

**REASONING AND PROCESSING OF BEHAVIOURAL AND
CONTEXTUAL INFORMATION: INFLUENCES ON PRE-
JUDGEMENT REASONING, POST-JUDGEMENT INFORMATION
SELECTION AND ENGAGEMENT, AND MORAL BEHAVIOUR**

by

NINA LAUREL POWELL

(B.A. The University of North Carolina at Greensboro)

A thesis submitted to the University of Birmingham

for the degree of

DOCTOR OF PHILOSOPHY

School of Psychology
College of Life and Environmental Sciences
The University of Birmingham
June 2013

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

ABSTRACT

Recent research on moral judgements tends to emphasise the role of intuition, emotion and non-deliberative gut-reactions to moral violations. The aim of this thesis was to investigate instances during the judgement process and on resulting behaviour when deliberative consideration and processing of behavioural and contextual information (i.e., information beyond initial gut-reactions and intuitions) occurred. Specifically, this thesis examined the effects of reasoning about behavioural and contextual information pre-judgement, the desires and needs for and engagement with behavioural information and the effects of behavioural and contextual information on eliciting moral behaviour. Across seven experiments, I demonstrated (1) that age-related changes in the ability to reason about the means through which a negative outcome occurred influenced attributions of blameworthiness, (2) that post-judgement information selection and engagement differed depending on the moral violation judged, emotions elicited from the violations and the amount of reported epistemic certainty, and (3) that the presence of information about the outcome of a morally virtuous act influenced later helping behaviour. These findings suggest that deliberative reasoning and processing of behavioural and contextual information can occur and influence judgements and behaviour at different stages in the judgement process.

ACKNOWLEDGEMENTS

First of all, I would like to thank Kim Quinn and Sarah Beck for their support and guidance, without which I would not have had the opportunity to conduct this research. I would like to thank Kim for making me a well-rounded thinker and researcher, and her encouragement to always “think like a lecturer”; and Sarah for unwavering faith in my ideas, and her ability to consistently challenge and motivate me. I owe my deepest gratitude to both Kim and Sarah for giving me the opportunity to work with them during the summer of 2008 and the winter of 2009; you both have inspired me.

This thesis would not have been possible without the studentship I was awarded for my PhD research provided by the University of Birmingham, School of Psychology. I would like to thank the studentship committee for their decision to fund my research proposal.

I am grateful to Rob Guttentag who has been a mentor, collaborator, teacher and friend. Rob’s encouragement and faith in me during my time as an undergraduate inspired me to pursue a PhD in this topic. Rob planted the seed for the opportunity to study and work at the University of Birmingham. I am eternally grateful to Rob for providing me with unique opportunities to teach, present and conduct research at such an early stage, and especially for the weekly meetings in Tate Street Coffee.

Finally, I would like to thank my family; my mom, dad, John, Laura, and Grandma Isabel for their kindness, love, support and encouragement over the past several years. I would particularly like to thank my dad for always encouraging me to grab every opportunity to learn and to be a “citizen of the world”. I would like to thank Kim Sein for his kindness, support and generosity during the difficult period of writing-up. And last but not least, I would like to thank Stuart Derbyshire for his guidance and faith in me; knowing Stuart has made me a better thinker, researcher and person.

TABLE OF CONTENTS

CHAPTER 1 THE ROLE OF REASONING AND INTUITION IN MORAL JUDGEMENTS.....	12
1 General Introduction.....	13
1.1 The Shift from Reasoning to Intuition.....	13
1.2 The Moral Emotions.....	22
1.2.1 Characteristics and Groupings of Moral Emotions.....	22
1.2.2 When Information and Reasoning Can Override Emotion.....	23
1.2.3 Dual-Process Theory of Moral Judgement.....	25
2 Core Hypotheses.....	26
CHAPTER 2 THE INFLUENCE OF BEHAVIOURAL AND CONTEXTUAL INFORMATION PRE-JUDGEMENT: EVIDENCE FROM REGRET AND BLAME FOR DEVELOPMENTAL CHANGES IN REASONING AND PROCESSING OF INFORMATION.....	32
1 Introduction.....	33
2 Experiment 1.....	41
2.1 Method.....	42
2.1.1 Participants and design.....	42
2.1.2 Materials and procedure.....	43
2.2 Results.....	45
2.2.1 Commission/omission stories.....	46
2.2.2 Atypical/typical stories.....	47
2.3 Discussion.....	48
3 Experiment 2.....	49
3.1 Method.....	50
3.1.1 Participants and design.....	50
3.1.2 Materials and procedure.....	50
3.2 Results.....	52
3.3 Discussion.....	54
4 General Discussion.....	55
CHAPTER 3 ENGAGEMENT WITH BEHAVIOURAL AND CONTEXTUAL INFORMATION POST-JUDGEMENT: MORAL CONDEMNATION OF HARM AND PURITY VIOLATIONS AND DESIRE TO KNOW MORE.....	59
1 Moral Judgements, Attitudes, and Information Engagement.....	60
1.1 Differences between Attitudes and Moral Principles.....	61
1.2 Attitudes and Additional Information.....	64
1.2.1 Attitude Certainty.....	69
1.3 Summary.....	70

2 Post-Judgement Selection and Engagement with Behavioural and Contextual Information	71
3 Experiment 3	79
3.1 Method	79
3.1.1 Participants and design	79
3.1.2 Materials	79
3.1.3 Procedure	80
3.2 Results and Discussion	81
3.2.1 Moral appropriateness, judgement certainty, and emotion intensity	81
3.2.2 Desire for information	82
3.2.3 Usefulness of information	84
3.2.4 Summary	84
4 Experiment 4	85
4.1 Method	86
4.1.1 Participants and design	86
4.1.2 Materials and procedure	86
4.2 Results and Discussion	87
4.2.1 Moral appropriateness, judgement certainty, and emotion intensity	87
4.2.2 Desire for additional information	88
4.2.3 Functions of post-judgement reasoning	89
4.2.4 Summary	90
5 Experiment 5	91
5.1 Method	91
5.1.1 Participants and design	91
5.1.2 Materials	92
5.1.2.1 Moral violations	92
5.1.2.2 Additional information	93
5.1.3 Procedure	93
5.2 Results and Discussion	96
5.2.1 Set 1 violations	97
5.2.1.1 Engagement with additional information	97
5.2.1.2 Moral wrongness	97
5.2.1.3 Judgement certainty	101
5.2.1.4 Emotion ratings	103
5.2.1.5 Facial expression ratings	104
5.2.2 Set 2 violations	105
5.2.2.1 Engagement with additional information	105
5.2.2.2 Moral wrongness	105
5.2.2.3 Judgement certainty	110
5.2.2.4 Emotion ratings	111
5.2.2.5 Facial expression ratings	112
5.2.3 Summary	113
6 Experiment 6	116
6.1 Method	117
6.1.1 Participants and design	117

6.1.2 Materials.....	117
6.1.2.1 Moral violations.....	117
6.1.2.2 Additional information.....	117
6.1.3 Procedure.....	118
6.2 Results and Discussion.....	119
6.2.1 Engagement with additional information.....	119
6.2.2 Moral wrongness.....	119
6.2.3 Judgement certainty.....	120
6.2.4 Emotion ratings.....	120
6.2.5 Facial expression ratings.....	121
6.3 Summary.....	123
7 General Discussion.....	124
CHAPTER 4 THE INFLUENCE OF BEHAVIOURAL AND CONTEXTUAL INFORMATION ON BEHAVIOUR: MORAL ELEVATION WITHOUT OUTCOME INFORMATION MAY NOT PROMOTE HELPING BEHAVIOUR.....	135
1 Introduction.....	136
2 Experiment 7.....	141
2.1 Method.....	141
2.1.1 Participants and design.....	141
2.1.2 Stimuli.....	142
2.1.3 Measures.....	143
2.1.3.1 Emotion ratings.....	143
2.1.3.2 Memory.....	143
2.1.2 Procedure.....	143
2.2 Results.....	145
2.2.2 Emotion.....	145
2.2.2.1 Comparisons between video conditions.....	146
2.2.2.2 Comparisons within video condition.....	146
2.2.1 Helping behaviour.....	146
2.2.1.1 Time spent completing anagrams task.....	146
2.2.1.2 Willingness to volunteer for future research.....	147
2.3 General Discussion.....	147
3 Conclusion.....	150
CHAPTER 5 GENERAL DISCUSSION.....	152
1 Background and Aims of the Thesis.....	152
2 Summary of Findings.....	153
2.1 Sophistication of Reasoning Ability in Processing Information.....	154
2.2 Different Moral Violations and Emotions Lead to Different Engagement with Information.....	157
2.2.1 Limitations.....	160

2.3 Impact of Behavioural and Contextual Information on Behaviour	162
3 Implications for Models of Moral Judgement	163
4 Future Directions	168
4.1 Asymmetry between moral condemnation and moral praise.	168
5 Conclusion	171
REFERENCES	173
APPENDICES	185

LIST OF FIGURES

Figure 1. The social intuitionist model (Haidt, 2001)	19
Figure 2. Mean ratings of desire for wrongness-confirming and wrongness-contradicting information as a function of violation type, Experiment 4.....	88
Figure 3. Sample facial expressions taken from (Beaupré, Cheung, & Hess, 2000) used in emotion measure questionnaire, Experiments 5 and 6.	95
Figure 4. Interaction between the times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and the violation type for Set 1, Experiment 5.....	98
Figure 5. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 1, Experiment 5.	99
Figure 6. Interaction between times of judgement (initial wrongness-judgement vs. final wrongness judgement post-information), violation type (harm vs. purity) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 1, Experiment 5.	100
Figure 7. Interaction between times of judgement (initial certainty judgement vs. final certainty judgement post-information) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 1, Experiment 5.	102
Figure 8. Interaction between violation type (harm vs. purity) and facial expression ratings (angry face vs. disgusted face) for Set 2, Experiment 5.....	113
Figure 9. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and violation type (harm vs. purity) for Set 2, Experiment 5.	106
Figure 10. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and information congruency (wrongness-confirming vs. wrongness contradicting) for Set 2, Experiment 5.....	107
Figure 11. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information), violation type (harm vs. purity) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 2, Experiment 5.	108
Figure 12. Interaction between times of judgement (initial certainty judgement vs. final certainty judgement post-information) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 2, Experiment 5.	110
Figure 13. Interaction between violation type (harm vs. purity) and emotion ratings (anger vs. disgust), Experiment 6.....	121

Figure 14. Interaction between violation type (harm vs. purity) and facial expression ratings (angry face vs. disgusted face), Experiment 6..... 122

Figure 15. Mean self-reported emotion scores as a function of emotion type and video condition. 145

Figure 16. Mean time (in minutes) spent completing anagrams task as a function of video condition. 147

LIST OF TABLES

Table 1. Mean Number (and Standard Errors) of Target Character Responses for Commission/Omission Stories and Atypical/Typical Stories for Both Regret and Blame Questions as a Function of Age, Experiment 1	466
Table 2. Percentage of Non-Target, Target, and Same Character Responses for Atypical/Typical Stories and Commission/Omission Stories for Regret and Blame Question as a Function of Age, Experiment 1	466
Table 3. Mean Ratings (and Standard Deviations) in Response to Harm versus Purity Violations, Experiment 3.....	82
Table 4. Number of Participants Expressing Desire for Additional Information as a Function of Probe Question and Violation Type, Experiment 3	833
Table 5. Mean Number (and Standard Deviations) of Responses to Probes Regarding Desire for Additional Information as a Function of Response Type and Violation Type, Experiment 3	833
Table 6. Mean Ratings (and Standard Deviations) in Response to Harm versus Purity Violations, Experiment 4.....	877
Table 7. Correlations between Desire for Information and Responses to the Moral Violations as a Function of Violation Type and Information Congruity, Experiment 4	90
Table 8. Mean Number of Lines of Additional Information Read (and Standard Errors) Depending on the Information Congruency Depending and Violation Type for Set 1, Experiment 5	97
Table 9. Mean Ratings (and Standard Errors) of Emotion Depending on Violation Type for Set 1, Experiment 5	104
Table 10. Mean Number of Lines of Additional Information Read (and Standard Errors) Depending on the Information Congruency Depending and Violation Type for Set 2, Experiment 5	10505
Table 11. Mean Ratings (and Standard Errors) of Emotion Depending on Violation Type for Set 2, Experiment 5	112

LIST OF APPENDICES

<i>Appendix A:</i> Materials from Experiments 1 and 2, Chapter 2.....	185
<i>Appendix B:</i> Response Scales Used in Experiment 2, Chapter 2	193
<i>Appendix C:</i> Stories Used for Training for Response Scales, Experiment 2, Chapter 2	194
<i>Appendix D:</i> Stories and Lines of Additional Information, Experiment 5, Chapter 3.....	197
<i>Appendix E:</i> Stories and Lines of Additional Information, Experiment 6, Chapter 3.....	207
<i>Appendix F:</i> 10-item Memory Questionnaires, Experiment 7, Chapter 4.....	211

CHAPTER 1

THE ROLE OF REASONING AND INTUITION IN MORAL JUDGEMENTS

The study of morality is the study of when, why and how people perform and assess good and bad actions on an individual and societal level. Morality arguably derives from a set of deontological principles, the weighing of positive and negative outcomes, natural intuitions and/or relativistic social norms (Haidt, 2008). One important role in morality is the role of the third-party; a disinterested person or group who maintain moral order in society through their judgement of others' actions (DeScioli & Kurzban, 2009). Judgements are argued to be made through a process of conscious, deliberative reasoning or unconscious intuition and emotion. The moral judgement literature has largely presented the two processes as competing processes, with more recent theories (e.g., social intuitionist model (Haidt, 2001)) favouring the predominant role of intuition and emotion. There have been efforts to bridge this dichotomy of reason versus intuition. Although this thesis does not provide any direct comparisons of reason versus intuition in the moral judgement process, this thesis examines how moral judgements from disinterested, third-party judges are affected by the consideration or presence of behavioural and contextual information at different points during the judgement processes and subsequent moral behaviour. The examination of how moral judgements and elicited moral behaviour are affected by behavioural and contextual information may shed light on other factors that influence moral judgements besides emotion and intuition, and may contribute to a more nuanced understanding of how emotion and some level of reasoning about information are influential.

1 General Introduction

Moral psychology has largely focused on determining the origin of our moral judgements. In doing so, reasoning and emotion have been identified as being influential processes in making moral judgements. The precise causal role reasoning and emotion have in moral judgement, however, is arguable. The present research aimed to investigate conditions under which processing of behavioural and contextual information affects moral judgements at different stages: pre-judgement, post-judgement, and during behaviour. With evidence to support the role of both reasoning and emotion, I turn to the moral psychology literature for insight into how moral and judgements are formed and affected by affect and reasoning at different stages.

With respect to the moral psychology literature, I will first explore the shift from the belief that moral judgements are largely rooted in causal reasoning (Kohlberg, 1971, 1973, 1976; Nucci & Turiel, 1978) to the current belief that moral judgements are primarily the result of our natural, intuitive responses to violations of basic moral foundations (e.g., Haidt, 2001). Furthermore, I will explore the more recent evidence to support the role of moral emotions in forming moral judgements alongside evidence to suggest that information, knowledge and reasoning can, nonetheless, mediate the influence of judgements made on the basis of emotional responses. The aim of this thesis is to provide a comprehensive and contextual approach to the process of moral judgement that includes an examination of when (e.g., pre-judgement, post-judgement, and during behaviour) and how behavioural and contextual information is useful, important and/or desirable.

1.1 The Shift from Reasoning to Intuition

During the cognitive revolution of the late twentieth century, evidence was presented for a cognitive-developmental theory of moral judgement (Kohlberg, 1971). Adults and

children were presented with moral dilemmas, such as whether it was permissible for someone to steal in order to save another life. Based on how participants reasoned through these dilemmas and made their judgements of rightness/wrongness, Kohlberg identified six stages through which people move to make moral judgements, corresponding largely to Piaget's (1965) stages of cognitive development. Young children started making moral judgements in a stage that involved making egoistic moral judgements on the basis of the good or bad outcome for the actor. As children's perspective-taking ability improves, children move into stages of moral development that involve considering other nuanced features of the dilemma (e.g., motivation and mitigation for the actions). Kohlberg's theory stated that the mechanisms behind moral judgements were based in cognition, language and conscious deliberation (Kohlberg, 1973, 1976).

Further to Kohlberg's cognitive-developmental theory of moral judgement is additional evidence for the influence of social information and understanding on how we resolve moral dilemmas. A social interactionist model, presented by Nucci and Turiel (1978), explained the development of moral judgement as deriving from children's peer-to-peer interactions that involve negotiating and reflecting on harmful consequences of one's own and others' behaviours. The social interactionist model focused on people's consideration of the outcomes of actions when judging whether something is a moral transgression. For instance, when harm results from an action, both the actions and outcome are part of the moral judgement, and lead to principled and universal rules that aim to prevent harm from reoccurring. Rules that are negotiated on the basis of social norms and conventions, however, may not involve harm resulting from an action. The social interactionist model, like Kohlberg's theory, is grounded in our conscious deliberation on rules, norms and consequences that develop alongside cognition, theory of mind, reasoning, and language.

Contrary to the cognitive-developmental accounts are intuition- and emotion-based accounts. Although much of the cognitive-developmental research (e.g., Turiel, Hildebrandt, & Wainryb, 1991) has shown that people's reasoning about harm, rights and justice inform and ultimately determine their moral position on a social issue, those who argue for the causal role of intuition in moral judgements believe that rightness and wrongness can, in some instances, be determined without solely relying on the consideration of harm, rights, and justice. If wrongness can be determined in the absence of reasoning about harmful or unfair consequences (e.g., when considering the wrongness of a violation of purity or a sexual taboo), then there may be something else underlying our sense of wrongness.

Haidt, Koller, and Dias (1993) conducted a study whereby participants were presented with victimless actions that were likely to offend or evoke a negative affective response in both Western (in this case, American) and non-Western (in this case, Brazilian) cultures. The victimless scenarios involved a woman who cuts the American or Brazilian flag into pieces and uses the pieces to clean her kitchen floor, a son who promises to visit his dying mother's grave often but doesn't keep the promise, a family who eat their dead dog, a brother and sister who kiss, and a man who has sex with a dead chicken before cooking and eating it. American and Brazilian children (ages 10-12) and adults, from high- and low-socioeconomic (SES) backgrounds, were presented with moral violations that depicted either actions that led to harmful outcomes or actions that were offensive but did not lead to harmful outcomes. The results demonstrated that Western participants, compared to non-Western participants, found the victimless violations to be less disgusting, disrespectful, and immoral (although this effect was somewhat weaker in adults than children). High-SES adult participants, compared to low-SES adult participants, similarly found the violations less disgusting, disrespectful, and immoral; interestingly, for high-SES participants, perceived harm predicted perceived

wrongness, whereas for low-SES participants, perceived offensiveness predicted perceived wrongness. This study provides some preliminary evidence in contrast to the emphasis originally placed on reasoning to suggest that some groups of people (in this case, non-Westerners and low-SES adults) may consider the wrongness of moral violations on the basis of a feeling or something more visceral than just a reasoned calculation of resulting harm.

Further evidence for reasoning without consideration of harm involved in determining moral wrongness comes from research that examined reactions to sexual taboo acts that did not involve any resulting harm (Haidt & Hersh, 2001). American participants who identified as either liberals or conservatives judged scenarios depicting acts of homosexuality, incest and masturbation where there was no resulting harm. Participants judged these scenarios on the basis of how wrong the action was, why the action was/was not wrong, whether or not anyone was harmed, whether the actor should be stopped, how the action made the participant feel, whether the participant experienced any feelings of contamination from the action, and whether the act would be considered universally wrong across cultures and norms. The main findings of the study were that liberals were driven mainly by a concern for harm, whereas conservatives were more concerned with violations of community and divinity¹. More importantly, however, the best predictor of participants' moral judgements for the victimless and harmless actions was their emotional response to the scenario—evidence against a reasoned basis for moral judgement. Additionally, participants often struggled to justify their judgements of wrongness, engaging instead in a circular and shallow reasoning process referred to as moral dumbfounding (Haidt, Bjorklund, & Murphy, 2000). These studies in

¹ The ethics of community and divinity are two of the three ethics of the CAD triad hypothesis (Rozin, Lowery, Imada, & Haidt, 1999) where the ethic that is violated in a given moral situation determines the moral standards used and evaluations made. Community refers to a violation of communal and social duties and respect. Divinity refers to a violation of something sacred or pure.

combination suggest that moralising and judging wrongness can occur in some instances without reasoning about resulting harm from an action, and can be driven by affect and non-conscious, inaccessible reactions to a violation other than harm.

In response to research that suggested that moral judgements could and would be made in the absence of reasoning about harmful outcomes, the social intuitionist model was proposed (Haidt, 2001). The social intuitionist account (see Figure 1) puts intuition at the forefront of the moral judgement causal chain, and reasoning second. According to Haidt's model, people rely on intuitive reactions, gut feelings and affective responses to make intuitive judgements. Haidt does not see intuition as separate from cognition, but rather a part of cognition that is automatic where the reasons underlying the immediate, snap intuitions are less accessible to the person. Reason, according to this model, serves the intuitions typically for the purpose of allowing people to persuade others that their judgement is appropriate and justifiable; that is, reasoning is post-hoc and does not necessarily modify the initial intuition.

Reasoning is not only for the social purpose of persuading and convincing others, however; reasoning can also occur privately. Haidt acknowledges that some people in some circumstances are able to override their intuitions in favour of reasoning and logic. Although Haidt acknowledges this possibility, there is little research to date which explains when people override intuitions and why. According to Haidt, however, these circumstances are rare, and may in some cases just reflect the fact that a person is able to express their reasoned judgement, but may still hold onto the intuition on some other, deeper level. Additionally, people are able to reason in a way that involves overriding one intuition with a different intuition that leaves the person with either competing intuitions or a re-negotiated intuition. Haidt suggests that people do so in a way similar to what the cognitive-developmentalists suggest in that people consider others' perspectives. Whereas cognitive-developmentalists

such as Kohlberg (1969) would consider this a development in the reasoning process for moral judgement, Haidt (2001) suggests that perspective-taking and empathising with others allows us to hold new intuitions. The social intuitionist account (Haidt, 2001) rejects the causal role of deliberative reasoning in moral judgement. The evidence to support the lack of deliberative reasoning comes from people's ability to make moral judgements when consequential harm is absent, the predictive relationship between emotional reactions and moral judgements, the presence of moral dumbfounding when justifying judgements, and the social motivations for justifying judgements.

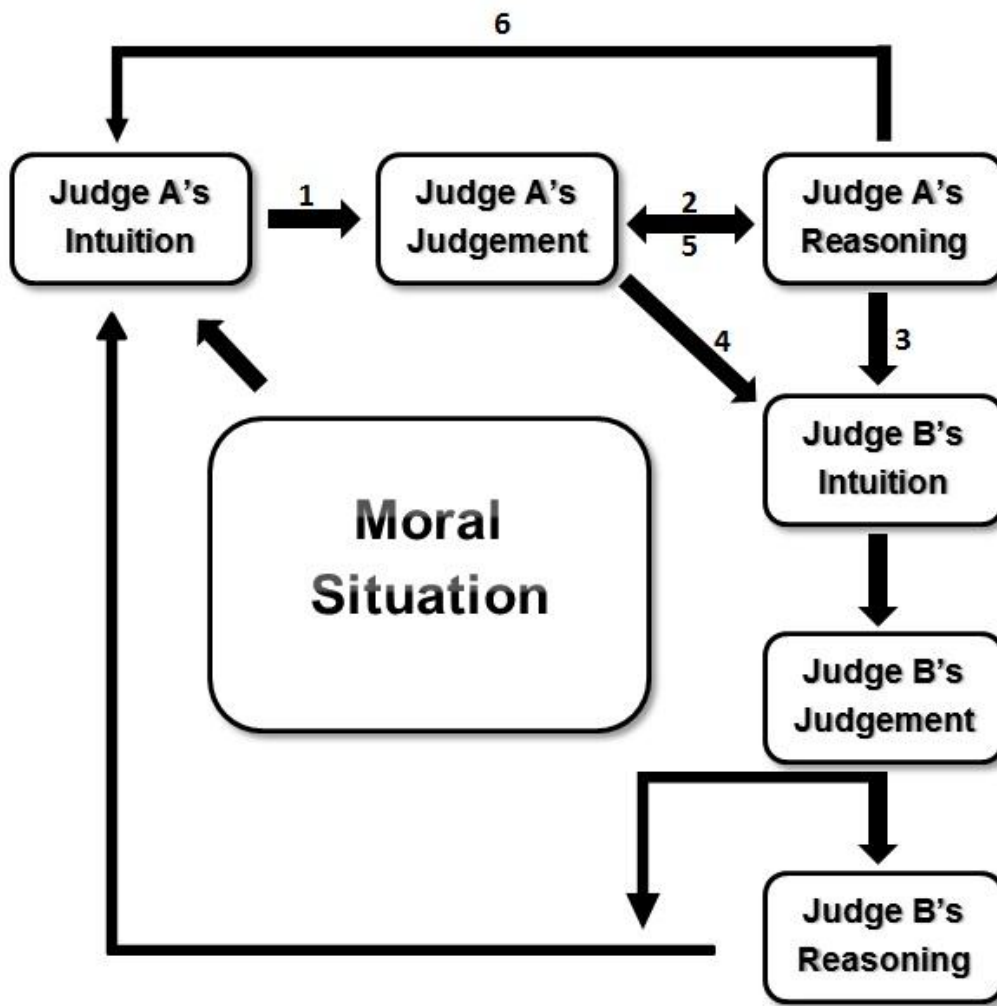


Figure 1: The social intuitionist model (Haidt, 2001). For Judge A’s moral judgement process, the numbered links represent the following process: Link 1 is the intuitive judgement link whereby Judge A’s intuition informs their judgement. Link 2 is the post hoc reasoning link whereby reasoning occurs after the judgement has been made. Link 3 is the reasoned persuasion link whereby one person’s post hoc reasoning influences another person’s intuition. Link 5 is the reasoned judgement link whereby a person’s reasoning about a moral situation leads to their judgement. Link 6 is the private reflection link whereby a person’s reasoning influences their intuition that then leads on to their judgement.

Further research has examined the role of emotions and intuitions in moral judgements. Wheatley and Haidt (2005) conducted two studies that showed the prevalence of intuition and affective responses in judgements of wrongness. Participants who were classed as being highly hypnotisable were hypnotised to feel flashes of disgust in response to neutral words such as “often” and “take”. The neutral words were then included in scenarios that depicted moral violations. The participants who had been hypnotised to feel flashes of disgust in response to the words that were included in the moral scenarios made more severe moral judgements, suggesting that the emotional experience of disgust was at least in part accounting for the judgements of wrongness. Additionally, some participants continued to condemn actions that were neutral and non-moral because of the inclusion of the words which had been associated with disgusting feelings during an earlier hypnotic state. Those participants also engaged in post-hoc reasoning that supported their emotion-laden intuition by manufacturing negative characteristics about the actor in the neutral scenario. The findings from these studies suggest that emotional reactions and intuitions can lead to judgements of wrongness.

Further evidence for the role of intuition and emotion as a driving-force of moral judgement comes from research demonstrating that exposure to disgust in several forms can lead to more severe moral judgements (Schnall, Haidt, Clore, & Jordan, 2008). Across several studies, participants were exposed to either a disgusting smell or a disgusting environment (Experiments 1 and 2), were asked to remember their own disgusting experience (Experiment 3), or were exposed to a video designed to elicit disgust (Experiment 4). After being primed with the experience of disgust, participants judged the wrongness of several moral violations depicted in vignettes. Participants primed with feelings of disgust made more severe moral judgements in response to the moral violations than did control participants who were not

exposed to disgusting stimuli or who were exposed to stimuli intended to elicit feelings of sadness. The findings from these studies further suggest the role of unconscious, non-deliberative intuition and affect in judging moral wrongness.

In addition to evidence from people's verbal moral judgements, evidence from neuronal activity further emphasises the role of emotion and intuition in moral judgement. Greene, Sommerville, Nystrom, Darley and Cohen (2001) gave participants moral dilemmas where there was personal involvement that led to the harmful action, and impersonal involvement that led to a harmful action; they also read a non-moral control dilemma. For instance, Greene et al. (2001) used the classic trolley problem as one such dilemma. The impersonal version of this dilemma involves pulling a switch that will divert a runaway trolley that is heading towards five people that would be killed by the trolley, onto a different track where the trolley would head towards only one person that would be killed. The dilemma is classed as impersonal because it does not involve any personal connection to the harm or the person being killed. The personal version of this dilemma involves physically pushing a man to his death in order to stop the trolley from hitting and killing five other people. Participants' neuronal activity and response times demonstrated that personal versions of moral dilemmas such as the trolley problem activated regions of the brain associated with emotion, such as the medial prefrontal cortex, posterior cingulate/precuneus, and superior temporal sulcus/temporo-parietal junction, and are judged to be more morally wrong and faster than impersonal and non-moral dilemmas. This study further demonstrates that emotional snap judgements are involved in at least some types of moral judgements.

With evidence to support the role of intuition and emotion in moral judgement, empirical and theoretical research has aimed to classify and distinguish moral emotions from other non-moral emotions. The aim of understanding moral emotions is to better understand

what elicits a moral judgement. If it is the case that emotion is a large part of the underlying causal process for moral judgement, as the social intuitionist model would suggest (Haidt, 2001), then understanding what moral emotions are and how moral emotions are elicited is important for understanding what modifies moral judgements and behaviours.

1.2 The Moral Emotions

1.2.1 Characteristics and Groupings of Moral Emotions. Haidt has identified two features that distinguish moral emotions from other emotions (Haidt, 2003). While basic emotions such as anger or sadness often involve a direct relationship between the person experiencing the emotion and the event that led to the emotion, moral emotions can involve disinterested elicitors. This means that people can respond to a moral violation or good deed that does not impact them directly with an emotional reaction. For instance, a person may experience feelings of disgust in response to reading about a brother and sister who have a sexual relationship. The story does not necessarily have any personal impact on the person exposed to it, yet as a removed third party, the person exposed to the story can still respond to the violation by experiencing a moral emotion.

The second feature that Haidt argues classifies an emotion as a moral emotion is the behaviour that the emotion elicits. Certain moral emotions have been shown to elicit behavioural responses, such as anger leading to revenge tendencies (Shweder, 1994). Moral emotions, however, may be distinct from other emotions in that they usually encourage social behaviour, whether that behaviour is for punishing transgressors or rewarding moral excellence for the betterment and functionality of society. One example of this comes from research investigating the effects of exposure to morally elevating stimuli on women's interactions with their infants (Silvers & Haidt, 2008). Silvers and Haidt found that exposing women to a video clip that depicted altruistic behaviour led to feelings of moral elevation and

being uplifted, relative to a video clip designed to elicit feelings of mirth. Women who observed altruistic behaviour and experienced moral elevation were subsequently more likely to nurse and hug their infants, compared to women in the control condition. Although not all—and maybe even few—moral emotions are positive, there may be a relationship between prosocial action tendencies that distinguish moral emotions from other emotions.

Haidt (2003) has grouped moral emotions into categories based on their shared properties. The other-condemning emotion family includes the emotions contempt, anger and disgust. The self-conscious emotions include shame, embarrassment and guilt. The other-suffering emotion family includes sympathy and empathy. Finally, the other-praising family includes awe and elevation. Haidt theorises about what moral situations evoke each of the above moral emotions and what resulting behavioural responses follow, and there have been several recent empirical studies that have further investigated moral emotions (e.g., Chapman, Kim, Susskind, & Anderson, 2009; Gutierrez & Giner-Sorolla, 2007; Russell & Giner-Sorolla, 2011, 2013). The overriding argument when researching moral emotions is that moral emotions are largely the driving force behind moral judgement. While this may be the case, additional research has demonstrated conditions under which it is possible to downplay the role of emotion or intuition, as well as factors that can influence a moral judgement when they are considered.

1.2.2 When Information and Reasoning Can Override Emotion. Some research has demonstrated that information about the actor's belief about causing harm and desire to cause harm can affect how wrong we judge an action to be and how much blame we attribute to the actor (Cushman, 2006, 2008). Knowledge that harm resulted from a bystander not intervening (passive harm) leads to the action being judged less morally wrong than equivalent active and direct harm (Spranca et al., 1991). Additionally, if the action leading to

harm involves physical contact with the victim, as seen in the footbridge variation of the trolley problem (Foot, 1978), then the harm is judged to be more morally wrong than equivalent harm that does not involve physical contact (Greene et al., 2001).

Further research on responses to the classic trolley problem by Valdesolo and DeSteno (2006) demonstrated that participants who are exposed to a brief comedy video clip designed to elicit positive affect were more likely to make a reasoned decision in the personal and emotion driven version of the trolley problem. That is, participants judging the footbridge version of the trolley problem—where participants decide that it is morally impermissible to push a large man off of a footbridge in front of the trolley to his death to stop a runaway trolley from killing five people—were more likely to override the typical emotional response for a more reasoned response when they had previously been exposed to video clip that elicited positive affect. Participants who were given the opportunity to react emotionally (prior to judging the trolley problem), even in response to different stimulus (video clip), were able to put their initial emotional response aside in favour of reasoning. The emotional response typically leads the majority of participants to decide that pushing the large man is wrong, whereas the reasoned utilitarian response that Valdesolo and DeSteno (2006) evoked in their study makes pushing the large man to save five other people more acceptable. This study suggests that when there is an outlet for emotion that comes from another source, there may be room for more reasoned appraisals of moral dilemmas.

A more recent series of studies supports the notion that moral judgements are less influenced by emotion than the social intuitionist account would suggest (Feinberg, Willer, Antonenko, & John, 2012). The first study measured participants' tendency to reappraise their emotions. Reappraisal is a technique for regulating emotion that involves constructing or portraying an emotionally charged situation in a way that lessens the emotional reaction to the

situation (Gross, 1998, 2002, 2007). Participants who were more likely to reappraise their emotional responses to a situation were able to judge two non-victim scenarios that have previously demonstrated the role of emotion in judging wrongness (incest, sex with a dead chicken; Haidt et al., 1993) as less morally wrong. This study suggests that participants who re-evaluate and reappraise their emotional reactions can override their emotional reactions in a moral context, thereby affecting their perception of wrongness. The second study confirmed the influence of reappraising emotions in overriding emotional responses to judging moral wrongness. Participants who reappraised their emotional responses to the moral scenarios while judging the moral scenarios took longer to make their judgement and also found the scenarios to be less morally wrong. In the final study, participants watched a video clip designed to evoke sadness; participants were either asked to reappraise their feelings of sadness or to just watch the video. Participants then read and judged three moral dilemmas. Those participants who had engaged in the process of reappraising their sadness in response to the video clip reported milder emotional reactions to the moral dilemmas and were more likely to change their initial moral judgement and judge the moral dilemmas as less wrong overall than participants who had not engaged in emotion reappraisal. These studies demonstrate that there are conditions under which the role of emotion and intuition in moral judgement can be dampened to allow for more reasoned responses to moral scenarios.

1.2.3 Dual-Process Theory of Moral Judgement. A dual-process theory of moral judgement (Greene, 2007, 2009; Greene, Morelli, Lowenberg, Nystrom, & Cohen, 2008; Greene, Nystrom, Engell, Darley, & Cohen, 2004; Greene, Sommerville, Nystrom, Darley, & Cohen, 2001) provides an alternative explanation to moral judgement processing that goes beyond using either emotion and intuition or reasoning and deliberation. This account suggests that moral judgements based on deontology—a concern with absolute and universal

principles, duties and rights, such as the belief that killing another person is morally wrong—rely on more automatic processing and are based on emotional reactions and gut-feelings (Greene, 2009). Moral judgements that are utilitarian—involving weighing costs and benefits, and attempting to maximise utility and minimise suffering—tend to be reasoned and involve controlled processing. Although the majority of the research supporting a dual-process account for moral judgement pertains to moral dilemmas such as the trolley dilemma and footbridge dilemma where one decides to kill one person to save five others by virtue of pulling a switch or pushing a person to their death (Greene et al., 2001; Mikhail, 2000; Petrinovich, O'Neill, & Jorgensen, 1993), there may be different factors, such as the type of violation, in other moral situations that lead to controlled processing versus emotional reaction and vice versa. This thesis aims to investigate moral judgements made in response to moral situations as opposed to moral dilemmas, and to examine under which contexts, for which violations and at which point behavioural and contextual information is considered, desired, useful and influential. I address the dual-process theory of moral judgement further in the final chapter, and how the findings presented in this thesis can be understood within this framework.

2 Core Hypotheses

Taken together, research from the psychological study of moral judgements has demonstrated that depending on the context, reasoning and emotion can both moderate moral judgements. The current moral judgement literature tends to focus largely on the role of intuition and emotion in making non-deliberative snap judgements. Although the social intuitionist model of moral judgements (Haidt, 2001) does not entirely negate the role of processing information or deliberative reasoning, the majority of recent research has directly examined the role of emotions and intuition over the influence of behavioural and contextual

information, or the interaction between reason, emotion and the processing of behavioural and contextual information.

The present research aimed to investigate the role of deliberative processing of behavioural and contextual information that has the potential to modify moral judgements and elicited behaviour. The effect of behavioural and contextual information is examined at three distinct stages: pre-judgement, post-judgement, and during behaviour. The pre-judgement stage refers to how behavioural and contextual information presented prior to making a moral judgement is considered and affects the subsequent judgement. The post-judgement stage refers to how behavioural and contextual information is selected and engaged with when presented after the initial moral judgement. Behaviour refers to how consideration of behavioural and contextual information can influence behaviour that follows from witnessing a moral act. Overall, I predicted that moral judgements would be influenced by the processing of behavioural and contextual information differently across the three stages and according to the following mediators: age-related changes in reasoning ability, violation and information type, and judgement valence. The empirical research is divided by chapter across the three stages of judgement with various mediators: Chapter 2 examines the role of behavioural and contextual information; specifically, different means that eventually lead to something negative occurring and the consideration of hypothetical alternative outcomes pre-judgement. Chapter 3 examines the desire and engagement with post-judgement mitigating or inflammatory information. Both Chapters 2 and 3 examine judgements of wrongness, condemnation or blame. Chapter 4 examines the role of outcome information in behaviour that follows from experiencing a positive moral emotion (i.e., elevation), and is the only chapter to examine the role of information in positive moral acts and emotion.

I chose moral stimuli in each experiment that have been demonstrated to evoke specific moral emotions: judgements of violations of purity that evoke feelings of disgust (e.g., Gutierrez & Giner-Sorolla, 2007), judgements of harm that evoke feelings of anger and disgust (e.g., Russell & Giner-Sorolla, 2011a), judgements of morally uplifting acts of virtue that evoke feelings of elevation (Schnall, Roper, & Fessler, 2010) and judgements of blame and regret that result in part from the consideration of counterfactuals and consideration of how a story character feels about a negative outcome (e.g., Weisberg & Beck, 2010). The reason behind choosing moral stimuli that evoke emotion is to directly address the social intuitionist model of moral judgement. I aimed to examine how the role of emotion and intuition in moral judgement can be moderated by consideration and processing of behavioural and contextual information at different stages of judgement. The social intuitionist model does not provide an understanding of how emotion and intuition might be more or less accountable for moral judgements depending on the stage of judgement (e.g., prior to a judgement or following a judgement). Additionally, the social intuitionist model does not provide an understanding of how emotion and intuition might be more or less accountable for a judgement depending on age-related changes in reasoning ability, differences in violation type and differences in information types.

Chapter 2 addresses the role of information regarding the means through which a negative outcome occurs prior to making a negative moral judgement, specifically a judgement of blame. There is a significant body of evidence to suggest that pre-judgement information will yield different moral judgements. For instance, knowledge about an actor's desire and intention to cause harm will impact the judgement (e.g., no desire to cause harm is less bad than a desire to cause harm; Cushman, 2008). Furthermore, knowledge that an actor acted to cause harm rather than was a bystander to harm will impact the amount of

wrongness, blame and punishment attributed to that actor (Cushman, Young, & Hauser, 2006; Spranca, Minsk, & Baron, 1991). Although this evidence does suggest that processing of information pre-judgement is influential to moral judgements, this evidence is often discussed in terms of operating on the basis of an unconscious, non-deliberative set of intuitive moral principles. In Chapter 2, I examine the role of pre-judgement information with specific regard to reasoning by testing participants of different ages who have different levels of cognitive sophistication when it comes to processing and reasoning about information. This will enable an understanding of whether pre-judgement information is considered similarly at all ages (evidence for non-conscious, non-deliberative operation of intuitive moral principles), or whether the use of pre-judgement information depends on age-related changes in reasoning ability.

Chapter 3 addresses participants' engagement with behavioural and contextual information that has the potential to mitigate versus inflame a negatively valenced judgement of wrongness post-judgement. There is a lack of research in the moral judgement literature on post-judgement processing—how people consider information that has the potential to change their initial judgement. The social intuitionist model argues that post-judgement reasoning is typically done in an effort to convince others to adopt one's moral judgement (Haidt, 2001). I predict that moral judgements can also guide post-judgement information selection on the basis of epistemic and personal motivations relating to desires to approach and avoid moral situations and emotions, and desires to enhance and boost epistemic certainty. I predict that these motivations are for the individual judge, not only for convincing others as described by the social intuitionist account (Haidt, 2001) (i.e., for the purpose of convincing others to share one's moral judgement).

Chapter 4 addresses how elicited moral behaviour following exposure to a positive moral act is affected by information. While the research in this thesis does not provide a direct comparison between positive versus negative moral judgements, I believe it is important to begin to investigate how information might affect positive moral judgements. Moral judgement research has largely focused on negative moral judgements (e.g., blame, condemnation, wrongness and punishment) and few moral emotions have been identified (Haidt, 2003). This is an imbalance that I seek to redress, and begin to do so in Chapter 4, as there may be differences in the role of reasoning, emotion and information processing depending on whether moral judgements are in response to positive moral acts or negative moral acts.

