C. F., & Phelps, E. A. (2009). Thinking like a trader selectively reduces individuals' loss aversion. *Proceedings of the National Academy of Sciences*, *106*(13), 5035-5040.
- Steingroever, H., Wetzels, R., Horstmann, A., Neumann, J., & Wagenmakers, E. J. (2013). Performance of healthy participants on the Iowa Gambling Task. *Psychological assessment*, *25*(1), 180.
- Stephan, E., Liberman, N., & Trope, Y. (2011). The effects of time perspective and level of construal on social distance. *Journal of experimental social psychology*, *47*(2), 397-402.
- Strick, H.K. (2007). George-Louis Leclerc Buffon (1707-1788). Retrieved 29.09.2014 from <http://www.mathematics-in-europe.eu/de/startseite/76-enjoy-maths/strick/654-georges-louis-leclerc-buffon-1707-1788-by-heinz-klaus-strick-germany>.

- Stroebe, W., Postmes, T., & Spears, R. (2012). Scientific misconduct and the myth of self-correction in science. *Perspectives on Psychological Science*, 7(6), 670-688.
- Sundar, S. S., & Kalyanaraman, S. (2004). Arousal, memory, and impression-formation effects of animation speed in web advertising. *Journal of Advertising*, 33(1), 7-17.
- Tamir, M., & Robinson, M. D. (2007). The happy spotlight: Positive mood and selective attention to rewarding information. *Personality and Social Psychology Bulletin*, 33(8), 1124-1136.
- Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: The mobilization-minimization hypothesis. *Psychological Bulletin*, 110, 67-85.
- Terasawa, Y., Fukushima, H., & Umeda, S. (2013). How does interoceptive awareness interact with the subjective experience of emotion? An fMRI Study. *Human brain mapping*, 34(3), 598-612.
- Thibaut, J., & Kelly, H. (1959): *The Social Psychology of Groups*. New York.
- Thorndike, E.L. (1935). *The psychology of wants, interest, and attitudes*. New York: Appleton-Century-Crofts.
- Tiedens, L. Z. (2001). Anger and advancement versus sadness and subjugation: the effect of negative emotion expressions on social status conferral. *Journal of Personality and Social Psychology* (Vol. 80, pp. 86-94).
- Tolin, D. F., Lohr, J. M., Sawchuk, C. N., & Lee, T. C. (1997). Disgust and disgust sensitivity in blood-injection-injury and spider phobia. *Behaviour Research and Therapy*, 35(10), 949-953.
- Tomkins, S. S. (1970). Affect as the primary motivational system. *Feelings and emotions*, 101-110.

- Trope, Y., & Liberman, N. (2000). Temporal construal and time-dependent changes in preference. *Journal of personality and social psychology*, 79(6), 876.
- Trope, Y., & Liberman, N. (2003). Temporal construal. *Psychological review*, 110(3), 403.
- Trope, Y., & Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological review*, 117(2), 440.
- Tukey, J. W. (1969). Analyzing data: Sanctification or detective work?. *American Psychologist*, 24(2), 83.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, 5(2), 207-232.
- Tversky, A., & Kahneman, D. (1981). The framing of decisions and the psychology of choice. *Science*, 211(4481), 453-458.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4), 297-323.
- Verplanken, B. (1997). The Effect of Catastrophe Potential on the Interpretation of Numerical Probabilities of the Occurrence of Hazards¹. *Journal of Applied Social Psychology*, 27(16), 1453-1467.
- Von Winterfeldt, D., & Edwards, W. (1986). *Decision analysis and behavioral research* (Vol. 604). Cambridge: Cambridge University Press.
- Vosgerau, J. (2010). How prevalent is wishful thinking? Misattribution of arousal causes optimism and pessimism in subjective probabilities. *Journal of Experimental Psychology. General*, 139(1), 32-48.