Some moral emotions have been demonstrated to lead to behavioural outcomes, such as cooperation (de Hooge, Zeelenberg, & Breugelmans, 2007). The experience of moral elevation has also been demonstrated to lead to prosocial behavioural outcomes (e.g., Schnall et al., 2010). Chapter 4 discusses how behavioural outcomes resulting from the experience of moral elevation may be moderated by information pertaining to different and specific facets of the moral situation (e.g., the outcome of the helpful act for the beneficiary). The potential influence of information on behaviour has interesting implications for the social intuitionist model. The behavioural outcomes of moral elevation are typically discussed as resulting from the emotional experience of elevation; this would support the social intuitionist model of moral judgement in resting the importance of moral judgement with emotion. If, however, behavioural and contextual information about the moral situation can mediate the behavioural outcome, then the experience of a moral emotion may not be solely responsible for the behavioural outcome, suggesting a role of information processing and reasoning. Chapter 4

examines this hypothesis, and hypotheses regarding the role of information in determining and responding to moral virtue are discussed further in Chapter 4.

By the conclusion of this thesis, I aim to have added to the moral judgement literature through an examination of how moral judgements and behaviour are influenced by the availability and processing of behavioural and contextual information. Additionally, I aim to have an understanding of the potential influence of behavioural and contextual information at three specific stages: prior to making a moral judgement, following a moral judgement and on elicited moral behaviour. This thesis does not address whether behavioural and contextual information is processed deliberately rather than intuitively or vice versa, but does demonstrate the potential relevance and mediation of something more than affect and unconscious knee-jerk reactions to moral situations.

CHAPTER 2

**THE INFLUENCE OF BEHAVIOURAL AND CONTEXTUAL
INFORMATION PRE-JUDGEMENT: EVIDENCE FROM REGRET
AND BLAME FOR DEVELOPMENTAL CHANGES IN REASONING
AND PROCESSING OF INFORMATION**

Some moral judgements are believed to originate from intuitive, innate principles, although other research has demonstrated that moral judgements can be influenced by more deliberative reasoning and processing of information. In this chapter, I turn to developmental literature and methodology to determine whether certain moral judgements can be influenced by age-related changes in reasoning, and whether attention to behavioural and contextual information (in this case, information regarding active or passive and atypical or typical means leading to a bad outcome) provided pre-judgement influences the judgement and subsequent reasoning. In two studies, I examined the developmental change in the influence of behavioural and contextual information on attributions of blame and judgements of others' experience of regret. I tested 6-year-olds, 8-year-olds and adults, and found that 6-year-olds did not reason counterfactually about behavioural and contextual information provided pre-judgement when considering others' regret or blameworthiness. Eight-year-olds, however, reasoned counterfactually and exhibited evidence of omission bias when making judgements about both others' regret and blameworthiness. Adults showed a similar pattern to that of 8-year-olds when judging regret but not when making judgements of blameworthiness (recognizing that all bad outcomes were accidental). The findings from Experiment 2 demonstrated that both 6-year-olds and 8-year-olds attributed greater regret and more blame

*to characters when the negative outcome affected a greater number of characters. Our findings suggest that blame attributions, like judgements of regret, are influenced by developmental change in reasoning and processing of behavioural and contextual information provided pre-judgement.*²

1 Introduction

In this thesis, I examine the effects of information and reasoning on how we judge and process certain components of morality at different stages: pre-judgement, post-judgement and during behaviour. The role of behavioural and contextual information pre-judgement is important to consider given the debate regarding the role of reasoning and intuition in moral judgements. As Chapter 1 detailed, moral psychology was largely dominated by cognitive developmentalists during the 20th century who argued that moral judgements were made on the basis of reasoning about harm (Kohlberg, 1969; Piaget, 1932/1965; Turiel, 1983). According to this perspective, moral judgement changed developmentally in stages alongside other forms of reasoning. With evidence against the need for harm to be present when judging wrongness, the inability for people to properly justify their judgements and the overall relationship between emotional reactions and judgements (Haidt, 2001; Haidt, Bjorklund, & Murphy, 2000; Haidt, Koller, & Dias, 1993; Prinz, 2007), more recent research has turned to less reasoning-based explanations for how we judge morality. The social intuitionist account for moral judgement (Haidt, 2001) suggests that very little influences moral judgements beyond emotional, intuitive, gut-reactions (Haidt, 2001). Although the social intuitionist

² The research in this chapter is in the stages of revision for a submission to *Cognitive Development* as the following manuscript: Powell, N. L., Guttentag, R. E., Quinn, K. A., & Beck, S. R. (in revision). *The role of reasoning in moral judgments: Evidence from regret and blame.*

account posits that intuitions are not separate from cognition, the model suggests that intuitions that form our moral judgements are neither deliberative nor immediately accessible.

This chapter includes two experiments that examined age-related changes in how behavioural and contextual information and reasoning about behaviours that lead to negative outcomes influence how we blame others. If moral judgements are largely grounded in emotional and intuitive reactions, then it is unlikely that the development of the ability to reason about more sophisticated and nuanced pieces of information will have a strong influence on our ability to make moral judgements. Additionally, given that much of the research on moral intuition pertains to how we judge moral wrongness, these two experiments examined the role of intuition and reason in judgements of a person's blameworthiness.

Moral judgements about wrongness, blame, and punishment may involve some reasoning when, for instance, constructing an argument post-hoc for why something is wrong (Haidt, 2001). Many authors argue that moral judgements originate from intuitive, gut-feelings (e.g., Gutierrez & Giner-Sorolla, 2007; Haidt, 2001; Russell & Giner-Sorolla, 2011). If this is the case, then questions emerge about the role of cognitive mechanisms in moral judgement and the developmental origins of moral understanding.

In this chapter, I look at an arguably intuitive moral principle: the omission bias (Baron & Ritov, 2004; Spranca, Minsk, & Baron, 2005). The omission bias refers to the tendency to judge harm caused by action as worse than that caused by inaction (Baron & Ritov, 2004; Spranca, Minsk, & Baron, 1991). For example, it is worse to steal than to refrain from donating to help the poor (Spranca et al., 1991). I compare the potential influence of reasoning when attributing blame—one facet of morality (Cushman, 2008)—on judgements of others' experience of regret, which are known to be influenced by reasoning (e.g., Burns, Riggs, & Beck, 2012). The apparent inflexibility of moral judgements (e.g., Russell & Giner-

Sorolla, 2011a) and the apparent lack of cognitive influence or insight for these judgements (Haidt, Bjorklund, & Murphy, 2000) have been used to support a strong position that moral judgements reflect the activity of an inherent moral faculty that does not depend upon conscious deliberative reflection, but instead involves a set of unconscious, intuitive processes based upon a set of fundamental principles, one of which is the omission bias (Cushman, Young, & Hauser, 2006).

With regard to development, if there are fundamental moral principles that influence moral judgements through an intuitive non-conscious mechanism, then we might expect that the kinds of age-related changes in conscious moral reasoning that have been identified by Kohlberg and others (Kohlberg 1986; Murphy & Gilligan, 1980) should not affect the degree to which children of different ages adhere to moral principles. Indeed, evidence that supports the idea of basic intuitive moral principles being available from a very young age comes from studies showing the early emergence of sensitivity to aspects of moral judgement such as fairness and intention. For instance, Sloane, Baillargeon, and Premack (2012) have found that 19-month-old infants look longer at people and events that violate their expectations of fairness—for example, when an experimenter behaves inequitably as opposed to equitably to another person, or when an experimenter rewards two individuals who failed to equally distribute the workload needed to complete a task compared to when the experimenter rewards two individuals who worked together equally to complete a task. This evidence is used to support the notion that infants possess some basic moral principles such as fairness—that is, that some moral principles are intuitive and natural and do not require complex reasoning. The difficulty with this interpretation is that an understanding of the concept of fairness can be sophisticated (e.g., holding mental representations of action and intention while evaluating cost and benefit), and is therefore unlikely to be present other than in a crude

form in infants because infants are incapable of sophisticated reasoning. This raises questions about how meaningful a crude understanding of fairness is, and if a crude understanding is comparable to a later-developing understanding of fairness.

In addition, the “snap” judgements that adults make in response to moral scenarios involving differentiations of action- versus inaction-produced harm (Baron & Ritov, 2004; Cushman, Young, & Hauser, 2006; Spranca, Minsk, & Baron, 1991)), in combination with their inability to access justifications for their responses (Haidt, Bjorklund, & Murphy, 2000), suggests that moral judgements rely on intuitive processes (Cushman et al., 2006; Haidt, 2001). Evidence suggests that people find it difficult to access and report justifications for strong moral beliefs about wrongness, especially when the situation does not produce any evident harm—a further case for the operation of intuitive moral principles (Haidt et al., 2000). Furthermore, research has demonstrated that the operation of some basic moral principles, such as believing action is worse than inaction, is strong enough to outweigh a rational decision made on the basis of calculating harm. Spranca et al. (1991) found that people are more accepting of inaction- than action-produced harm even if the outcome of inaction-produced harm involves a greater degree of harm than action-produced harm.

The alternative to the intuitionist position is that making moral judgements relies on the ability to reason about harm. In this case, we should expect to see age-related changes in how children of different ages adhere to moral principles when making judgements of harm. Preliminary research suggests that some of the basic harm principles respond to age-related changes in cognitive development, such as the understanding of cost and benefit (Powell, Derbyshire, & Guttentag, 2012). Two experiments examined biases in children’s (5/6-year-olds and 7/8-year-olds) and adults’ moral judgements. Participants at all ages judged that it was worse to produce harm only when harm occurred: (1) through action rather than inaction

(Omission Bias), (2) when physical contact with the victim was involved (Physical Contact Principle), and (3) when the harm was produced as a direct means to an end rather than as an unintended but foreseeable side effect of the action (Intention Principle). The youngest participants, however, did not consider benefit when making judgements about situations in which harm to one individual resulted in benefit to five individuals. Older participants showed some preference for benefit resulting from action (commission) as opposed to inaction (omission). The findings support the theory that moral judgements result, in part, from the operation of an inherent, intuitive moral faculty compared to the theory that moral judgements require development of necessary cognitive abilities (Hauser, 2006; Hauser, Young, & Cushman, 2008; Kohlberg, 1969; Piaget, 1932/1965).

Potentially relevant to our understanding of the process of reasoning about negative moral judgements is recent research exploring the role of reasoning in children's reactions to negative events and their causal links, such as judgements about another person's experience of regret (Amsel & Smalley, 2000; see Weisberg & Beck, 2010, 2012, for discussion of the difference between understanding and experiencing regret). Although regret is an affective state, it is dependent on being able to reason about alternative outcomes—that is, to reason about what might have been. For example, imagine I go out to a party the night before a big exam, drink too much and fail the exam. The failure itself will lead to a sense of sadness. However, the experience of regret is prompted not just by the negative outcome that actually occurred but also by counterfactual reasoning, that is, by the generation of an alternative possible world in which I did not go out (or perhaps stayed sober) and the comparison of that alternative outcome with reality (see Burns, Riggs, & Beck, 2012; Zeelenberg, van Dijk, Manstead, & van der Pligt, 1998).

Being able to reason about when others experience regret is relatively late-developing. Amsel and Smalley (2000) showed that 3- to 5-year-olds playing a card game could report how another player would feel if a different (better) card had been picked; that is, they created a counterfactual world. However, they did not attribute regret in cases in which the card not chosen was better than the card actually chosen. That is, they did not realise that a player who *could have* picked a better card would feel worse *about reality* than someone who did not have such a counterfactual. Beck and Crilly (2009) also showed that understanding regret lagged behind developments in children's counterfactual thinking and was relatively late-developing. Participants in this study were able to answer an open counterfactual question about whether the character in a story could have done something different more easily than answering a question about who felt worse (a measure of regret), suggesting that understanding regret requires not just the consideration of an alternative action, but comparing that alternative action with the actual action and the outcome.

In a related study, Guttentag and Ferrell (2004) assessed the emergence of children's understanding of the conditions under which others would feel regret. In their tasks, children read stories about two protagonists who undertook different actions but experienced similar negative outcomes. One of the factors included in the study was the difference between the effects of action versus inaction, that is, the omission bias. Adults judge negative outcomes that result from an action (commission) to be more wrong and deserving of punishment than negative outcomes that result from inaction by a bystander (omission). The bias towards less harsh judgements and punishment for omission is argued to be so strong that people will favour outcomes that result in a greater number being harmed as long as the harm occurs as a result of inaction by a bystander as opposed to active harm that results in fewer people being harmed. Additional research with adults has found that judgements of regret are higher for

those who acted and experienced a negative outcome compared to those who failed to act and experienced a similar negative outcome (Gilovich, Medvec, Husted, & Kahneman, 1998; Zeelenberg, van den Bos, van Dijk, & Pieters, 2002). Guttentag and Ferrell also examined the effects of the typicality of an action on children's regret judgements; previous research with adults has found that adults attribute greater regret to someone whose actions led to a negative outcome if the action was unusual or atypical for that individual than if it was a usual or typical course of action (Kahneman & Miller, 1986).

Guttentag and Ferrell (2004) found that children at the age of 7 years reasoned like adults about alternative outcomes; that is, they attributed more regret to the actor as opposed to the bystander and more regret to an individual engaging in an atypical than a typical action. Importantly, these findings demonstrate that the 7-year-olds were making judgements consistent with the omission bias and with the use of counterfactual reasoning by weighing possible hypothetical outcomes against actual outcomes—a cognitive reasoning ability that is not typically present in younger children (Amsel, Robbins, Tumarkin, Janit, Foulkes, & Smalley, 2003).

Given the current research on intuitive and non-deliberative moral judgements (Cushman, 2008; Haidt, 2001) moral intuitionists would see making blame judgements as a facet of intuitive moral judgement. It is evident that understanding and experiencing regret by definition must involve reasoning about counterfactual worlds, and therefore must involve developmental changes in reasoning ability. Intuitionists would not expect, however, to see the same developmental relationship as seen with judgements of regret for judgements of blame; blame judgements are a type of moral judgement (Cushman, 2008) and therefore exist independently of developmental changes in reasoning ability given that they are based on intuitive and non-deliberative reactions (Haidt, 2001). Therefore, intuitionists would most

likely expect an understanding of how to attribute blame to emerge earlier than an understanding of regret since moral judgements do not rely on the development of sophisticated reasoning. If, however, blame judgements in fact rely on reasoning rather than intuition, then we might expect to see judgements of blame and regret involving the same factors (e.g., action vs. inaction) to emerge together.

In the present studies, I sought to test these two accounts of the relationship between judgements of blame and regret. In two experiments, I explored children's judgements of regret and blame in response to situations that depict different behaviours that lead to the same negative outcomes, using Guttentag and Ferrell's (2004) procedure and comparing omission/commission to explore the omission bias principle. For the purpose of consistency with Guttentag and Ferrell, I also compared atypical and typical actions. Omission/commission stories involved target characters whose action leads to negative outcomes (commission) compared to non-target characters whose inaction leads to the same negative outcomes (omission). Atypical/typical stories involved target characters whose unusual behaviour leads to negative outcomes (atypical) compared to non-target characters whose normal behaviour leads to the same negative outcomes (typical). Age was used as a blocking factor for testing whether the omission bias occurs regardless of age-related changes in sophisticated reasoning ability.

I hypothesized that if blame judgements are based upon intuitions of harm rather on more conscious moral reasoning, then I would find evidence of omission bias (and perhaps sensitivity to action typicality) at a younger age than the age at which the omission/commission and typical/atypical factors affect children's judgements of regret. Alternatively, if the omission bias affects blame judgements through more reflective moral

reasoning, then judgements of blame and judgements of regret would be expected to exhibit a similar and linked pattern of age-related change.

2 Experiment 1

The aim of Experiment 1 was to determine whether blame judgements would respond to age-related changes in reasoning similar to regret, counter to the intuitionist account. I tested children whom I expected to understand when others felt regret (8-year-olds) and those whom I expected to lack this understanding due to immature counterfactual reasoning (6-year-olds). Both age groups of children and a sample of adults made blame judgements in response to a negative outcome that occurred for a character in a story as a result of the character acting (commission) or failing to act (omission). Children and adults also made judgements of regret for omission versus commission similar to Guttentag and Ferrell (2004). Additionally, children and adults made regret and blame judgements in response to atypical and typical actions. Age was used as a blocking factor for testing whether the omission/commission bias occurs regardless of reasoning ability.

I predicted that regret responses would replicate the findings from Guttentag and Ferrell (2004), such that 6-year-olds would fail to reason counterfactually about a character's experience of regret and would not choose the target character (the character who acted or engaged in an atypical action) as experiencing more regret than the non-target character (the character who failed to act or engaged in a typical action). Eight-year-olds and adults, however, were predicted to reason counterfactually and deduce that target characters would experience more regret than the non-target characters.

Blame judgements, however, may show a different pattern from regret if responses to the omission bias—an arguably early-emerging, natural and intuitive moral principle—are directly driving the judgements of blame, without the need for counterfactual reasoning (as in

regret). If, however, age-related changes emerge in blame judgements, similar to regret judgements, then this would argue against an account whereby judgements of blame are intuitive rather than reason-based.

When judging regret, atypical/typical judgements have been shown to respond similarly to changes in reasoning across ages as omission/commission judgements (Guttentag & Ferrell, 2004). Therefore, if judging blame is based on similar reasoning, then we would expect atypical/typical judgements to be similar to omission/commission judgements. However, if the omission bias is a special early-emerging moral principle, we might expect differences between omission/commission and atypical/typical judgements.

Additionally, I manipulated the number of characters affected by the negative outcome. Participants either responded to stories where the negative outcome affected only the character or affected the character and bystanders. Previous research with adults has demonstrated that the number of people affected by a negative outcome influences how wrong that action is perceived to be: harm to many is judged to be worse than harm to few (e.g., Spranca et al., 1995). I predicted that the intensity of judgements about regret or blame should be stronger when the negative outcome affected a greater number of people.

2.1 Method

2.1.1 Participants and design. The participants were 40 six-year-old children (26 girls; mean age = 6 years 9 months (6;1), range = 6;3–7;3), 40 eight-year-old children (20 girls; mean age = 8 years 9 months (8;9), range = 8;3–9;3), and 59 adults. The children were recruited from and tested at two schools in the United Kingdom. The adults were first-year students enrolled in an online introductory psychology course at a public university in the United States who participated for extra credit towards their final course grade. I did not collect further demographic data from the adult participants.

Participants took part in a 2 (Outcome: character, character and bystanders) x 2 (Action: atypical/typical, commission/omission) x 3 (Age Group: 6-year-olds, 8-year-olds, adults) mixed-model design where action was manipulated within participants. Age and outcome were between-participants factors.

2.1.2 Materials and procedure. Children were read four stories aloud (fully counterbalanced) and adult participants read the same four stories themselves (see Appendix A). For children, each story was accompanied by one 21cm x 30cm sheet depicting the sequence of events in each story. The experimenter pointed to the picture during each story. The pictures did not show emotional responses of any characters in the story. After each story, the experimenter recapped the main points (e.g., what the negative outcome was and how it occurred). Children were then asked to choose which character felt worse (regret) and why, and which character was more to blame and why.

There were four different sets of stories. In two, the negative outcome always affected only the acting character; in the other two, the outcome affected other people as well. Crossed with this, in half the stories the target character was always presented first and in half the non-target character was first. The presentation of the stories and questions was fully counterbalanced.

The stories were modelled after the stories used in Guttentag and Ferrell (2004). Each story involved two characters who did different things but encountered the same negative outcome (e.g., spilling soda pop or choosing unpleasant candies). As in Guttentag and Ferrell, two of the stories described an action that was atypical for one character (e.g., sitting in a different seat at the cinema) and typical for the other character (e.g., sitting in the same seat as always at the cinema). The other two stories described one character's commission (e.g., changing their original candy choice) and the other character's omission (e.g., keeping the

original candy choice). Stories either described negative consequences for only the acting character or negative consequences for both the acting character and bystanders. Once children had heard each story they were asked both a regret question and a blame question (counterbalanced between children). To use the cinema soda story as an example, the regret question was:

Do you think one boy will feel worse than the other about having decided to sit in the back row that day where they got covered by spilt sticky fizzy pop? David, who always sits in the back row and sat in the back row today where he got covered in spilt sticky fizzy pop? Or Charles, who always sits in the front row, but just today decided to sit in the back row where he got covered in spilt sticky fizzy pop? Or the two boys will feel the same. Why?

The blame question was:

Who do you think is more to blame for getting himself covered in spilt sticky fizzy pop? David, who always sits in the back row and sat in the back row today where he got covered in spilt sticky fizzy pop? Or Charles, who always sits in the front row, but just today decided to sit in the back row where he got covered in spilt sticky fizzy pop? Or are they equally to blame? Why?

Children had a choice of three possible options: target characters were those who acted (commission) in the commission/omission stories, and characters whose behaviour was atypical in the atypical/typical stories. Non-target characters were characters who failed to act (omission) in the commission/omission stories, and characters whose behaviour was typical in the atypical/typical stories. Some responses indicated that participants believed the target and non-target characters felt the same (regret) or were equally to blame (blame).

Adults accessed stories online through an introductory psychology course webpage hosted by BlackBoard[®]. For adults, the presentation of the target character was the only item counterbalanced across stories. The blame question preceded the regret question and separate samples of adults had access to the stories involving negative outcomes for one character versus many characters. Adults responded by typing the character choice and justification for both regret and blame questions into a text field. Adults were not given accompanying pictures with the stories.

2.2 Results

Participants were given a score of 1 for each story when the target character was chosen and a score of 0 for any other response (non-target or same) for regret and blame questions. Total scores for regret and blame were calculated for each participant to reflect the total number of target choices for commission/omission stories and for atypical/typical stories; that is, each participant had four separate scores out of 2. A preliminary multivariate analysis of variance (ANOVA) revealed that outcome (whether the negative outcome affected only the character or more people) did not yield any differences (lowest $p = .06$); therefore, responses were collapsed across outcome. The additional factor of outcome may have decreased the number of participants within each condition, and reduced the ability to detect some of the interactions; however, I collapsed across this factor in an effort to reduce this possibility. Mean commission/omission and atypical/typical scores for the three age groups are reported in Table 1. The distribution of responses (target, non-target and same) for both commission/omission and atypical/typical stories are reported in Table 2.

Table 1

Mean Number (and Standard Errors) of Target Character Responses for Commission/Omission Stories and Atypical/Typical Stories for Both Regret and Blame Questions as a Function of Age, Experiment 1

	Commission/Omission		Atypical/Typical	
	Regret	Blame	Regret	Blame
6-year-olds (n = 80)	.75 (.13)	.53 (.12)	.57 (.12)	.53 (.12)
8-year-olds (n = 80)	1.10 (.13)	1.13 (.12)	1.30 (.12)	1.00 (.12)
Adults (n = 118)	1.20 (.10)	.24 (.10)	1.54 (.10)	.49 (.10)

Note. Possible range = 0 to 2.

Table 2

Percentage of Non-Target, Target, and Same Character Responses for Atypical/Typical Stories and Commission/Omission Stories for Regret and Blame Question as a Function of Age, Experiment 1

		Commission/Omission			Atypical/Typical		
		Non-Target	Target	Same	Non-Target	Target	Same
Regret	6-year-olds (n = 80)	18	38	45	26	29	45
	8-year-olds (n = 80)	21	53	26	16	65	19
	Adults (n = 118)	1	60	34	0	77	21
Blame	6-year-olds (n = 80)	19	26	55	13	26	61
	8-year-olds (n = 80)	20	56	24	18	50	33
	Adults (n = 118)	0	12	85	0	25	74

2.2.1 Commission/omission stories. A 2 (regret question vs. blame question) x 3 (6-year-olds vs. 8-year-olds vs. adults) mixed-model ANOVA for commission/omission stories revealed a significant main effect of age group, $F(2, 136) = 6.84, p = .001, \eta^2_p = .091$. There was also a significant main effect of question, $F(1, 136) = 21.87, p < .001, \eta^2_p = .14$. There was a significant Age Group \times Question interaction, $F(2, 136) = 16.91, p < .001, \eta^2_p = .199$.

To break down this interaction I ran two univariate ANOVAs with age group as a between-participants factor for regret judgements, $F(1, 139) = 3.85, p = .024, \eta^2_p = .054$, and for blame judgements, $F(1, 139) = 19.79, p < .001, \eta^2_p = .23$. I made a Bonferroni correction for 9 post hoc t -tests across the two ANOVAs ($\alpha = .006$). Six-year-olds showed a borderline difference to adults on regret questions $t(97) = 2.75, p = .07$. Both 6-year-olds and adults

differed from 8-year-olds on the blame question, $t(78) = 3.64, p < .001$ and $t(97) = 6.22, p < .001$ respectively. Only the adult group showed a difference in their scores on regret and blame judgements, $t(58) = 8.73, p < .001$. There were no other differences revealed by the post hoc t -tests (lowest $p = .083$).

In response to the regret question, most 6-year-olds (45%) selected “same”, while the majority of 8-year-olds (52%) and adults (60%) selected the target character. In response to the blame question, the majority of 6-year-olds (55%) and adults (85%) selected “same”, while the majority of 8-year-olds (56%) selected the target character.

I looked to see whether the justifications given were in keeping with the forced choice (target, non-target, or the same). To do this I focused on the most common responses at each age. Blame justifications were coded as either making reference to the target character (Target), the non-target character (Non-target), the actions and/or outcomes being the same (Same) or the outcome being unintentional, unforeseen, chance or no-one’s fault (Rational). Only justifications for the majority character choice of each age group were coded. The majority of 6-year-olds gave a “same” response, and 92% of those justifications were coded as Same. The majority of 8-year-olds gave a target response, and 91% of those justifications were coded as Target. The majority of adults gave a “same” response, and 55% of those justifications were coded as Rational.

2.2.2 Atypical/typical stories. In response to the regret question, most 6-year-olds (45%) selected “same”, while the majority of 8-year-olds (65%) and adults (77%) selected the target character. In response to the blame question, the majority of 6-year-olds (61%) and adults (74%) selected “same”, while the most common response from 8-year-olds (50%) was the target character.

A 2 (regret question vs. blame question) x 3 (6-year-olds vs. 8-year-olds vs. adults) mixed-model ANOVA was conducted for mean atypical/typical scores. There was a significant main effect of age group, $F(2, 136) = 11.20, p < .001, \eta^2_p = .14$. There was also a significant main effect of question, $F(1, 136) = 30.87, p < .001, \eta^2_p = .19$. There was a significant Age Group \times Question interaction, $F(2, 136) = 14.37, p < .001, \eta^2_p = .17$.

Post-hoc tests with a Bonferroni correction revealed that for the regret question the 6-year-olds differed from both the 8-year-olds and the adults (both $ps < .001$), but that the 8-year-olds did not differ from adults ($p = .38$); 8-year-olds and adults had higher mean target scores than 6-year-olds. For the blame question, 6-year-olds and adults performed similarly ($p > .999$), but both differed from 8-year-olds ($p < .02$). Eight-year-olds had higher mean target scores than 6-year-olds and adults.

As with Commission/omission, blame justifications were coded as either making reference to the target character (Target), the non-target character (Non-target), the actions and/or outcomes being the same (Same) or the outcome being unintentional, unforeseen, chance or no-one's fault (Rational). The majority of 6-year-olds gave a "same" response, and 86% of those justifications were coded as Same. The majority of 8-year-olds gave a target response, and 88% of those justifications were coded as Target. The majority of adults gave a "same" response, and 70% of those justifications were coded as Rational.

2.3 Discussion

The findings from Experiment 1 replicated the findings from Guttentag and Ferrell (2004) with respect to children's and adults' judgements of another character's experience of regret. When they judged the consequent regret, 6-year-olds did not differentiate between characters who acted and those who were passive, or between characters who acted atypically and typically. Eight-year-olds, like adults, made a distinction on the basis of active/passive

and atypical/typical means resulting in the negative outcome by attributing more regret to the active and atypical characters than the passive and typical characters. Thus, in line with Guttentag and Ferrell (2004) I interpret these findings as showing that only after the age of about 7 do children understand the circumstances that lead others to feel regret.

Blame judgements in Experiment 1, however, did not follow the same developmental trajectory as regret judgements. Six-year-olds, similarly to their regret judgements, did not differentially attribute blame on the basis of the means through which the negative outcome occurred. Eight-year-olds, however, blamed characters that acted and behaved atypically rather than characters that were passive and typical, demonstrating sensitivity to the means through which the negative outcome occurred. The findings that the younger children did not exhibit evidence of the omission bias with regard to judgements of either regret or blame, suggests that these particular moral judgements were based upon reasoning rather than intuition. At first glance, adults' attributions of blame are surprising as they seem not to be sensitive to the omission bias. However, their justifications revealed that they thought that blaming either character was inappropriate because the negative outcome was unavoidable and accidental. So, their responses differed from that of both groups of children, and importantly, the moral decision to blame both characters equally was firmly rooted in reasoning rather than being intuitive.

3 Experiment 2

In Experiment 1, I demonstrated that judgements of blame as well as regret developed in parallel and concluded that both were based on counterfactual reasoning. Recall that I also manipulated the number of people affected by the negative outcome but no significant effects of this between-participants factor were found. One likely explanation for this lack of an effect of the number of individuals harmed is that in Experiment 1, I asked participants to

make forced choice (Who feels worse?/ Who is more to blame?) judgements rather than assessments of the magnitude of blame or regret felt by each character. Accordingly, Experiment 2 was designed solely to test the hypothesis that attributions of target characters' regret and blameworthiness would vary as a function of the number of people affected by a negative outcome. To measure attributions of characters' regret and blame, participants were read the same stories used in Experiment 1. Unlike in Experiment 1, however, participants rated regret and blameworthiness on a rating scale for each character. I predicted that the use of parametric ratings would demonstrate that attributions of characters' regret and blameworthiness would be higher overall when the negative outcome affected both the character and bystanders. I predicted that the use of parametric ratings would allow for participants to attribute regret and blameworthiness in less absolute terms, allowing participants to allocate more or less regret or blame on the basis of how many were affected by the negative outcome.

3.1 Method

3.1.1 Participants and design. The participants were 24 six-year-old children (11 girls; mean age = 7;4, range = 6;9–7;8), 24 eight-year-old children (10 girls; mean age = 9;3, range = 8;8–9;7). The children were recruited from and tested at two schools in the United Kingdom. A 2 (Outcome: negative outcome for character, negative outcome for character and bystanders) \times 2 (Action: atypical/typical, commission/omission) \times 2 (Age Group: 6-year-olds, 8-year-olds) mixed-model design was used similar to Experiment 1.

3.1.2 Materials and procedure. The only change to the procedure in Experiment 2 was the dependent measure. Participants rated how negatively each character in each story felt on a 5-point “feels bad” scale (1 = *does not feel bad at all*, 5 = *feels very, very bad*). Participants also rated how much each character was to blame for the negative outcome on a

5-point scale “blame scale” (1 = *not at all to blame*, 5 = *loads and loads to blame*).

Participants gave justifications for each character rating.

The “feels bad” and “blame” scales were illustrated with pictures of separate blue boxes that increased in size along the page (see Appendix B). The beginning and end points were labelled with text (e.g., *feels very, very bad*). Participants were trained to use the scales with three “feels bad” and three “blame” stories (see Appendix C). All six training stories were unrelated to the experimental stories in content; the stories were set at a playground in a sandbox where the various actions of a main character led to either good, mildly negative or very negative outcomes (e.g., A boy playing in the sandbox stomps on another boy’s sandcastle and makes the boy cry). The three “feels bad” training stories depicted negative outcomes for the main character (e.g., “Sally wants to make a sandcastle. All of the toys to build the sandcastles are being used by the other boys and girls, so she can’t build her sandcastle like she wanted”). The three “blame” stories depicted negative outcomes for another character that resulted from the main character’s action (e.g., “Dan is cross, so he goes over to the other boy’s sandcastle and stomps on it. Dan ruins the other boy’s sandcastle and doesn’t feel sorry at all”).

The first training story corresponded to a rating of 1 on the scale; a character that felt happy and a character that was not at all to blame for something bad happening. The second training story corresponded to a rating between 2 and 4 on the scale; a somewhat happy character and a character that was partially to blame for something bad happening. The third training story corresponded to a rating of 5 on the scale; a very unhappy character and a character that was very much to blame for something bad happening. Participants were given standardized feedback after each rating for correct and incorrect ratings. An example of feedback for a correct response to a blame story would read:

Yes, I think so too. He is to blame loads and loads because he stomped on the other boy's sandcastle.

From the same story, an example of feedback for an incorrect response would read:

Well, I think he is to blame loads and loads because he stomped on the other boy's sandcastle, so I'd pick this one [pointing to 4 on the scale] or this one [pointing to 5 on the scale].

3.2 Results

The aim of experiment 2 was to solely focus on the outcome paradigm tested in Experiment 1, whereby I predicted that participants would attribute greater amounts of regret and blame in situations where a greater number of characters were affected by the negative outcome. Given the only aim of experiment 2 was to further test the effect of magnitude of the negative outcome, I do not present the analyses of the commission/omission and atypical/typical factors separately, but instead focus on the effects of magnitude of outcome on the parametric ratings of blame and regret. I had no reason to predict differences between ratings of the magnitude of outcome for both blame and regret questions to differ between commission/omission stories and atypical/typical stories. For reasons of consistency with Experiment 1, however, I included Story Type (commission/omission, atypical/typical) as a factor in the overall analysis. Ratings were averaged across both stories to give four blame ratings: commission/omission story target character, commission/omission story non-target character, atypical/typical story target character, atypical/typical story non-target, and the equivalent regret ratings. I conducted separate ANOVAs for blame and regret ratings.

I conducted a 2 (Story Type: commission/omission, atypical/typical) \times 2 (Character Rating: target, non-target) \times 2 (Outcome: character affected, character and bystanders affected) mixed-model ANOVA for blame ratings \times 2 (Age group: 6-year-olds, 8-year-olds).

There was a significant effect of character, $F(1, 44) = 18.72, p < .001, \eta^2_p = .30$, with targets rated more harshly ($M = 2.65, SE = .16$) than non-targets ($M = 2.21, SE = 1.5$). There was a main effect of outcome $F(1,44) = 5.17, p = .028$, partial $\eta^2 = .11$. More blame was attributed when other people were affected ($M = 2.76, SE = .20$) than only the protagonist ($M = 2.10, SE = .20$). There was an interaction between character and age, $F(1, 44) = 5.25, p = .027$, partial $\eta^2 = .107$. There were no other main effects or interactions.

I examined the interaction using t tests with Bonferroni corrections ($\alpha = .0125$). The younger children did not differentiate between the target and non-target in their blame ratings $p = .12$. The older children blamed target characters ($M = 2.58, SD = 1.13$) more than non-target ($M = 1.91, SD = .80$), $t(23) = 4.40, p < .001$. Comparing across the age groups, there was no difference in the blame ratings of target characters, but there was a trend for older children to be more lenient in their ratings of the non-targets, $t(46) = 2.02, p = .049$.

I conducted a parallel analysis for regret ratings. There was a significant effect of character, $F(1, 44) = 4.68, p = .036$, partial $\eta^2 = .10$, with targets once again rated more harshly ($M = 4.30, SE = .10$) than non-targets ($M = 4.17, SE = .12$). There was an interaction between Character and Age, $F(1, 44) = 4.68$, partial $\eta^2 = .10$ and an interaction between Story Type and Character, $F(1, 44) = 4.22, p = .046$, partial $\eta^2 = .09$.

I examined the interactions using t-tests with Bonferroni corrections for each set of multiple comparisons ($\alpha = .0125$). For the interaction between character and age, the younger children did not differentiate between the target and non-target in their regret ratings ($p > .999$). The older children judged that targets ($M = 4.33, SD = .78$) would feel greater regret than non-targets ($M = 4.07, SD = .88$), $t(23) = 2.75, p = .011$. Comparing across the age groups, there was no difference between regret ratings for either target or non-target characters.

For the interaction between story type and character, there was a difference between target ($M = 4.45$, $SD = .76$) and non-target ($M = 4.23$, $SD = .85$) in regret ratings on the typical stories, $t(47) = 3.01$, $p = .004$, and between targets in the commission stories ($M = 4.15$, $SD = .81$) and typical stories ($M = 4.45$, $SD = .76$), $t(47) = 3.26$, $p = .002$. The effect of character on regret ratings was greater for the atypical/typical stories; differences between target and non-target characters may be driven by the atypical/typical stories.

3.3 Discussion

The findings from Experiment 2 demonstrate that younger children do not differentiate how they rate regret and blame among target and non-target characters. Older children attribute more blame and more regret to target characters than non-target characters. Thus, older but not younger children behave as if they are influenced by commission and typicality when judging blame and regret. We can be confident that children at both ages were using the scales appropriately because the blame ratings were sensitive to outcome (i.e., characters were judged to be more blameworthy if more people were affected by the bad outcome). It is of note that this manipulation did not influence judgements of regret.

The findings suggest that blame ratings are sensitive to the number of people involved in the negative outcome at both ages. Regret ratings, however, seem to be uninfluenced by the number of people affected by a negative outcome. This differentiation between regret and blame ratings can be explained by the fact that regardless of the number of characters affected by the negative outcome, the character who felt regret was always regretful about their negative experience regardless of the impact the action had on other people. Regret pertains to the negative experience of the character (i.e., an internal mental and emotional state) whereas blame requires consideration of the amount of total damage done in order to determine the level of individual responsibility. Therefore, the regret question is responding to the

character's negative experience, which is negative regardless of who else was also affected. Blame, however, is assessing whether the character *did* something wrong, which seems to be responsive to considering the global cost and is therefore affected by the number of people experiencing that cost, above and beyond the character's own experience. We should also note that there is some evidence that the different stories have an influence on how children judge regret. Evidently, characters that did something atypical were judged to feel the most regret. As this did not interact with age, we did not consider it further.

4 General Discussion

The findings from these experiments demonstrate that both judgements about regret and blame can be influenced by counterfactual reasoning; we saw no evidence for an intuitive omission bias influencing judgements of harm in the youngest children and instead it appeared that blame judgements, like regret judgements, showed a reasoning-based change at around 8 years of age. In both Experiments 1 and 2, six-year-olds did not incorporate counterfactual reasoning when attributing regret and blameworthiness to another person; moreover, they did not focus on the means by which a negative outcome occurred but, instead, focused solely upon the magnitude of the negative outcome. Older children's and adults' judgements of regret and blame did incorporate counterfactual reasoning, but to different extents when attributing blame: Whereas the older children's blame judgements closely paralleled their judgements of regret, the adults in Experiment 1 focused more than did the children on the uncontrollable nature of the negative outcomes for the story protagonists when making judgements of blame but not regret. These findings are in keeping with the reasoning account of judgements of blame and counter to the intuitionist account.

Younger children, even at 6 years old, focused on the negative outcome in their moral judgements; older children, in contrast, took into account other factors beyond the number of

people harmed by the negative outcome, such as the means through which the negative outcome occurred. For children to make judgements on the basis of the means through which a negative outcome occurs, and to consider distinctions of active versus passive harm, the ability to reason about counterfactuals (including the ability to compare counterfactuals with reality) needs to be in place. Whether what younger children lack is a competence in reasoning ability or a lack of knowledge regarding which factors (e.g., active/passive) impact the experience of regret and judgements of blame is unknown. However, the present findings suggest that at least one of the proposed intuitive fundamental principles suggested by Cushman and colleagues (2006) does not impact judgements of blame prior to the age at which children tend to incorporate counterfactual reasoning into the blame judgement process.

There is some evidence that infants are sensitive to certain basic moral principles such as harm and fairness (e.g., Sloane et al., 2012). One possibility is that there are a few early developing moral principles, in which case our evidence indicates a need to differentiate which principles are early developing and which, like the omission bias, are not—or at least are not early developing for judgements of blame. Alternatively, some sensitivity to some aspects of morality may appear in a crude form at an early age, but undergo substantial change before becoming fully-fledged, normative moral judgements. This raises the question of whether what we see in infancy is rightly termed morality.

Our evidence from adults' blame judgements supports this second view. Adults chose not to blame either individual in our forced choice question (claiming characters were equally to blame). This suggests that adults were bringing additional kinds of information and mitigating factors in to their reasoned decision making, compared to 8-year-olds. One question for further research is when this shift (or shifts) takes place.

With both experiments, I aimed to address the overarching question of how reasoning influences moral judgements, specifically judgements of blame. The current dominant theory in research with adults is that moral judgements are largely driven by emotion and intuition, and the influence of some basic moral principles should not be influenced by age-related changes in reasoning ability (Haidt, 2001; Cushman et al., 2006). Both experiments presented here provide evidence that there are developmental changes in blame attributions similar to age-related changes in other judgements about negative events (e.g., regret) that involve reasoning about the negative outcome, the means by which the negative outcome occurs, counterfactual alternatives, and accidental or unintended negative outcomes. The evidence found here for protracted development of blame judgements runs counter to the view that these judgements are based primarily upon a small set of intuitively understood and applied principles.

In support of previous research (e.g., Cushman et al., 2006) and my hypotheses regarding the role of behavioural and contextual information in moral judgements, these studies demonstrate that reasoning about behavioural and contextual information pre-judgement, specifically about active versus passive behaviours and atypical versus typical behaviours leading to a negative outcome, is influential. Information regarding the means, specifically the action a character takes that leads to a negative outcome, and a comparison of the actual and hypothetical outcomes, is influential only when there is an age-related sophistication in reasoning. For young children who lack sophisticated reasoning abilities, information about the means through which an action occurs is not taken into account resulting in a lack of distinction between actions that lead to a negative outcome. Once the sophisticated reasoning is in place, however, information about the comparison of

hypothetical and actual outcomes, and active and atypical behaviour, influences whether and the extent to which someone is blamed.

CHAPTER 3

ENGAGEMENT WITH BEHAVIOURAL AND CONTEXTUAL INFORMATION POST-JUDGEMENT: MORAL CONDEMNATION OF HARM AND PURITY VIOLATIONS AND DESIRE TO KNOW MORE

This chapter investigates selection of and engagement with post-judgement behavioural and contextual information. This chapter is examining engagement with various types of behavioural and contextual information following judgement of moral violations relating to the two moral foundations of harm (rape) and purity (incest). In Experiment 3, participants' ratings of moral appropriateness, judgement certainty, and emotion intensity did not differ as a function of violation type, nor did their ratings of the potential utility of additional behavioural and contextual information; however, participants wanted more information after reading about a harm violation than a purity violation. In Experiment 4, participants in the harm-violation condition expressed equivalent interest in wrongness-confirming and -contradicting information, whereas participants in the purity-violation condition expressed more interest in wrongness-contradicting than -confirming information. Experiments 3 and 4 suggest that post-judgement information engagement in the moral domain is not driven by a simple desire for epistemic certainty or by the perceived utility of the information. Instead, my findings suggest that post-judgement information desire differs as a function of the moral domain. These findings are consistent with a motivation to regulate affect when faced with purity violations and a motivation to maintain certainty about harm violations. Experiments 5 and 6, however, did not yield evidence to support my predictions: These experiments investigated how participants actually engaged with additional information, and demonstrated that the majority of participants were interested in wrongness-confirming

*information regardless of their exposure to harm or purity violations. Additionally, the amount of information participants chose to engage with did not differ as a function of violation. Overall, the experiments in this chapter demonstrate that behavioural and contextual information play at least some role post-judgement that differs depending on the type of information available and the type of moral violation that is being judged.*³

1 Moral Judgements, Attitudes, and Information Engagement

The overarching goal of this thesis is to examine the impact of behavioural and contextual information that goes beyond the “gut” reactions that are believed to be the basis for moral judgement. In this chapter, I aim to investigate engagement with additional information post-judgement. I took as my starting point a substantive literature on how attitudes influence informational engagement. Importantly, both moral principles and non-moral attitudes are inherently evaluative (right versus wrong, positive versus negative), and both direct behaviour, suggesting that much about the impact of moral judgement on post-judgement processing can be gleaned from the attitudes literature. Even the differences between moral principles and attitudes tend to relate to the notion that moral principles are enhanced, emotional, universal and more certain attitudes, which has important consequences for post-judgement information search. Therefore, both the similarities and differences between attitudes outlined in this chapter provide reasons to examine post-judgement information selection and engagement with respect to epistemic certainty.

³ The research from Experiments 3 and 4 of this chapter is in revision for resubmission to the *Journal of Experimental Social Psychology* as the following manuscript: Powell, N. L., Beck, S. R., Giner-Sorolla, R., & Quinn, K. A. (in revision). *Reactions to harm and purity violations: Moral condemnation and desire to know more.*

1.1 Differences between Attitudes and Moral Principles

Attitudes are defined as positive and negative evaluations a person makes regarding a certain attitude object (Skitka, 2010). Moral judgements are also evaluative, but are framed in the context of determining what is right and wrong as opposed to positive and negative (Skitka et al., 2005). Research by Skitka et al. (2005) looked directly at the similarities and difference between attitude strength and moral conviction. Moral conviction is a strong sense that something is morally right or wrong (Skitka & Mullen, 2002). There is some inconsistency across disciplines of moral philosophy, moral psychology, and attitude research as to whether moral convictions should be treated separately from other non-moral attitudes. In other words, there may be reason to believe that moral judgements potentially have some shared processes to attitudes. That being said, however, moral judgements may have some unique qualities that dissociate them from attitudes, specifically with respect to how moral judgements arguably rely on emotional processing, a lack of deliberative processing, circular post-hoc reasoning, and disinterested third-party judges (Haidt, 2001).

An argument in support of classing moral beliefs (i.e., beliefs that arise from making judgements of right and wrong) as separate from attitudes is that moral beliefs and moral principles are seemingly universal. Many moral beliefs, such as the belief that doing harm to someone is wrong, hold true across culture, religiosity, and context (Skitka & Mullen, 2002). Determining what is right and wrong is often experienced as standardised and universal (Shweder, 2002; Turiel, 2002). This is not to say that moral beliefs are experienced in the same way as non-moral facts (i.e., facts about the natural world that do not arise from evaluative judgements of right and wrong). For example, a moral belief that harming others is wrong is not the same as the fact that human beings on earth experience a gravitational pull. The reason moral convictions feel like universal fact has been a topic of speculation. The

Humean distinction, for instance, between non-moral and moral beliefs is that morals have a motivational quality that dictates how people should behave. Moral beliefs include an element of “should” or “ought”—for example, one should not cause harm. Beliefs or facts about the world do not involve any implication for how people or society ought to do things; a belief that blue is a nice colour does not necessarily suggest anything about the way things should be. Non-moral attitudes may also be easier to explain than moral beliefs because moral beliefs are less deliberative, especially when providing justifications for moral beliefs (Haidt, 2001). When left with the justification “it just is” for why something is immoral, it suggests that morals are self-evident facts which do not require justification.

In addition to the universality and fact-like nature of morality, the emotional component of moral conviction may be another case for arguing that non-moral attitudes are distinctly different. Haidt (2003) has identified moral emotions that are associated with different foundations of morality, for example, feelings of disgust being associated with violations of purity. The social intuitionist model of moral judgement posits that morality is largely an emotional and intuitive process (Haidt, 2001). Some theories regarding attitude suggest that non-moral attitudes are also emotional. Moral convictions, however, often deal with raw affective experiences that result in snap-judgements when the judge is a third-party disinterested elicitor, whereas non-moral attitudes can be comprised of personal opinions and preferences— potentially less emotionally salient and less prescriptive for how people should and should not behave (Skitka, Bauman, & Sargis, 2005).

A series of studies by Skitka et al. (2005) examined the distinction between moral convictions and strong non-moral attitudes by testing what is called the moral mandate hypothesis. An implication of the moral mandate hypothesis is that people will place a greater significance on other people sharing moral beliefs than sharing strong but non-moral attitudes.

Skitka et al. (2005) first asked participants to think of the most significant issues facing the world and to rate the issues according to their significance. Assessments were given on the strength of the attitude towards the issue, the degree to which the issues were related to moral principles, and the responses to hypothetical interactions with people who disagreed with the issue. Attitude strength was measured through ratings of how strongly participants felt about their chosen issues. Moral conviction was measured through participants' ratings of the relatedness of each issue to participants' core moral beliefs. Skitka et al.'s (2005) findings supported the moral mandate hypothesis: A significant number of participant-generated pressing issues facing the world stemmed from a core moral belief. Issues that stemmed from strong moral convictions were also more likely to yield greater desirability for social distance from individuals who disagreed than issues that were loosely or not at all related to moral convictions. In other words, world issues are likely to be moral in nature, and those who disagree with a person's moral issue are more undesirable than those people who just disagree with a non-moral issue.

Skitka et al. (2005) further examined the moral mandate hypothesis with respect to implications for behaviour. Two additional studies confirmed the initial findings from the first study and also demonstrated that participants chose to physically distance themselves significantly further away from those with an opposing moral attitude on the issue of abortion rights than those who opposed a strong non-moral attitude. Further, when participants were given the task of solving a moral or non-moral conflict, the interpersonal interactions were less cooperative and cohesive, particularly with respect to resolving group disagreement, for those resolving moral conflicts than those resolving non-moral conflicts. These series of studies provide compelling evidence for the argument that in some contexts moral beliefs evoke different reactions from strongly held non-moral beliefs.

Notwithstanding the differences outlined above, it is nonetheless true that both non-moral attitudes and moral beliefs are inherently evaluative (right versus wrong, positive versus negative) and that both direct behaviour (Haidt, 2001; Skitka & Mullen, 2002). Similarly, judgements relating to both moral and non-moral attitudes can be made more automatically than deliberately (Haidt, 2001). It may be the case that moral attitudes are much stronger and less context-sensitive than non-moral attitudes (to the extent that they are perceived as universally true and as worthy of defence at any cost; Skitka & Mullen, 2002), but these differences appear to speak more to motivational intensity than actual processes and mechanisms. For that reason, I began with the loose assumption that moral attitudes could be conceived as a class of particularly strong attitudes, and looked to the literature to examine attitude strength and informational engagement.⁴

Attitude strength—that is, how resistant an attitude is to change, how long an attitude lasts, and how much an attitude affects judgements and behaviour— has been shown to be affected by different attributes that rely on deliberative processing and information (Krosnick & Petty, 1995; Skitka, Bauman, & Sargis, 2005). Some of the attitude-strength-related attributes that rely on processing and information are knowledge, certainty, intensity, and extremity (Skitka et al., 2005).

1.2 Attitudes and Additional Information

The greater the amount of psychological significance a person attaches to an attitude, the more influential and robust that attitude becomes, which in turn impacts how that person selects and attends to information that might pertain to the attitude (Berent, Krosnick, &

⁴ Although the research on attitudes has informed my hypotheses regarding the influence of reasoning and information on moral judgements, this thesis does not make any direct empirical or theoretical comparisons between moral judgements and attitudes.

Boninger, 1993; Boninger, Krosnick, & Berent, 1995). The leading contributors to attitude importance include believing that the object of the attitude is related to one's self-interest (Boninger, Krosnick, & Berent, 1995), relating to others who share the same connection between the object of the attitude and one's self-interest (Key, 1961; Modigliani & Gamson, 1979), and identifying a close kinship between personal attitudes and values (Johnson & Eagly, 1989; Katz, 1960; Rosenberg, 1956). Other research has identified additional contributing factors to attitude importance, such as the accessibility of the attitude (Fazio, 1995) and the ability to easily retrieve information pertaining to the attitude from memory (Haddock et al., 1996, 1999; Pelham, 1991).

Research examining attitude importance often involves asking participants to report how important something is, followed by reports of how much that thing is personally cared for by the participant. The importance participants place on something has been strongly correlated with the subsequent importance of the attitude participants have towards that something (Boninger et al., 1995). Boninger et al. (1995) found this effect in activists' attitude about things such as animal rights; animals, the object of the attitude, were considered highly important to the activist, and the attitude towards animal rights activism, the attitude, was also considered to be highly important. Some research in the attitudes domain has shown that the importance of an attitude can increase the desire to seek additional, relevant information (Berent & Krosnick, 1993). Additionally, the importance of an attitude can also impact trait inferences (Judd & Johnson, 1981) and increase the expression of attitudes and attitude-related behaviours (Krosnick & Telhami, 1995). Attitude importance plays a useful, motivating role that encourages acquisition of more information. Attitude importance is suggested to cause the acquisition of attitude-related knowledge; attributing a great deal of

importance to the subject of an attitude generates the desire to acquire knowledge about that subject (Visser et al., 2006).

Several factors have been identified that make an attitude to a subject important. For example, Visser et al. (2004) demonstrated that when the issue of legalised abortion is highly relevant to the values of similar others, the attributed importance of the issue increases. Specifically, participants attributed greater importance to their own attitudes toward legalised abortion when they had knowledge that the issue of legalised abortion was important to other relevant groups of people. The importance other people placed on legalised abortion did not, however, have an effect on participants' amount of attitude-relevant knowledge. Participants' attitude-relevant knowledge was primarily a result of exposure to news media; news media did not relate to the amount of importance participants placed on their attitudes towards legalised abortion. This study demonstrated that relevance of an attitude to participants and the relevance to their like-minded social groups' values most influenced their attitude importance and related to their degree of attitude-relevant knowledge.

Relating back to the influence of personal importance to moral judgements, liberals living in the United States who show a greater intuitive sensitivity to violations of care (i.e., the harm foundation) are more likely to judge harm as being morally wrong above and beyond other moral violations, such as violations of authority (Haidt, 2012). Attitude importance shares several properties with moral judgements. According to the social intuitionist model (Haidt, 2001; Haidt & Graham, 2007; Haidt & Joseph, 2004), moral judgements are based on intuitions that are based on a set of moral foundations: so-called "individualising" foundations of harm and fairness, and "binding" foundations of loyalty, authority, and purity. The moral foundations are more or less accessible or important to people on the basis of their educational background, culture, or political viewpoint (Haidt,

2012; Haidt et al., 1993). For instance, research has demonstrated that in the United States, people who identify as liberal are more sensitive to the individualising moral foundation above and beyond the binding foundations, whereas people who identify as conservative are sensitive to all of the moral foundations in relatively equal measure (Graham, Haidt, & Nosek, 2009).

People's knowledge about the subject of their attitude can facilitate flexibility to the input of new, potentially updated information (Davidson, 1995). On the other hand, however, having a great deal of knowledge about the subject of an attitude can cause attitude resistance, preventing attitude change (Wood et al., 1995). Attitude resistance provides protection against useless and unfounded information, such as propaganda, but may also prevent empathy towards another person's legitimate attitude. Moral judgements may be similarly influenced by knowledge in that a person who has a great deal of knowledge about a moral scenario may be resistant to additional information. The way in which we gather information depends on the importance of an issue and the knowledge we have of an issue. For instance, research on attitudes towards various political candidates asked participants to learn about the candidates' positions (Visser et al., 2004) and found that attitude importance affected what participants learned. Specifically, the participants who attached more importance to the issue of capital punishment, for example, asked for information about the candidates relating to that topic significantly more often than the other areas of the candidates' positions. If the subject of an attitude is important to individuals, they will be selective in their information gathering and would like to deal with topics that can relate back to their own, already-established attitudes.

The consideration of information is important to the moral judgement process when we consider the role of reasoning in moral judgements. For instance, if moral judgements are believed to be based largely on unconscious intuitions guided by a set of natural moral

principles and foundations, then moral judgements may not be largely influenced by knowledge and information. Additionally, if moral judgements are similar to attitudes, people could lack interest in additional information that could mitigate a judgement when people have made a moral judgement they believe is important and for which they believe they have a great deal of knowledge. Perhaps if people are given the opportunity to learn new information about a moral situation after making a judgement they believe was made with a great deal of knowledge and is important, people may choose to seek information that is relevant to maintaining their important and previously established moral judgement.

One hypothesis I aimed to explore in the following set of studies relates to epistemic certainty, and how our own need for certainty can influence our need and desire for further information and knowledge. The literature on attitudes has demonstrated that certainty about one's own attitude can influence information selection and attention, as well as influence the resulting attitude itself. Here I will summarise some of the relevant findings regarding attitude certainty that have informed my hypotheses about moral judgements and epistemic certainty, and how epistemic certainty motives can influence selection of and engagement with post-judgement information. The similarities between non-moral attitudes and moral beliefs (i.e., the evaluative nature of these judgements, and the implications for behaviour) would suggest that epistemic certainty about moral judgements can influence our attention to and selection of further information in a similar way to attitudes. Likewise, the differences between non-moral attitudes and moral beliefs would suggest to an even greater extent that moral judgements and information selection engagement will be influenced by epistemic certainty needs, given that the differences relate to moral judgements yielding an even greater level of epistemic certainty and confidence than non-moral attitudes (Skitka & Mullen, 2002).

Therefore, epistemic certainty motives are likely to play a role in selection to and engagement with post-judgement information.

1.2.1 Attitude Certainty

One of the most influential factors that have been demonstrated to influence attitudes is certainty. If it is the case that moral principles are enhanced attitudes, held with greater conviction, certainty and intensity, it is worthwhile to consider how attitudes are affected by certainty before examining how post-judgement information selection and engagement are affected by certainty in the moral realm. Research by Visser et al. (2003) explored attitude certainty and importance with respect to people's preference for a political candidate. Participants rated various political preferences in terms of their importance, and reported their confidence in their political attitudes. The more important the political issue was rated, and the higher degree of certainty, the more attached participants became to their selected political candidate. Additional measures revealed that high certainty and high importance increased the likelihood that participants would feel upset if the candidate were not chosen and would encourage others to adopt their attitude and preference. Additionally, high certainty and high importance led to an increased likelihood that the attitude would inform subsequent behaviour (e.g., voting).

A great deal of the literature on attitude strength deals with factors which are likely to increase an attachment to an attitude, such as importance, knowledge, information and certainty. It is therefore important to consider how these factors are related to attitude resistance and the contexts under which we can flexibly and comfortably move from one attitude to another, such as when a strong attitude is able to resist change in the face of a counter-attitude. This has implications for moral judgements and the role of information and reasoning when making moral judgements. For instance, there are opportunities after making

a moral judgement when behavioural and contextual information might be available that could modify our moral judgement.

Some research supports the idea that the more often we resist attitude or judgement change, the more confidence we develop in our attitudes, the more valid we believe our attitudes to be, and the more we will be comfortable resisting other forms of attitude persuasion in future situations (Fazio & Zanna, 1978). This could be particularly important to moral judgements and our engagement with information and reasoning. If we make moral judgements on the basis of emotion and intuition, we may resist further behavioural and contextual information that has the potential to change our moral judgements, which in turn reinforces our confidence in our moral judgements. This means that in the absence of comprehensive knowledge, the confidence in our judgements and imperviousness to other information increases while the motivation to address and acquire potentially influential information decreases.

1.3 Summary

At least in the domain of non-moral attitudes, factors such as importance, knowledge, and certainty can motivate further engagement with certain types of information—and the commonalities between moral and non-moral attitudes (e.g., the evaluative component, their impact on behaviour) suggest that moral judgements facilitate engagement with behavioural and contextual information on the basis of similar factors such as knowledge and certainty. This allows me to use the findings from the attitudes literature to inform my hypotheses about post-judgement behavioural and contextual information engagement. The attitudes literature provided an analysis of how attitudes maintain and react to certain types of information (e.g., attitude-congruent information or attitude-incongruent information). I turned to the research on attitudes for the hypotheses in this chapter because attitudes can lead to selective exposure

of information that have different consequences for the resulting strength and maintenance of an attitude (Canon, 1964; Jonas, Greenberg, & Frey, 2003).

Research on attitudes has shown the importance of knowledge and certainty in leading to different outcomes for information engagement (Kanwar, Grund, & Olson, 1990; Krosnick, Boninger, Chuang, Berent, Carnot, 1993; Krosnick & Shuman, 1988). Additionally, research on attitudes has demonstrated that information search can be manipulated. Jonas, Graupmann, and Frey (2006) found that participants who were manipulated to feel in a negative mood had a preference for information that supported their attitude whereas participants who were manipulated to feel in a positive mood took a more balanced approach to both supporting and contradicting information. This chapter will address how epistemic certainty and affect regulation influence post-judgement information engagement, specifically behavioural and contextual information.

This chapter investigates how participants engage with behavioural and contextual information after making a moral judgement. Specifically, the research in this chapter aimed to investigate whether participants' judgement certainty influenced their desire and willingness to engage with post-judgement behavioural and contextual information. The research on attitudes provides reason to believe that, given the similarities between moral judgements and attitudes, a great deal of epistemic certainty about one's moral judgement might lead people to find additional behavioural and contextual information that could change a judgement to be irrelevant and uninteresting.

2 Post-Judgement Selection and Engagement with Behavioural and Contextual Information

This thesis aims to investigate how and when people engage with behavioural and contextual information that either precedes a judgement or follows a judgement, and how

behavioural and contextual information affects behaviour. The overarching goal of this thesis is to provide some evidence that while judgements may result from intuitions and emotions (Haidt, 2001), some degree of deliberative processing can take place and be influential. Chapter 2 demonstrated that consideration of behavioural and contextual information that is available pre-judgement can be influential to the amount of blame that is attributed to an actor when participants are capable of engaging in sophisticated counterfactual reasoning processes. This chapter aims to investigate how participants choose to engage with behavioural and contextual information that follows their moral judgement, and how that choice to engage with information differs depending on various factors such as the type of moral violation that is being judged, the level of certainty participants have about their own judgement and the type of information that is available for selection. This chapter addresses the overarching goal of the thesis by demonstrating that while initial moral judgements may depend on intuitive gut reactions and emotions, participants are capable of engaging in more deliberative information selection processes after making a judgement on the basis of minimal information.

Judging the moral worth of others' actions, even on the basis of minimal information, feels effortless—and, indeed, one prevailing view is that moral judgement is automatic and intuitive (Haidt, 2001). What *follows* moral judgement, however, has been studied less extensively, reflecting perhaps the assumption that post-judgement processing serves primarily to persuade others of the correctness of the judge's moral principles (Haidt, 2001). In the current research, I questioned whether post-judgement processing might serve functions other than other-directed persuasion—specifically, to establish epistemic certainty or to regulate affect.

According to the social intuitionist model, individuals make “snap” judgements about moral behaviour on the basis of minimal information because these types of judgements rely on an intuitive and affective response to the behaviour (Haidt, 2001). Any subsequent need or desire for additional information comes after the fact and is not part of the moral judgement itself. That is, reasoning in the moral domain is assumed to be post hoc and directed primarily toward providing social (i.e., other-directed) justifications for emotion- or intuition-driven judgements—that is, to persuade others to accept and adhere to the judge’s moral principles. Although persuading others to join our moral crusade when making a judgement of wrongness can be a motivating factor in post-judgement processing, such processing might also serve more self-directed motives. Most obviously, post-judgement processing may be driven by epistemic need for certainty—that is, to convince ourselves of the correctness of our own moral judgements, as well as to maintain levels of certainty post-judgement. In addition, however, post-judgement reasoning may be driven by our desire to approach or avoid the affective consequences of our judgements—perhaps reflecting the extent to which we want to engage with or indulge our emotional reaction to a violation.

According to moral foundations theory (Haidt & Joseph, 2004), moral violations can be divided into five different types, which elicit different (albeit sometimes overlapping) emotions (Haidt, 2001, 2003; Haidt & Graham, 2007; Haidt & Joseph, 2004; Rozin, Lowery, Imada, & Haidt, 1999). For the purpose of the present research, I explored how moral foundation type mediates post-judgement processing, focusing specifically on the foundations of harm and purity. Violations of these foundations have been shown to elicit differing affective responses—namely, anger and disgust, respectively (Gutierrez & Giner-Sorolla, 2007; Horberg, Oveis, Keltner, & Cohen, 2009; Rozin et al., 1999).

Of importance to the current research, recent evidence demonstrates not only that anger and disgust tend to be elicited by different types of moral violations, but also that anger can be more flexible than disgust in response to a moral wrongdoing. For example, Russell and Giner-Sorolla (2011a) found that when prompted, participants were able to generate more hypothetical mitigating factors to explain a harm violation (which elicited anger) than a purity violation (which elicited disgust). Further, when participants were asked to imagine that the hypothetical mitigating factors had actually occurred, those who experienced anger in response to a harm violation show a greater shift in moral judgement and emotion than did those who experienced disgust in response to a purity violation. In other words, participants' ability to generate information about possible mitigating circumstances surrounding different kinds of moral violations predicted their subsequent emotional states.

Russell and Giner-Sorolla (2011b) have also shown that explanations for judgements related to moral anger versus moral disgust differ qualitatively. When participants were asked to explain why they felt moral anger, they favoured "elaborated" reasons that included possible causes and consequences of moral actions. In contrast, when participants were asked to explain why they felt moral disgust, they favoured non-elaborated reasons that did little more than reiterate subjective evaluations and emotions. Participants were able to generate more complex justifications for feeling angry versus disgusted.

Further work by Russell and Giner-Sorolla (2011c) provides evidence that moral anger, but not moral disgust, is sensitive to mitigating information—in this case, actor intentionality. Participants read about a bodily norm violation (cannibalism) that was either intentional or unintentional, completed measures of anger and disgust, and appraised the wrongness of the behaviour. The results demonstrated that, controlling for disgust, participants' anger was predicted by their assessment of the actor's intentionality but not by

participants' assessment of the extent to which the behaviour violated the cannibalism taboo; controlling for anger, participants' disgust was predicted by their assessment of the extent to which the behaviour violated the cannibalism taboo but not their assessment of the actor's intentionality.

Thus, different moral violations and different moral emotions are associated with different patterns of reasoning—both in terms of flexibility and sensitivity to additional behavioural and contextual information, but also in terms of the kinds and complexity of post-judgement explanations that can be generated. In the current research, I investigated whether differences in post-judgement reasoning in particular arise for motivational reasons. Specifically, I considered two possible functions of post-judgement processing.

First, individuals may be driven by the motive to achieve epistemic certainty—that is, to convince themselves (rather than or in addition to convincing others) of the rightness of their moral judgements—and this need may vary as a function of violation/emotion type. Indeed, the greater sensitivity and flexibility of moral anger than moral disgust reported by Russell and Giner-Sorolla (2011a, 2011b, 2011c) might suggest that moral anger is associated with less certainty than moral disgust. Outside the moral emotion literature, however, appraisal theories of emotion (e.g., Smith & Ellsworth, 1985) suggest that, if anything, anger is associated with greater certainty than is disgust. The status of epistemic certainty motives following moral judgement thus remains unclear.

Alternatively, individuals may be differentially willing to engage with behavioural and contextual information related to different moral violations because of the motivations that they elicit. Anger has been designated an approach-oriented emotion (e.g., Carver & Harmon-Jones, 2009). It is assumed to result from an interrupted approach towards a goal (Depue & Iacono, 1989; Fox, 1991) and to produce action tendencies that attempt to remove barriers to

goal pursuit (Frijda, 1986), to change others' behaviour (Fischer & Roseman, 2007), or to punish offending others (Shaver, Schwartz, Kirson, & O'Connor, 1987). In contrast, disgust is typically experienced as an avoidant feeling of physical revulsion, along with a strong desire to withdraw from the elicitor (Rozin, Haidt, & McCauley, 2000). Indeed, disgust may have evolved as a response that acts as a disease-avoidance mechanism (Rubio-Godoy et al., 2007; Oaten, Stevenson, & Case, 2009; Schaller & Duncan, 2007; Stevenson & Repacholi, 2005). Individuals may thus be more motivated to engage with anger- than disgust-inducing moral violations.

Although it may seem counterintuitive to suggest that individuals would seek to indulge the negative emotion of anger, it is worth noting that I am referring to moral rather than personal anger. Elsewhere, it has been argued that anger at the violation of a moral standard should be distinguished from anger at personal harm (Batson et al. 2007). I suggest that because moral anger is evoked by conditions that have no direct implications for the perceiver's own outcomes, while being approach- or action-oriented, it may actually have positive implications for the perceiver's sense of self. In recognizing that another individual has violated a moral standard, moral anger can invoke a sense of moral righteousness (Haidt, 2003) that affirms the perceiver's moral integrity. Disgust's avoidant tendencies, however, are unlikely to afford the same opportunity. Moral anger may show an aspect of approach and engagement that moral disgust does not have.

The present research investigated whether different types of moral violations elicit different patterns of post-judgement processing that are consistent with these motives. I conducted four experiments. In the first two, participants judged the wrongness of either a harm or purity violation, then reported their certainty in their moral judgement and the intensity of their emotional reaction. They then rated desire for and usefulness of additional

information. In the second experiment, I also varied whether participants believed they would have the opportunity to view wrongness-confirming or -contradicting information. In the third and fourth experiments, I tested participants' self-reported willingness to engage with behavioural and contextual information by seeing how participants would actually engage with additional behavioural and contextual information designed to either inflame or mitigate their wrongness judgements and elicited emotional reactions. This was to determine how participants believed they would engage with behavioural and contextual information if information were made available when deliberating on their post-judgement information interests. I asked participants particularly about their interest in mitigating or inflammatory information because mitigating information would be information that cooled their initial emotional judgement and reaction by allowing them to consider information that directly contradicts their initial judgement, and inflammatory information would be information that maintained or enhanced their initial emotional judgement and reaction by reinforcing and confirming what led to their initial judgement. Therefore, I deduced that an interest in mitigating information would be demonstrative of some degree of interest in deliberating and reasoning while an interest in inflammatory information would be demonstrative of some degree of interest in maintaining an intuitive and emotional approach. I considered two possible patterns of evidence for why participants might choose the types of information available on the basis of what they judged:

(1) Based on the literature on attitudes, there is evidence to suggest that certainty about an attitude can interact with information engagement, such as research which suggests a person's certainty about their own judgement can increase after resisting counter-attitudinal information (e.g., Tormala & Petty, 2002; Tormala & Petty, 2004). Given the similarities between moral beliefs and attitudes described in the beginning of this chapter, there is reason

to believe that certainty of one's own belief or judgement (i.e., epistemic certainty) will interact with information engagement. Assuming differences in *epistemic certainty motives* as a function of moral violation type, I would predict that judgements of harm violations, relative to judgements of purity violations, would be associated with lower ratings of judgement certainty, as well as stronger desire for and greater perceived usefulness of additional judgement-relevant information, which could be used to establish certainty. Harm violations are violations in which there is a clear victim and transgressor. Violations involving a victim are ones where society demonstrates the need to establish justice, sometimes by punishing the transgressor. Therefore, people may be motivated to achieve certainty that the act is wrong because of the implications for upholding justice and determining appropriate punishment. Purity violations, however, are often the result of the violation of a moral norm or taboo, and lack identifiable, if any, victim and transgressor. Therefore, post-judgement processes following purity violations may be mediated less by the need for achieving certainty than with harm violations, because there is less, if any, concern about upholding justice and punishing a transgressor.

(2) Assuming differences in *emotion regulation motives* as a function of moral violation type, I would predict that judgements of harm violations, relative to judgements of purity violations, would be associated with a stronger desire for additional information of unspecified valence (indicating an approach tendency towards their emotional reaction). These judgements would not be associated with lower ratings of certainty or greater perceived usefulness of additional information, however, as this hypothesis would simply predict differences in participants' general orientation toward new information. Additionally, I would predict differences in the *kind* as well as amount of information desired following judgements of harm versus purity violations. Judgements of harm violations should also be associated

with a relatively stronger desire for wrongness-*confirming* information; this pattern of information desire would be consistent with a motivation to indulge and approach their feelings of moral outrage and maintain certainty. Judgements of purity violations, in contrast, should be associated with a stronger desire for wrongness-*contradicting* information; this pattern would be consistent with a motivation to escape their feelings of moral disgust. This emotion regulation account would not predict differences in ratings of certainty or greater perceived usefulness of additional information as a function of moral violation type.

3 Experiment 3

Experiment 3 examined post-judgement desire for and perceived utility of additional judgement-relevant information in the context of moral anger and moral disgust. Participants judged the (in)appropriateness of a moral transgression and rated their certainty in their judgement and the intensity of their emotion prior to rating their desire for additional information and the usefulness of additional information.

3.1 Method

3.1.1 Participants and design. One hundred and twenty-one undergraduate psychology students from the University of Birmingham (107 women; $M_{\text{age}} = 19.2$ years, $SD = 2.22$) participated for course credit. Participants were randomly assigned to one of two conditions of a single-factor (violation type: harm, purity) between-participants design.

3.1.2 Materials. The moral violation scenarios were three-line vignettes describing either the moral violation of a woman who is raped while on a date (harm violation) or a brother and sister who have sex (purity violation). The incest scenario was taken from Gutierrez and Giner-Sorolla (2007), with much of the original passage that contained potentially judgement-mitigating information removed, leaving only a brief explanation of who was involved and what took place:

Julie and Mark are brother and sister. They are travelling together in France on a summer vacation from college. One night they are staying alone in a cabin near the beach. They decide it would be interesting and fun if they tried making love.

The date rape scenario was generated to match the incest scenario in length and valence:

Becky is setup on a blind date by her friends with a man named Luke. After going to a house party together, they end up in one of the bedrooms in the house. The next morning, Becky wakes up and goes to the police because she was raped by Luke that night.

3.1.3 Procedure. All materials were presented in an online survey hosted on SurveyMonkey™. Participants completed the online survey remotely without the presence of an experimenter. Participants learned that the experimenter was interested in people's judgements of the moral appropriateness of certain behaviours, as well as people's perceptions of their own judgements (e.g., confidence in judgements, ability to explain judgements, interest in learning additional information relevant to judgements). Participants then read either the harm-violation scenario or the purity-violation scenario, and responded to a series of questions about the scenario.

First, participants responded to the scenario itself. Participants rated the moral appropriateness of the story character's action on a 7-point scale anchored by *not at all appropriate* and *very appropriate*. Participants then rated their own certainty about having made the right judgement of appropriateness on a 7-point scale anchored by *not at all certain* and *very certain*. Finally, participants rated the extent to which they felt the emotion that the scenario was designed to elicit. Specifically, participants who read the date rape scenario rated how angry the scenario made them feel on a 7-point scale anchored by *not at all angry* and *very angry*, whereas participants who read the incest scenario rated how disgusted the

scenario made them feel on a 7-point scale anchored by *not at all disgusted* and *very disgusted*.

Next, participants made three-alternative forced-choice responses to three questions about their interest in additional information about the scenario they had read: “Would you need additional information to be sure that your judgement is correct?”, “Would additional information help you to understand/explain why your judgement is correct?”, and “Would you like additional information about the scenario you just judged?”. For each question, participants indicated whether more information was needed/helpful/desired, did not care either way to have more information, or more information was not needed/helpful/desired.

Finally, participants rated how useful they would find additional information. They provided separate ratings for how useful they would find additional information about the actor’s disposition/mental state at the time of the act, the actor’s personality in general, and the situation the actor was in at the time of the act. Participants made these ratings on a 7-point scale anchored by *not at all useful* and *very useful*.

3.2 Results and Discussion

3.2.1 Moral appropriateness, judgement certainty, and emotion intensity. Ratings of moral appropriateness were coded along a -3 to +3 scale; ratings of judgement certainty and emotion intensity were coded on a 1 to 7 scale. A series of univariate ANOVAs demonstrated that none of the ratings differed as a function of violation type, all $t < 1.01$, all $ps > .31$. As can be seen in Table 3, however, all ratings tended to be very close to the scale endpoints, demonstrating that participants had very strong reactions to both scenarios.

Table 3
Mean Ratings (and Standard Deviations) in Response to Harm versus Purity Violations, Experiment 3

	Harm Violation (n = 59)	Purity Violation (n = 62)
Moral appropriateness	-2.49 (1.28)	-2.68 (1.14)
Judgement certainty	6.22 (1.31)	6.47 (1.40)
Emotion intensity	6.25 (1.21)	6.29 (1.46)

Note. Possible ranges: moral appropriateness, -3 to +3; judgement certainty and emotion intensity, 1 to 7.

3.2.2 Desire for information. The number of participants responding “yes,” “no,” and “don’t care” to the three questions about their interest in additional information as a function of violation type are presented in Table 4. A chi-squared test of significance revealed a significant difference between the harm and purity violation with respect to reported “need” for information, $\chi^2(2, N = 121) = 26.204, p < .001$, reported “want” for information, $\chi^2(2, N = 121) = 23.78, p < .001$, and reported “like” for information, $\chi^2(2, N = 121) = 12.44, p = .002$. To simplify the analysis, and because the pattern of responding was similar across the three questions (suggesting that needing and wanting additional information were treated similarly by our participants), I created a summary score for each participant that reflected the number of “yes,” “no,” and “don’t care” responses (see Table 5) and conducted a series of Bonferroni-corrected *t*-tests.

Table 4***Number of Participants Expressing Desire for Additional Information as a Function of Probe Question and Violation Type, Experiment 3***

	Yes	Don't Care	No
Need more information to be certain of judgement			
Harm violation	31	25	3
Purity violation	8	36	18
Want more information to explain judgement			
Harm violation	36	19	4
Purity violation	12	33	17
Would like more information			
Harm violation	45	12	2
Purity violation	28	27	7

Note. Harm violation, $n = 59$; purity violation, $n = 62$.

Table 5***Mean Number (and Standard Deviations) of Responses to Probes Regarding Desire for Additional Information as a Function of Response Type and Violation Type, Experiment 3***

	Harm Violation ($n = 59$)	Purity Violation ($n = 62$)
Yes	1.90 (1.20)	0.77 (0.86)
Don't care	0.95 (1.12)	1.55 (0.95)
No	0.15 (0.45)	0.68 (0.95)

Note. Possible range = 0 to 3.

I found that participants were more likely to respond “yes,” that they desired more information, in the harm than purity violation condition, $t(119) = 5.95$, $p < .001$, $d = 1.08$, and to respond “no,” that they did not desire more information, in the purity than harm violation condition, $t(119) = 3.84$, $p < .001$, $d = 0.71$. Although not predicted, I also found that participants were more likely to respond that they did not care either way about the additional information in the purity than harm violation condition, $t(119) = 3.17$, $p = .002$, $d = .58$. Given the strength of participants’ reactions to the scenarios, this apparent lack of interest is surprising—but perhaps, rather than lack of interest, this response reflects another form of disengagement (i.e., refusing to even consider whether additional information is desired).

3.2.3 Usefulness of information. Ratings of information usefulness were coded along a 1 to 7 scale. Collapsing across the three types of information (actor mental state, actor personality, actor situation; Cronbach's alpha = .87), I found that participants rated the additional information to be equal in usefulness for harm and purity violations ($M_s = 4.70$ and 4.55 , respectively), $t(119) = 0.55$, $p = .59$, $d = .10$. In general, participants rated additional information to be only moderately useful (as evidenced by mean ratings near the scale midpoint). In both conditions, participants were less interested in character information than either mental state or situational information. Although interesting, this pattern did not interact with the other factors and was not directly relevant to the questions of interest, and so I elected to remove the information type factor from the analysis.

3.2.4 Summary. In this experiment, participants rated the harm and purity violations as equally—and very—morally inappropriate. They also rated their judgement certainty and respective emotional response intensity (anger or disgust) as equally (and, again, very) high. Their reactions to the availability of additional violation-relevant information, however, showed a noteworthy difference: Although participants who were exposed to harm versus purity violations did not differ in how useful they expected additional information would be, participants exposed to harm violations expressed a desire for additional information—and participants exposed to purity violations expressed a lack of desire to read additional information. As participants in the harm- versus purity-violation conditions did not differ in their moral judgements or their views on how useful additional information would be, then the greater desire for additional information among participants in the harm-violation condition cannot be attributed to an epistemic certainty motive. Instead, I argue, the differences in desire for information derive from the need to maintain and bolster their certainty rather than to achieve epistemic certainty in the first place.

4 Experiment 4

The goal of Experiment 4 was to further explore the role of certainty maintenance and emotion regulation motives with regards to post-judgement desire for different types of additional information. I presented participants with the same moral violations as in Experiment 3. Instead of asking participants whether they desired additional information of an unspecified nature, however, I told them that they were going to be required to read additional information and asked them how much they would desire a specific type of additional information. Importantly, for half of the participants, I asked about their desire for wrongness-confirming information, and for the remaining participants, I asked about their desire for wrongness-contradicting information.

For participants exposed to the harm violation, the emotion regulation hypothesis would predict a greater interest in wrongness-confirming than -contradicting information; seeking such information would be consistent with a motivation to approach the harm violation and moral anger. For participants exposed to the purity violation, the emotion regulation hypothesis would predict a greater interest in wrongness-contradicting information and a negative correlation; seeking such information would be consistent with a motivation to escape the purity violation and moral disgust.

I also examined the relationships between participants' ratings of how desirable wrongness-confirming and -contradicting information would be and their reactions to the moral violations. For participants given the opportunity to learn additional wrongness-*confirming* information, the emotion regulation hypothesis would predict a positive relationship between reactions to the moral violation and interest in learning additional wrongness-confirming information for participants in the harm-violation condition, but a negative relationship for participants in the purity-violation condition. The certainty

maintenance hypothesis would suggest that to the extent that participants were certain of their judgement in response to the moral violation, they should embrace additional information to boost or maintain their certainty.

For participants given the opportunity to learn additional wrongness-*contradicting* information, the certainty maintenance hypothesis would predict that participants should generally resist additional information that might change their judgement and threaten their certainty. The emotion regulation hypothesis would instead predict that participants in the *purity*-violation condition, when reporting feelings of disgust towards the moral violation, should be motivated to embrace additional information that had the potential to change their judgement, thus repair their mood, and allow them to avoid further engagement with the moral violation. Under this hypothesis, though, participants in the *harm*-violation condition should be motivated to reject additional information that has the potential to change their condemning judgement, in order to maintain their feelings of moral outrage.

4.1 Method

4.1.1 Participants and design. Eighty undergraduate psychology students from the University of Birmingham (78 female; $M_{\text{age}} = 19.4$ years) participated for course credit. Participants were randomly assigned to one of four conditions in a 2 (violation type: harm, purity) \times 2 (information congruity: wrongness-confirming, wrongness-contradicting) between-participants design.

4.1.2 Materials and procedure. The moral violation scenarios used in Experiment 4 were identical to Experiment 3. All materials were presented in an online survey hosted on SurveyMonkeyTM. Participants read a scenario depicting either a harm violation (date rape) or a purity violation (incest). As in Experiment 3, participants made judgements of appropriateness (this time, right/wrong), judgements of their own certainty about the

appropriateness, and emotion intensity judgements (again for the target emotion only); all ratings were made along 7-point scales.

Participants then responded to three questions probing their desire for additional information; depending on random assignment, they were asked about their desire for either confirming or contradicting information. Participants were asked, “If you had to read additional information, how much would you like to read information that confirms [contradicts] (1) what you think about the actor’s disposition/mental state at the time of the act, (2) what you think about the actor’s personality in general, and (3) what you think about the situation the actor was in at the time of the act?” Participants made their ratings on a 7-point scale anchored by *not at all* and *very much*.

4.2 Results and Discussion

4.2.1 Moral appropriateness, judgement certainty, and emotion intensity.

Replicating the effects of Experiment 3, ratings of moral appropriateness (right/wrong) and emotion intensity did not differ as a function of violation type, both $t < 1.09$, both $p > .28$ (see Table 6). Unlike in Experiment 3, however, ratings of judgement certainty were higher in the purity-violation condition than in the harm-violation condition, $t(78) = 2.40$, $p = .02$, $d = .54$.

Table 6
Mean Ratings (and Standard Deviations) in Response to Harm versus Purity Violations, Experiment 4

	Harm Violation (n = 39)	Purity Violation (n = 41)
Right/wrong	-2.74 (0.44)	-2.59 (0.81)
Judgement certainty	5.90 (1.21)	6.27 (1.34)
Emotion intensity	5.64 (1.65)	5.83 (1.67)

Note. Possible ranges: right/wrong, -3 to +3; judgement certainty and emotion intensity, 1 to 7.