- Wagenmakers, E. J., Wetzels, R., Borsboom, D., & Van Der Maas, H. L. (2011). Why psychologists must change the way they analyze their data: the case of psi: comment on Bem (2011).
- Wakslak, C. J. (2012). The where and when of likely and unlikely events. *Organizational Behavior and Human Decision Processes*, *117*(1), 150-157.
- Wakslak, C. J., Trope, Y., Liberman, N., & Alony, R. (2006). Seeing the forest when entry is unlikely: probability and the mental representation of events. *Journal of Experimental Psychology. General*, *135*(4), 641-53.
- Wakslak, C., & Trope, Y. (2009). The effect of construal level on subjective probability estimates. *Psychological Science*, *20*(1), 52-58.
- Wallsten, T. S., & Budescu, D. V. (1983). State of the art—Encoding subjective probabilities: A psychological and psychometric review. *Management Science*, *29*(2), 151-173.
- Warr, P., Barter, J., & Brownbridge, G. (1983). On the independence of positive and negative affect. *Journal of Personality and Social Psychology*, *44*, 644-651.
- Watson, D., Clark, L. A., & Tellegen, A. (1984). Cross-cultural convergence in the structure of mood: A Japanese replication and a comparison with U.S. findings. *Journal of Personality and Social Psychology*, *47*, 127-144.
- Watson, D., & Tellegen, A. (1985). Toward a consensual structure of mood. *Psychological Bulletin*, *98*, 219-235.
- Windschitl, P. D., Smith, A. R., Rose, J. P., & Krizan, Z. (2010). The desirability bias in predictions: Going optimistic without leaving realism. *Organizational Behavior and Human Decision Processes*, *111*(1), 33-47.

- Webb, W. B. (1985). Volume 22 MARCH 1985 Number 2, 22(2).
- Weber, E. U. (1994). From subjective probabilities to decision weights: The effect of asymmetric loss functions on the evaluation of uncertain outcomes and events. *Psychological Bulletin*, 115(2), 228–242.
- Weber, E. U., & Hilton, D. J. (1990). Contextual effects in the interpretations of probability words: Perceived base rate and severity of events. *Journal of Experimental Psychology: Human Perception and Performance*, 16(4), 781.
- Weick, M., & Guinote, A. (2008). When subjective experiences matter: power increases reliance on the ease of retrieval. *Journal of Personality and Social Psychology*, 94(6), 956–70.
- Weinberger, D. (1990). The construct validity of the repressive coping style. In Singer, J.L., editor. *Repression and dissociation: implications for personality theory, psychopathology, and health* (pp. 337–386). Chicago, IL: University of Chicago Press.
- Weinberger, D. A., & Davidson, M. N. (1994). Styles of inhibiting emotional expression: distinguishing repressive coping from impression management. *Journal of Personality*, 62(4), 587–613.
- Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of Personality and Social Psychology*, 39, 806–820.
- Weinstein, N. D. (1982). Unrealistic optimism about susceptibility to health problems. *Journal of Behavioral Medicine*, 5, 441–460.
- Weinstein, N. D. (1984). Why it won't happen to me: Perceptions of risk factors and susceptibility. *Health Psychology*, 3, 431–457.

- Weinstein, N. D., & Klein, W. M. (1995). Resistance of personal risk perceptions to debiasing interventions. *Health Psychology, 14*, 132–140.
- Weinstein, N. D., Marcus, S. E., & Moser, R. P. (2005). Smokers' unrealistic optimism about their risk. *Tobacco control, 14*(1), 55-59.
- Weisz, J. R., Rothbaum, F. M., & Blackburn, T. C. (1984). Standing out and standing in: The psychology of control in America and Japan. *American Psychologist, 39*(9), 955.
- Werner, N. S., Jung, K., Duschek, S., & Schandry, R. (2009). Enhanced cardiac perception is associated with benefits in decision-making. *Psychophysiology, 46*(6), 1123–9.
- Werner, N. S., Mannhart, T., Reyes Del Paso, G. A., & Duschek, S. (2014). Attention interference for emotional stimuli in cardiac interoceptive awareness. *Psychophysiology, 51*(6), 573-578.
- Werner, N. S., Schweitzer, N., Meindl, T., Duschek, S., Kambeitz, J., & Schandry, R. (2013). Interoceptive awareness moderates neural activity during decision-making. *Biological Psychology, 94*(3), 498–506.
- Whitehead, W. E., Drescher, V. M., Heiman, P., & Blackwell, B. (1977). Relation of heart rate control to heartbeat perception. *Biofeedback and Self-Regulation, 2*(4), 371-392.
- Wiens, S., Mezzacappa, E. S., & Katkin, E. S. (2000). Heartbeat detection and the experience of emotions. *Cognition & Emotion, 14*(3), 417-427.
- Wright, W. F., & Bower, G. H. (1992). Mood effects on subjective probability

assessment. *Organizational behavior and human decision processes*, 52(2), 276-291.

Zeelenberg, M., Nelissen, R. M., Breugelmans, S. M., & Pieters, R. (2008). On emotion specificity in decision making: Why feeling is for doing. *Judgment and Decision making*, 3(1), 18-27.

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.

Zevon, M. A., & Tellegen, A. (1982). The structure of mood change: An idiographic/nomothetic analysis. *Journal of Personality and Social Psychology*, 43, 111-122.

Zillmann, D. (1971). Excitation transfer in communication mediated aggressive behaviour. *Journal of Experimental Social Psychology*, 4, 419-434.

Zillmann, D., Mody, B., & Cantor, J. R. (1974). Empathetic perception of emotional displays in films as a function of hedonic and excitatory state prior to exposure. *Journal of Research in Personality*, 8(4), 335-349.

Appendix

Appendix A

Items of the Consciousness of Body – Private scale

-
1. I am sensitive to internal bodily tensions.
 2. I know immediately when my mouth or throat gets dry.
 3. I can often feel my heart beating.
 4. I am quick to sense the hunger contractions of my stomach.
 5. I am very aware of changes in my body temperature.
-

Answers were provided on a 6-point Likert scale from 1 (disagree strongly) to 6 (agree strongly).

Appendix B

Items of the Body Awareness Questionnaire

-
1. I notice differences in the way my body reacts to various foods.
 2. I always know when I've exerted myself to the point where I'll be sore the next day.
 3. I am always aware of changes in my energy level when I eat certain foods.
 4. I know in advance when I'm getting the flu.
 5. I know I'm running a fever without taking my temperature.
 6. I can distinguish between tiredness because of hunger and tiredness because of lack of sleep.
 7. I can accurately predict what time of day lack of sleep will catch up with me.
 8. I am aware of a cycle in my activity level throughout the day.
 9. I don't notice seasonal rhythms and cycles in the way my body functions.
 10. As soon as I wake up in the morning, I know how much energy I'll have during the day.
 11. I can tell when I go to bed how well I will sleep that night.
 12. I notice distinct body reactions when I am fatigued.
 13. I notice specific body responses to changes in the weather.
 14. I can predict how much sleep I will need at night in order to wake up refreshed.
 15. When my exercise habits change, I can predict very accurately how that will affect my energy level.
 16. There seems to be a "best" time for me to go to sleep at night.
 17. I notice specific bodily reactions to being overhungry.
 18. I can always tell when I bump myself whether or not it will become a bruise.
-

Answers were provided on a 6-point Likert scale from 0 (never) to 5 (always).

Appendix C

Items of the Noticing subscale from the Multidimensional Assessment of Interoceptive Awareness

-
1. I notice how my body changes when I am angry.
 2. When something is wrong in my life I can feel it in my body.
 3. I notice that my body feels different after a peaceful experience.
 4. I notice that my breathing becomes free and easy when I feel comfortable.
 5. I notice how my body changes when I feel happy/joyful.
-

Answers were provided on a 6-point Likert scale from 0 (never) to 5 (always).

Appendix D

Items of the Body Listening subscale from the Multidimensional Assessment of Interoceptive Awareness

-
1. I listen for information from my body about my emotional state.
 2. When I am upset, I take time to explore how my body feels.
 3. I listen to my body to inform me about what to do.
-

Answers were provided on a 6-point Likert scale from 0 (never) to 5 (always).

Appendix E

Items of the Sense of Power scale

-
1. Even if I voice them, my views have little sway.
 2. I can get others to do what I want.
 3. I can get people to listen to what I say.
 4. I think I have a great deal of power.
 5. My wishes don't carry much weight.
 6. My ideas and opinions are often ignored.
-

Answers were provided on a 7-point Likert scale from 1 (disagree strongly) to 7 (agree strongly).