4.2.2 Desire for additional information. Participants' ratings of how much they wanted additional information were coded along a 1 to 7 scale. I collapsed across the three types of information (actor mental state, actor personality, actor situation; Cronbach's alpha = .78) and subjected the scores to a 2 (violation type: harm, purity) \times 2 (information congruity: wrongness-confirming, wrongness-contradicting) between-participants ANOVA. The analysis yielded only a Violation Type \times Information Congruity interaction, $F(1, 76) = 4.47, p = .04, \eta_p^2 = .06$. As depicted in Figure 2, participants in the harm-violation conditions perceived confirming and contradicting information as equally desirable, $t(37) = 0.62, p = .75, d = .20$, whereas participants in the purity-violation condition perceived contradicting information to be more desirable than confirming information, $t(39) = 2.50, p = .02, d = .78$.

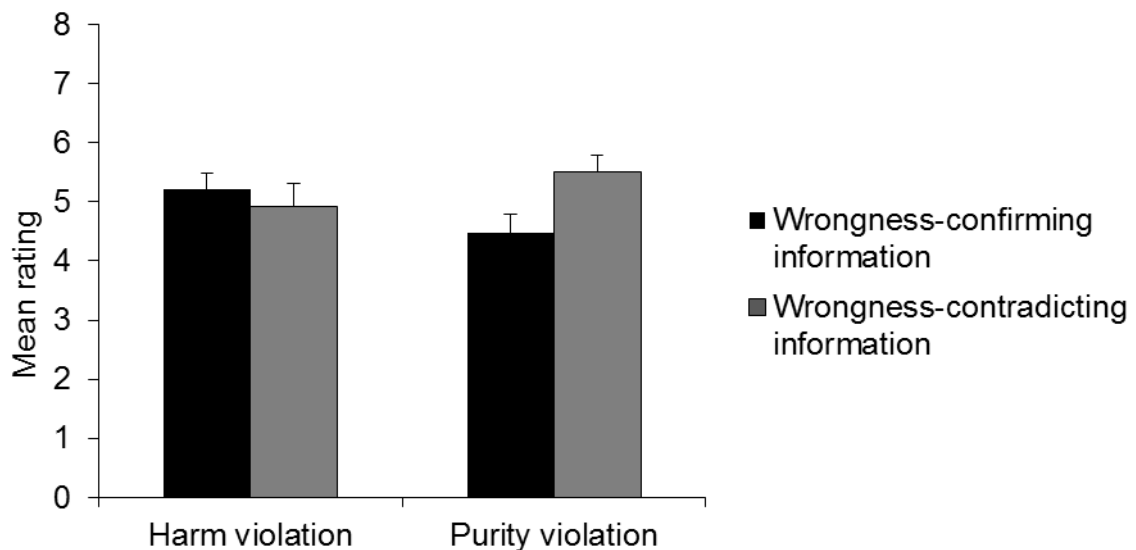


Figure 2. Mean ratings of desire for wrongness-confirming and wrongness-contradicting information as a function of violation type, Experiment 4.

Note. Error bars represent standard error.

4.2.3 Functions of post-judgement reasoning. To explore further the possible functions served by post-judgement processing, I calculated correlations between participants' ratings of how desirable additional information would be with their reactions to the moral scenarios; this was done separately for each Violation Type \times Information Congruity condition. As depicted in Table 7, the analysis yielded no reliable correlations in the contradictory-information conditions. In the confirming-information conditions, however, the analysis revealed differences between the harm- and purity-violation conditions: To the extent that participants in the harm-violation condition were certain of their judgement, they expressed a desire to read *more* wrongness-confirming information; to the extent that participants in the purity-violation condition were disgusted by the scenario, however, they expressed a desire to read *less* wrongness-confirming information.

Table 7
Correlations between Desire for Information and Responses to the Moral Violations as a Function of Violation Type and Information Congruity, Experiment 4

	Harm Violation		Purity Violation	
	Confirming (n = 20)	Contradicting (n = 19)	Confirming (n = 21)	Contradicting (n = 20)
Right/wrong	-.041 (.868)	.073 (.758)	.262 (.264)	.213 (.354)
Judgement certainty	.592 (.008)	-.187 (.430)	-.045 (.851)	-.083 (.722)
Emotion intensity	.214 (.379)	-.233 (.322)	-.444 (.050)	-.249 (.277)

Notes. Two-tailed correlations. Numbers in parentheses are *p*-values.

4.2.4 Summary. As in Experiment 3, participants' ratings of moral appropriateness and emotional response intensity did not differ as a function of whether they were exposed to a harm or purity violation. In Experiment 4, however, participants expressed less judgement certainty in response to harm than purity violations. Moreover, participants who read about harm violations reported equivalent interest in additional information regardless of whether that information was presented as wrongness-confirming or -contradicting. Information search about the harm violation, rather than being negatively correlated with prior certainty, was actually positively correlated with it; to the extent that participants who read about harm violations were *certain* of their judgement, they wanted *more* wrongness-confirming information. This argues against epistemic certainty motives as an explanation for the bias. Information search, rather than being directed toward (re-)establishing absent certainty, was most directed toward bolstering certainty when certainty was already present. Also, participants who read about purity violations reported greater interest in additional information when that information was presented as wrongness-contradicting rather than wrongness-confirming, supporting the emotion regulation account over the epistemic certainty account, which would have only predicted interest in wrongness-confirming information. This may stem from the presence of disgust; participants might have been

expressing a desire to alleviate their feelings of disgust as well as to avoid the situation. Moreover, to the extent that participants who read about purity violations were disgusted by the scenario, they also wanted *less* wrongness-confirming information, suggesting a desire to escape moral disgust.

5 Experiment 5

Experiments 3 and 4 examined people's self-reported desire and need for additional behavioural and contextual information following exposure to either harm or purity violations. The findings from Experiments 3 and 4 suggest overall that people report a greater desire for additional information following exposure to a harm violation compared to a purity violation. In Experiment 5, I aimed to investigate participants' actual engagement with additional information rather than their self-reported engagement with additional information. I predicted that participants who were exposed to harm violations information would engage with more information when given wrongness-confirming than wrongness-contradicting information, because they would be motivated to conserve and boost their judgement certainty. Additionally, I predicted that participants who were exposed to purity violations would engage with information less overall, especially with wrongness-confirming information, reflecting an effort to alleviate and avoid their emotional experience of disgust.

5.1 Method

5.1.1 Participants and design. Eighty-five undergraduate psychology students from the University of Birmingham (73 women; $M_{\text{age}} = 18.8$ years, $SD = 1.26$) participated for course credit. Participants were randomly assigned to one of four conditions of a 2 (violation type: harm, purity) \times 2 (information congruency: wrongness-confirming, wrongness-contradicting) between-participants design.

5.1.2 Materials.

5.1.2.1 Moral violations. The moral violations were short vignettes describing one of four moral violations: two purity violation scenarios and two harm violation scenarios. The harm violation scenarios either depicted a woman raped on a date or a woman being permanently tattooed by her boyfriend while unconscious. The purity violation scenarios either depicted a brother and sister having sex or a man consuming the meat from his dead dog. The incest scenario was taken from Experiments 3 and 4:

Julie and Mark are brother and sister. They are travelling together in France on a summer vacation. They are staying in a cabin near the beach. One night, Julie and Mark have sex.

The dead dog scenario was modified from a scenario used to elicit feelings of disgust in a harmless-outcome taboo violation from Haidt, Koller, and Dias (1993):

Alex's family dog is dead and buried in the garden. One day, Alex digs up the dead family dog, cooks it on the barbecue and eats it for dinner.

The date rape scenario was generated to match the incest scenario in length and valence, and depicted a harmful act involving both a transgressor and victim:

Becky is setup on a blind date with a man named Luke. They go to a house party together for their date. The next day, Becky wakes up and goes to the police because Luke raped her.

The tattoo scenario was modified from a scenario used to create socio-moral and bodily moral violations from Gutierrez, Giner-Sorolla, and Vasiljevic (2012):

Mike and his girlfriend Sarah go on a trip. Mike thinks Sarah would look really good with a tattoo. One night when Sarah's unconscious, Mike gets a permanent tattoo put on her body.

5.1.2.2 Additional information. Thirty lines of additional information were generated for each moral violation scenario (see Appendix D). The information was designed to either confirm or contradict a judgement of wrongness, depending on the condition. All lines of information for all moral violation scenarios included equal amounts of information about the character and the situation depicted in the scenario. Lines of information were presented one-at-a-time, line-by-line, with one line per screenshot. Lines of information were matched for length and presented in a pseudo-random order.

5.1.3 Procedure. All materials were presented in a laboratory on a computer in the program MediaLab™ (Empirisoft, 2006). Participants completed the experiment individually in a laboratory cubicle without the presence of an experimenter. Participants learned that the experimenter was interested in people's judgements of the moral appropriateness of certain behaviours, people's perceptions of their own judgements (e.g., confidence in judgements, ability to explain judgements, interest in learning additional information relevant to judgements), and emotional reactions to different moral violations. Participants in all conditions read a scenario (either the date rape or incest scenario), answered follow-up questions, and then read as many lines of either wrongness-contradicting or wrongness-confirming information as they wished. The order of the information that was presented in each condition was fixed for reasons that I address in the results section. Participants were asked to read as many lines of information as they felt they should in order to make a final judgement. They then completed the Need for Cognition scale (Cacioppo, Petty, Feinstein, & Jarvis, 1996) as a filler task before reading a second scenario (either the tattoo or dead dog scenario, always the same violation type as for Set 1) and completing the same measures as for the first scenario.

For each scenario, participants first rated the wrongness of the story character(s)'s action on a 7-point scale anchored by *not at all wrong* and *very wrong*. Participants then rated their own certainty about having made the right judgement of wrongness on a 7-point scale anchored by *not at all certain* and *very certain*. Finally, participants rated the extent to which they felt angry, disgusted, frustrated, sad and afraid in response to the scenario they read. These ratings were all made on 7-point scales anchored by *not at all [emotion]* and *very [emotion]*, and were presented in a randomised order.

Participants then rated their emotional response on a pencil-paper questionnaire taken from Gutierrez et al. (2011). Each participant saw two sets of photographs (see Figure 3) from the MFSDE set of emotional facial expressions (Beaupré, Cheung, & Hess, 2000). One set of facial expressions expressed anger and the other set expressed disgust. Participants were first asked to "Select one set of faces that best describes your feelings about the story." Participants were then asked to indicate for both the anger and disgust faces, "Now, select which set of faces best describes your feelings towards (insert character(s) name) in the story. Once you've selected the set of faces that best describes how you feel, rate BOTH sets for how much you feel the emotion depicted in the faces". Participants' ratings were made on a 9-point scale anchored by *not at all* and *extremely*.


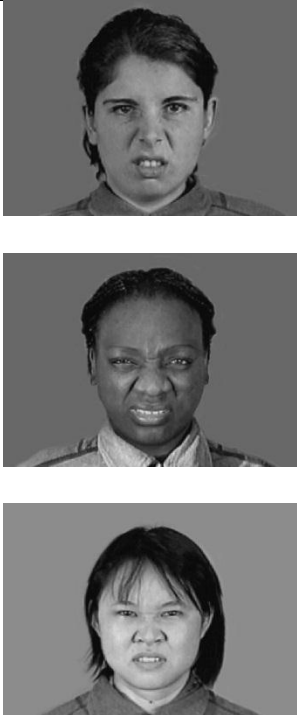
 <p>I feel more like this.</p> <p><input type="checkbox"/></p>	 <p>I feel more like this.</p> <p><input type="checkbox"/></p>
<p>How much do you feel this way towards (insert character(s) name)?</p> <p>1 2 3 4 5 6 7 8 9</p> <p>1 = not at all; 9 = extremely</p>	<p>How much do you feel this way towards (insert character(s) name)?</p> <p>1 2 3 4 5 6 7 8 9</p> <p>1 = not at all; 9 = extremely</p>

Figure 3. Sample facial expressions taken from (Beaupré, Cheung, & Hess, 2000) used in emotion measure questionnaire, Experiments 5 and 6.

Next, participants learned that they would read additional information about the character and the situation that occurred in the scenario. Participants were instructed to read each line of information until they felt "... they had enough information and know enough about the situation to make a good judgement". Participants were instructed to press the "continue" button on the bottom right-hand side of the screen if they wanted to read an additional line of information, and to press the button at the bottom of the screen labelled "escape – I have had enough information" followed by the "continue" button on the bottom right-hand side of the screen when they had enough information. When participants pressed the "continue" button, a new screen with a different line of information would appear with the same response options as before (continue or escape). When participants chose to "escape" from a line of information, the program redirected them to a set of final questions. Participants re-rated the wrongness of the story character(s)'s action on a 7-point scale anchored by *not at all wrong* and *very wrong*, and then re-rated their own certainty about having made the right judgement of wrongness on a 7-point scale anchored by *not at all certain* and *very certain*.

5.2 Results and Discussion

The harm and purity violation scenarios from Experiments 3 and 4 (incest and date rape) were analysed separately from the new harm and purity violation scenarios (dead dog and tattoo). The reasoning behind having separate analyses was to be able to further test our hypotheses from Experiments 3 and 4 by having identical stimuli while at the same time introducing new stimuli to determine whether findings from Experiments 3 and 4 could be replicated with new scenarios. Separate analyses meant that the responses to new scenarios would not affect how the new measures added in this experiment (i.e., amount of information read) influenced responses to the original scenarios that were used in Experiments 3 and 4. I

will present the main results that directly addressed my hypotheses, along with further analyses of manipulation checks and unexpected results by set (i.e., the stories that were used in Experiments 3 and 4 (set 1), and the new stories that were used only in experiments 5 and 6 (set 2).

5.2.1 Set 1 violations. Set 1 violations depicted either a woman raped on a date (harm) or siblings having sex (purity).

5.2.1.1 Engagement with additional information. The number of lines of information that participants chose to read were analysed in a 2 (violation type: harm, purity) \times 2 (information congruency: wrongness-confirming, wrongness-contradicting) mixed-model ANOVA. The analysis revealed no significant main or interaction effects, all $p > .25$ (see Table 8):

Table 8
Mean Number of Lines of Additional Information Read (and Standard Errors) Depending on the Information Congruency Depending and Violation Type for Set 1, Experiment 5

	Harm violation (n=43)	Purity violation (n=42)
Wrongness-confirming information	23.20 (1.81)	22.71 (1.81)
Wrongness-contradicting information	26.50 (1.77)	23.57 (1.81)

Note. Possible ranges 1 to 7.

5.2.1.2 Moral wrongness. Ratings of moral wrongness were coded along a 1 to 7 scale. A 2 (violation type: harm, purity) \times 2 (information congruency: wrongness-confirming, wrongness-contradicting) \times 2 (time: pre-information, post-information) mixed-model ANOVA revealed a significant main effect of time, $F(1, 85) = 12.84, p = .001, \eta^2_p = .14$, such that judgements of wrongness were significantly higher following versus preceding exposure to additional information ($M = 6.31, SE = .16$, versus $M = 5.68, SE = .12$, respectively). There was also a main effect of information congruency, $F(1, 85) = 36.100, p < .001, \eta^2_p = .31$, such that judgements of wrongness were higher overall for those participants in the wrongness-

confirming condition than the wrongness-contradicting condition ($M = 6.66$, $SE = .16$, versus $M = 5.33$, $SE = .16$, respectively).

The time of judgement interacted with the violation type, $F(1, 85) = 29.79$, $p < .001$, $\eta^2_p = .27$ (see Figure 4).




Figure 4. Interaction between the times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and the violation type for Set 1, Experiment 5.

Post-hoc comparisons revealed an unexpected significant effect of time of judgement in the harm-violation condition, $t(83) = 2.40$, $p = .019$, such that judgements of wrongness for the harm violation decreased from initial to final judgement.

The analysis also revealed a significant interaction between time and information congruency, $F(1, 85) = 31.27$, $p < .001$, $\eta^2_p = .28$ (see Figure 5).

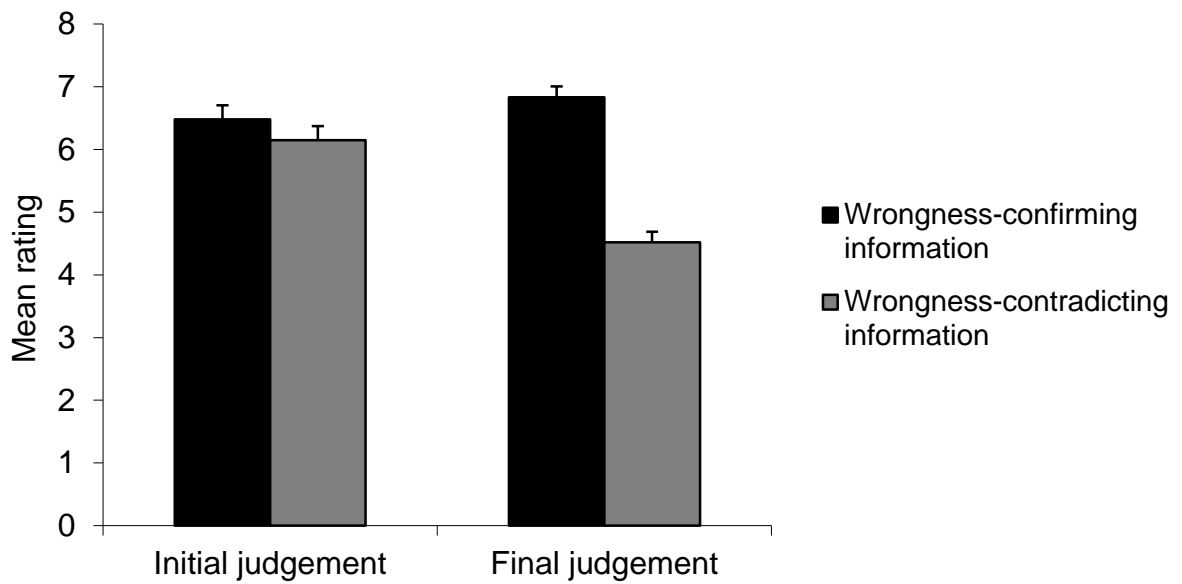


Figure 5. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 1, Experiment 5.

Post-hoc comparisons revealed only a significant effect of information congruency on the final judgement of wrongness, $t(83) = 7.72, p < .001$, such that final judgements of wrongness were higher after receiving wrongness-confirming information than wrongness-contradicting information.

The analysis also revealed a significant interaction between time, violation type and information congruency, $F(1, 85) = 22.93, p < .001, \eta^2_p = .22$ (see Figure 6). The three-way interaction was broken down into separate 2 (information congruency) \times 2 (time) mixed-model ANOVAs for harm and purity violations.

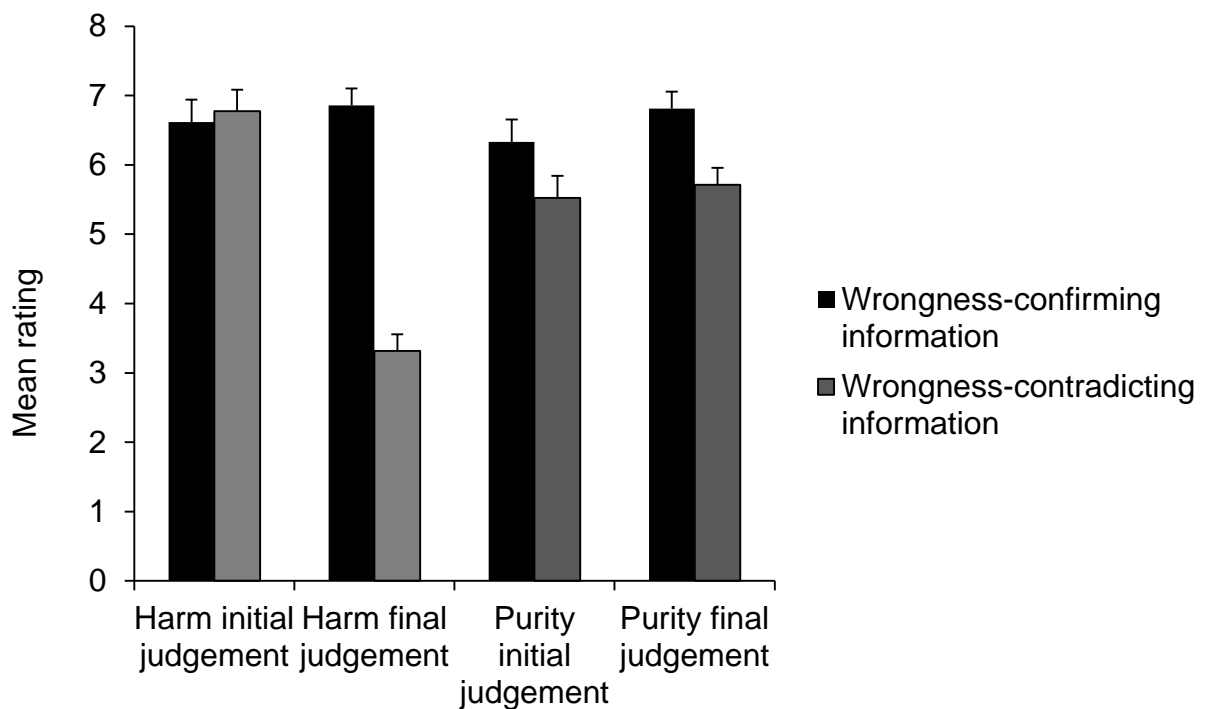


Figure 6. Interaction between times of judgement (initial wrongness-judgement vs. final wrongness judgement post-information), violation type (harm vs. purity) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 1, Experiment 5.

For the harm violation, the 2 (information congruency) \times 2 (time) mixed-model ANOVA revealed a significant main effect of time, $F(1, 43) = 74.39, p < .001, \eta^2_p = .65$ (initial judgement $M = 6.70, SE = .13$, final judgement $M = 5.09, SE = .17$) and a significant main effect of information congruency, $F(1, 43) = 53.109, p < .001, \eta^2_p = .56$ (wrongness-confirming information $M = 6.74, SE = .17$, wrongness-contradicting information $M = 5.05, SE = .16$). There was also a significant interaction between time and information congruency, $F(1, 43) = 98.045, p < .001, \eta^2_p = .705$.

This interaction was broken down as a manipulation check to determine whether the post-judgement information that was presented adequately affected judgements following the

presentation of information. I would expect that final judgements of wrongness would be lower than initial judgements of wrongness following exposure to wrongness-contradicting information. Likewise, I would expect that final judgements of wrongness would be higher than initial judgements of wrongness following exposure to wrongness-confirming information. Paired samples *t*-tests revealed a significant difference between initial and final judgements of wrongness after exposure to wrongness-contradicting information, $t(21) = 10.78$, $p < .001$; judgements of wrongness were lower after exposure to wrongness-contradicting information, but not as a function of exposure to wrongness-confirming information, $p = .204$.

For the purity violation, the 2 (information congruency) \times 2 (time) mixed-model ANOVA revealed only a main effect of information congruency, $F(1, 42) = 6.39$, $p = .016$, $\eta^2_p = .14$ (wrongness-confirming information $M = 6.57$, $SE = .27$, wrongness-contradicting information $M = 5.62$, $SE = .27$). Unexpectedly, exposure to additional information had no effect (main effect for time, $p = .28$). This result supports previous research suggesting that responses to purity violations are more inflexible than responses to harm violations, potentially as a result of the elicited experience of disgust in purity violations (Russell & Giner-Sorolla, 2011a).

5.2.1.3 Judgement certainty. Ratings of judgement certainty were coded along a 1 to 7 scale. The analyses reported here are to further check the manipulation of post-judgement information; wrongness-contradicting information should decrease participants' final judgement certainty. A 2 (violation type: harm, purity) \times 2 (information congruency: wrongness-confirming, wrongness-contradicting) \times 2 (time: pre-information, post-information) mixed-model ANOVA revealed a marginally significant main effect of time, $F(1, 85) = 3.63$, $p = .060$, $\eta^2_p = .043$. Certainty ratings were higher overall when made prior to

receiving additional information than after receiving additional information ($M = 6.06$, $SE = .15$, versus $M = 5.78$, $SE = .15$, respectively). There was a significant main effect of information congruency, $F(1, 85) = 9.88$, $p = .002$, $\eta^2_p = .11$, such that certainty ratings were higher after wrongness-confirming information than after wrongness-contradicting information ($M = 6.33$, $SE = .18$, versus $M = 5.51$, $SE = .18$, respectively).

The analysis also revealed a significant interaction between time and information congruency, $F(1, 85) = 23.23$, $p < .001$, $\eta^2_p = .22$ (see Figure 7).




Figure 7. Interaction between times of judgement (initial certainty judgement vs. final certainty judgement post-information) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 1, Experiment 5.

Independent samples t -tests revealed only a significant difference between final ratings of certainty based on information congruency, $t(83) = 4.95$, $p < .001$, such that final ratings of certainty were higher after exposure to wrongness-confirming versus wrongness-contradicting information ($M = 6.56$, $SE = .16$, versus $M = 5.00$, $SE = .26$, respectively).

5.2.1.4 Emotion ratings. Ratings of emotion were coded along a 1 to 7 scale. The analyses reported here were to determine whether the violations elicited the predicted emotion (i.e., harm eliciting anger and purity eliciting disgust) in contrast to the emotion counterpart (anger vs. disgust) and in contrast to other control emotions (e.g., sadness). A 2 (Violation Type: harm, purity) \times 5 (Emotion Type: fear, anger, disgust, frustration and sadness) mixed-model ANOVA revealed a significant main effect of emotion, $F(1, 85) = 51.037, p < .001, \eta^2_p = .38$ (fear $M = 3.17, SE = .19$, anger $M = 4.01, SE = .19$, disgust $M = 5.56, SE = .19$, frustration $M = 3.64, SE = .21$, sadness $M = 4.03, SE = .20$). There was also a main effect of violation type, $F(1, 85) = 28.400, p < .001, \eta^2_p = .26$, such that emotion intensity ratings overall were higher in the harm violation condition versus the purity violation condition ($M = 4.93, SE = .22$, versus $M = 3.24, SE = .23$, respectively).

There was a significant interaction between emotion and violation type, $F(1, 85) = 6.96, p < .001, \eta^2_p = .077$ (see Table 9). Independent samples t -tests revealed a significant difference in ratings of anger based on the violation type, $t(83) = 6.15, p < .001$, such that participants' ratings of anger were higher in the harm violation condition versus the purity violation condition ($M = 5.19, SE = .26$, versus $M = 2.83, SE = .28$, respectively). Ratings of disgust did not differ significantly on the basis of violation type, $p = .089$. Ratings of frustration were significantly different based on the violation type, $t(83) = 4.006, p < .001$, such that participants' ratings of frustration were higher in the harm than purity violation condition ($M = 4.47, SE = .29$, versus $M = 2.81, SE = .29$, respectively). Ratings of sadness were significantly different based on violation type, $t(83) = 5.55, p < .001$, such that participants were more sad in the harm than purity violation condition ($M = 5.12, SE = .26$, versus $M = 2.95, SE = .29$, respectively).

Table 9***Mean Ratings (and Standard Errors) of Emotion Depending on Violation Type for Set 1, Experiment 5***

	Harm violation (n=43)	Purity violation (n=42)
Fear	3.98 (.27)	2.36 (.27)
Anger	5.19 (.27)	2.83 (.27)
Disgust	5.88 (.26)	5.24 (.27)
Frustration	4.47 (.29)	2.81 (.29)
Sadness	5.12 (.27)	2.95 (.28)

Note. Possible ranges 1 to 7.

5.2.1.5 Facial expression ratings. The number of participants who selected the angry versus disgust expression as best representing their reaction to the scenario they read was analysed using a 2 (violation type: harm, purity) \times 2 (facial expression: anger, disgust) chi-square test. The test revealed no significant differences between the number of participants who selected the angry facial expression versus the disgusted facial expression as best representing their reaction to either the harm or purity violation, $\chi^2 = .331$, $df = 1$, $p = .56$.

Facial expression ratings were coded along a 1 to 9 scale. A 2 (violation type: harm, purity) \times 2 (facial expression: angry face, disgusted face) repeated measures ANOVA revealed no main effect of facial expression, $p = .41$, and no main effect of violation type, $p = .73$.

There was a significant Violation Type \times Facial Expression interaction, $F(1, 85) = 4.24$, $p = .043$, $\eta^2_p = .049$. Paired samples t -tests revealed that ratings did not differ as a function of facial expression in the harm violation condition, $p = .37$. There was a borderline difference between ratings as a function of facial expression in the purity violation condition, $t(41) = 1.98$, $p = .054$, such that participants' ratings of the extent to which their emotional reaction matched the disgusted face were higher than the extent to which their emotional

reactions matched the angry face ($M = 6.60$, $SE = .42$, versus $M = 5.40$, $SE = .40$, respectively).

5.2.2 Set 2 violations. Set 2 violations depicted either a man tattooing his girlfriend (harm) or a man eating his dead dog (purity).

5.2.2.1 Engagement with additional information. The number of lines of additional information read by participants (see Table 10) were analysed in a 2 (violation type: harm, purity) \times 2 (information congruency: wrongness-confirming, wrongness-contradicting) mixed-model ANOVA. The analysis revealed only a significant main effect of information congruency, $F(1, 85) = 4.77$, $p = .032$, $\eta^2_p = .056$, such that participants read more additional information when the information provided to them was wrongness-contradicting versus wrongness-confirming ($M = 24.79$, $SE = 1.34$, versus $M = 20.62$, $SE = 1.36$, respectively). There were no effects involving violation type, all $p > .95$.

Table 10
Mean Number of Lines of Additional Information Read (and Standard Errors) Depending on the Information Congruency Depending and Violation Type for Set 2, Experiment 5

	Harm violation (n=43)	Purity violation (n=42)
Wrongness-confirming information	20.57 (1.92)	20.67 (1.92)
Wrongness-contradicting information	24.73 (1.88)	24.86 (1.92)

5.2.2.2 Moral wrongness. Ratings of moral wrongness were coded along a 1 to 7 scale. A 2 (violation type: harm, purity) \times 2 (information congruency: wrongness-confirming, wrongness-contradicting) \times 2 (time: pre-information, post-information) mixed-model ANOVA revealed a significant main effect of time, $F(1, 85) = 69.35$, $p < .001$, $\eta^2_p = .46$, such that judgements of wrongness were significantly higher preceding versus following exposure to additional information ($M = 6.43$, $SE = .103$, versus $M = 5.57$, $SE = .104$, respectively). There was also a main effect of information congruency, $F(1, 85) = 61.89$, $p < .001$, $\eta^2_p = .43$,

such that judgements of wrongness were higher for those participants in the wrongness-confirming condition than the wrongness-contradicting condition ($M = 6.702$, $SE = .13$, versus $M = 5.29$, $SE = .13$, respectively).

Time of judgement interacted with violation type, $F(1, 85) = 17.79$, $p < .001$, $\eta^2_p = .18$ (see Figure 9).

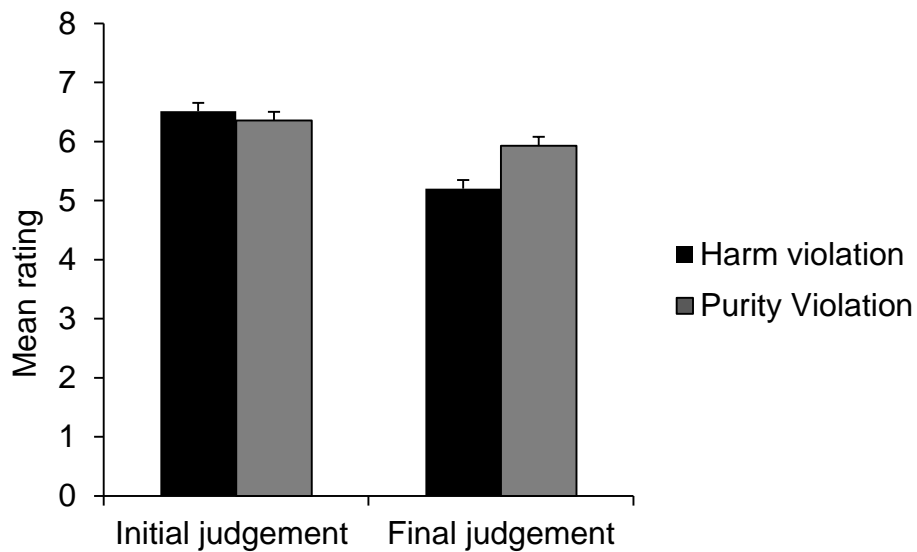


Figure 9. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and violation type (harm vs. purity) for Set 2, Experiment 5.

Post-hoc comparisons revealed that in the harm violation condition, judgements of wrongness were higher at the initial stage than the final stage, $t(42) = 4.25$, $p < .001$ ($M = 6.53$, $SE = .107$, versus $M = 5.16$, $SE = .307$, respectively). In the purity violation condition, judgements of wrongness also differed significantly between initial judgement and final judgement, $t(41) = 2.508$, $p = .016$, such that judgements were higher at the initial stage than the final stage ($M = 6.36$, $SE = .18$, versus $M = 5.93$, $SE = .19$, respectively).

The analysis also revealed a significant interaction between time and information congruency, $F(1, 85) = 68.87, p < .001, \eta^2_p = .65$ (see Figure 10).

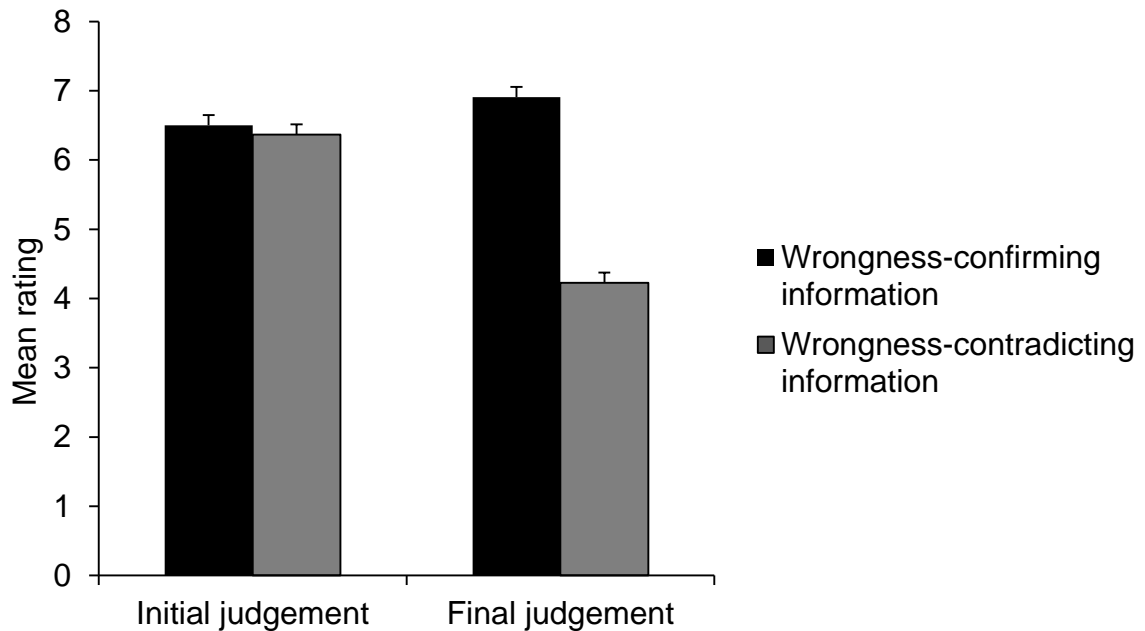


Figure 10. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information) and information congruency (wrongness-confirming vs. wrongness contradicting) for Set 2, Experiment 5.

This interaction depicts a manipulation check to determine whether post-judgement information affected judgements depending on the type of information presented. I expected that final judgements of wrongness would be lower than initial judgements of wrongness following the presentation of wrongness-contradicting information. Post-hoc comparisons revealed only a significant effect of information congruency on the final judgement of wrongness, $t(83) = 11.51, p < .001$, such that final judgements of wrongness were higher after receiving wrongness-confirming information than wrongness-contradicting information.

The analysis also revealed a significant interaction between time, violation type and information congruency that further checked the manipulation of the two types of information

(wrongness-confirming, wrongness-contradicting), $F(1, 85) = 19.77, p < .001, \eta^2_p = .20$ (see Figure 11). The three-way interaction was broken down into separate 2 (information congruency) \times 2 (time) mixed-model ANOVAs for harm and purity violations.

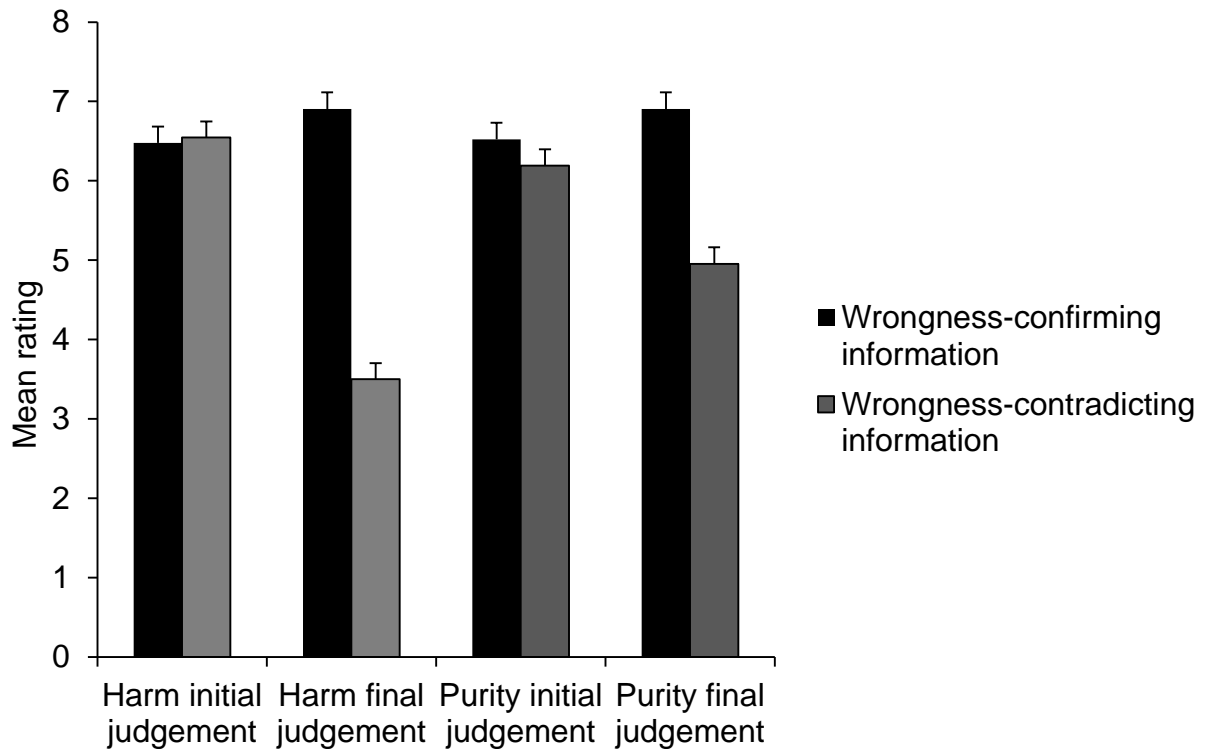


Figure 11. Interaction between times of judgement (initial wrongness judgement vs. final wrongness judgement post-information), violation type (harm vs. purity) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 2, Experiment 5.

For harm violations, the 2 (information congruency) \times 2 (time) ANOVA revealed a significant main effect of time, $F(1, 43) = 57.904, p < .001, \eta^2_p = .59$ (initial judgement $M = 6.51, SE = .11$, final judgement $M = 5.20, SE = .16$), and a significant main effect of information congruency, $F(1, 43) = 61.33, p < .001, \eta^2_p = .60$ (wrongness-confirming information $M = 6.70, SE = .15$, wrongness-contradicting information $M = 5.02, SE = .15$).

There was also a significant interaction between time and information congruency, $F(1, 43) = 102.05, p < .001, \eta^2_p = .71$.

Paired samples t -tests revealed that judgements of wrongness were lower after exposure to wrongness-contradicting information, $t(21) = 9.98, p < .001$. ($M = 3.500, SE = .307$, versus $M = 6.55, SE = .14$, respectively). Judgements of wrongness were also higher after exposure to wrongness-confirming than -contradicting information, $t(20) = 2.905, p = .009$.

For the purity violation condition, the 2 (information congruency) \times 2 (time) ANOVA revealed a significant main effect of time, $F(1, 42) = 13.56, p = .001, \eta^2_p = .25$ (initial judgement $M = 6.36, SE = .18$, final judgement $M = 5.93, SE = .13$). There was a significant main effect of information congruency, $F(1, 42) = 15.706, p < .001, \eta^2_p = .28$ (wrongness-confirming information $M = 6.71, SE = .204$, wrongness-contradicting information $M = 5.57, SE = .204$).

There was also a significant interaction between time and information congruency, $F(1, 42) = 48.37, p < .001, \eta^2_p = .55$.

Paired samples t -tests revealed a significant difference between initial and final judgements after receiving wrongness-contradicting information, $t(20) = 6.38, p < .001$, such that judgements of wrongness were lower after exposure to wrongness-contradicting information ($M = 4.95, SE = .25$, versus $M = 6.19, SE = .31$, respectively). Additionally, there was a significant difference between initial and final judgements after receiving wrongness-confirming information, $t(20) = 2.96, p = .008$, such that judgements of wrongness were higher after exposure to wrongness-confirming information ($M = 6.91, SE = .066$, versus $M = 6.52, SE = .16$, respectively).

5.2.2.3 Judgement certainty. Ratings of judgement certainty were coded along a 1 to 7 scale. A 2 (violation type: harm, purity) \times 2 (information congruency: wrongness-confirming, wrongness-contradicting) \times 2 (time: pre-information, post-information) mixed-model ANOVA revealed a significant main effect of information congruency, $F(1, 85) = 18.97, p < .001, \eta^2_p = .19$ (wrongness-confirming information $M = 6.39, SE = .15$, wrongness-contradicting information $M = 5.50, SE = .15$).

There was a significant interaction between time and information congruency, $F(1, 85) = 38.13, p < .001, \eta^2_p = .32$ (see Figure 12).




Figure 12. Interaction between times of judgement (initial certainty judgement vs. final certainty judgement post-information) and information congruency (wrongness-confirming vs. wrongness-contradicting) for Set 2, Experiment 5.

This interaction was explored to further check the manipulation of post-judgement information type. Post-hoc comparisons revealed that final ratings of certainty differed significantly as a function of information congruency, $t(83) = 8.03, p < .001$, such that participants' ratings of certainty were higher after exposure to wrongness-confirming

information than wrongness-contradicting information ($M = 6.76$, $SE = .13$, versus $M = 4.95$, $SE = .19$, respectively).

5.2.2.4 Emotion ratings. Ratings of emotion were coded along a 1 to 7 scale. Emotion ratings are explored to determine whether anger and disgust differed as expected for the respective violation (i.e., more anger than disgust for harm violations and more disgust than anger for purity violations). Additionally, these analyses compared the emotions of interest with other control emotions (e.g., sadness). A 2 (violation type: harm, purity) \times 5 (emotion type: fear, anger, disgust, frustration and sadness) repeated measures ANOVA revealed a significant main effect of emotion, $F(1, 85) = 58.74$, $p < .001$, $\eta^2_p = .41$ (fear $M = 2.79$, $SE = .21$, anger $M = 4.304$, $SE = .22$, disgust $M = 5.43$, $SE = .17$, frustration $M = 3.60$, $SE = .22$, sadness $M = 4.22$, $SE = .20$). There was no main effect of violation type, $p = .31$.

There was a significant interaction between emotion and violation type, $F(1, 85) = 10.45$, $p < .001$, $\eta^2_p = .11$ (see Table 11). Ratings of anger did not differ significantly as a function of violation type, $p = .46$. Ratings of disgust did differ significantly as a function of violation type, $t(83) = 2.94$, $p = .004$, such that ratings of disgust were higher in the purity versus harm violation condition ($M = 5.93$, $SE = .22$, versus $M = 4.93$, $SE = .26$, respectively). Ratings of frustration did not differ significantly as a function of violation type, $p = .12$. Ratings of sadness did differ significantly as a function of violation type, $t(83) = 3.072$, $p = .003$, such that ratings of sadness were higher in the purity versus harm violation condition ($M = 4.83$, $SE = .29$, versus $M = 3.605$, $SE = .28$, respectively).

Table 11
Mean Ratings (and Standard Errors) of Emotion Depending on Violation Type for Set 2, Experiment 5

	Harm violation (n=43)	Purity violation (n=42)
Fear	2.56 (.29)	3.024 (.29)
Anger	4.47 (.31)	4.14 (.31)
Disgust	4.93 (.24)	5.93 (.24)
Frustration	3.93 (.302)	3.26 (.31)
Sadness	3.61 (.28)	4.83 (.28)

Note. Possible ranges 1 to 7.

5.2.2.5 Facial expression ratings. A 2 (violation type: harm, purity) \times 2 (facial expression: angry expression, disgusted expression) chi-square test revealed no significant differences between the number of participants who selected the angry facial expression versus the disgusted facial expression as best representing their reaction to either the harm or purity violation, $\chi^2 = 2.08$, $df = 1$, $p = .15$.

Facial expression ratings were coded along a 1 to 9 scale. A 2 (violation type: harm, purity) \times 2 (facial expression: angry face, disgusted face) repeated-measures ANOVA revealed no main effect of facial expression, $p = .16$, and no main effect of violation type, $p = .76$.

There was a significant Violation Type \times Facial Expression interaction, $F(1, 85) = 8.82$, $p = .004$, $\eta^2_p = .096$ (see Figure 8).

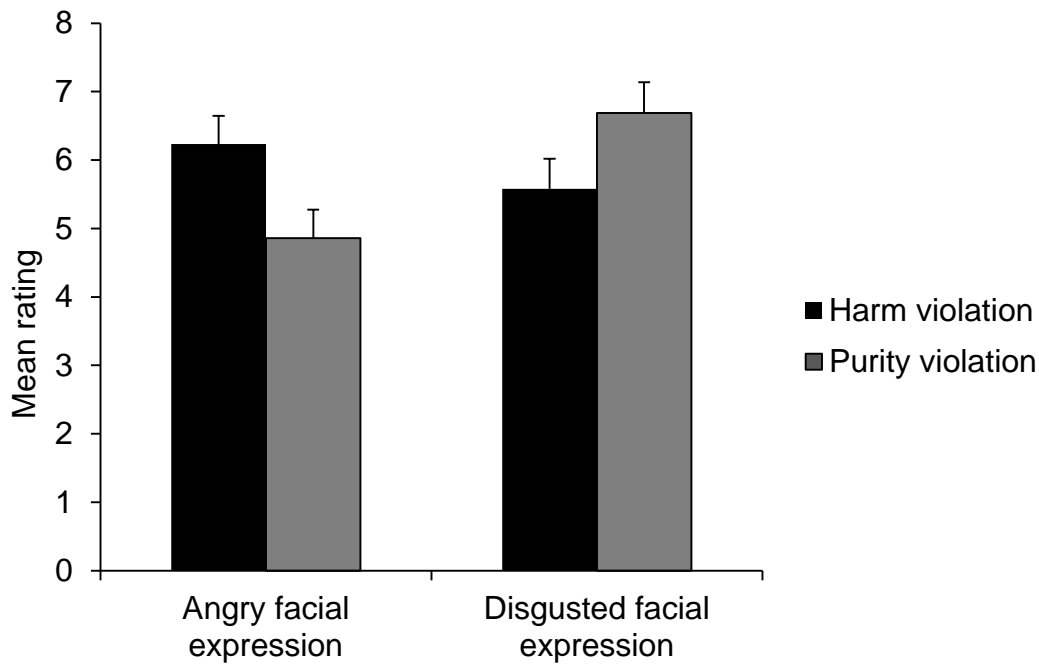


Figure 8. Interaction between violation type (harm vs. purity) and facial expression ratings (angry face vs. disgusted face) for Set 2, Experiment 5.

Paired samples *t*-tests revealed that ratings did not differ as a function of facial expression in the harm violation condition, $p = .28$. Ratings did differ as a function of facial expression in the purity violation condition, $t(41) = 3.13$, $p = .003$, such that ratings were higher for the disgusted face than the angry face ($M = 6.69$, $SE = .404$, versus $M = 4.86$, $SE = .46$, respectively).

5.2.3 Summary. The primary goal of Experiment 5 was to examine how participants' behaviour with regard to information selection differed depending on their exposure to either harm or purity moral violations. I predicted that, on the basis of findings from Experiments 3 and 4, participants in the harm violation condition would select more lines of wrongness-confirming information than wrongness-contradicting information in an effort to boost or maintain their epistemic certainty. Conversely, I predicted that participants in the purity

violation condition would select more lines of wrongness-contradicting information than wrongness-confirming information in an effort to reduce or alleviate feelings of disgust associated with the purity violation. Experiment 5 replicated the high reports of wrongness and judgement certainty from Experiments 3 and 4. However, the findings from Experiment 5 did not support the prediction that participants would engage with information post-judgement differently depending on the type of violation they judged.

There are two possible explanations for the findings in Experiments 5, the second of which is explored further in Experiment 6. First, findings from Experiments 3 and 4 result from participants' self-reported, introspective and hypothetical desires and needs to engage with different types of post-judgement information—information, however, that was never presented. Experiment 5 explored participants' *actual* engagement with additional information, and therefore was not measuring participants' introspection, suggesting a potential discrepancy between the findings in Experiments 3 and 4, and Experiment 5 that is due to differences between hypothetical and actual information engagement. Furthermore, these self-report measures from Experiments 3 and 4 were also gathered from responses to an online survey, whereas responses in Experiment 5 were collected in the lab with the presence of an experimenter. Therefore, there may be reason to believe that the lack of continued support for the findings in Experiments 3 and 4 is due to the change in methodology; an online survey that is conducted privately, compared to an in-lab computer task where an experimenter who may be perceived to be judgemental is present. This is not to say that the findings from the online measures used in Experiments 3 and 4 should be dismissed, but rather suggests that decisions about engaging with information about a highly emotionally charged moral violation may be sensitive to the presence of another, potentially judgemental person (e.g., the experimenter). Additionally, a discrepancy between self-reported

hypothetical engagement with information and actual engagement with information does not necessarily suggest a weakness in the findings from Experiments 3 and 4, but may demonstrate the limits of introspection (Nisbett & Wilson, 1977). Nisbett and Wilson (1977) suggest that participants who attempt to report on their own cognitive processes may not always be doing so on the basis of genuine and accurate introspection. Therefore, participants in Experiments 3 and 4 may be reporting a limited understanding of their own cognitive processes with respect to how they would hypothetically select and engage with additional behavioural and contextual information.

Secondly, in part due to the potentially more serious and demanding setting of Experiment 5 (in a lab cubicle on a university campus with the occasional presence of an experimenter), participants may have felt that the appropriate response was to read all available information. As the results from Experiment 5 demonstrated, the majority of participants read nearly all 30 lines of available information regardless of condition (see Tables 7 and 9). Participants may have felt that reading fewer than all available lines of information was in some way taking the experiment less seriously, quitting too early and therefore giving a negative impression. Relatedly, the instruction in Experiment 5 that participants read prior to receiving additional information made reference to reading as much as necessary before making a final judgement. The use of the word “necessary” implies that participants *should* read a certain number of items as opposed to reading on the basis of the number desired. Experiments 3 and 4 tested participants’ self-reported *desire* for additional information.

Finally, it is worth noting that the emotion measures did not reflect clear differences between anger and disgust for either harm or purity violations. Given the conclusions that were drawn from Experiments 3 and 4 about the role of disgust in yielding a lack of desire for

additional information overall, and a specific desire for disgust-alleviating information, the goal of Experiment 5 was to directly measure differences in elicited emotions. The lack of a clean difference between anger and disgust, particularly in response to purity violation, made it difficult to draw any conclusions about the direct role of either anger or disgust in yielding different information desires. Additionally, the second set of violations yielded the cleanest distinction between anger and disgust to their respective and predicted violations; the second set of violations was therefore used in Experiment 6.

6 Experiment 6

Experiment 6 aimed to further investigate what types of information participants would choose to engage with and the extent to which participants would engage with the information. Several modifications were made to the procedure in this experiment. First, Experiment 5 did not yield clear differences with respect to emotion ratings on the basis of the violation type. The clearest trend towards differences between anger and disgust associated with their respective violations (harm, purity) was found in the second set of stories (tattoo and dead dog) Therefore, I included only the two new scenarios from the second set in Experiment 6. Additionally, I emphasised precisely what I meant by anger and disgust in order to allow for more dissociation between the two, and to control for reports of general negative affect as opposed to specific feelings of anger and disgust. I also allowed participants to select the type of information they wanted to read, and modified the information so as to either conserve or alleviate the participants' emotional experience in response to the violation.

As with Experiments 3, 4 and 5, I predicted that participants exposed to the harm violation would choose to read more emotion-conserving information and would engage with more of the information in an attempt to maintain or boost their certainty. I also aimed to explore whether participants who experienced anger in response to a harm violation would

conserve their feelings of anger by choosing to read wrongness-confirming information, as the information following both violations was tailored to give both more or less wrongness *and* more or less anger/disgust to the violation. I predicted further that participants exposed to the purity violation would choose to read wrongness-contradicting information and would engage with more wrongness-contradicting information in an attempt to alleviate and escape their feelings of disgust following their judgement.

6.1 Method

6.1.1 Participants and design. Forty-six undergraduate psychology students from the University of Birmingham (31 women; $M_{\text{age}} = 20.7$ years, $SD = 3.73$) participated for course credit. Participants were randomly assigned to one of two conditions of a single-factor (violation type: harm, purity) between-participants design.

6.1.2 Materials.

6.1.2.1 Moral violations. The moral violation scenarios were the same short vignettes from Set 2 in Experiment 5. The harm violation scenario depicted a woman being permanently tattooed by her boyfriend while unconscious. The purity violation scenario depicted a man consuming the meat from his dead dog.

6.1.2.2 Additional information. Thirty new lines of additional information were generated for each moral violation scenario (see Appendix E). The information was designed to either conserve or alleviate feelings of disgust in the case of the purity violation, or feelings of anger in the case of the harm violation. As in Experiment 5, lines of information included equal amounts of information about the character, and the situation surrounding the violation, and were matched for length and theme across wrongness-confirming and wrongness-contradicting content. Lines of information appeared one-at-a-time, line-by-line, with one line per screenshot for a period of three seconds. Response options for “escape” or “continue”

were presented on a separate screen at the end of the three seconds. There were two counterbalanced pseudo-random presentation orders for both conditions.

6.1.3 Procedure. The procedure for Experiment 6 was similar to Experiment 5 with some alterations. First, participants only read one moral violation and therefore did not complete a filler task between violations as in Experiment 5. Additionally, participants completed the facial expressions questionnaire used in Experiment 5 immediately after the presentation of the moral violation scenario, and before rating the moral wrongness of the violation, their certainty about their judgement of wrongness, and their feelings of anger and disgust in response to the moral violation. All ratings were made on the same 7-point scales used in Experiment 5. Additionally, participants were given different instructions regarding the additional information available to them:

Now you're going to read more about this story. What kind of information would you like to read more about? Would you like information that supports your judgements of [story character's name] and his actions, or information that might challenge these judgements? In other words, would you like to read negative (supporting) or positive (challenging) information about [story character's name] and his/her actions?

Participants were then given the option to select a response button that redirected them to either the lines of wrongness-confirming or wrongness-contradicting information. Participants were given the following instructions once redirected and prior to receiving the first line of additional information:

You have chosen to read more negative [positive] information about [story character's name] and his/her actions. There is going to be a lot of information that you can read. There will be more information than there is time to complete this study, so don't feel as though you need to keep reading until you reach the end. Instead, we'd

like you to just read as much or as little as you want. That means the amount of information you read is completely up to you. For example, you can read 2 lines or you can read 20 lines; it's up to you.

After a line of information was presented for three seconds, a screen appeared that provided the same response button options from Experiment 5 to “escape” or “continue”. Participants were led to additional information if “continue” was selected or to a demographics questionnaire if “escape” was selected or if all 30 lines of information were viewed.

6.2 Results and Discussion

The results of Experiment 6 are presented in the same way as the results in Experiment 5 for the two violations used in Experiment 6. The results to the main question of interest are presented first, followed by the results of the manipulation checks for both the post-judgement information, and the elicited violation-specific emotions.

6.2.1 Engagement with additional information. A Fisher's exact test revealed no significant differences between conditions (harm violation versus purity violation) with respect to the selection of wrongness-confirming versus wrongness-contradicting information, $p = .18$. In both conditions, the majority of participants chose wrongness-contradicting information (96% in harm condition, 82% in purity condition). Post-hoc comparisons revealed no significant differences between the number of lines read on the basis of the violation type, $p = .506$. Due to the fact that the majority of participants selected wrongness-contradicting information, I am unable to make a comparison between the numbers of lines of information read on the basis of information congruency.

6.2.2 Moral wrongness. Ratings of wrongness were coded along a 1 to 7 scale. A 2 (violation type: harm, purity) \times 2 (counterbalance order: 1, 2) between-participants ANOVA revealed no significant effect of counterbalance order, $p = .17$. There was a significant main

effect of violation type, $F(1, 46) = 11.43, p = .002, \eta^2_p = .21$, such that judgements of wrongness were higher in the harm versus purity violation condition ($M = 6.909, SE = .16$, versus $M = 6.14, SE = .17$, respectively).

6.2.3 Judgement certainty. Ratings of certainty were coded along a 1 to 7 scale. A 2 (counterbalance order: 1, 2) \times 2 (violation type: harm, purity) ANOVA revealed no significant effect of counterbalance order, $p = .52$. Additionally, there were no significant main effects or interactions, all $p > .15$. Participants' ratings of judgement certainty were equivalent and high in both the harm and purity violation conditions ($M = 6.44, SE = .23$, and $M = 5.96, SE = .24$, respectively).

6.2.4 Emotion ratings. Emotion ratings were coded along a 1 to 7 scale. A 2 (counterbalance order: 1, 2) \times 2 (violation type: harm, purity) \times 2 (emotion: anger, disgust) mixed-model ANOVA revealed no significant main effect of counterbalance order, $p = .29$. There was a main effect of emotion, $F(1, 46) = 6.19, p = .017, \eta^2_p = .13$, such that overall ratings of disgust were higher than ratings of anger ($M = 5.39, SE = .18$, and $M = 4.73, SE = .24$, respectively).

There was also a significant interaction between violation type and emotion, $F(1, 46) = 22.29, p < .001, \eta^2_p = .35$ (see Figure 13).




Figure 13. Interaction between violation type (harm vs. purity) and emotion ratings (anger vs. disgust), Experiment 6.

Post-hoc comparisons revealed that for the purity violation condition, ratings of disgust were higher than ratings of anger, $t(21) = 4.98, p < .001$ ($M = 5.86, SE = .19$, and $M = 3.95, SE = .39$, respectively). For the harm violation condition, emotion ratings did not differ significantly, $p = .096$, although the pattern suggests that ratings of anger were higher than ratings of disgust.

6.2.5 Facial expression ratings. A Fisher's exact test revealed a significant difference between participants' selection of the set of angry facial expressions versus the set of disgusted facial expressions as a function of the violation type, $p = .038$. Twenty-eight per cent of participants in the harm condition chose the angry face and 24% of participants in the harm condition chose the disgusted face in the harm condition; 37% of participants chose the disgusted face and 11% of the participants chose the angry face in the purity condition.

The ratings of each facial expression set were along a 1 to 9 scale. A 2 (counterbalance order: 1, 2) x 2 (violation type: harm, purity) x 2 (facial expression: angry, disgusted) repeated measures ANOVA revealed no significant main effect of counterbalance order, $p = .71$. There were no additional significant main effect, lowest $p = .14$.

There was a significant interaction between violation type and facial expression rating, $F(1, 46) = 5.28, p = .027, \eta^2_p = .11$ (see Figure 14).




Figure 14. Interaction between violation type (harm vs. purity) and facial expression ratings (angry face vs. disgusted face), Experiment 6.

Post-hoc comparisons revealed that in the purity violation condition, there was a significant difference between emotion, $t(21) = 2.59, p = .017$, such that ratings of disgust were higher than ratings of anger, $t(21) = 2.59, p = .017$. In the harm violation condition, there was no significant difference between ratings of anger and disgust, $p = .55$.

6.3 Summary

In Experiment 6, my goal was to measure post-judgement information preferences in two ways: the selection of wrongness-confirming or wrongness-contradicting information, and the number of lines read of the selected information type. Additionally, the wording of the instructions changed from Experiment 5 to emphasise that participants should choose to read lines of information on the basis of their *desire* for more information as opposed to on the basis of acquiring a sufficient amount of information to make a judgement.

Participants in Experiment 6 were given the opportunity to select the type of information that they were interested in reading. The findings suggest that the overwhelming majority of participants in both harm and purity conditions chose to read positive information; information designed to alleviate or reduce the emotions elicited from the violation, and to mitigate a judgement of wrongness. In Experiment 5, the information options were designed purely to mitigate or inflame a judgement of wrongness. In Experiment 6, however, the information options also focused on information designed to elicit either more or less anger in response to a harm violation, or more or less disgust in response to a purity violation. This was done in an effort to directly address the role of emotion in guiding post-judgement information choices that were discussed in Experiments 3, 4 and 5. The second set of violations from Experiment 5 was used in Experiment 6. With respect to emotion, the findings suggest that participants did report more anger than disgust in the harm violation condition, and more disgust than anger in the purity violation condition.

The harm and purity violations elicited their respective predicted emotions. Emotion, however, did not guide information selection in the way that was predicted. I predicted that participants who experienced anger in response to a harm violation would be motivated to conserve their anger and therefore choose to continue reading wrongness-confirming

information. To support the findings from Experiments 3, 4 and 5, I also predicted that participants in the harm condition would seek additional wrongness-confirming information for the purposes of boosting or maintaining epistemic certainty; hypotheses regarding certainty maintenance or boosting were not directly tested in this experiment. I also predicted that participants who experienced disgust in response to a purity violation would want to avoid and alleviate their disgusting feelings by choosing to read wrongness-contradicting information. Although participants in the purity condition did report high levels of disgust and chose wrongness-contradicting information, I do not have evidence to support the claim that the disgust and purity violation were driving the choice of information, given that participants in the harm condition who reported feeling angry also chose wrongness-contradicting information.

7 General Discussion

The research presented in this chapter examined post-judgement processing with respect to hypothetical and actual engagement with additional behavioural and contextual information following exposure to moral violations of harm or purity. Experiments 3 and 4 examined self-reported hypothetical desires and needs for certain types of post-judgement information. I found that participants in the harm violation condition reported a greater desire for additional information than participants in the purity violation condition. Additionally, the desire for either wrongness-confirming or wrongness-contradicting information was influenced by the amount of epistemic judgement certainty participants reported in the harm violation condition, and by the emotion experienced in the purity violation condition. Participants who reported more epistemic judgement certainty in the harm condition were more interested in wrongness-confirming information more than wrongness-contradicting information, and participants who reported more intense disgust in the purity condition were

more interested in wrongness-contradicting information than wrongness-confirming information. Experiments 5 and 6 examined participants' actual engagement with additional post-judgement behavioural and contextual information and found that, overall, there was a high degree of engagement with all types of information regardless of the type of violation that was judged, and an overwhelming preference for wrongness-contradicting information when participants had the opportunity to select either wrongness-contradicting or wrongness-confirming information.

Experiments 3 and 4 provided evidence that self-reported desire for additional information differs depending on the type of violation being judged and the type of additional information that is available. In these two experiments, participants who made judgements after reading a harm violation reported a desire for more information. In contrast, participants who made judgements after reading a purity violation reported *not* wanting, needing or liking more information. This difference emerged even though participants in the harm- and purity-violation conditions did not differ in the extent to which they believed that the additional information would be useful; indeed, participants in both conditions reported that additional information would be equally *useless* for informing their judgement. Evidence for the further role of the type of additional information came from Experiment 4. Specifically, participants in the purity-violation condition were more interested in information contradicting than supporting their initial judgements of wrongness.

Two additional experiments were conducted to further test these findings with respect to actual behaviour. The findings from Experiments 5 demonstrated that participants engaged with the majority of available information regardless of the violation that was judged or the type of information available. Furthermore, the findings from Experiment 6 demonstrated that

participants chose to engage with positive, wrongness-contradicting information after exposure to either a harm or purity violation.

To frame my research, I suggested two accounts for differences in post-judgement processing as a function of moral violation type. First, I suggested that post-judgement processing might reflect epistemic certainty motives. In Experiment 3, however, I demonstrated that participants' ratings of their judgement certainty were high and did not differ as a function of violation type. In other words, participants in both the harm- and purity-violation conditions reported feeling very certain that both violations were very wrong, so that differences in certainty were not a viable explanation for the different need for information across conditions. Further evidence against the epistemic certainty account came from a more subtle, indirect measure of epistemic certainty: participants' ratings of the usefulness of more information. If epistemic certainty was motivating post-judgement information search, then I would expect that information to be perceived as useful, the utility being to satisfy participants' epistemic certainty of their judgements. In Experiment 3, however, participants reported that additional information would be useless for informing their judgements regardless of whether they were exposed to a harm violation or a purity violation.

Experiment 4 did find differing certainty as a function of violation type, such that participants were more certain of their judgement when judging the purity violation; however, among participants in the harm-violation condition, greater certainty actually predicted interest in wrongness-*confirming* information. Although participants were equally likely to choose information that confirmed or contradicted their wrongness judgement (i.e., did not show a preference for wrongness-confirming information that would maintain their mood), their certainty ratings correlated positively with their desire for wrongness-confirming information. Thus, the more certain participants were of their wrongness judgement, the

greater their interest in *maintaining* and *boosting* certainty by choosing to engage with information that sustains certainty by confirming their judgement of wrongness. Possibly, this relationship between certainty and biased search reflects a desire to maintain certainty once achieved, due to the positive subjective properties of the certain state. Achieving epistemic certainty in the face of uncertainty is likely to serve the function of helping people to check and confirm their judgement of wrongness. Maintaining and boosting already existing certainty, however, is likely to serve the function of allowing people to indulge in condemnation. Further research should explore the extent to which this need to maintain certainty in the face of harm violations is a result of people upholding their sense of moral righteousness and indulging in the experience of condemning. Importantly, this bias could present problems when legitimately mitigating information is ignored that could dampen or reverse a judgement.

With regards to how participants actually engaged with available additional behavioural and contextual information, participants' certainty in Experiments 5 and 6 differed only as a function of the time of their judgement (pre-information vs. post-information), and differed as a function of the type of information that followed their initial rating of epistemic certainty. Unsurprisingly, participants were overall more certain of their final judgement than their initial judgement following the presentation of information that was wrongness-confirming. In conjunction with the above account for a need and desire to maintain and boost epistemic certainty, this pattern of findings suggests that certainty can be maintained and, in fact, boosted with the presentation of information that supports initial judgements of wrongness. Experiments 5 and 6 did not find differences in epistemic certainty as a function of the violation that was judged (harm violations versus purity violations).

The second account that I considered was that motivational tendencies to approach versus avoid different types of information in the effort to regulate emotional experience would drive post-judgement information search. This account predicts that participants exposed to a harm violation would approach and engage with information whereas participants exposed to a purity violation would avoid information, based on evidence that anger is an approach emotion (Carver & Harmon-Jones, 2009) but disgust is an avoidance emotion (Rozin et al., 2000). The findings from Experiment 3 were consistent with this account with respect to purity violations. Participants exposed to a purity violation, associated with feelings of disgust and avoidance tendencies, reported having no desire for additional information. Experiment 4, which measured choice between wrongness-confirming versus -contradicting information as a function of violation type, did not present a direct test of this account; however, given that wrongness-confirming information in purity violations was relatively avoided, this preference could also reflect a desire to avoid information and therefore avoid further feelings of disgust.

Harm violations are known to elicit approach-oriented feelings of anger (Gutierrez & Giner-Sorolla, 2007). I predicted that harm violations would lead to a desire for more information, and specifically more wrongness-confirming information that would maintain the feeling of anger. Because both anger and disgust were reported equally in reaction to harm violations, it is difficult to reconcile post-judgement responses following harm violations with this emotion-based account. However, it is possible that the correlation of certainty with biased search in Experiment 4's harm condition reflects a kind of self-regulation more focused on keeping alive the positive feelings of confidence and moral righteousness than the negative feeling of anger. On the other hand, my findings for purity violations were more supportive of initial predictions based on emotion regulation. Purity violations were predicted

to lead to desire to avoid information or, if forced to engage with information, a preference for wrongness-contradicting information that would repair the emotion and take them affectively further from their experience of disgust.

While ratings of disgust did differ significantly from ratings of anger in the purity violation condition in Experiments 5 and 6, I did not find any differences in the actual engagement with additional information in Experiment 5, and found an overwhelming preference for wrongness-contradicting information in Experiment 6. Therefore, Experiment 5 does not support an emotion regulation account with regards to actual engagement with information. Experiment 6, however, does demonstrate that participants who judged both purity and harm violations did report high levels of disgust, while sometimes in conjunction with anger, and did choose to engage with emotion avoiding and alleviating information (i.e., wrongness-contradicting information). While I expected this to only occur for those participants who judged purity violations, the fact that disgust was reported for both violations (and higher levels of disgust than anger were reported in the purity violation condition) may suggest that disgust is driving a preference to engage with emotion-avoiding and -alleviating post-judgement information.

The findings from Experiment 3 support an emotion regulation account in that participants in the purity-violation condition expressed a clear lack of desire for more information. Further support for this account was found in Experiment 4 and Experiment 6. After judging purity violations, participants chose information that contradicted their wrongness judgement, consistent with the notion that they sought to repair their feelings of disgust. Moreover, in Experiment 4, their emotion intensity ratings correlated positively with their desire for wrongness-contradicting information, suggesting that the greater their feeling of disgust, the greater their interest in disgust-alleviating information. When the opportunity

to actually engage with post-judgement information was presented, participants in Experiment 6 who reported both feelings of anger and disgust chose to engage with emotion avoiding and alleviating information.

As the social intuitionist model (Haidt, 2001) suggests that emotion is the driving force behind moral judgement, perhaps emotion and certainty regulation are also the driving forces behind post-judgement information search. In violations where harm involves a clear victim and transgressor and leads to concerns regarding justice, maintaining certainty may drive post-judgement information search bias. People may feel the need to maintain certainty in order to induce feelings of moral righteousness and establish confidence when condemning or punishing a transgressor. When moral violations involve the violation of a taboo or norm without an obvious victim and transgressor, the emotion tends to be the driving force in post-judgement information desire. An emotion-regulation account would suggest that disgust in the purity-violation condition is driving post-judgement decisions to disengage with, and avoid, information. Disgusted participants chose to not engage with information post-judgement, and to engage with information that was positive and disgust-alleviating when forced to engage with information. Consistent with the hypothesis that disgusted participants seek to repair their mood, previous research has shown that disgust is functional by keeping people away from things that could lead to contamination or disease (Park, Faulkner, & Schaller, 2003; Rozin, Millman, & Nemeroff, 1986).

With support for two of the predicted accounts for post-judgement information selection, Experiments 5 and 6 aimed to solidify these accounts with respect to how participants actually engaged with post-judgement information. Unfortunately, in Experiment 5 there were no differences in information engagement as a function of the type of violation judged, the emotions elicited or the amount of epistemic certainty. There is some evidence

from the findings of Experiment 5 that certainty can be maintained and boosted following judgement-congruent and wrongness-confirming information. This adds to the accounts from Experiments 3 and 4 that maintaining and boosting epistemic certainty can motivate participants' hypothetical self-reported information engagement selections. Due to limitations in methodology discussed in the summary of Experiment 5, differences in information engagement were difficult to disentangle as a majority of participants read the majority of additional information made available regardless of whether they were exposed to a harm or purity violation. Reports of anger were not higher than reports of disgust in response to the harm violation; likewise, reports of disgust were not higher than reports of anger in response to the purity violation. Although the lack of support for the previous findings may be due to the discrepancy between hypothetical self-reports (Experiments 3 and 4) and actual behaviour, some modifications were made in Experiment 6 to test this hypothesis further.

Experiment 6 addressed some of the methodological limitations of Experiment 5, but supported an emotion-avoiding and -alleviating account for post-judgement information engagement. Although I aimed to demonstrate differences on the basis of the violation that was judged (harm vs. purity), Experiment 6 only demonstrated that, regardless of the violation judged, the majority of participants chose to engage with wrongness-contradicting information. The initial hypothesis was that, given the avoidance tendencies associated with feelings of disgust, participants who felt disgust, presumably only in response to a purity violation, would choose to engage with wrongness-contradicting information in an effort to avoid and alleviate their emotional response to the violation. While this was not the case for purity violations as opposed to harm violations as initially predicted, reports of disgust were present in response to both harm and purity violations, and therefore may be driving the decision to engage with disgust alleviating and avoiding post-judgement information.

Although Experiment 6 provided evidence for actual information selection, in order to further investigate information engagement (i.e., the amount of the type of information selected was actually read), I would need to have proportionate numbers of participants choosing the information types that are available. Further experiments should either allocate participants to the information types as was done in Experiment 5, or data collection should be on-going until there are comparable numbers of participants reading each of the information types.

Experiments 5 and 6 did not provide entirely consistent accounts for post-judgement information engagement with the findings from Experiments 3 and 4. One possible methodological explanation for this discrepancy is the difference in hypothetical self-reported post-judgement information desire and *actual* post-judgement information desire and engagement. Participants in Experiments 3 and 4 reflected on how they *would* engage with different information following harm and purity violations, whereas participants in Experiments 5 and 6 actually chose the information that they would engage with and made real decisions about how much or little information they wanted. The inconsistency in these findings may reflect a limitation of self-report (Nisbett & Wilson, 1977).

Furthermore, Experiments 5 and 6 were conducted in such a way that there was an opportunity for participants to feel self-conscious in engaging with additional post-judgement information. Particularly for those participants exposed to a harm violation, my prediction would have required participants to indulge in the reading of negative, inflammatory information. Although people may engage with conserving anger and condemnation, there may be certain conditions that are necessary for that to take place. For instance, people may not want to feel judged in displaying a desire for more negative, condemning information that serves no purpose. The presence of the experimenter may have hindered this need for being in a non-judgemental environment, a situation that was available in the online surveys used in

Experiments 3 and 4. Additionally, if reading information that perpetuates condemnation just for the sake of condemning is undesirable, perhaps this paradigm should be tested with situations where boosting and maintaining certainty serve a purpose, such as when deciding to punish on the basis of a wrongdoing.

The findings from Experiments 3 and 4 clearly suggest that hypothetical self-reported interest in post-judgement behavioural and contextual information about moral violations differs on the basis of violation type. Two additional experiments suggest that actual behaviour of participants differs from their self-reported post-judgement interest in additional information in that all participants engage with most available information and choose positive information over negative information. Of the two predicted accounts for self-reported post-judgement information search, epistemic certainty maintenance seems to best fit the data with regards to harm violations. Emotion regulation accounts seem to best fit the data with regards to purity violations for hypothetical self-reported information selection, and both harm and purity violations eliciting feelings of disgust with regards to actual information selection. Certainty-maintenance and emotion-regulation motivations have important implications for the role of post-judgement information, and how people who condemn different violations choose to expose themselves to potentially important information.

Overall, the findings taken together suggest preferences for post-judgement information selection and engagement, whether or not that role is dependent on the type of violation that is judged. Post-judgement information is important to consider for two reasons. Firstly, if self-reported desires are reflective of how people will treat post-judgement information, then when there is information that could mitigate a wrongness judgement, it may be unlikely that a person who is largely driven to maintain their certainty about their initial wrongness judgement will be motivated to look for information that could dampen their

judgement and influence punishment of the transgressor. People who experience disgust in response to a moral violation, however, may be quick to look for mitigation in an effort to repair and avoid their emotional state. If, however, the final two experiments are reflective of how people will treat post-judgement information, then there may be a positivity bias that leads people to mitigate their judgements of wrongness or emotional state, in some cases before reaching information that has the opportunity to change a judgement.

Secondly, these findings are important for considering the role of emotion, intuition and reasoning in moral judgements. Judgements may be made on the basis of emotion and intuition, as posited by the social intuitionist account. If, however, there is a case for the importance and influence of post-judgement information, there may be reason to suggest that people could engage with information and reasoning more so when there is an understanding that information following a judgement is available. So far, I have presented evidence that both information pre-judgement (information about actions) and information post-judgement (information that has the potential to mitigate or inflame a judgement) can be considered and influential. In Chapter 4, I will address the influence of behavioural and contextual information, specifically information about the outcomes of a helping act, on moral behaviour.

CHAPTER 4

**THE INFLUENCE OF BEHAVIOURAL AND CONTEXTUAL
INFORMATION ON BEHAVIOUR: MORAL ELEVATION WITHOUT
OUTCOME INFORMATION MAY NOT PROMOTE HELPING
BEHAVIOUR**

The majority of this thesis thus far has examined how the processing and consideration of behavioural and contextual information influences moral judgements prior to the moral judgement being made, and how behavioural and contextual information is selected and engaged with following the moral judgement. This chapter⁵ will examine how the consideration of different types of information can affect moral behaviour. The experiment in this chapter specifically examined conditions that lead from witnessing acts of moral beauty to engaging in helping behaviour. Participants were provided with the opportunity to engage in helping behaviour after viewing morally virtuous behaviour; critically, some participants learned about the outcome for the actor, others learned about the outcome for the beneficiary, and others learned nothing about the outcome. A significant increase in helping was observed only when the morally elevating event was accompanied by positive outcome information—when participants learned that the recipient benefited from the help or observed the helper being praised. This pattern of helping behaviour emerged despite the fact that participants in all three helping conditions reported high levels of moral elevation (and significantly more

⁵ The research from Experiment 7 of this chapter is under revision for resubmission to the *Journal of Experimental Social Psychology* as the following manuscript: Powell, N. L., Zumbé, S., Beck, S. R., & Quinn, K. A. (in revision). *Is It Worth the Effort? Moral Elevation without Outcome Knowledge Does Not Promote Helping*.

elevation than control participants). Moral elevation on its own might not be sufficient to prompt helping. The goal of this thesis is to examine how the consideration of behavioural and contextual information alongside intuition and emotion can moderate moral judgement, and this final chapter will provide a preliminary investigation into how the consideration of behavioural and contextual information can also moderate moral behaviour.

1 Introduction

The majority of this thesis thus far has examined the influence of behavioural and contextual information and reasoning on negative moral judgements, specifically condemnation, wrongness and blame. Although behaviour is separate from judgement, behaviour is important to consider in the moral judgement process because moral behaviour, specifically in response to witnessing another moral act, does involve some degree of evaluation, and can be the result of the experience of an elicited emotion. Additionally, this thesis has largely focused on the experience of negative moral emotions, such as anger and disgust, which follow from judgements of wrongness and condemnation. The focus in the moral psychology literature has overwhelmingly been on moral condemnation, judgements of wrongness and blame, and negative moral emotions. For instance, a search of the current PsycINFO database yields 245 hits for moral anger, but only 48 for moral elevation. Contemporary moral psychology can thus be characterised more as the study of moral condemnation than praise. I will begin to address the role of behavioural and contextual information and evidentiary standards for judging virtue and experiencing positive moral emotions. More specifically, I am interested in how behavioural and contextual information, and consideration of specific types of information, influence positive moral judgements and

emotions with regards to the behaviours that result from experiencing these positive moral emotions and evaluating virtuous acts.

Negative moral emotions dominate the identified moral emotion groups. Of the four identified groups of moral emotions, the positive moral emotions largely belong to the category known as the other-praising family, with the exception of pride (Haidt, 2003). There are many positive emotions, but the positive moral emotions, like the negative moral emotions, must meet the criteria of having an uninvolved judger (i.e., disinterested elicitor) and leading to prosocial and/or some moral behaviour. These emotions include awe, gratitude and elevation, and are associated with different behavioural tendencies from the negative, condemnatory moral emotions. The positive moral emotions, like other positive emotions, are useful for creating social cohesion and allow people to make self-improvements for creating sustainable social bonds later on (Fredrickson, 1998). Of the positive moral emotions, the majority of empirical research has focused on the study of moral elevation.

With respect to examining the effects of behavioural and contextual information on moral behaviour, I chose to examine how the consideration of different aspects of a witnessed act of virtue can moderate the likelihood that a person will engage in prosocial behaviour as a result of experiencing positive moral elevation. The current experiment will add to the overall aim of the thesis by exploring how consideration of different types of information, in addition to intuition and emotion, can lead to different outcomes for moral behaviour. I aimed to investigate conditions that might account for acts of altruism that result from exposure and judgement of morally good and virtuous acts. Specifically, I examined the conditions that lead from witnessing acts of moral beauty to engaging in helping behaviour. Witnessing or recognition of an act of “moral beauty” performed by another in the aid of a third party is said to lead to the experience of “moral elevation” (Haidt, 2003a, 2003b), characterised by feelings

of uplift and warmth in the chest. Algoe and Haidt (2009) have demonstrated that moral elevation is distinct from other forms of positive affect and motivates the desire to be helpful for a specific other or in general (see Cox, 2010; Freeman, Aquino, & McFerran, 2009; Silvers & Haidt, 2008). Further evidence comes from research by Silvers and Haidt (2008).

Silvers and Haidt (2008) conducted a study whereby female participants were videotaped as part of a study meant to observe their behaviour towards their infants in private; participants also wore nursing pads to measure the amount of breast milk produced. All participants began by watching a short video that was designed to be emotionally neutral and rated the video clip on several dimensions. Participants were then divided into separate video conditions such that the participants in the elevation condition watched a video clip whereby a man is praised for having mentored a poor, disadvantaged young boy by teaching him how to play a musical instrument. To control for the effects of general positive affect, participants in the control condition watch a video clip of a comedian. The mothers who had watched the morally elevating video clip reported a high amount of feelings of elevation, warmth and uplift, and in turn leaked more breast milk into their breast pads and were more likely to nurse their children in the room and show affection to their children, compared to the mothers who watched the amusing video clip. This study provided preliminary evidence that feelings of moral elevation in response to witnessing another's prosocial behaviour can lead to the witness engaging in more positive behaviours themselves.

Additional research by Freeman, Aquino, and McFerran (2009) has demonstrated that witnessing morally virtuous behaviour can have an effect of the witness' engagement in subsequent morally virtuous behaviour. In one study, participants watched video clips depicting either White people helping Black people or just generic helping that was unrelated to race. White participants who witnessed these video clips depicting acts of moral virtue and

experienced feelings of moral elevation reported being more willing to donate to a charity that was focused towards helping Black people than participants in the control condition who did not watch a video clip and therefore did not experience feelings of moral elevation; this pattern was true even participants who scored high on a measure of group-based dominance, which is usually associated with a reluctance to engage in prosocial behaviour with an out-group. A follow-up study confirmed that self-reports of willingness to donate to charity were representative of people's actual behaviour; White people were more likely to donate to a predominately Black charity after experiencing feelings of moral elevation after reading a story that depicted morally virtuous behaviour. These findings suggest that experiencing moral elevation after witnessing virtuous act will lead to an engagement in prosocial behaviour.