Appendix F

Items of the Positive and Negative Affect scale

1. Interested	11. Irritable
2. Distressed	12. Alert
3. Excited	13. Ashamed
4. Upset	14. Inspired
5. Strong	15. Nervous
6. Guilty	16. Determined
7. Scared	17. Attentive
8. Hostile	18. Jittery
9. Enthusiastic	19. Active
10. Proud	20. Afraid

Participants indicated to what extent they felt the emotion at the present moment (very slightly or not at all, a little, moderately, quite a bit, extremely).

Appendix G

Items of the short version of the State Trait Anxiety Inventory

1. I feel calm.	4. I feel relaxed.
2. I feel tense.	5. I feel content.
3. I feel upset.	6. I am worried.

Participants indicated to what extent they felt the emotion at the present moment (not at all, somewhat, moderately, very much).

Appendix H

Items of the BIS-BAS questionnaire

-
1. My family is the most important thing in my life.
 2. Even if something bad is about to happen to me, I rarely experience fear or nervousness.
 3. I go out of my way to get things I want.
 4. When I'm doing well at something I love to keep at it.
 5. I'm always willing to try something new if I think it will be fun.
 6. How I dress is important to me.
 7. When I get something I want, I feel excited and energized.
 8. Criticism or scolding hurts me quite a bit.
 9. When I want something I usually go all-out to get it.
 10. I will often do things for no other reason than that they might be fun.
 11. It's hard for me to find the time to do things such as get a haircut.
 12. If I see a chance to get something I want I move on it right away.
 13. I feel pretty worried or upset when I think or know somebody is angry at me.
 14. When I see an opportunity for something I like I get excited right away.
 15. I often act on the spur of the moment.
 16. If I think something unpleasant is going to happen I usually get pretty "worked up."
 17. I often wonder why people act the way they do.
 18. When good things happen to me, it affects me strongly.
 19. I feel worried when I think I have done poorly at something important.
 20. I crave excitement and new sensations.
 21. When I go after something I use a "no holds barred" approach.
 22. I have very few fears compared to my friends.
 23. It would excite me to win a contest.
 24. I worry about making mistakes.
-

Participants indicated to what extent the statement was true for them at the present moment (very true for me, somewhat true for me, somewhat false for me, very false for me). Note that items 1, 6, 11 and 17 are fillers.

Appendix I

Items of the modified BIS-BAS questionnaire

-
1. At the moment my family is the most important thing in my life.
 2. Even if something bad was about to happen to me now, I would rarely experience fear or nervousness.
 3. Right now, I would go out of my way to get things I want.
 4. Currently, when I would do well at something I would love to keep at it.
 5. I would try something new right now if I think it would be fun.
 6. How I dress is important to me.
 7. If I would get something I want now, I would feel excited and energized.
 8. If I would get criticized now, it would hurt me quite a lot.
 9. If I wanted something at the moment, I would go all-out to get it.
 10. Right now, I would do something for no other reason than that it might be fun.
 11. Nowadays, it's hard for me to find the time to do things such as get a haircut.
 12. If I would see the chance to get something I want at this moment, I would move on it right away.
 13. If currently I would think or know that someone is angry at me, I would feel pretty worried or upset.
 14. I would get excited right away if I saw an opportunity for something I would like now.
 15. Right now I would act on the spur of the moment.
 16. If I would currently think that something unpleasant is about to happen I would get pretty "worked up".
 17. I wonder nowadays why people act the way they do.
 18. If something good would happen to me right now, it would affect me strongly.
 19. If I would think I would have done poorly at something important now I would feel worried.
 20. I crave excitement and new sensations at the moment.
 21. If I would go after something now I would use a "no holds barred" approach.
 22. I have very few fears right now.
 23. If I could currently win something nice, then that would strongly draw my attention.
 24. At the moment I worry about making mistakes.
-

Participants indicated to what extent the statement was true for them at the present moment (very true for me, somewhat true for me, somewhat false for me, very false for me). Note that items 1, 6, 11 and 17 are fillers.