Further evidence to support the behavioural effects of experiencing moral elevation comes from research that examined the role of elevation-induced helping behaviour for people's engagement with volunteering (Cox, 2010). Participants in this study were undergraduate student volunteers on a service trip during their spring break who volunteered their time to help with the Nicaraguan Orphan Fund. Participants' self-reported experience of moral elevation in response to their volunteering was assessed at various times during the trip. Reports of moral elevation were predictive of later volunteering on the trip, and specifically the volunteering and service engagement that was related to the experience of moral elevation. In other words, the experience of moral elevation from volunteering in a specific context perpetuated further volunteering in that context. The effects of self-reported moral elevation were not mediated by any other factors such as personality (e.g., openness to experience) or reports of general positive affect.

The most direct support for the relationship between elevation and prosocial behaviour is reported by Schnall, Roper, and Fessler (2010). In the first experiment, participants either watched the morally elevating video clip from Silvers and Haidt (2008) or a control clip with no emotional charge. Participants watched the video clips and then recalled as much about the video as they could, as participants were told that they were participating in a memory recall experiment. Participants were then asked whether they would be willing to sign up for future unpaid research as a measure of willingness to help the experimenter at a later date. Participants in the elevation video condition were more likely to report that they would sign up for future unpaid research than participants in the control condition. In the second experiment, participants were allocated to either the elevation video clip condition, the neutral video clip condition, or a condition whereby participants watch a video clip designed to elicit feelings of mirth. Participants then completed manipulation checks that included emotion measures regarding feelings of elevation and general positive affect. Believing that the experiment was a test of episodic memory, participants completed a memory recall task. The experimenter then feigned technical problems loading the next part of the experiment and told participants that they could not complete the full experiment and were free to leave. As participants were leaving, and as an apparent afterthought, the experimenter asked whether participants would be willing to complete an unrelated and boring task to help the experimenter with a second project. Participants who were in the elevation condition and reported feelings of elevation beyond general positive affect were more likely to help the experimenter with the additional task and spent a longer amount of time on the task than participants in the neutral video and amusing video conditions. Thus, there appears to be direct support for the relationship between elevation and prosocial behaviour.

Nonetheless, research to date has not examined whether witnessing an act of moral beauty and feeling moral elevation is *sufficient* to elicit helping behaviour. Applying a crude distinction between selfless altruism and selfish egoism (e.g., Batson, 1991; Batson, Duncan, Ackerman, Buckley, & Birch, 1981; Cialdini, Darby, & Vincent, 1973; Cialdini & Kenrick, 1976) would suggest either that observers might be inspired directly by the elevation they experience when they observe helping behaviour, or that they might be sensitive to other information pertaining to the helping act—for example, cues that helping was effective or that the helper received praise.

With this consideration in mind, I investigated the impact of outcome information on participants' feelings of moral elevation and subsequent willingness to help. Participants were provided with the opportunity to engage in helping behaviour after viewing an act of moral beauty; critically, some participants learned about the outcome for the actor, others learned about the outcome for the beneficiary, and others learned nothing about the outcome. I examined the impact of outcome information on participants' feelings of moral elevation and subsequent willingness to help.

2 Experiment 7

2.1 Method

2.1.1 Participants and design. Eighty students from the University of Birmingham (71 female; $M_{age} = 19.7$ years, range = 18 to 26 years) completed the experiment in exchange for course credit. Data from one participant was omitted; this participant had difficulty following instructions due to a poor comprehension of the English language. The study used a single factor (video condition: no-outcome, actor-outcome, beneficiary-outcome, control) between-participants design.

2.1.2 Stimuli. The video stimuli used in three of the four experimental conditions were edited extracts from a 7-minute excerpt of the *Oprah Winfrey Show*, previously used by Schnall et al. (2010).

In the *no-outcome* condition, the video told the story of a troubled young man, Fernando, who grew up in conditions of extreme poverty and consequently believed that his only prospects lay in a life of crime. The video then described how Fernando had a chance meeting with Marcellus Brown, a college band director who heard him play the trumpet and offered him a scholarship. Marcellus acted as a mentor to Fernando, nurturing him and giving him the opportunity to build a better life for himself. Crucially, the video did not reveal whether the help given to Fernando yielded a positive outcome. The video was 1:23 seconds in duration.

In the *beneficiary-outcome* condition, the excerpt included the same background footage as in the *no-outcome* condition but also contained supplementary footage which revealed that the help Marcellus provided was extremely beneficial to Fernando; the viewer learned that Fernando forged a successful career as a musician. The video was 2:30 seconds in duration.

In the *actor-outcome* condition, the excerpt also contained the same background footage plus supplementary footage that showed the helper, Marcellus Brown, receiving praise and thanks from Fernando and being applauded by the studio audience. The video was 2:14 seconds in duration.

The control video was an excerpt from the television sitcom *Friends* taken from the episode, “*The one where Ross gets high.*” This stimulus was designed to elicit mirth and thus provide a control for general positive affect.

2.1.3 Measures.

2.1.3.1 Emotion ratings. A ratings scale was constructed to assess participants' emotional responses to the video stimuli, adapted from the scale devised by Schnall et al. (2010). Participants rated the extent to which they were currently experiencing each of 12 mood states along a scale anchored by 1 (*didn't feel at all*) and 9 (*felt very strongly*). Six items (taken from Schnall et al., 2010) assessed elevation: *moved, uplifted, optimistic about humanity, warm feeling in chest, want to help others, and want to be a better person*. An additional six items were included to control for general positive affect: *happy, excited, cheerful, energized, contented, and joyful*.

2.1.3.2 Memory. To retain the cover story that the study was investigating episodic memory recall, 10-item questionnaires (see Appendix F) were devised to test participants' recall of events in the videos that they watched. The questionnaires for the *no-outcome*, *beneficiary-outcome*, and *actor-outcome* conditions tested memory for events in the *Oprah Winfrey Show*. The questionnaire given to participants in the control condition was based upon events that took place on the *Friends* episode. Participants answered questions on visual properties of the videos (e.g., the colour of actors' clothing) and episodic information (e.g., the instrument that musician Fernando learned to play).

2.1.2 Procedure. The procedure was based on Schnall et al. (2010, Experiment 2). Participants completed the study individually. Upon arrival at the laboratory, participants were greeted by the experimenter and learned that they were to take part in a study investigating episodic memory recall within different media contexts. They learned that they would watch two short video clips and then answer questions on their content. Following on-screen instructions, participants watched one of the four videos. They then completed the 12-item emotion measure and the 10-item memory questionnaire.

Participants were then led to believe that they would watch a second video and complete two corresponding questionnaires. In view of the participant, the researcher clicked a folder on the computer screen labelled “Task 2”, and a (fictitious) error message flashed on screen. The experimenter feigned surprise, exclaiming, “Oh no, what’s going on here?”, and clicked on the icon a second time, with the same result. The experimenter explained that “the computer file seems to have corrupted.” The experimenter then indicated that the experiment must be stopped, that the participants could leave and would still be given full credit for their participation.

As participants began gathering their belongings, however, the experimenter interjected, as an apparent afterthought, by saying, “Actually, I am also running a pilot study at the moment. Would you mind helping me with that for a little while?” The experimenter then gave participants an 8-page list of anagrams and asked them to solve as many as they wanted. Participants were told that any amount of the task that they completed would be extremely helpful to the experimenter and that they could finish at any time. The experimenter then left the room, but stood outside and used a stopwatch to surreptitiously record the amount of time that the participants spent completing the anagrams task.

When the participants indicated that they had done as much of the task as they were motivated to complete, the experimenter administered an exit questionnaire; this questionnaire asked for basic demographic information, but also included a second measure of helping. Specifically, participants were asked whether they would be willing to volunteer for future unpaid research in the lab, and if so, to provide an email address through which they could be contacted.

Finally, participants were probed for suspicion, debriefed as to the true nature of the experiment, and thanked for their cooperation.

2.2 Results

2.2.2 Emotion. Separate scores were calculated for elevation and general positivity by averaging across six items for each (Cronbach's alphas = .874 and .899, respectively). These scores were submitted to a 2 (emotion: elevation, positivity) \times 4 (video: no-outcome, beneficiary-outcome, actor-outcome, control) mixed-model ANOVA with emotion as a within-participants factor.

The analysis showed a significant main effect of emotion, $F(1, 76) = 7.35, p = .008, \eta^2_p = .088$. The analysis also yielded a significant interaction between emotion and condition, $F(3, 76) = 61.03, p < .001, \eta^2_p = .71$; interaction means are depicted in Figure 15.

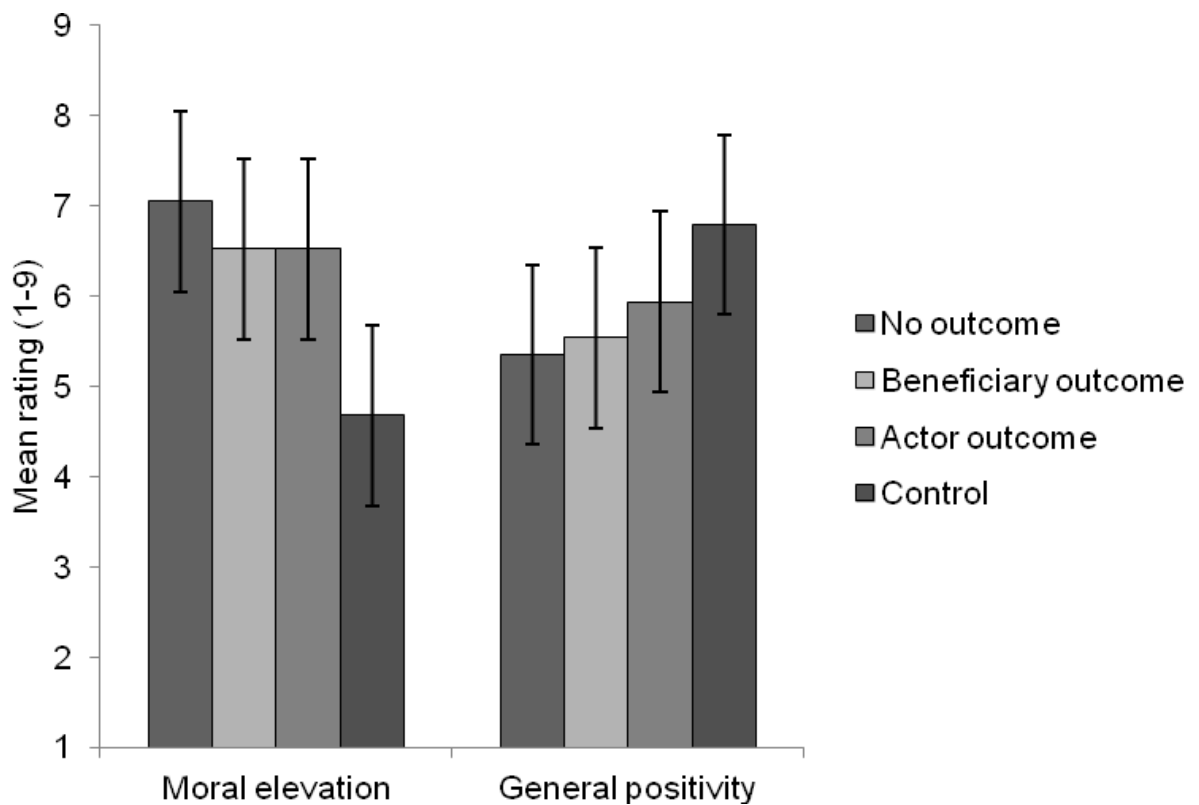


Figure 15. Mean self-reported emotion scores as a function of emotion type and video condition.

Note. Error bars represent standard error.

2.2.2.1 Comparisons between video conditions. Dunnett's comparisons were conducted to determine whether levels of self-reported emotion in the experimental conditions differed from the control condition. Participants reported significantly higher levels of elevation in the *no-outcome*, *actor-outcome*, and *beneficiary-outcome* conditions than did participants in the control condition (all $ps < .001$). Furthermore, participants in the *no-outcome* and *beneficiary-outcome* conditions reported significantly lower levels of general positivity than did participants in the *control* condition (both $p < .012$); general positivity did not differ in the *actor-outcome* versus *control* conditions ($p = .11$).

2.2.2.2 Comparisons within video condition. Bonferroni-corrected t -tests were conducted to analyse differences in self-reported levels of elevation and positive affect within each condition. Participants in the *no-outcome* ($t(19) = 6.32, p < .001$), *beneficiary-outcome* ($t(19) = 4.66, p < .001$) and *actor-outcome* ($t(19) = 4.14, p = .001$) conditions all reported higher levels of elevation than general positive affect. Conversely, participants in the control condition reported higher levels of positive affect than elevation ($t(19) = 10.92, p < .001$).

2.2.1 Helping behaviour.

2.2.1.1 Time spent completing anagrams task. A one-way between-participants ANOVA revealed a significant main effect of video type on the amount of time participants spent completing an anagrams task, $F(3, 76) = 2.73, p = .05, \eta^2_p = .097$; condition means are presented in Figure 16. Post hoc Dunnett's comparisons revealed that participants in the *actor-outcome* and *beneficiary-outcome* conditions spent significantly longer completing an anagrams task than participants in the control condition ($ps = .037$ and $.049$, respectively). Importantly, participants in the *no-outcome* condition did not spend significantly more time on the anagrams task than did participants in the control condition ($p = .36$).

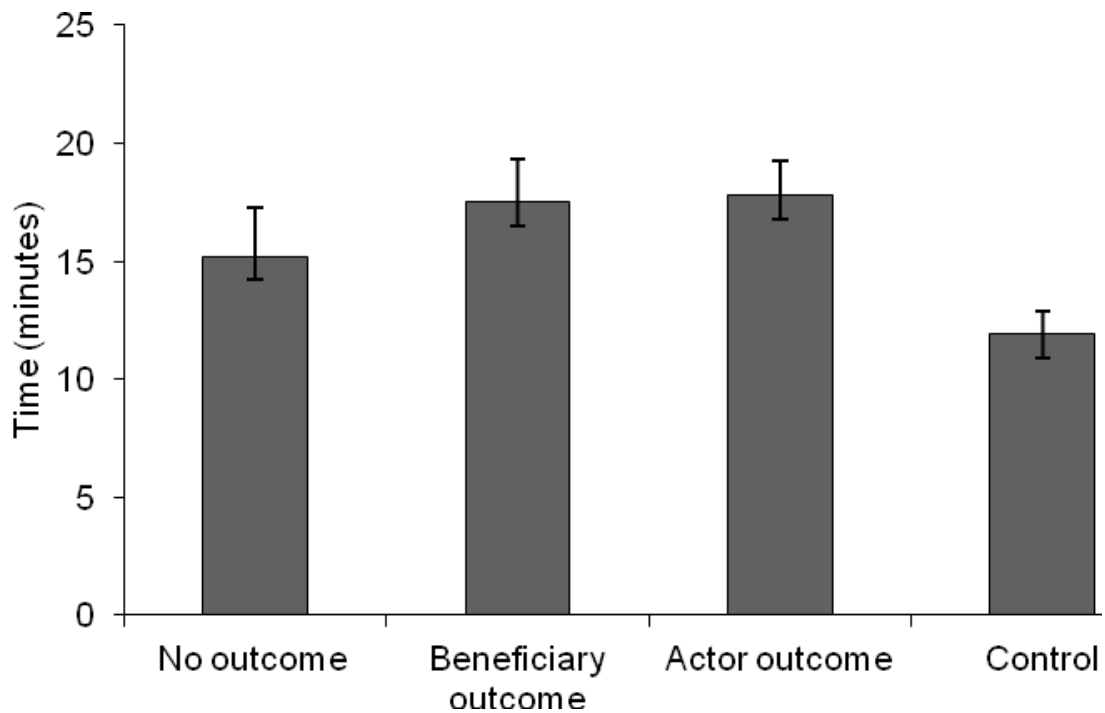


Figure 16. Mean time (in minutes) spent completing anagrams task as a function of video condition.

Note. Error bars represent standard error.

2.2.1.2 Willingness to volunteer for future research. A chi-square test revealed no significant relationship between the video participants watched and their subsequent willingness to volunteer for future unpaid research in the lab ($X^2 = 2.83, N = 80, p = .42$).

2.3 General Discussion

My goal in the current research was to differentiate between the experience of moral elevation in response to observing virtuous behaviour and the outcome knowledge that accompanies these observations. In so doing, I sought to provide greater insight into the relationship between observing others' morally virtuous acts, feeling elevated and deciding to behave virtuously. Our findings confirm those of Schnall et al. (2010) that witnessing acts of

moral beauty motivate subsequent helping behaviour. Information about the outcomes of these acts, however, moderated the effect of observation on helping. Specifically, watching a morally elevating video with no outcome information did not produce a significant increase in helping that was different from what occurred in the control condition; the control condition measures helping behaviour that was elicited after watching a neutral video clip that did not elicit feelings of moral elevation. A significant increase in helping was observed only when the morally elevating event was accompanied by positive outcome information—when participants learned that the recipient benefited from the help (in the *beneficiary-outcome* condition) or observed the helper being praised (in the *actor-outcome* condition). This pattern of helping behaviour emerged despite the fact that participants in all three helping video conditions reported equal levels of moral elevation (and significantly more elevation than control participants). Thus, moral elevation on its own is not sufficient to prompt helping.

Returning to the egoism–altruism debate (Batson, 1991), these results are consistent with the argument that helping behaviour is motivated by the opportunity for personal reward. It was only when participants could imagine positive outcomes (e.g., successfully helping another, earning praise)—and presumably put themselves in the helper’s shoes—that they engaged in more helping than control participants. Exploration of participants’ framing of helping (e.g., helping as an end in itself versus a means to a positive end for the helper) would provide important insight into this relationship.

Elevation without any outcome information does not demonstrate a greater increase in helping behaviour than the control, non-moral, non-elevation inducing condition. That being said, the fact that the two outcome-information conditions did not yield differences in the amount of helping behaviour may suggest that information about praise is unnecessary; participants did not need to know that the actor was praised in order for the participant to

engage in as much helping behaviour as when the participants knew the actor had been praised. This finding may suggest that while some outcome information is important for eliciting helping behaviour, information that leads the participant to believe that helping will result in benefits for the helper is not necessary for eliciting helping behaviour. Information about praise for the helper's good deeds is the more selfish of outcome information, and the lack of importance of praise information may suggest that true altruism is occurring.

It is important to note two limitations in the current experiment. Firstly, our second measure of helping behaviour—participants' willingness to volunteer for a future unpaid study—did not vary as a function of video condition. I interpret this null effect as reflecting the fact that participants had satisfied their helping goal during completion of the earlier helping task (Förster, Liberman, & Friedman, 2009; Förster, Liberman, & Higgins, 2005). Future follow-up studies should also counterbalance the two helping measures in order to determine if participants always satisfy their desire to help in the first helping task that is completed; this would solidify my explanation for why the second measure of helping was unsuccessful⁶.

Secondly, the video used in the no-outcome condition, where I did not observe increased helping, was substantially shorter than the videos used in the other two elevation conditions. I felt it was critical to use the same stimuli that had yielded the help-inducing effects in past research (e.g., Schnall et al., 2010), and so did not have the flexibility to control video length. More importantly, however, participant in the *no-outcome* condition reported the highest levels of elevation of all participants, again underscoring our argument that moral elevation is not sufficient to prompt helping. The fact that the no-outcome condition did not

⁶ I would like to thank the anonymous reviewer of my manuscript for this study for this useful suggestion.

differ from the control condition does suggest that once a virtuous act has been stripped of its positive outcome information, it will lead to no more helping than the control condition where no morally elevating content is present. Finally, a comparison of both positive and negative outcome information would be useful for determining whether witnessing a morally virtuous act without outcome information is better or worse for yielding helpful behaviour than negative outcome information.

3 Conclusion

Moral elevation leads to the desire to “be a better person” and to emulate others’ virtuousness (Algoe & Haidt, 2009), but moral elevation on its own might not be sufficient to prompt helping. Instead, it requires a “boost”—for example, from additional information that extends beyond the helping behaviour itself (in this case, outcome information). As in other helping contexts (e.g., Clary & Orenstein, 1991), the desire to help that accompanies the experience of moral elevation might translate better into helping behaviour when such behaviour is deemed worth the effort. These findings reveal the importance of witnessing moral virtue *and* its positive effects in motivating helping behaviour.

The question I aimed to answer with this experiment was about the effects of behavioural and contextual information on behaviour, an aspect of morality that has evaluative and affective components similar to moral judgements. Specifically, I was interested in understanding the effect of behavioural and contextual information on the only empirically studied positive moral emotion of moral elevation. Moral elevation served as a useful moral emotion for the study of behavioural and contextual information effects of behaviour as this emotion has consistently been demonstrated to mediate helping behaviour. Information about others’ helping behaviour influences one’s own helping behaviour, and the current study examined the specific qualifiers of information that were necessary for

eliciting behaviour. I have demonstrated that behavioural and contextual information, and the consideration of specific types of information, are influential pre-judgement and post-judgement in moral decision making, and this chapter demonstrates further the importance of behavioural and contextual information, and the specific types of information, on behaviour. Additionally, this chapter begins to redress the largely imbalanced study of how behavioural and contextual information and reasoning affect moral condemnation by demonstrating that behavioural and contextual information and reasoning are influential when making positive moral judgements.

This thesis has thus far examined how behavioural and contextual information affects moral judgements. I have provided evidence that consideration and reasoning about behavioural and contextual information prior to a making a moral judgement is influential to a moral judgement of blame, and that different types of information are considered after making moral judgements of wrongness depending on the type of moral violation that is judged. This chapter has added support for the overarching hypothesis that the consideration of behavioural and contextual information is valuable to how we construct our evaluations of morality by demonstrating that information regarding outcomes of virtuous behaviour has implications for moral behaviour. Additionally, this chapter has preliminary implications for how we consider the judgement process for positive moral judgements—an area that needs further exploration. The social intuitionist model of moral judgement (Haidt, 2001) may need to be closely examined for the role of reasoning (specifically, reasoning that accompanies the consideration of behavioural and contextual information) when making positive moral judgements, as well as the judgement process at different stages (pre-judgement, post-judgement and in behaviour).

CHAPTER 5

GENERAL DISCUSSION

1 Background and Aims of the Thesis

The social intuitionist account of moral judgement has focused largely on the role of intuition and emotion in determining our moral judgements (Haidt, 2001). Although this account does not negate the role of reasoning or the potential for conscious deliberation when making moral judgement (e.g., post-hoc reasoning), the literature would benefit from an examination of some of the instances when potentially deliberative processing beyond the “gut” reaction about behavioural and contextual information, can occur. For the purposes of this thesis, I have defined behavioural and contextual information as being information that, although not necessarily devoid of affective qualities, is primarily information that moves people from their initial intuitions and gut-reactions in response to the initial features of a moral violation (e.g., the harmful outcome). This thesis aimed to examine how people process and engage with different types of behavioural and contextual information at different points during the judgement process, and the effects on subsequent moral behaviour. In particular, this thesis aimed to examine how people can reason about and engage with behavioural and contextual information, and reasoning, in moral scenarios with respect to different:

- (a) points during the judgement process (i.e., pre-judgement, post-judgement),
- (b) processes and outcomes (i.e., pre-judgement reasoning, post-judgement information selection and engagement, post-evaluative behaviour),
- (c) moral scenarios (based on the moral foundations theory; Haidt & Graham 2007; Haidt & Joseph, 2004) and their respective eliciting moral emotions (Haidt, 2003),
- (d) moral evaluations (i.e., judgements of blame, wrongness and virtue), and

(e) types of information (e.g., wrongness-confirming, wrongness-contradicting information).

Considering the role of reasoning and evaluation of behavioural and contextual information is important for several reasons. Firstly, the debate about whether moral judgements are more intuitive than reasoned will benefit from a more nuanced understanding of the conditions under which potentially deliberative processing of behavioural and contextual information can occur and influence judgements and behaviour. Secondly, comparisons between the different facets of moral situations as listed above will allow for an overall understanding of how moral judgements respond differently to information and reasoning in different contexts. Finally, given motivation and context can play a role in moral judgement (e.g., Cushman, 2008; Greene & Haidt, 2002), this dissection of moral information into behavioural and contextual components will allow for further understanding of when people attend to this information, what types of information specifically people select and engage with, and how information and reasoning affect judgements and behaviour.

2 Summary of Findings

Chapter 1 outlined the shift in moral psychology from reasoning about harmful outcomes to emotional and intuitive snap judgements in response to non-deliberative “gut” feelings. The majority of current research on moral judgements tends to suggest that moral judgements are more reliant on gut-reactions, intuitions and emotions than conscious deliberative reasoning (Cushman, Young, & Hauser, 2006; Greene & Haidt, 2002; Haidt, 2001; Haidt, Bjorklund, & Murphy, 2000; Pizarro & Bloom, 2003). I aimed to identify a need to further explore more precisely the role of reasoning and deliberative processing of and engagement with information that goes beyond a person’s initial gut-reactions and emotions. The following chapters investigated the role of reasoning and deliberative evaluations of

behavioural and contextual information at different stages of the judgement process in seven experiments. I predicted that deliberative engagement with behavioural and contextual information would occur at the three stages that I investigated (pre-judgement, post-judgement, and in behaviour), but would be mediated by (1) the sophistication of reasoning ability, (2) the type of moral violation and elicited emotional reaction, (3) the type of information (e.g., outcome information, wrongness-confirming information) and (4) the effects of behavioural and contextual information on elicited moral behaviour that follows from witnessing a moral act.

2.1 Sophistication of Reasoning Ability in Processing Information

As discussed in Chapter 1, moral psychology was initially dominated by cognitive developmentalists who argued that morality developed in stages across the lifespan similarly to the development of other forms of reasoning. Following the influential paper by Haidt (2001) proposing a new model of moral judgement that emphasised the role of intuition ahead of reasoning, most research has found evidence to support the hypothesis that adults' moral judgements are driven largely by their emotional responses (Greene, Sommerville, Nystrom, Darley, & Cohen, 2001; Haidt, 2001; Pizarro, 2000; Prinz, 2006; Schnall, Haidt, & Clore, & Jordan, 2008; Wheatley & Haidt, 2005). It is important, however, to continue to understand the precise role of reasoning and intuition by examining how moral judgements change developmentally. If moral judgements are largely intuitive and affect-driven, then I would expect that moral judgements would not necessarily be influenced by cognitive sophistication. In other words, if we rely on a set of unconscious, intuitive and innate moral principles to guide our “snap” judgements, then we would not expect to see judgements changing dramatically on the basis of our cognitive ability to reason about different types of information. In two experiments, Chapter 2 tested whether moral judgements regarding

blaming another person for an accidental negative outcome would be influenced by the development of sophisticated reasoning about hypothetical alternative outcomes and consideration of the means through which a negative outcome occurs. Additionally, I tested whether blame judgements were influenced by changes in ability to reason about greater or fewer numbers harmed in the negative outcome. Blame judgements were compared alongside judgements of another person's experience of regret because judging another's experience of regret has been shown to be influenced by the ability to reason about hypothetical alternative outcomes (i.e., counterfactual reasoning) (e.g., Guttentag & Ferrell, 2004).

The findings from both experiments, alongside additional research that demonstrates that consideration of factors such as an actor's intention and desire to do harm are influential (Cushman, 2008), suggest that blame judgements—including blame judgements about active versus passive harm, one of the intuitive, natural moral principles—are largely influenced by consideration of behavioural and contextual information at the pre-judgement stage. The results from these the two studies presented in Chapter 2 demonstrated that information about the means through which a negative outcome occurred (i.e., actively versus passively; atypically versus typically) influenced blame judgements, as well as judgements of another's experience of regret, but only at the point during which participants were able to reason counterfactually—a form of sophisticated reasoning (Amsel & Smalley, 2000; Guttentag & Ferrell, 2004). This ability to incorporate counterfactual reasoning with judgements of blame did not occur in participants younger than 7 years of age. Younger children did not reason counterfactually about the means through which the negative outcome occurred, and therefore blamed characters who behaved differently but experienced the same negative outcome equally. Interestingly, blame judgements continued to change with age to the extent that adults reasoned about counterfactuals *and* about the accidental, unintended nature of the

negative outcome. Adults, therefore, blamed the characters equally, but unlike younger children, did so because they were reasoning about the fact that accidental negative outcomes are not blameworthy. Blame judgements were affected by reasoning and consideration of information that was provided pre-judgement to the extent that judgements of blame differed at all three ages that were tested.

Although there is substantial evidence to suggest that moral judgements are mainly driven by affect and intuition at the pre-judgement stage (Hodson & Costello, 2007; Haidt, 2001; Inbar, Pizarro, Knobe, & Bloom, 2009; Miller, 1997; Nussbaum, 2001), Chapter 2 provided experimental evidence to suggest that age-related changes in cognitive reasoning abilities affect how people process information, and that once able to reason, certain pre-judgement information (e.g., active versus passive harm) is influential to the resulting moral judgement. For instance, pre-judgement information about behaviours leading to a negative outcome is not considered by 6-year-olds attributing blame when those 6-year-olds are unable to reason about the information and focuses solely on harmful outcomes. Pre-judgement information, however, is influential for blame judgements made by 8-year-olds who factor in the more nuanced behavioural and contextual information into their judgement. Behavioural and contextual pre-judgement information continues to be important but is considered alongside further information (e.g., understanding about accidents) into adulthood. The experiments in Chapter 2, alongside additional developmental evidence suggesting age-related changes in the processing of moral principles (Powell et al., 2012), provided evidence that a developmental approach to moral judgements highlights the significance of the use of reasoning and deliberative consideration of information more so than the social intuitionist theory suggests.

2.2 Different Moral Violations and Emotions Lead to Different Engagement with Information

Following on from the social intuitionist theory and the findings that suggest moral judgements can occur in the absence of harmful outcomes, several moral foundations have been identified (Graham, Haidt, & Nosek, 2009; Haidt & Graham, 2007; Haidt & Joseph, 2004). The moral foundations are the organisation of people's moral intuitions into categories such as care/harm, fairness/cheating, liberty/oppression, loyalty/betrayal, authority/subversion and sanctity/degradation (Haidt, 2012). These moral foundations are believed to be innate (Haidt, 2012), yet have been demonstrated to differ on the basis of culture, political affiliation and socioeconomic status (Graham et al., 2009; Haidt & Graham, 2007). The moral foundations are considered to be similar to taste buds whereby everyone has the potential to experience sensitivity to each of the moral foundations, but some people on the basis of experience and environment may be more sensitive to some foundations at the cost of others (Haidt, 2012). For instance, research has demonstrated that people who come from a low socioeconomic background are more sensitive to the purity/sanctity foundation than people who come from a high socioeconomic background (Haidt, Koller, & Dias, 1993). When asked about the wrongness of purity violation where a harmful outcome was not present, people from the low socioeconomic background expressed disgust reactions that led to high wrongness judgements followed by circular, dumbfounded justifications for why the act was wrong. Princeton undergraduate students, however, who came from a high socioeconomic background, were more likely to abandon their initial emotional response in favour of reasoning about the lack of harm done and the private nature of the act which led to lesser wrongness judgements.

Moral violations associated with the harm and purity foundations have been shown to differ with respect to how people judge wrongness of the violations, their justifications for wrongness and the eliciting emotions. For instance, purity violations typically do not involve a victim or a harmful outcome, whereas harm violations do. Purity violations are typically associated with feelings of disgust whereas harm violations are typically associated with feelings of anger (e.g., Gutierrez & Giner-Sorolla, 2007). Additionally, wrongness judgements in response to harm violations can be mediated by information regarding the actor's intentionality, and in general are more responsive to mitigating factors than are purity violations (Russell & Giner-Sorolla, 2011a).

Chapter 3 provided four experiments that further investigated how harm and purity violations differ with respect to responsiveness to information at the post-judgement stage. Chapter 3 provided evidence that there are differences in post-judgement information selection and engagement on the basis of the type of violation that was judged; for instance, violations that elicit feelings of disgust (e.g., purity violations) lead to disgust-avoidance tendencies for information selection, and violations that result in high levels of judgement certainty (e.g., harm violations) lead to a desire to engage with information that maintains or boosts epistemic certainty. Post-judgement information selection and engagement is important to understand because it is often the case that new information will be brought to light after an initial judgement of wrongness has been made on the basis of little information. For instance, we might read a headline in a newspaper that gives the gist of the violation, including information about the harmful or impure outcome. There may be additional information to learn about regarding more nuanced aspects of the act or actor, such as what the actor intended, desired and believed about what s/he does that can play an important role in making an accurate judgement of wrongness, blame and subsequent punishment (e.g.,

Cushman, 2008). Chapter 3 explored what people report to want from potentially available information that can either mitigate or enhance the degree of wrongness of an act, and what people actually do with the information, on the basis of the type of moral violation that is being judged (harm versus purity).

Experiments 3 and 4 found that people's self-reported willingness and desire to engage with additional information differed as a function of the type of violation that was being judged. Participants reported a desire and need for additional behavioural and contextual information after judging a harm violation as opposed to a purity violation; purity violations yielded reports of not wanting or needing additional information. More specifically for harm violations, patterns of self-reported hypothetical information selection were consistent with a need to maintain and boost certainty for those participants who reported high levels of judgement certainty. Feelings of disgust appeared to drive participants' self-reported desire to disengage and avoid further post-judgement information. Disgust has been demonstrated to be associated with avoidance tendencies and therefore I argued that participants in Experiments 3 and 4 were likely reporting a lack of desire for post-judgement information because of a desire or need to repair feelings of disgust (e.g., Stevenson & Repacholi, 2005).

Experiments 5 and 6 aimed to further test how people actually engaged with post-judgement information. There were some inconsistencies between the self-reported desire and need to engage with post-judgement information reported in Experiments 3 and 4 and participants' actual selection and engagement reported in Experiments 5 and 6. In Experiment 5, the actual pattern of engagement with post-judgement information reflected no differences between violation types, such that the majority of participants read the majority of available information regardless of the violation type or the information's congruency with

participants' judgements (i.e., wrongness-confirming information versus wrongness-contradicting information). In support of an epistemic certainty maintenance explanation for information engagement, Experiment 5 did find that participants' ratings of judgement certainty were higher following engagement with wrongness-confirming information.

Experiment 6 gave participants the opportunity to select the type of post-judgement information with which they wished to engage. When participants were in control of the type of post-judgement information they could read, both those participants who were exposed to harm and those who were exposed to purity violations chose to engage with wrongness-contradicting information. This pattern supports an emotion regulation account for those choosing to engage with wrongness-contradicting information, as participants reported disgust in response to both harm and purity violations, and disgust has been demonstrated to lead to avoidance and alleviating tendencies (Gutierrez & Giner-Sorolla, 2007).

2.2.1 Limitations. I offer an explanation for the potential inconsistency between self-reported and actual post-judgement information engagement for those who judged harm violations, and a general limitation with the methodology used for exploring information selection and engagement specifically for harm violations. Firstly, participants in Experiments 3 and 4 who gave self-reports of their desire and willingness to engage with different types of information did so online. This means that participants were able to engage with the moral violation that they were exposed to with complete anonymity without any consideration of how they may be judged for reporting wanting or not wanting information about a violation. Participants in Experiments 5 and 6, in contrast, made their choices to engage with information in a lab setting with the presence of an experimenter. There may be a stigma associated with choosing to perpetuate a condemnation of another person's actions for no reasons associated with punishment or justice. Therefore, people may be less willing to report

wanting to engage with wrongness-confirming information that would indicate a desire to perpetuate their condemnation with the hope of appearing egalitarian and fair. In other words, it may not be desirable to appear to be the type of person who is interested in pursuing condemnation. The participants in Experiments 3 and 4, however, did not have this concern as there was no experimenter present (i.e., no one to judge) when participants reported their desire, or lack thereof, for additional information. The testing environment in Experiments 3 and 4 may reflect a similar environment to the environment when people are privately condemning someone in the news who has committed a moral violation.

One possible explanation for the lack of willingness to engage with post-judgement information following judgement of harm violations may be related to the private nature of engaging with wrongness-confirming information. According to the social intuitionist model, one of the reasons people engage in reasoning, albeit post-hoc, is to rally others to join their moral crusade (Haidt, 2001). Participants in Experiments 5 and 6 had the opportunity to engage with post-judgement wrongness-confirming information, as opposed to self-report their hypothetical preferences as in Experiments 3 and 4. For actual post-judgement information engagement, there was no opportunity for participants to feel encouraged to rally with others in support of their wrongness judgement or a need for participants to gain others' support, as they made their choices privately apart from the presence of a neutral experimenter. Anecdotally, there are cases of people perpetuating their condemnation in groups when gossiping about others' wrongdoings, such as when people rally in support of punishing a transgressor (e.g., groups of people online expressing support of Tiger Woods losing his endorsement following a sexual infidelity scandal). Social support from likeminded others may be the most likely situation to lead people to have a desire to condemn others in a harm violation, and Experiments 5 and 6 did not provide a conducive setting.

In essence, information appears to have an important post-judgement function that does vary on the basis of the violation type. The extent to which information impacts judgements on the basis of the violation remains unclear. Future research should investigate how people engage with wrongness-confirming and wrongness-contradicting information in private, low-individual-cost settings (i.e., when judging for the sake of judging as opposed to judging for the sake of convicting or punishing) where there is an opportunity to condemn without public scrutiny.

2.3 Impact of Behavioural and Contextual Information on Behaviour

The majority of experimental research presented in this thesis focused on the processing of and engagement with behavioural and contextual information pre-judgement and post-judgement, and when considering moral transgressions (i.e., blame and condemnation). The final experimental chapter in this thesis examined the effects of processing behavioural and contextual information on behaviour, and for the consideration of acts of moral good deeds that elicit positive feelings of moral elevation.

The majority of research on evaluations of moral good deeds has focused largely on the experience of moral elevation in response to witnessing moral good deeds (Algoe & Haidt, 2005; Englander, Haidt, & Morris, 2012; Haidt, 2000, 2003; Schnall & Roper, 2012; Schnall, Roper, & Fessler, 2010). One of the main functions of moral elevation that has been identified in the literature is to elicit further moral good deeds. Schnall, Roper and Fessler (2010) demonstrated that participants who watched a video depicting a helping act and who reported feelings of moral elevation in response to witnessing the helping act were more likely to engage in a future unrelated act of helping than those participants who watched a control video depicting a humorous situation designed to elicit feelings of mirth. As one of the goals of this thesis has been to examine the role of processing behavioural and contextual

information on behaviour, I explored the effects of witnessing a helpful act on subsequent behaviour by dissecting the helpful act used in Schnall, Roper and Fessler (2010) into separate components based on behavioural and contextual information presented.

Experiment 7 examined the helpful act presented in Schnall, Roper and Fessler (2010) to investigate whether separating the information about the outcome of the helping yielded differences in participants' later willingness to engage in their own helpful act. The results suggest that participants who were only exposed to the helpful act that did not contain information regarding the outcome (either success for the beneficiary or praise for the helper) were no more likely to help the experimenter than those participants who were exposed to the control video that did not depict any moral or helpful content. Therefore, I argued that this provides preliminary evidence to suggest that some degree of processing of outcome information is needed for subsequent helpful behaviour to occur. Behavioural and contextual information does appear to be considered pre- and post-judgement, but also to have important implications for eliciting moral behaviour. Additional studies are needed, however, to determine whether outcome information regarding the efficacy of a helping act compared to outcome information regarding praise and reward for the helper are equally or differentially important for eliciting a further act of helping.

3 Implications for Models of Moral Judgement

The research presented in this thesis provides evidence to suggest a more nuanced understanding of how consideration of information and reasoning can affect moral judgements. Additionally, this thesis aimed to dissect the process of moral judgement into separate parts: pre-judgement reasoning, post-judgement information search and behaviour. The dominant theory in current moral psychology—the social intuitionist model—suggests that emotions and intuitions drive judgements, with reasoning playing a secondary role that

serves the purpose of convincing others to join our moral crusade (Haidt, 2001). Although the evidence presented in this thesis does not directly contradict the social intuitionist account, it does highlight the importance of taking information into account as opposed to relying solely on emotion. For instance, Chapter 2 suggests that reasoning about hypothetical outcomes and actions does impact moral judgements. Chapter 3 suggests that although initial judgements can be made on the basis of emotion and intuition, information following judgements does serve a function related to the self (e.g., maintaining certainty). Chapter 4 demonstrated that emotions that are elicited from witnessing virtue may not be useful for motivating behaviour unless accompanied by additional information. Taken together, these findings may challenge or highlight limitations of the social intuitionist account of moral judgements.

Firstly, the social intuitionist model should take into account the points during the judgement process when deliberation can occur at least in addition to, but perhaps in favour of, intuition and gut-reactions. Empirical evidence from Experiments 1 and 2 suggest that with the appropriate cognitive reasoning sophistication, people can attend to behavioural and contextual information that goes beyond initial intuitions and affective responses, and rely on reasoning about the features of a scenario prior to making a judgement about who to blame for a negative outcome. People who have developed the cognitive sophistication to attend to features of a situation, such as whether harm is active or passive, and to compare hypothetical alternative behaviours/outcomes with actual behaviours/outcomes will attribute blame differently than those without that cognitive sophistication. Furthermore, blame judgements continue to change after the cognitive maturity of late childhood, as adults apply a rational correction to their blame judgements that accounts for accidental, unintended, unforeseen negative outcomes.

Experiments 3 and 4 provided preliminary evidence that the desire and need for further post-judgement information is driven by motivations for the self (i.e., regulating affect and maintaining certainty) and may differ on the basis of the type of violation that is judged. If people have a desire and motivation to seek information post-judgement, then perhaps the typical reliance on emotion and intuition is due to a lack of consideration about potential additional information, a lack of motivation to search for additional information when information is not provided, or a lack of available post-judgement information. This does not necessarily refute the social intuitionist model, but rather emphasises that people *can* use conscious deliberation when considering post-judgement information. Additionally, some of the preliminary findings from Chapter 3 suggest that people may have internal, self-related reasons for seeking information: Information maintains judgement certainty in the context of judging harmful violations, or alleviates an undesirable emotional experience in the context of judging disgusting violations. Furthermore, Experiment 7 provided evidence to suggest that consideration of certain types of behavioural and contextual information (i.e., information about the outcomes of a helpful act) affects behaviour following judgement. Emotion may not be the primary motivator of behaviour in some contexts (e.g., judging and witnessing virtuous acts) if the emotion is not accompanied by other relevant information. The process of considering behavioural and contextual information may be a deliberative and conscious process more than originally posited by the social intuitionist account (Haidt, 2001). If so, however, there needs to be further research to determine whether the presence and consideration of information in the contexts described above are truly deliberative and conscious.

Secondly, the social intuitionist model should take into account that its proposal for the moral judgement process might differ with respect to differing moral foundations. The

majority of research in support of the social intuitionist account for moral judgement has focused on negative moral judgements of wrongness, condemnation and blame. Moral elevation research did provide some support for the role of emotion following judgements of virtue and in eliciting behaviour. The research on moral elevation, however, may need to consider that there may be deliberative and reasoned elements to how we judge acts of helping and virtue, as evidenced from Experiment 7. Emotion may be acting in combination with other more deliberative processes for judgements of morally good and virtuous behaviours, especially when that morally good and virtuous behaviour leads to further good and virtuous behaviours. Furthermore, there is a need to consider how deliberative processing of information affects not just evaluations of moral wrongdoings, but also evaluations of morally good deeds and acts of virtue. Evaluations of moral good deeds and acts of virtue may interact differently with behavioural and contextual information than evaluations of moral wrongdoings. There is no direct comparison between how deliberative, reasoned or intuitive judgements of moral wrongdoings are compared to judgements of moral good deeds. Drawing on other literature (e.g., how we attribute honesty), I will present a case for a potential asymmetry in judging moral good deeds and moral wrongdoings that would suggest judging moral wrongdoings relies on a less reasoned process than judging moral good deeds, or at least that the evidentiary standards for the two types of judgements may differ.

While there is compelling evidence to suggest that moral judgements are often made on the basis of emotion and intuition, the literature would benefit from a more contextual approach to understanding and modelling the process of moral judgements that takes into account the different stages and context-specific processes of moral judgements, and their relationship to available and potential information. Furthermore, there may be different motivational factors for reasoning and information search underpinning the different stages;

pre-judgement reasoning can be informative for judgement accuracy, post-judgement information search can be desirable for satisfying self-related motives, and behaviour can be modified. In support of a motivational basis for moral judgements, recent research has identified that although emotions play a role in moral judgements, their effect on moral judgements can be determined by their motivational qualities (Ugazio, Lamm, & Singer, 2012). If we are motivated to escape from a moral situation or the elicited emotion following our judgement of that situation, then we might rely heavily on our emotional reaction such as disgust. Disgust can be a useful and effective avoidance mechanism (e.g., Chapman, Kim, Susskind, & Anderson, 2009). If, however, we are motivated to be helpful after witnessing others' helping behaviour, we might rely heavily on having information about how others' helpfulness was rewarded and/or successful. In essence, moral psychology should not be a study of how we always use intuition over reason, but rather a study of when, in what contexts, and for what purpose we use intuition, and when we use reason, deliberative processing and additional information.

Overall, the framework for the research presented in this thesis may fit best with a dual-process account for moral judgements. The dual-process theory for moral judgements (Greene, Nystrom, Engell, Darley, & Cohen, 2004) suggests that moral judgements involve both deliberative reasoning and intuition because the mechanisms behind moral judgements are both affective and cognitive; the affective is automatic and intuitive whereas the cognitive component is reasoned and controlled. Research in support of this theory has also identified that different judgements may rely on different systems, some that are affective more than reasoned and vice versa (Cushman, Young, & Greene, 2010). The research presented in this thesis suggests that although emotion and intuition can generate moral judgements, reasoning

and consideration of information are also influential to moral judgements, and potentially dependent on the type of moral violation and moral judgement.

4 Future Directions

The majority of this thesis has examined how information affects judgements of wrongness and condemnation. The majority of the moral psychology literature also focuses on wrongness and condemnation. Chapter 4 aimed to apply my hypotheses regarding the importance of information to more than just condemning moral judgements. The findings from Experiment 7 suggest that information about positive outcomes of others' helping behaviour may be useful for encouraging others to help after witnessing someone else's helpful act. In doing so, I have speculated about the potential for an asymmetrical relationship between positive and negative moral judgements that I aim to investigate.

4.1 Asymmetry between moral condemnation and moral praise.

The majority of research in moral psychology has focused on condemnation over praise: A search of the current PsycINFO database yields 245 hits for moral anger, for example, but only 48 for moral elevation. Although excellent progress has recently been made in understanding positive moral emotions (e.g., Algoe & Haidt, 2009), contemporary moral psychology can nonetheless be characterised more as the study of moral condemnation than praise. There are reasons why condemnation might outweigh praise, both as a topic of study and as a driving force in moral judgement. First, there is a well-documented "negativity bias" in human information processing (Baumeister, Bratslavsky, Finkenauer, & Vohs, 2001; Jordan, 1965), which might support greater attention to moral vice than moral virtue (and which might have even prompted more research interest in condemnation than praise). Relatedly, a wealth of evidence suggests that people are averse to loss (e.g., Kahneman & Tversky, 1979; Tversky & Kahneman, 1986), to the extent that they are even more motivated

to avoid a loss than acquire a gain (e.g., Tversky & Kahneman, 1992), which suggests that individuals should generally be more motivated to avoid condemnation than earn praise.

Beyond these general biases, however, there might also be a morality-specific reason to focus on condemnation over praise: Cosmides and Tooby (2005) have argued that the human mind has evolved a neurocognitive system that is specialised for reasoning about social exchange, characterised by a “subroutine” for detecting cheaters. That is, they have posited a morality-specific negativity bias, driven by the importance of identifying individuals who cannot be trusted. In brief, we may be “wired” to process negative information, making vice and condemnation likely candidates as the core of morality.

In addition to a system evolved for detecting cheaters, people might also have different standards of evidence for “good” versus “bad” moral character. Personality traits differ in the extent to which they are “hierarchically restrictive”—that is, the extent to which one end of a trait continuum is associated with a wider range of possible behaviours than is the other end of that continuum (Devine, Hirt, & Gehrke, 1990)—and honesty is one such trait. People view dishonest behaviour as more diagnostic of personality than honest behaviour: Dishonest people can engage in both honest and dishonest behaviour, but honest people can only engage in honest behaviour; one instance of dishonest behaviour effectively makes a dishonest person of an honest one (McGraw, 1985; Reeder & Coovert, 1986; Skowronski & Carlston, 1989). Honesty can be viewed as a “moral” trait, suggesting that moral judgement might be subject to the same hierarchically restrictive representation. That is, it might be easier to consider the presence versus absence of evidence for morally “bad” than “good” behaviour: In judging the moral character of actors, one instance of harm might undermine many instances of care, one instance of betrayal might undermine many instances of loyalty, and so on.

Research in Chapter 4 provided some preliminary evidence to suggest that moral behaviour resulting from the experience of moral emotions may require a knowledge/information base to encompass information such as whether or not a helpful act was successful and/or whether or not the helper was praised for helping. Research has demonstrated that feelings of elevation in response to witnessing virtuous behaviour can lead to helping behaviour (Schnall, Roper, & Fessler, 2010). Experiment 7, however, found that helping behaviour may not solely result from witnessing the virtuous behaviour in isolation, but may require having knowledge of either a positive outcome for the help recipient or praise of the helper. This evidence suggests that the positive effects of witnessing morally virtuous behaviour are reliant on more information than was previously thought, suggesting that people are in need of extensive proof and confirmation of virtue in order for that virtuous act to inspire a further helpful behaviour. Additionally, this adds to the imbalanced understanding of how information can play a hugely influential role in positive moral judgements as well as negative moral judgements.

Reactions to “moral excellence” can enhance human functioning. For example, compelling evidence indicates that gratitude strengthens social connectedness and resilience. In a series of diary studies, Emmons and McCullough (2003) found that participants who were assigned to write about people and events that made them feel grateful, compared to participants who wrote about neutral or negative life events, showed higher well-being (e.g., coping strategies, life satisfaction, health behaviours). Additionally, studies by Algoe and Haidt (2009) found that other-praising emotions, including gratitude and elevation, lead to increased pro-social and affiliative behaviour, improved relationships, and striving for self-improvement. Thus, where condemnation protects us from those who pose threat (Cosmides & Tooby, 2005), praise and gratitude connect us to those who provide strength.

In this sense, both condemnation and praise are useful to human functioning, albeit at different levels. Condemnation may be particularly functional at the societal level, ensuring that norms supporting cooperation and inhibiting cheating are enforced (Cosmides & Tooby, 2005). The societal value of condemnation, however, may come at the cost of individual functioning by leading to a focus on negative emotion. Research on human flourishing suggests that individuals need to experience approximately three times as much positive as negative affect to live optimally (Fredrickson & Losada, 2005). This suggests that praise may be more crucial than condemnation at the individual level, in promoting personal wellbeing and self-improvement.

Future research should redress the imbalance in moral psychology by investigating whether/how praise differs from condemnation. Based on negativity bias and the potential for different evidentiary standards for “goodness” versus “badness”, future research should explore the general hypothesis that people are more motivated and able to condemn than praise—that they are more interested in vice than virtue, that vice is processed more fluently than virtue. Additionally, future research should conduct an investigation of the impact of praise versus condemnation on people’s wellbeing.

5 Conclusion

Theoretical research on moral psychology has identified the potential oddness of moral decision-making coming from the ability for unaffected and uninvolved people to make confident judgements on the basis of little information (DeScioli & Kurzban, 2009). Further research suggests that attitudes in the moral realm may be different from other non-moral attitudes in part due how conviction and universality make moral attitudes become moral mandates (Mullen & Skitka, 2006; Skitka & Mullen, 2002). The literature on attitudes, and in part moral judgement, suggests that information and knowledge can mediate the strength of

an attitude and its impact on information search and engagement. Moving away from cognitive developmental theories of moral reasoning, the moral psychology literature has focused largely on moral judgements being made on the basis of allegedly innate intuitions and emotions that are only served by limited reasoning. This thesis aimed to break down the content and process of moral judgement so as to more closely examine how reasoning and information can be influential. The findings presented in this thesis demonstrated that reasoning and information can be influential, desirable and potentially motivating depending on the type of information and moral violation. The effects of reasoning and information are variable and contextual, but are evident in pre-judgement reasoning, post-judgement information search, and behaviour following judgement.

REFERENCES

- Albright, L., Kenny, D. A., & Malloy, T. E. (1988). Consensus in personality judgments at zero acquaintance. *Journal of Personality and Social Psychology*, *55*, 387–395.
- Algoe, S. B., & Haidt, J. (2009). Witnessing excellence in action: The ‘other praising’ emotions of elevation, gratitude, and admiration. *Journal of Positive Psychology*, *4*, 105–127.
- Ambady, N., & Rosenthal, R. (1993). Half a minute: Predicting teacher evaluations from thin slices of behavior and physical attractiveness. *Journal of Personality and Social Psychology*, *64*, 431–441.
- Amsel, E., & Smalley, D. (2000). Beyond really and truly: Children’s counterfactual thinking about pretend and possible worlds. In K. Riggs & P. Mitchell (Eds.), *Children’s reasoning and mind* (pp. 99–134). Brighton, England: Psychology Press.
- Amsel, E., Robbins, M., Tumarkin, T., Janit, A., Foulkes, S., & Smalley, J.D. (2003). *The card not chosen: The development of judgments of regret in self and others*. Unpublished manuscript, Weber State University.
- Baron, J., & Ritov, I. (2004). Omission bias, individual differences, and normality. *Organizational Behavior and Human Decision Processes*, *94*, 74–85.
- Batson, C. D. (1991). *The altruism question: Toward a social-psychological answer*. Hillsdale, NJ: Erlbaum.
- Batson, C. D., Duncan, D., Ackerman, P., Buckley, T., & Birch, K. (1981). Is Empathic Emotion a Course of Altruistic Motivation? *Journal of Personality and Social Psychology*, *40*, 290–302.
- Batson, C. D., Kennedy, C. L., Nord, L., Stocks, E. L., Fleming, D. A., Marzette, C. M., Lishner, D. A., Hayes, R. E., Kolchinsky, L., M. & Zerger, T. Anger at unfairness: is it moral outrage? *European Journal of Social Psychology*, *37*, 1272–1285.
- Baumeister, R. F., Bratslavsky, E., Finkenauer, C., & Vohs, K. D. (2001). Bad is stronger than good. *Review of General Psychology*, *5*, 323–370.
- Beck, S.R. & Crilly, M. (2009). Is understanding regret dependent on developments in counterfactual thinking? *British Journal of Developmental Psychology*, *27*, 505–510.
- Berent, M. K., & Krosnick, J. A. (1993). Comparison of Party Identification and Policy Preferences: The Impact of Survey Question Format. *American Journal of Political Science*, *37*, 941–964.
- Berent, M. K., Krosnick, J. A., & Boninger, D. S. (1993). *Attitude importance and memory for attitude-relevant information*. Unpublished manuscript, Ohio State University, Columbus, Ohio.

- Biek, M., Wood, W., & Chaiken, S. (1996). Working Knowledge, Cognitive Processing, and Attitudes: On the Determinants of Bias. *Personality and Social Psychology Bulletin*, *22*, 547-556.
- Boninger, D. S., Krosnick, J. A., & Berent, M. K. (1995). Origins of attitude importance: Self-interest, social identification, and value relevance. *Journal of Personality and Social Psychology*, *68*, 61-80.
- Boninger, D. S., Krosnick, J. A., & Berent, M. K. (1995). The origins of attitude importance: Self-interest, social identification, and value relevance. *Journal of Personality and Social Psychology*, *68*, 61– 80.
- Burns, P., Riggs, K. J., & Beck, S. R. (2012). Executive control and the experience of regret. *Journal of Experimental Child Psychology*, *111*, 501–515.
- Cacioppo, J. T., Petty, R. E., Feinstein, J. A., & Jarvis, W. B. G. (1996). Dispositional Differences in Cognitive Motivation: The Life and Times of Individuals Varying in Need for Cognition. *Psychological Bulletin*, *119*, 197-253.
- Canon, L. K. (1964). Self-confidence and selective exposure to information. *Conflict, decision, and dissonance*, *1*, 83-95.
- Carver, C. S., & Harmon-Jones, E. (2009). Anger is an approach-related affect: Evidence and implications. *Psychological Bulletin*, *135*, 183–204.
- Chapman, H. A., & Anderson, A. K. (2011). Varieties of moral emotional experience. *Emotion Review*, *3*, 255-257.
- Chapman, H. A., Kim, D. A., Susskind, J. M., & Anderson, A. K. (2009). In bad taste: Evidence for the oral origins of moral disgust. *Science*, *323*, 1222–1226.
- Cialdini, R. B., & Kenrick D. T. (1976). Altruism as hedonism: a social development perspective on the relationship of negative mood state and helping. *Journal of Personality and Social Psychology*, *34*, 907-914.
- Cialdini, R. B., Darby, B. L., & Vincent, J. E. (1973). Transgression and altruism: A case for hedonism. *Journal of Experimental Social Psychology*, *9*, 502–516.
- Clary, E. G., & Orenstein, L. (1991). The amount and effectiveness of help: The relationship of motives and abilities to helping behavior. *Personality and Social Psychology Bulletin*, *17*, 58–64.
- Cosmides, L., & Tooby, J. (2005). Neurocognitive adaptations designed for social exchange. In D. M. Buss (Ed.), *The handbook of evolutionary psychology* (pp. 584-627). New York, NY: John Wiley and Sons.
- Cox, K. S. (2010). Elevation predicts domain-specific volunteerism 3 months later. *Journal of Positive Psychology*, *5*, 333–341.

- Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analysis in moral judgment. *Cognition*, *108*, 353–380.
- Cushman, F., Young, L., & Hauser, M. D. (2006). The role of conscious reasoning and intuition in moral judgment: Testing three principles of harm. *Psychological Science*, *17*, 1082–1089.
- Cushman, F. A., Young, L., & Greene, J. D. 2010: Multi-system moral psychology. In J. Doris et al. (Eds.), *The Oxford Handbook of Moral Psychology* (pp. XX–XX). City, State/Country: Oxford University Press.
- Davidson, A. R. (1995). From attitudes to actions to attitude change: The effects of amount and accuracy of information. In R. E. Petty, & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 315-336). City, NJ: Lawrence Erlbaum Associates.
- de Hooge, I. E., Zeelenberg, N. & Breugelmans, S. M. (2007). Moral sentiments and cooperation: Differential influence of shame and guilt. *Cognition & Emotion*, *21*, 1025-1042.
- Depue, R. A., & Iacono, W. G. (1989). Neurobehavioral aspects of affective disorders. *Annual Review of Psychology*, *40*, 457–492.
- DeScioli, P., & Kurzban, R. (2009). Mysteries of morality. *Cognition*, *112*, 281–299.
- Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, *56*, 5-18.
- Devine, P. G., Hirt, E. R., & Gehrke, E. M. (1990). Diagnostic and confirmation strategies in trait hypothesis testing. *Journal of Personality and Social Psychology*, *58*, 952-963.
- Eagly, A. H., & Chaiken, S. (1995). Attitude Strength, Attitude Structure, and Resistance to Change. In R. E. Petty, & J. A. Krosnick (Eds.), *Attitude Strength: Antecedents and Consequences* (pp. 413-432). Mahwah, NJ: Lawrence Erlbaum Associates.
- Emmons, R. A., & McCullough, M. E. (2003). Counting Blessings Versus Burdens: An Experimental Investigation of Gratitude and Subjective Well-Being in Daily Life. *Journal of Personality and Social Psychology*, *84*, 377-389.
- Fazio, R. H., & Zanna, M. P. (1978). On the predictive validity of attitudes: The roles of direct experience and confidence. *Journal of Personality*, *46*, 228-243.
- Feinberg, M., Willer, R., Antonenko, O., & John, O. P. (2012). Liberating reason from the passions: Overriding intuitionist moral judgments through emotion reappraisal. *Psychological Science*, *23*, 788-795.
- Fischer, A. H., & Roseman, I. J. (2007). Beat them or ban them: The characteristics and social functions of anger and contempt. *Journal of Personality and Social Psychology*, *93*, 103–115.

- Foot, P. (1978). *Virtues and vices*. Oxford: Blackwell.
- Förster, J., Liberman, N., & Friedman, R. S. (2009). What Do We Prime? On Distinguishing Between Semantic Priming, Procedural Priming, and Goal Priming. In E. Morsella, J. A. Bargh., & P. M. Gollwitzer (Eds.), *Oxford Handbook of Human Action* (pp. 173-193). Oxford: Oxford University Press.
- Förster, J., Liberman, N., & Higgins, T. (2005). Accessibility from active and fulfilled goals. *Journal of Experimental Social Psychology, 41*, 220-239.
- Fox, N. A. (1991). If it's not left, it's right: Electroencephalogram asymmetry and the development of emotion. *American Psychologist, 46*, 863-872.
- Fredrickson, B. L. (1998). What good are positive emotions? *Review of General Psychology, 2*, 300-319.
- Fredrickson, B. L., & Losada, M. F. (2005). Positive Affect and the Complex Dynamics of Human Flourishing. *American Psychologist, 60*, 678-686.
- Freeman, D., Aquino, K., & McFerran, B. (2009). Overcoming beneficiary race as an impediment to charitable donations: Social dominance orientation, the experience of moral elevation and donation behavior. *Personality and Social Psychology Bulletin, 35*, 72-81.
- Frijda, N. H. (1986). *The emotions*. London, England: Cambridge University Press.
- Gilovich, T., Medvec, V. H., & Kahneman, D. (1998). Varieties of regret: A debate and partial resolution. *Psychological Review, 105*, 602-605.
- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology, 96*, 1029-1046.
- Greene, J. D. (2007). The secret joke of Kant's soul. In W. Sinnott-Armstrong (Ed.), *Moral psychology* (Vol. 3, pp. 35-79). Cambridge, MA: MIT Press.
- Greene, J. D. (2009). Dual-process morality and the personal/impersonal distinction: A reply to McGuire, Langdon, Coltheart, and Mackenzie. *Journal of Experimental Social Psychology, 45*, 581-584.
- Greene, J. D., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive load selectively interferes with utilitarian moral judgment. *Cognition, 107*, 1144-1154.
- Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The Neural Bases of Cognitive Conflict and Control in Moral Judgment. *Neuron, 44*, 389-400.
- Greene, J., & Haidt, J. (2002). How (and where) does moral judgment work? *Trends in Cognitive Sciences, 6*, 517-523.

- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI investigation of emotional engagement in moral Judgment. *Science*, *293*, 2105-2108.
- Gross, J. J. (1998). The emerging field of emotion regulation: An integrative review. *Review of General Psychology*, *2*, 271 – 299.
- Gross, J. J. (2002). Emotion regulation: Affective, cognitive, and social consequences. *Psychophysiology*, *39*, 281 – 291.
- Gross, J. J. (Ed.). (2007). *Handbook of Emotion Regulation*. Guilford: New York.
- Gutierrez, R. & Giner-Sorolla, R. (2011). Disgusting but harmless moral violations are perceived as harmful due to the negative emotions they elicit. *Revista de Psicología Social*, *26*, 141-148.
- Gutierrez, R., & Giner-Sorolla, R. S. (2007) Anger, disgust, and presumption of harm as reactions to taboo-breaking behaviors. *Emotion*, *7*, 853–868.
- Gutierrez, R., Giner-Sorolla, R., & Vasiljevic, M. (2012). Just an anger synonym? Moral context influences predictors of disgust word use. *Cognition & Emotion*, *26*, 53-64.
- Guttentag, R. E., & Ferrell, J. (2004). Reality compared with its alternatives: Age differences in judgments of regret and relief. *Developmental Psychology*, *40*, 764–775.
- Haddock, G., Rothman, A. J., & Schwarz, N. (1996). Are (some) reports of attitude strength context dependent? *Canadian Journal of Behavioural Science*, *28*, 313-316.
- Haidt, J. (2001). The emotional dog and its rational tail. *Psychological Review*, *108*, 814–834.
- Haidt, J. (2003a). Elevation and the positive psychology of morality. In C. L. M. Keyes & J. Haidt (Eds.), *Flourishing: Positive psychology and the life well-lived* (pp. 275–289). Washington DC: American Psychological Association.
- Haidt, J. (2003b). The moral emotions. In R. J. Davidson, K. R. Scherer, & H. H. Goldsmith (Eds.), *Handbook of affective sciences* (pp. 852–870). Oxford: Oxford University Press.
- Haidt, J. (2008). Morality. *Perspectives on Psychological Science*, *3*, 65–72.
- Haidt, J. (2012). *The Righteous Mind: Why Good People Are Divided by Politics and Religion*. Pantheon: New York.
- Haidt, J., & Graham, J. (2007). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, *20*, 98–116.
- Haidt, J., & Hersh, M. (2001). Sexual morality: The cultures and emotions of conservatives and liberals. *Journal of Applied Social Psychology*, *31*, 191-221.

- Haidt, J., & Joseph, C. (2004). Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus*, 133(4), 55–66.
- Haidt, J., Bjorklund, F., & Murphy, S. (2000). Moral dumbfounding: When intuition finds no reason. Unpublished manuscript, University of Virginia.
- Haidt, J., Koller, S., & Dias, M. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology*, 65, 613-628.
- Hauser, M. D. (2006). *Moral Minds: How Nature Designed Our Universal Sense of Right and Wrong*. Ecco/Harper Collins: New York.
- Hauser, M., Cushman, F., Young, L., Kang-Xing, R. J., & Mikhail, J. (2007). A Dissociation Between Moral Judgments and Justifications. *Mind & Language*, 22, 1-21.
- Hodson, G., & Costello, K. (2007). Interpersonal disgust, ideological orientations, and dehumanization as predictors of intergroup attitudes. *Psychological Science*, 18, 691-698.
- Horberg, E. J., Oveis, C., Keltner, D., Cohen, A.B. (2009). Disgust and the moralization of purity. *Journal of Personality and Social Psychology*, 97, 963–976.
- Inbar, Y., Pizarro, D. A., Knobe, J., & Bloom, P. (2009). Disgust Sensitivity Predicts Intuitive Disapproval of Gays. *Emotion*, 9, 435-439.
- Jonas, E., Graupmann, V., & Frey, D. (2006) The Influence of Mood on the Search for Supporting Versus Conflicting Information: Dissonance Reduction as a Means of Mood Regulation? *Personality and Social Psychology Bulletin*, 32, 3-15.
- Jonas, E., Greenberg, J., & Frey, D. Connecting Terror Management and Dissonance Theory: Evidence that Mortality Salience Increases the Preference for Supporting Information after Decisions. *Personality and Social Psychology Bulletin*, 29, 1181-1189.
- Jordan, N. (1965). The asymmetry of liking and disliking. A phenomenon meriting further reflection and research. *Public Opinion Quarterly*, 29, 315–322.
- Judd, C. M., & Johnson, J. T. (1981). Attitudes, polarization and diagnosticity: Exploring the effect of affect. *Journal of Personality and Social Psychology*, 41, 26-36.
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93, 136–153.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, 47, 263-291.
- Kanwar, R., Grund, L., & Olson, J. C. (1990). When Do the Measures of Knowledge Measure What We Think They Are Measuring? In M. E. Goldberg, G. Gorn, & R. W. Pollay (Eds.), *Advances in Consumer Research* (pp. 603-608). Duluth, MN: Association for Consumer Research.

- Kohlberg, L. (1969). *Stages in the development of moral thought and action*. New York: Holt, Rinehart & Winston.
- Kohlberg, L. (1973). The Claim to Moral Adequacy of a Highest Stage of Moral Judgment". *Journal of Philosophy*, 70, 630–646.
- Kohlberg, L. (1976). Moral stages and moralization: The cognitive-developmental approach. In T. Lickona (Ed.) *Moral development and behavior: Theory, research, and social issues* (pp. 31-53). New York: Holt, Rinehart and Winston.
- Kohlberg, L. (1986). A current statement on some theoretical issues. In S. Modgil & C. Modgil (Eds.) *Lawrence Kohlberg: Consensus and controversy* (pp. 485–546). Philadelphia: Falmer.
- Kohlberg, L. (1971). From is to ought: How to commit the naturalistic fallacy and get away with it in the study of moral development. In T. Mischel (Ed.) *Cognitive development and epistemology* (pp. 151-235). New York: Academic Press.
- Krosnick, J. A., & Petty, R. E. (1995). Attitude strength: An overview. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 413–432). Mahwah, NJ: Erlbaum.
- Krosnick, J. A., & Shuman, H. (1988). Attitude intensity, importance, and certainty and susceptibility to response effects. *Journal of Personality and Social Psychology*, 54, 940-952.
- Krosnick, J. A., Boninger, D. S., Chuang, Y. C., Berent, M. K., & Carnot, C. G. (1993). Attitude strength: One construct or many related constructs? *Journal of Personality and Social Psychology*, 65, 1132-1151.
- Krosnick, J., & Telhami, S. (1995). Public Attitudes Toward Israel: A Study of the Attentive and Issue Publics. *International Studies Quarterly*, 39, 535-554.
- Lord, C. G., Ross, L., & Lepper, M. R. (1979). Biased Assimilation and Attitude Polarization: The Effects of Prior Theories on Subsequently Considered Evidence. *Journal of Personality and Social Psychology*, 37, 2098-2109.
- Martin, L. L., Ward, D.W., Achee, J.W., & Wyer, R. S., Jr. (1993). Mood as input: People have to interpret the motivational implications of their moods. *Journal of Personality and Social Psychology*, 64, 317–326.
- McGraw, M. (1985). Subjective probabilities in moral judgments. *Journal of Experimental Social Psychology*, 21, 501-518.
- McGuire, W. J. (1986). The vicissitudes of attitudes and similar representational constructs in twentieth-century psychology. *European Journal of Social Psychology*, 16, 89-130.
- Miller, N. (1944). Experimental studies of conflict. In J. M. Hunt (Ed.), *Personality and the behavior disorders* (Vol. 1, pp. 431–465). New York: Ronald Press.

- Miller, N. E., & Dollard, J. (1941) *Social learning and imitation*. Yale University Press.
- Miller, W. I. (1997). *The anatomy of disgust*. Cambridge, MA: Harvard University Press.
- Mullen, E., & Skitka, L. J. (2006). Exploring the psychological underpinnings of the moral mandate effect: Motivated reasoning, group differentiation, or anger? *Journal of Personality and Social Psychology*, *90*, 629-643.
- Murphy, J. M., & Gilligan, C. (1980). Moral development in late adolescence and adulthood: A critique and reconstruction of Kohlberg's theory. *Human Development*, *23*, 77–104.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, *84*, 231.
- Nucci, L., & Turiel, E. (1978). Social interactions and the development of social concepts in preschool children. *Child Development*, *49*, 400-407.
- Nussbaum, M. C. (2001). *Upheavals of thought: The intelligence of emotions*. New York: Cambridge University Press.
- Oaten, M., Stevenson, R. J., & Case, T. I. (2009). Disgust as a disease-avoidance mechanism. *Psychol Bull*, *135*, 303-321.
- Park, J. H., Faulkner, J., & Schaller, M. (2003). Evolved disease-avoidance processes and contemporary anti-social behavior: Prejudicial attitudes and avoidance of people with physical disabilities. *Journal of Nonverbal Behavior*, *27*, 65–87.
- Pelham, B. W. (1991). On confidence and consequence: The certainty and importance of self-knowledge. *Journal of Personality and Social Psychology*, *60*, 518-530.
- Petty, R. E., & Cacioppo, J. T. (1986). *Communication and persuasion: Central and peripheral routes to attitude change*. Springer-Verlag: New York.
- Petty, R. E., Cacioppo, J. T., & Goldman, R. (1981). Personal involvement as a determinant of argument-based persuasion. *Journal of Personality and Social Psychology*, *41*, 847-855.
- Piaget, J. (1932/1965). *The moral judgment of the child*. New York: Free Press. Pizarro, D. (2000). Nothing More than Feelings? The Role of Emotions in Moral Judgment. *Journal for the Theory of Social Behaviour*, *30*, 355-375.
- Pizarro, D. A., & Bloom, P. (2003). The intelligence of the moral intuitions: A reply to Haidt (2001). *Psychological Review*, *110*, 193–196.
- Powell, N. L., Derbyshire, S. W. G., & Guttentag, R. E. (2012). Biases in children's and adults' moral judgments. *Journal of Experimental Child Psychology*, *113*, 186–193.
- Prinz, J. J. (2006). The emotional basis of moral judgments. *Philosophical Explorations*, *9*, 29-43.

- Reeder, G. D., & Coovert, M. D. (1986). Revising an impression of morality. *Social Cognition, 4*, 1-17.
- Roberts, D. F., and Maccoby, N. (1985) *Effects of Mass Communication*. In Lindzey and Aronson (1985), pp. 539-98.
- Roese, N. J. (1997). Counterfactual thinking. *Psychological Bulletin, 121*, 133–148.
- Roseman, I. J. (1984). Cognitive determinants of emotion: A structural theory. *Review of Personality and Social Psychology, 5*, 11–36.
- Rozin, P., Haidt, J., & McCauley, C. R. (2000). Disgust. In M. Lewis & J. M. Haviland-Jones (Eds.), *Handbook of emotions* (2nd ed., pp. 637-653). New York: Guilford Press.
- Rozin, P., Lowery, L., Imada, S., & Haidt, J. (1999). The CAD triad hypothesis: A mapping between three moral emotions (contempt, anger, disgust) and three moral codes (community, autonomy, divinity). *Journal of Personality and Social Psychology, 76*, 574–586.
- Rozin, P., Millman, L., Nemeroff, C. (1986). Operation of the laws of sympathetic magic in disgust and other domains. *Journal of Personality and Social Psychology, 50*, 703–712.
- Rubio-Godoy, M., Aunger, R., & Curtis, V. (2007). Serotonin: A link between disgust and immunity? *Medical Hypotheses, 68*, 61–66.
- Rucker, D. D., & Petty, R. E. (2004). When resistance is futile: Consequences of failed counterarguing on attitude certainty. *Journal of Personality and Social Psychology, 86*, 219 –235.
- Russell, P. S., & Giner-Sorolla, R. (2011a). Moral anger, but not moral disgust, responds to intentionality. *Emotion, 11*, 233–240.
- Russell, P. S., & Giner-Sorolla, R. (2011b). Moral anger is more flexible than moral disgust. *Social Psychological and Personality Science, 2*, 360–364.
- Russell, P. S., & Giner-Sorolla, R. (2011c). Social justifications for moral emotions: When reasons for disgust are less elaborated than for anger. *Emotion, 11*, 637–646.
- Russell, P. S., & Giner-Sorolla, R. (2013). Bodily moral disgust: What it is, how it is different from anger, and why it is an unreasoned emotion. *Psychological Bulletin, 139*, 328–351.
- Schaller, M., & Duncan, L. A. (2007). The behavioral immune system: Its evolution and social psychological implications. In J. P. Forgas, M. G. Haselton, & W. von Hippel (Eds.), *Evolution and the social mind: Evolutionary psychology and social cognition* (pp. 293–307). New York: Psychology Press.
- Scherer, K. R. (1982). Emotion as a process: Function, origin, and regulation. *Social Science Information, 21*, 555–570.

- Schnall, S., Haidt, J., Clore, G., & Jordan, A. (2008). Disgust as embodied moral judgment. *Personality and Social Psychology Bulletin, 34*, 1096-1109.
- Schnall, S., Roper, J., & Fessler, D. M. T. (2010). Elevation leads to altruistic behavior. *Psychological Science, 21*, 315–320.
- Shaver, P., Schwartz, J., Kirson, D., & O'Connor, C. (1987). Emotion knowledge: Further exploration of a prototype approach. *Journal of Personality and Social Psychology, 52*, 1061–1086.
- Shaver, P., Schwartz, J., Kirson, D., & O'Connor, C. (1987). Emotion knowledge: Further exploration of a prototype approach. *Journal of Personality and Social Psychology, 52*, 1061–1086.
- Shweder, R. A. (1994). “You’re not sick, you’re just in love”: Emotion as an interpretive system. In P. Ekman & R. Davidson (Eds.), *The nature of emotions: Fundamental questions* (pp. 32-34). New York: Oxford University Press.
- Shweder, R. A. (2002). “What about female genital mutilation?” and why understanding culture matters in the first place. In R. A. Shweder, M. Minow, & H. R. Markus (Eds.), *Engaging cultural differences: The multicultural challenge in liberal democracies* (pp. 216–251). New York: Russell Sage Foundation.
- Silvers, J. A., & Haidt, J. (2008). Moral elevation can induce nursing. *Emotion, 8*, 291–295.
- Skitka, L. (2010). The psychology of moral conviction. *Social and Personality Psychology Compass, 4*, 267–281.
- Skitka, L. J., & Mullen, E. (2002). The dark side of moral conviction. *Analyses of Social Issues and Public Policy, 2*, 35–41.
- Skitka, L. J., Bauman, C. W., & Sargis, E. G. (2005). Moral conviction: Another contributor to attitude strength or something more? *Journal of Personality and Social Psychology, 88*, 895–917.
- Skowronski, J. J., & Carlston, D. E. (1989). Negativity and extremity biases in impression formation: a review of explanations. *Psychological Bulletin, 105*, 131–142.
- Sloane, S., Baillargeon, R., & Premack, D. (2012). Do infants have a sense of fairness? *Psychological Science, 23*, 196–204.
- Smith, C. A., & Ellsworth, P. C. (1985). Patterns of cognitive appraisal in emotion. *Journal of Personality and Social Psychology, 48*, 813–838.
- Smith, C. A., Haynes, K. N., Lazarus, R. S., & Pope, L. K. (1993). In search of the “hot” cognitions: Attributions, appraisals, and their relation to emotion. *Journal of Personality and Social Psychology, 65*, 916–929.
- Spranca, M., Minsk, E., & Baron, J. (1991). Omission and commission in judgment and choice. *Journal of Experimental Social Psychology, 27*, 76–105.

- Spranca, M., Minsk, E., & Baron, J. (2005). Omission and commission in judgment and choice. In M. Bazerman (Ed.), *Negotiation, decision making and conflict management* (pp. 463–492). Northampton, MA: Edward Elgar Publishing.
- Stevenson, R. J., & Repacholi, B. M. (2005). Does the source of an interpersonal odour affect disgust? A disease risk model and its alternatives. *European Journal of Social Psychology, 35*, 1–27.
- Tormala, Z. L., & Petty, R. E. (2002). What doesn't kill me makes me stronger: The effects of resisting persuasion on attitude certainty. *Journal of Personality and Social Psychology, 83*, 1298–1313.
- Tormala, Z. L., Petty, R. E., & Brinol, P. (2002). Ease of Retrieval Effects in Persuasion: A Self-Validation Analysis. *Personality and Social Psychology Bulletin, 28*, 1700-1712.
- Turiel, E. (1983). *The Development of Social Knowledge: Morality and Convention*. Cambridge: Cambridge University Press.
- Turiel, E. (2002). *The culture of morality: Social development, context, and conflict*. New York: Cambridge University Press.
- Turiel, E., Hildebrandt, C., & Wainryb, C. (1991). Judging social issues: Difficulties, inconsistencies, and consistencies. *Monographs of the Society for Research in Child Development, 56*, 1-103.
- Tversky, A., & Kahneman, D. (1986). Rational Choice and the Framing of Decisions. *The Journal of Business, 59*, 251-278.
- Tversky, A., & Kahneman, D. (1992). Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Journal of Risk and Uncertainty, 5*, 293-323.
- Ugazio, G., Lamm, C., & Singer, T. (2012). The role of emotions for moral judgments depend on the type of emotion and moral scenario. *Emotion, 12*, 579-590.
- Valdesolo, P., & DeSteno, D. A. (2006). Manipulations of emotional context shape moral judgment. *Psychological Science, 17*, 476–477.
- Visser, P. S., Krosnick, J. A., & Simmons, J. P. (2003). Distinguishing the cognitive and behavioral consequences of attitude importance and certainty: A new approach to testing the common-factor hypothesis. *Journal of Experimental Social Psychology, 39*, 118-141.
- Visser, P. S., & Mirabile, R. R. (2004). Attitudes in the social context: The impact of social network composition on individual-level attitude strength. *Journal of Personality and Social Psychology, 87*, 779-795.
- Visser, P. S., Bizer, G., & Krosnick, J. A. (2006). Exploring the latent structure of strength-related attitude attributes. In M. Zanna (Ed.), *Advances in Experimental Social Psychology* (Vol. 38, pp. 1-67). San Diego, CA: Academic Press.

- Weisberg, D. P., & Beck, S. R. (2010). Children thinking about their own and others' regret and relief. *Journal of Experimental Child Psychology, 106*, 184-191.
- Weisberg, D. P. & Beck, S. R. (2010). Children's thinking about their own and others' regret and relief. *Journal of Experimental Child Psychology, 106*, 184-191.
- Weisberg, D. P. & Beck, S. R. (2012). The development of children's regret and relief. *Cognition and Emotion, 26*, 820-835.
- Wheatley, T., & Haidt, J. (2005). Hypnotically induced disgust makes moral judgments more severe. *Psychological Science, 16*, 780-784.
- Wood, W. (1982). Retrieval of Attitude-Relevant Information From Memory: Effects on Susceptibility to Persuasion and on Intrinsic Motivation. *Journal of Personality and Social Psychology, 42*, 798-810.
- Wood, W., Rhodes, N., & Biek, M. (1995). Working knowledge and attitude strength: An information-processing analysis. In R. E. Petty, & J. A. Krosnick (Eds.), *Attitude strength: Antecedents and consequences* (pp. 283-313). New Jersey: Lawrence Erlbaum Associates.
- 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*, 314-327.
- Zeelenberg, M., van Dijk, W. W., Manstead, A. S. R., & van der Pligt, J. (1998). The experience of regret and disappointment. *Cognition and Emotion, 12*, 221-230.

APPENDICES

Appendix A: Materials from Experiments 1 and 2, Chapter 2

Stories Depicting Negative Outcomes for Target and Non-Target Characters

Omission/Commission Story

At school today, a teacher said that she was going to play a game. She had two boxes – a red box and a blue box. Inside one of the boxes was a sticker. Inside the other box there was nothing. One child in the class would get to pick a box, and if that box had the sticker inside it, that child would get a sticker. If that box had nothing, the child would not get anything.

Sarah was the one who got to choose the box. She picked the red box. Before she opened it, however, the teacher asked her if she wanted to keep that choice, or switch to the blue box. She said she wanted to stay with the red box. She opened it, and there was nothing inside. As a result, Sarah did not get a sticker.

That same teacher had another class of different children that day, and she played the same game with that class. Katie was the one who got to choose the box in that class. She picked the blue box. However, before she opened it, the teacher asked her if she wanted to keep that choice or switch to the red box. She decided to switch to the red box. She opened it, and there was nothing inside. As a result, Katie did not get a sticker.

Do you think one girl will feel worse than the other about having picked the red box – which meant she didn't get a sticker?

Sarah, who picked the red box (with nothing inside it) at first and didn't switch to the blue box (that had the sticker inside it).

Or

Katie, who picked the blue box (with the sticker inside it) at first, but then switched to the red box (which had nothing inside it)

Or

The two girls will feel the same

Why?

Who do you think is more to blame for not getting a sticker or are they equally to blame?

Sarah

Or

Katie

Or

The two girls are equally to blame

Why?

Omission/Commission Story

At a sweet shop, a shopkeeper had two types of new sweets. There were yellow sweets and green sweets that day. A boy named Ben went into the shop one day to choose one of the two types of new sweets for him to eat. Ben went into the shop and asked the shopkeeper for the yellow sweets.

Before the shopkeeper got him his sweets, the shopkeeper asked him if he wanted to keep that choice, or switch to the green sweets. Ben said he wanted to stay with the yellow sweets. The shopkeeper gave Ben the yellow sweets and Ben left. Ben tasted the yellow sweets and discovered that the yellow sweets did not taste nice at all.

On another day, a boy named Andy went into the same sweets shop to choose one of the two types of new sweets for him to eat. That day there were also yellow and green sweets. Andy went into the shop and asked the shopkeeper for the green sweets. Before the shopkeeper got him his sweets, the shopkeeper asked him if he wanted to keep that choice or switch to the yellow sweets. Andy decided to switch to the yellow sweets. The shopkeeper gave Andy the yellow sweets and Andy left. Andy tasted the yellow sweets and discovered that the yellow sweets did not taste nice at all.

Do you think one boy will feel worse than the other about having picked the sweets that did not taste nice – which meant they had bad tasting sweets?

Ben, who picked the yellow sweets (the nasty tasting sweets) at first and didn't switch

Or

Andy, who picked the green sweets (the nice tasting sweets) at first, but then switched to the yellow sweets (the nasty tasting sweets)

Or

The two girls will feel the same

Why?

Who do you think is more to blame for getting nasty tasting sweets or are they equally to blame?

Ben

Or

Andy

The two girls are equally to blame

Why?

Atypical/Typical Story

David and Charles are two boys who don't know each other and both like to go to the cinema a lot. Whenever David goes to the cinema, he ALWAYS sits in the back row. When Charlie goes to the cinema, he ALWAYS sits in the front row.

Today, David went to the cinema and sat in the back row the way he always does. David put his fizzy pop in a cup holder that he didn't know was nearly broken. Unfortunately, during the movie, David's cup holder broke and David's sticky fizzy pop spilled onto his seat. David got covered in very sticky, red fizzy pop which ruined his clothes and didn't feel nice.

Charles also went to the cinema today, but instead of sitting in the front where he usually sits, he decided for no special reason to sit in the back row. Charles put his fizzy pop in a cup holder he didn't know was nearly broken. Unfortunately, during the movie Charles' cup holder broke and Charles' sticky fizzy pop spilled onto his seat. Charles also got covered in very sticky, red fizzy pop which ruined his clothes and didn't feel nice.

Do you think one boy will feel worse than the other about having decided to sit in the back row that day where they got covered by spilt sticky fizzy pop?

David, who always sits in the back row and sat in the back row today where he got covered in spilt sticky fizzy pop

Or

Charles, who always sits in the front row, but just today decided to sit in the back row where he got covered in spilt sticky fizzy pop

Or

The two boys will feel the same.

Why?

Who do you think is more to blame for getting himself covered in spilt sticky fizzy pop?

David

Or

Charles

The two boys are equally to blame

Why?

Atypical/Typical Story

Anna and Laura are two girls who are both about 11 years old. They don't know each other, but they both like to go to football matches. The girls always ride their bikes to the football match.

When Anna goes to the football match, she has to ride around a large pond in a park, and there are two bike paths she can take. One path is green and the other path is yellow. The two paths are the same length and take the same amount of time on a bike.

Anna always rides on the yellow path and today she set off on the yellow path like she always does to get to the football match. Unfortunately, when she was part of the way around on the yellow path she saw there was a big fallen tree across the path, so she had to turn around and go back and then ride on the green path. Because of the big fallen tree, Anna didn't get to the football match.

Laura always rides on the green path but just that day Laura decided to set off on the yellow path to get to the football match. Unfortunately, just like with Anna, when Laura was part of the way around on the yellow path she saw there was a big fallen tree across the path, so she had to turn around and go back and then ride on the green path. Because of the big fallen tree, Laura didn't get to the football match.

Who do you think will feel worse about having missed the football match?

Anna – who always rides around the park on the yellow path like she always does where the big fallen tree way

Or

Laura – who always rides around the park on the green path, but just that day decided to ride on the yellow path where the big fallen tree way

Or

The two girls will feel the same

Why?

Who do you think is to blame for missing the football match?

Anna

Or

Laura

Or

The two girls are equally to blame

Why?

Stories Depicting Negative Outcomes for Target Characters, Non-Target Characters, and Bystanders

Omission/Commission Story

At school today, a teacher said that she was going to play a game. She had two boxes – a red box and a blue box. Inside one of the boxes was a sticker. Inside the other box there was nothing. One child in the class would get to pick a box, and if that box had the sticker inside it, everyone in the class would get a sticker. If that box had nothing, no one would get anything, even the child who made the choice.

Sarah was the one who got to choose the box. She picked the red box. Before she opened it, however, the teacher asked her if she wanted to keep that choice, or switch to the blue box. She said she want to stay with the red box. She opened it, and there was nothing inside. So, no one in his class got any stickers.

That same teacher had another class of different children that day, and she played the same game with that class. Katie was the one who got to choose the box in that class. She picked the blue box. However, before she opened it, the teacher asked her if she wanted to keep that choice or switch to the red box. She decided to switch to the red box. She opened it, and there was nothing inside. So, no one in his class got any stickers.

Do you think one girl will feel worse than the other about having picked the red box – which meant nobody got stickers?

Sarah, who picked the red box (with nothing inside it) at first and didn't switch to the blue box that had the sticker inside it)

Or

Katie, who picked the blue box (with the sticker inside it) at first, but then switched to the red box (which had nothing inside it)

Or

The two girls will feel the same?

Why?

Who do you think is more to blame for his class not getting any stickers or are they equally to blame?

Sarah

Or

Katie

The two girls are equally to blame

Why?

Omission/Commission Story

At a sweet shop, a shopkeeper had two types of new sweets. There were yellow sweets and green sweets that day. A boy named Ben had the job that day to go into the shop and choose one of the two types of new sweets for him and his friends to eat. Ben went into the shop and asked the shopkeeper for the yellow sweets.

Before the shopkeeper got him his sweets, the shopkeeper asked him if he wanted to keep that choice, or switch to the green sweets. Ben said he wanted to stay with the yellow sweets. The shopkeeper gave Ben the yellow sweets and Ben left. Ben and his friends discovered that the yellow sweets did not taste nice at all.

On another day, a boy named Andy had the job of going to the same sweets shop and to choose one of the two types of new sweets for him and his friends to eat. That day there were also yellow and green sweets. Andy went into the shop and asked the shopkeeper for the green sweets. Before the shopkeeper got him his sweets, the shopkeeper asked him if he wanted to keep that choice or switch to the yellow sweets. Andy decided to switch to the yellow sweets. The shopkeeper gave Andy the yellow sweets and Andy left. Andy and his friends discovered that the yellow sweets did not taste nice at all.

Do you think one boy will feel worse than the other about having picked the sweets that did not taste nice – which meant they and their friends had bad tasting sweets?

Ben, who picked the yellow sweets (the nasty tasting sweets) at first and didn't switch
Or

Andy, who picked the green sweets (the nice tasting sweets) at first, but then switched to the yellow sweets (the nasty tasting sweets)

Or

The two boys will feel the same

Why?

Who do you think is more to blame for getting bad tasting sweets for themselves and their friends or are they equally to blame?

Ben

Or

Andy

The two boys are equally to blame

Why?

Atypical/Typical Story

David and Charles are two boys who don't know each other and both like to go to the cinema a lot. Whenever David goes to the cinema, he ALWAYS sits in the back row. When Charles goes to the cinema, he ALWAYS sits in the front row.

Today, David went to the cinema and sat in the back row the way he always does. David put his fizzy pop in a cup holder that he didn't know was nearly broken. Unfortunately, during the movie, David's cup holder broke and David's sticky fizzy pop spilled onto him and everyone in the back row. David and everyone else got covered in very sticky, red fizzy pop which ruined their clothes and didn't feel nice.

Charles also went to the cinema today, but instead of sitting in the front where he usually sits, he decided for no special reason to sit in the back row. Charles put his fizzy pop in a cup holder he didn't know was nearly broken. Unfortunately, during the movie Charles' cup holder broke and Charles' sticky fizzy pop spilled onto him and everyone in the back row. Charles and everyone else also got covered in very sticky, red fizzy pop which ruined their clothes and didn't feel nice.

Do you think one boy will feel worse than the other about having decided to sit in the back row that day where they and everyone else got covered by spilt sticky fizzy pop?

David, who always sits in the back row and sat in the back row today where he got covered in spilt sticky fizzy pop

Or

Charles, who always sits in the front row, but just today decided to sit in the back row where he got covered in spilt sticky fizzy pop

Or

The two boys will feel the same.

Why?

Who do you think is more to blame for getting himself and everyone else covered in spilt sticky fizzy pop?

David

Or

Charles

The two boys are equally to blame

Why?

Atypical/Typical Story

Anna and Laura are two girls who are both about 11 years old. They don't know each other, but they both like to go to football matches. The girls always ride their bikes to the football match. They always meet their friends and bring their friends tickets to the match. When Anna goes to the football match, she has to ride around a large pond in a park, and there are two bike paths she can take. One path is green and the other path is yellow. The two paths take the same amount of time on a bike.

Anna always rides on the yellow path and today she set off on the yellow path like she always does to get to the football match. Unfortunately, when she was part of the way around on the yellow path she saw there was a big fallen tree on the path, so she had to turn around and go back and then ride on the green path. Because of the big fallen tree, Anna and her friends didn't get to see the football because Anna had their tickets.

Laura always rides on the green path but just that day Laura decided to set off on the yellow path to get to the football match. Unfortunately, just like with Anna, when Laura was part of the way around on the yellow path she saw there was a big fallen tree on the path, so she had to turn around and go back and then ride on the green path. Because of the big fallen tree, Laura and her friends didn't get to the football match because Laura had their tickets.

Who do you think will feel worse about having made herself and her friends miss the football match?

Anna – who always rides around the park on the yellow path like she always does where the big fallen tree was

Or

Laura – who always rides around the park on the green path, but just that day decided to ride on the yellow path where the big fallen tree was

Or

The two girls will feel the same

Why?

Who do you think is to blame for having made herself and her friends miss the football match?

Anna

Or

Laura


Or

The two girls are equally to blame


Why?

Appendix B: Response Scales Used in Experiment 2, Chapter 2

Regret Scale



Blame Scale



Appendix C: Stories Used for Training for Response Scales, Experiment 2, Chapter 2

Regret Scale Training Stories

Doesn't feel bad at all:

A little boy called James goes to the park with his mum. James plays in the sandbox all day with other boys and girls. James has a lot of fun that day playing in the sandbox.

How bad do you think James feels?

Why?

If chooses 1 on the scale (correct) say "Yes I think so too, he wouldn't feel bad at all, because he had lots of fun."

If chooses 2-5 on the scale (incorrect) say "Well, I think he wouldn't feel bad at all because he had lots of fun, so I'd pick this one (1)."

Feels a bit bad:

On a different day, a little girl called Sally goes to the park with her mum. Sally plays in the sandbox all day with other boys and girls. Sally wants to make a sandcastle. All of the toys to build the sandcastles are being used by the other boys and girls, so she can't build her sandcastle like she wanted. Sally decides to go play on the swings instead because she likes the swings, too.

How bad do you think Sally feels?

Why?

If chooses 1 on the scale (incorrect) say "Well, I think she would feel a little bit bad because she didn't get to make a sandcastle, so I'd pick this one (2) or this one (3)."

If chooses 2 or 3 on the scale (correct) say "Yes I think so too, she would feel a little bit bad because she didn't get to make a sandcastle but still had fun."

If chooses 4 or 5 on the scale (incorrect) say "Well, I don't think she would feel that bad because she did go on the swings, and she likes the swings, too, so I'd pick this one (2) or this one (3)."

Feels very, very bad:

On a different day, a little boy called Tommy goes to the park with his mum. Tommy plays in the sandbox all day with other boys and girls. Tommy builds a really big sandcastle. Tommy's sandcastle is really good and he likes his sandcastle very much. Just when Tommy finishes his sandcastle, another boy kicks Tommy's sandcastle ruins it. Tommy starts to cry.

How bad do you think Tommy feels?

Why?

If chooses 1-3 on the scale (incorrect) say “Well, I think he would feel really, really bad because he lost his sandcastle, so I’d pick this one (4) or this one (5).”

If chooses 4 or 5 on the scale (correct) say “Yes I think so too, he would feel really, really bad because he lost his sandcastle.”

Blame Scale Training Stories

Not at all to blame:

A little boy called John goes to the park with his mum. John plays in the sandbox all day with other boys and girls. John sees a boy in the park get pushed by someone.

How much do you blame John for the boy in the park getting pushed?

Why?

If chooses 1 on the scale (correct) say “Yes I think so too, he is not to blame at all because he didn’t do anything at all.”

If chooses 2-5 on the scale (incorrect) say “Well, I think he is not to blame at all because he didn’t do anything at all, so I’d pick this one (1).”

Somewhat to blame:

On a different day, a little girl called Gemma goes to the park with her mum. Gemma plays in the sandbox all day with other boys and girls. Gemma wants to make a sandcastle. She walks over to the place where the toys for building sandcastles are, and accidentally steps on a little boy’s sandcastle. Gemma feels very sorry for accidentally stepping on the sandcastle and ruining it.

How much do you blame Gemma for ruining the other boy’s sandcastle?

Why?

If chooses 1 on the scale (incorrect) say “Well, I think she is a little bit to blame because she did step on the little boy’s sandcastle, so I’d pick this one (2) or this one (3).”

If chooses 2 or 3 on the scale (correct) say “Yes I think so too, she is a little bit to blame because she did step on the little boy’s sandcastle.”

If chooses 4 or 5 on the scale (incorrect) say “Well, I don’t think she is that much to blame because she stepped on his sandcastle on accident, so I’d pick this one (2) or this one (3).”

Lots and lots to blame:

On a different day, a little boy called Dan goes to the park with his mum. Dan plays in the sandbox all day with other boys and girls. Dan builds a really big sandcastle. But, Dan sees another boy who has an even bigger and better sandcastle. Dan is cross, so he goes over to the other boy's sandcastle and stomps on it. Dan ruins the other boy's sandcastle and doesn't feel sorry at all.

How much do you blame Dan for ruining the other boy's sandcastle?

Why?

If chooses 1-3 on the scale (incorrect) say "Well, I think he is to blame loads and loads because he stomped on the other boy's sandcastle, so I'd pick this one (4) or this one (5)."

If chooses 4 or 5 on the scale (correct) say "Yes I think so too, he is to blame loads and loads because he stomped on the other boy's sandcastle."

Appendix D: Stories and Lines of Additional Information, Experiment 5, Chapter 3

Date rape story

Becky is setup on a blind date with a man named Luke. They go to a house party together for their date. The next day, Becky wakes up and goes to the police because Luke raped her.

Wrongness-confirming information:

- Luke shows very little regard for other people
- Luke has very few friends and does not stay in contact with his family.
- There was a party Luke really wanted to go to that night so he just decided to take Becky along.
- Luke is nasty to people and talks badly about people behind their backs.
- Becky was not interested in going to a party, and expressed her disinterest, but decided to try to be polite.
- When Luke and Becky arrived at the party, Luke immediately started taking shots with his friends and pressured Becky to drink more.
- Luke engages in a lot of risky behaviour like drug and excessive alcohol use.
- Luke has had a lot of sexual partners in the past.
- Luke called Becky lame and a prude in front of his friends because she wouldn't take more shots.
- Luke does not show respect to women, even those women he has sex with.
- After drinking a lot and ignoring Becky for most of the evening, Luke told her he wasn't going to take her home.
- At the end of the night, Becky decided to sleep alone in one of the bedrooms at the party.
- Luke often pretends to be a nice guy and show interest just to get sex.
- Luke drinks to the point of becoming aggressive with people around him.
- When Becky woke up the next morning, Luke was lying next to her naked.
- Luke thinks people are stupid and useless and deserve to be treated badly.
- Luke told Becky that he had sex with her after she had passed out because she was obviously too prudish to sleep with him when she was awake.
- Luke was fully aware of what had happened and showed no remorse for his actions even though Becky was upset.
- Luke is not a kind or generous person.
- Luke has no job or career plans and often cons his way into getting what he wants.
- Luke was setup on a blind date with Becky by one of his friends who was hoping Luke would meet a nice girl unlike the girls he usually spends time with
- People that know Luke think that he is mean spirited and not to be trusted.
- Luke was not particularly interested in Becky from the beginning, but wanted to shag that night.
- Luke knew that Becky did not want to sleep with him that night from the very beginning.

- Luke has lost many of his jobs in a short time period because his employers have found him stealing and engaging in poor conduct.
- Luke has stolen from his family and lied to them repeatedly.
- Luke planned to get Becky very drunk so that he could take advantage of her.
- Luke smokes, doesn't look after his hygiene and gambles excessively.
- Luke made Becky's drinks for her so that he could put extra shots of liquor in them to get her really wrecked.
- Luke conned Becky into trusting him by putting on a 'nice guy' persona that wasn't real.

Wrongness-contradicting information:

- Luke has been on a few dates when he can get a break in his schedule, but hasn't found the right person yet.
- Luke does not engage in risky behaviour and is happy to just go out with friends when he has time.
- Luke was very excited about his date with Becky.
- Luke does not have a temper, and is known for being very patient and friendly.
- Luke was very nervous because it was his first time being setup on a blind date.
- Luke took Becky to a friend's house party to take some of the pressure off of having a brilliant time one-on-one.
- Luke would like to settle down one day and get married and raise a family.
- Luke has never driven his car under the influence of alcohol.
- Becky seemed very keen to go on a date with Luke and was quite nervous.
- Luke has only had a few sexual partners.
- Becky and Luke had a drink in order to calm their nerves.
- Becky made strong hints that she was interested in Luke and was very flirtatious.
- Luke always uses protection when having sex.
- Luke is a firm believer in showing the utmost respect to women.
- Becky and Luke talked for a while and really enjoyed each other's company and had a lot of chemistry.
- Luke has a passion for helping others in his profession as well as in general, helping friends and family.
- At one point in the night, Becky and Luke started kissing with mutual interest quite intensely.
- The only thing Becky really recalled vividly was waking up in the bed the next morning next to Luke, both of them without their clothes on.
- Luke is a very shy and kind man who works very hard at his job as a medical intern.
- Luke is close to his family and friends and has been looking for someone special to spend his life with.
- Luke was surprised Becky didn't remember having sex because he thought she was into it.
- Friends of Luke describe him as smart, caring and responsible.
- Becky was very upset and Luke felt very bad and confused as to how he could've possibly done this to Becky.
- Luke believed that Becky was very interested in him in a romantic way.

- Luke's family are very proud of Luke.
- Luke takes very good care of his health and hygiene, eats right and exercises
- Luke wanted Becky to have a good date and a nice evening in his company.
- People at the party commented on how they thought it looked like Becky and Luke were really into each other.
- Becky was interested in Luke and was attracted to him.
- Luke is very easy to talk to and is an excellent listener.

Dead dog story

Alex's family dog is dead and buried in the garden. One day, Alex digs up the dead family dog, cooks it on the barbecue and eats it for dinner.

Wrongness-confirming information:

- Alex is not a kind or generous person.
- Alex is not responsible and does not take good care of his wife and young children.
- Alex didn't like the dog because he felt the dog got more attention from wife and kids than he did.
- Alex is disrespectful and rude, and his sense of humour involves making inappropriate comments and making fun of others.
- Alex had a good job and was financially secure at the time, and was not desperate for food.
- Alex was sick and tired of his wife and kids moaning about missing the dead dog.
- Alex does not invest a lot of his time or effort into being a good father, and does not provide for his children emotionally or financially.
- Alex has made many irresponsible decisions in the past including losing a lot of money at gambling.
- Alex wanted to teach his family a lesson for being so attached to a silly animal.
- Alex has been known to engage in drug use and excessive drinking, and spends his free time in pubs or casinos instead of with his kids.
- Alex believed that by digging up the dead dog and cooking it for dinner, he would really upset his family and teach them a lesson.
- Alex dug up the dead dog randomly and happily one day, without any provocation.
- Alex has not been able to maintain a steady job for years because of bad conduct at work.
- Alex is a slob who has poor hygiene, smokes, doesn't eat healthy and never looks after his fitness.
- Alex made it very public that he hated the dog and was glad to cook it up for dinner.
- Alex has never shown an interest in doing anything other than drinking, gambling and sitting in front of the telly for hours.
- Alex showed his wife and young kids the cooked dog on the barbeque and was really amused by their horrified reaction.
- Alex showed no remorse for cooking up the dead dog and eating it.
- Alex's wife has been unhappy for years because Alex treats her and the kids badly, but she tries to look after their family.

- Alex never picks his kids up from school, takes them to swimming or reads to them, and he never helps out with house cleaning.
- Alex thought that eating the dead family dog was perfectly acceptable, and a good form of revenge for his family being so attached to the dog.
- Alex dropped out of University and does not enjoy reading, but he likes to act as though he is the smartest person in the room.
- Alex told his neighbours and colleagues at work about what he did just to see their horrified reaction.
- When Alex's wife and kids begged him to not eat the family dog, he got even angrier and ate more of the dog.
- Alex litters his rubbish and cigarettes, and shows no care for the environment around him.
- Alex does not show respect or kindness to animals and teaches his kids to think of animals as just food products.
- Alex cut up the dead dog before placing the bits on the barbeque in the kitchen in front of his wife and young kids.
- Alex doesn't make an effort to keep in contact with friends or his family members.
- The family had the dog for 13 years and everyone, except for Alex, loved him very much.
- Alex tried to force his wife and young kids to eat the dog for dinner against their will and in spite of them being horrified and upset.

Wrongness-contradicting information:

- Alex is a very kind and generous person.
- Alex is very responsible and takes very good care of his wife and young children.
- Alex had recently lost his job and was going through a divorce, leaving him in financial ruins.
- Alex is very respectful and polite, but also has a very good sense of humour.
- Alex had barely enough money to buy food.
- Alex's family dog was extremely badly behaved and showed a lot of aggression.
- Alex always tries to provide a good life for his young kids, and is very devoted father and husband.
- Alex has never made irresponsible decisions in the past, financially or otherwise.
- Alex was forced to put the dog down by Animal Control only the day before he dug it back up for dinner.
- Alex does not engage in drug or excessive alcohol use, and his recreational time is spent playing and reading with his kids.
- The dog had randomly attacked his new-born baby girl and badly injured her.
- The dog was one of the biggest reasons for why his wife left him.
- Alex works very hard at his job so that he can afford to pay for the best education for his children.
- Alex has good hygiene and looks after his health and fitness routinely.
- Alex loved the dog, but knew he had to put it down.
- Alex likes to volunteer when he has time to run marathons for various charity organisations.

- The day Alex decided to dig up the dog; Alex had no food in his cupboards and no money to his name.
- When Alex asked a close friend the day he ate the dog what he should do for food, his close friend suggested he eat the dog.
- Alex's wife has been happily married to Alex for 10 years, and thinks that he is a very romantic, loyal partner and friend.
- Alex likes to do a lot of 'dad' activities like picking his kids up from school, taking them to swimming lessons and reading to them.
- Alex had not eaten in a couple of days before he finally dug up his dog to eat so ate the dog in desperation and hunger.
- Alex has a good university education and is very intelligent, but does not gloat or behave arrogantly.
- Alex never told anyone that he ate the dog because he didn't want to upset people.
- Alex did a lot of research about the safety and health risks of eating dog before he dug it up and cooked it.
- Alex recycles and believes in treating the environment with respect.
- Alex encourages his family to eat organic, free-range products.
- Alex was very depressed and confused about his life circumstances on the day he decided to eat the dog.
- Alex's friends and parents feel that they can always count on Alex for support and good advice.
- The dog had become a liability because it had previously attacked friends and strangers in the neighbourhood.
- Alex was familiar with some cultures where eating dog is acceptable.

Incest story

Julie and Mark are brother and sister. They are travelling together in France on a summer vacation. They are staying in a cabin near the beach. One night, Julie and Mark have sex.

Wrongness-confirming information:

- Julie and Mark have a sense of humour that involves making fun of other people and being rude at inappropriate times.
- When Julie or Mark has a friend or family member in need of help, they don't go out of their way to show support.
- Julie and Mark wanted to prove that they could do whatever they wanted to do.
- Julie and Mark both get poor marks at university and are known for being disrespectful and disruptive in lectures.
- Julie and Mark were both in serious relationships with other people.
- Julie and Mark were travelling without their partners and planned to have sex before going on holiday.
- Julie and Mark don't take their education seriously and may fail their degrees.
- Both Julie and Mark have engaged in drug use and risky sexual behaviour in the past.
- Julie and Mark decided they would lie to their partners about anything that happened on their trip.

- Julie and Mark's time is usually spent at the pub where they drink excessively and get into confrontations with friends or strangers.
- Julie and Mark had fooled around before this happened because they like doing whatever they want.
- Julie did not end up getting pregnant, but they didn't bother using protection because it was too much of a hassle.
- Julie and Mark don't follow social rules and try to upset social norms as much as possible.
- Julie and Mark don't believe in taking care of the things or people around them, and do whatever they please.
- Julie and Mark decided to tell random people about having sex with each other, except for their partners because they wanted to upset others.
- Julie and Mark take a lot of chances with their safety and well-being because they don't think life is very important.
- Julie and Mark enjoyed having sex and bragged about it.
- Julie and Mark had no regrets about what they did and would happily do it again.
- Julie and Mark are not kind or generous people.
- Julie and Mark show little respect or consideration for others.
- Julie and Mark had sex because they were both in the mood to do something reckless and stupid.
- Julie and Mark have lost many friends and jobs because of their bad attitude and inappropriate behaviour.
- Julie and Mark had been partying all night before having sex.
- Julie and Mark believed having sex with each other was wrong, but didn't care.
- Julie and Mark don't care about morals or social norms, and go out of their way to push boundaries.
- Julie and Mark smoke cigarettes and don't look after their hygiene or the environment around them.
- Julie and Mark wanted to prove to themselves that they didn't have to follow rules or moral codes.
- Julie and Mark have been known to be cruel to animals just for fun.
- Julie and Mark got immense pleasure out of having sex with each other because it felt wrong.
- Julie and Mark had sex for a long time and performed unusual sexual acts on each other.

Wrongness-contradicting information:

- Whenever Julie and Mark's friends are in need of help, they're always there to help.
- Julie and Mark are both very good students and take their studies very seriously at university.
- Most of Julie and Mark's friends at University had already had sex and talked openly about it, making Julie and Mark feel awkward.
- Julie and Mark both do very well at University but do not like to brag or gloat.
- Julie and Mark had low self-esteem because they had no experience with sex and felt out-of-place amongst their peers.

- Julie and Mark knew it would be weird but decided that if they were careful and only did it once it would be harmless.
- Julie and Mark both work hard but also have a rich social life and are well liked.
- Julie and Mark do volunteer work with the elderly and children in hospital.
- Julie and Mark only made the decision after talking and thinking it through a lot.
- Julie and Mark do not engage in risky behaviour and typically make sensible decisions.
- Julie and Mark did not treat the decision to have sex lightly.
- Julie was on the birth control pill.
- Neither Julie nor Mark has ever engaged in drug use or risky sexual behaviour.
- Julie and Mark drink socially in moderation and look after their health and fitness.
- Julie and Mark used a condom to be safe.
- Julie and Mark are dedicated students who value their friends and family.
- Julie and Mark decided not to tell anyone to avoid upsetting people because they were not proud of their actions.
- Julie and Mark decided never to have sex again and realised that it was not okay to do that kind of thing.
- Julie and Mark are both very kind and generous people.
- Julie and Mark both have a really good sense of humour and really entertaining when out with friends.
- Julie and Mark were both virgins before this night they decided to have sex together.
- Julie and Mark have a lot of respect for rules and norms.
- Julie and Mark felt that they were never going to have a positive sexual experience because of their inexperience with sex.
- Julie and Mark felt very badly about their actions, but also felt slightly better about having sex with a new partner in the future.
- Julie and Mark look after their family members and often hold family socials.
- Julie and Mark care a great deal for animals and the environment.
- Julie and Mark believed that having sex with protection only once would not harm anyone.
- Julie and Mark have not have many intimate relationships, but the ones they have had ended amicably and respectfully.
- Julie and Mark decided after they had sex that they shouldn't spend time together very often.
- Julie and Mark only had intercourse for less than two minutes and didn't enjoy it.

Tattoo story

Mike and his girlfriend Sarah go on a trip. Mike thinks Sarah would look really good with a tattoo. One night when Sarah's unconscious, Mike gets a permanent tattoo put on her body.

Wrongness-confirming information:

- Mike is not in love with his long-time girlfriend, Sarah, but doesn't have the courage to tell her.

- Mark has not kept a job consistently in years and is quite happy to live off of Sarah while he stays home and plays video games.
- Mike and Sarah were on a trip to try to save their relationship because it was going badly.
- Mike is not adventurous and doesn't like going new places or taking Sarah out.
- Sarah never wanted a tattoo, but Mike told her he wished she would get one.
- Sarah has made it really clear throughout their relationship and on the trip that she never wants a tattoo.
- Mike never cooks or cleans around the house he shares with Sarah, and doesn't have ambition to do anything useful.
- Mike is not respectful to women, and thinks of women as unintelligent sex objects.
- Mike told Sarah on the trip that she wasn't sexy because she didn't want or have a tattoo
- Mike doesn't have many friends because he's very antisocial, and is rude to Sarah's friends when they are around.
- Mike served Sarah her drink the night of the tattoo, and put extra shots in them without her knowing.
- Mike and Sarah were not getting along on the trip and Mike lashed out at Sarah because she wouldn't get the tattoo.
- Mike has a promiscuous sexual past and regularly cheats on his girlfriend, Sarah.
- Sarah cares for Mike a lot, and has tried to help him sort out his life, but doesn't like being treated badly.
- Sarah thought tattoos were unattractive and she didn't want to have something she didn't like permanently put on her body.
- Mike is not interested in being faithful to Sarah, and has lied to her about wanting to raise a family to get her to sleep with him.
- Sarah didn't know how drunk she was, and passed out in the hotel room.
- Mike knew Sarah didn't want a tattoo ever, but didn't care.
- Mike often drinks instead of working to the point of becoming aggressive and mean.
- Mike does drugs and smokes, but keeps it hidden from Sarah.
- When Sarah woke up with a tattoo on her body, she was furious and very upset.
- Mike does not exercise or eat healthy, and shows no regard for his hygiene.
- The tattoo was badly done by someone who Mike paid to come to the hotel room while Sarah was passed out.
- Sarah didn't remember anything about the tattoo and didn't know she drank so much.
- Mike is in a lot of debt and does not behave responsibly when it comes to finances, and likes to gamble.
- Mike did not go to University and shows no interest in bettering his education or knowledge, but likes to brag that he's the cleverest person in the room.
- Mike told the tattoo guy to have the tattoo say 'drunk slag' because he thought it was funny.
- Mike shows no regard for helping other people or his surroundings, and often is openly rude to people and litters.
- The man who gave Sarah the tattoo was not a professional tattoo artist, and just wanted to make some easy money.

- Mike was pleased that Sarah passed out from drinking so he could carry through his plan.

Wrongness-contradicting information:

- Mike is really in love with his long-time girlfriend, Sarah.
- Mike works full-time to provide for him and Sarah.
- Mike and Sarah were on a trip celebrating their 10 year anniversary and were at a great place in their relationship.
- Mike enjoys travelling to new, exotic places, and always takes Sarah with him.
- Sarah always wanted a tattoo, but was too afraid to do it every time she tried because of the pain.
- Mike and Sarah discussed her getting a tattoo on their trip at length beforehand.
- Mike is the kind of boyfriend who spontaneously brings home flowers for Sarah or cooks her a gourmet meal.
- Mike is very respectful to women and is a strong believer in feminism.
- Mike told Sarah that she was beautiful with or without a tattoo, but a tattoo would look great on her.
- Mike has a lot of friends who think very highly of him.
- Sarah got very drunk the night of the tattoo before the tattoo happened because she wanted to let loose and have fun.
- Mike and Sarah were having a fantastic time together on their trip and getting along really well.
- Even Mike's few ex-girlfriends have very good things to say about Mike always being a gentleman and kind-hearted.
- Sarah thinks Mike is a wonderful partner and thinks they have a future together.
- Sarah said she was going to finally get the tattoo she always wanted and asked Mike to make sure she followed through.
- Mike is very close to his family and hopes to one day have a family of his own and raise children.
- Sarah passed out while she was at the tattoo parlour because she drank so much.
- Mike believed that Sarah wanted the tattoo very badly and would be happy to have it done even though she'd passed out.
- Mike does not engage in risky behaviours and is a good citizen in his community.
- Mike does not engage in excessive drug or alcohol use, and has rarely shown aggression to anyone.
- Sarah was pleased when she woke up that she had the tattoo and was thankful Mike made sure it happened.
- Mike is a runner and looks after his fitness and diet rigorously.
- The tattoo Sarah got was of a heart she had designed previously for the day she finally got one.
- Sarah was so happy she didn't feel any pain when she got the tattoo since she'd accidentally passed out.
- Mike pays his bills on time and has never had any financial or legal troubles.
- Mike is very well-educated and intelligent, but does not brag or behave arrogantly.

- Sarah told Mike that she wanted the tattoo so badly but she needed to get drunk beforehand to go through with it.
- Mike volunteers at an animal shelter and battered women's shelter in his spare time.
- The man who gave Sarah the tattoo was not concerned about Sarah being unconscious while having the tattoo because he'd spoken to Sarah earlier.
- Mike wasn't thrilled about Sarah being unconscious, but knew she wanted the tattoo badly and that being unconscious would prevent any pain.

Appendix E: Stories and Lines of Additional Information, Experiment 6, Chapter 3

Dead dog story

Alex's family dog is dead and buried in the garden. One day, Alex digs up the dead family dog, cooks it on the barbecue and eats it for dinner.

Wrongness-confirming information:

- Alex ate the dog meat nearly raw.
- The dog meat was nearly rotten.
- The dog meat was covered in maggots.
- Alex has fantasies about eating dead dog.
- Alex can go weeks without showering.
- Alex eats the bloody dog meat with his fingers.
- Alex licks the bloody meat off the plate when he finishes.
- Alex eats every drop of meat from the dog.
- Alex leaves some of the dog fur on the meat to eat.
- Alex tears into the dog meat with his bare hands.
- The meat smelled rancid and was covered in flies.
- Alex gets bits of dog meat stuck between his teeth and doesn't mind.
- Alex is so excited to eat dead dog because he fancies freshly dead animals that have been buried.
- Alex thinks he might try other meat like cat meat after he finishes eating the dog.
- Alex goes weeks without brushing his teeth.
- Alex likes the taste of dog meat.
- Alex finds hair in the meat, but continues eating.
- Alex thinks eating dead dog is erotic and quite sexy.
- Alex lets the blood from the dog meat get all over his face.
- Alex sucks some of the dog meat off of the bones.
- Alex eats all parts of the dog, including eyes and tongue.
- Alex gets bits of dog meat in his fingernails.
- Alex often found himself salivating over the dog when it was alive.
- Alex lives in a filthy house that is overrun by animals and their waste.
- Alex lets bits of the dog meat fall on the floor and doesn't clean up.
- Alex really likes the look of the dead dog whilst it's cooking.
- Alex didn't use any utensils when preparing the dog, just his hands.
- Alex thinks that the taste of dead dog is better than any other meat he's had.
- Alex finds himself getting turned on by the sight of the dead dog.
- Alex decided to try to find other dead animals to eat.

Wrongness-contradicting information

- Alex made sure the meat was well cooked.
- The meat was in good condition.

- The meat had been carefully wrapped in plastic and preserved.
- Alex had never considered eating dog before.
- Alex has very good personal hygiene.
- Alex eats the meat like other food, with a knife and fork and the table set.
- Alex finishes eating and does the washing up.
- Alex takes only a few bites and then stops.
- Alex was very careful about making sure the meat was prepared well.
- Alex prepares the meat and then sits down to have dinner alone once it's cooked.
- The meat was suitable to eat and looked like a piece of steak you might find at a restaurant.
- Alex eats the meat like a civilised person.
- Alex is interested in trying something different because he knows other cultures find dog acceptable to eat.
- Alex decides not to eat dog meat again, and will stick to standard chicken and beef.
- Alex takes very good care of his teeth very often.
- Alex decides he does not like the taste of dog meat.
- Alex is meticulous about preparing the meat and makes sure it's flawless.
- Alex thinks eating dead dog is strange, but knows that other cultures do it.
- Alex takes small bites of the well cooked meat and uses a napkin.
- Alex only tries a couple of bites of the part of dog that is acceptable to eat.
- Alex researches the parts of dog that are acceptable for consumption, and only eats those parts.
- Alex is very meticulous about not handling the dog meat with his hands.
- Alex was a very good pet owner and never considered eating the dog while he was alive.
- Alex keeps a very clean and tidy house, hoovering a few times a week.
- Once Alex finishes eating, he cleans the space where he cooked the meat thoroughly.
- Alex doesn't find anything exciting about the process of cooking the meat.
- Alex is a culinary expert, and so uses his finest tools for meat preparation.
- Alex decides that chicken and steak are much nicer meats.
- Alex wasn't very attached to the dog because the dog was a terrible and vicious pet.
- Alex decided that he would never tell anyone nor do it again.

Tattoo story

Mike and his girlfriend Sarah go on a trip. Mike thinks Sarah would look really good with a tattoo. One night when Sarah's unconscious, Mike gets a permanent tattoo put on her body.

Wrongness-confirming information:

- Mike got a tattoo that he knew Sarah would hate.
- Mike doesn't have friends because he's very antisocial, and rude.
- Mike had the tattoo put on Sarah's arm so that she would get into trouble at work.
- Mike knew Sarah never, ever wanted a tattoo.
- Mike knew that Sarah's family has forbidden their children to get tattoos.
- Mike hired a guy who does prison tattoos to come do Sarah's tattoo.

- Mike spiked Sarah's drink with the date-rape drug to get her to pass out before the tattoo.
- The tattoo was done so quickly and badly that it got badly infected.
- Mike refused to take care of Sarah's tattoo.
- Mike got Sarah a tattoo of a girl with her breasts exposed with the words, "Property of Mike".
- When Sarah woke up to find the horrible tattoo, Mike just laughed and called her a prude.
- Sarah never in her life wanted to get a tattoo and had constantly told Mike this when he asked.
- Mike and Sarah have been together for a week.
- Mike thinks women are fun toys.
- Mike spends his free time drinking and sleeping with random women.
- Mike doesn't actually like Sarah, he just wants sex.
- Secretly, Mike had been planning for weeks to have her drugged and tattooed.
- Mike told the guy who did the tattoo to make the tattoo sloppy and leave a scar.
- Mike took photos of Sarah with the tattoo to put on the internet.
- Sarah was horrified when she woke up that she had been drugged and tattooed.
- Mike told Sarah many times that she looked ugly and should get a tattoo.
- Mike made sure the tattoo was enormous.
- Mike cheats on Sarah a lot.
- Mike lies to Sarah about almost everything.
- Mike decided he was going to dump Sarah the next day.
- While Sarah is getting tattooed, Mike gets bored and leaves to have sex with someone else.
- Mike stole money from Sarah's purse to pay for the tattoo.
- Mike thinks it's so funny that Sarah is getting a tattoo she hates.
- Mike doesn't keep track of how much he used to drug Sarah.
- Mike has an extensive criminal record.

Wrongness-contradicting information:

- Mike got the tattoo that Sarah always wanted.
- Mike has a lot of friends who think very highly of him.
- Mike got the tattoo put on a place that was very discrete so she could keep it private.
- Mike knew Sarah has always, always wanted a tattoo.
- Mike made sure Sarah's family were okay with her getting a tattoo first.
- Mike hired a well-known professional tattoo artist to come do Sarah's tattoo.
- Sarah took a safe amount of drowsy medication prescribed by her doctor.
- The tattoo came out really well and healed very nicely.
- Mike took very good care of the tattoo while it healed.
- Sarah got the tattoo she designed herself a few weeks before.
- When Sarah woke up, Mike offered to pay for the removal if she decided she didn't want to keep the tattoo.
- Sarah really wanted a tattoo but was so afraid of the possible pain.
- Mike and Sarah have been together for 10 years.

- Mike strongly believes in feminism.
- Mike volunteers at a battered women's shelter in his spare time.
- Mike is very much in love with Sarah and wants to spend his life with her.
- Mike and Sarah discussed her needing to be unconscious to go through with the tattoo.
- Mike told the tattoo artist to be very careful and delicate with the tattoo.
- Mike promised to keep the tattoo a secret as long as Sarah wanted.
- Sarah was so happy and relieved that she hadn't felt any pain while having the tattoo done.
- Mike told Sarah many times that he loved her the way she was and didn't have to get a tattoo for him.
- Mike made sure the tattoo was small.
- Mike has always been faithful to Sarah.
- Mike is completely honest and upfront with Sarah.
- Mike was planning on proposing to Sarah in the next few days.
- While Sarah is getting tattooed, Mike stays by her side the entire time to make sure it goes well.
- Mike paid for the tattoo himself as a gift to Sarah.
- Mike thinks it's unfortunate that Sarah is so desperate for a tattoo, but wants her to be happy.
- Mike makes sure that Sarah only has the smallest amount of drowsy medication to work.
- Mike has never done anything illegal in his life.

Appendix F: 10-item Memory Questionnaires, Experiment 7, Chapter 4

Questionnaire used for Oprah clips:

What instrument did Fernando play?
What colour of shirt was Oprah wearing?
What is one of Fernando's earliest memories?
How and where did Fernando first meet Marcelus?
What was Marcelus's job?
Were there more men or women in the audience?
What TV shows did Fernando perform on?
What colour was Marcelus's tie?
At what school was Fernando appointed band director?
Where did Fernando move to pursue a career as a musician?
You have now completed the first half of the study. Please inform the experimenter.

Questionnaire used for Control clip:

What dessert was Rachel trying to make?
What colour of shirt was Ross wearing?
What memory prompted Rachel to use beef in her dessert?
Why did Rachel leave the apartment?
What was the man drinking from the can?
Was there recorded laughter during this clip?
What colour was Chandler's shirt?
What two pages of the cook book were stuck together?
How hot (i.e. temperature) did Rachel claim Joey's apartment was?
How many people were in the room?
You have now completed the first half of the study. Please inform the experimenter.