

Investigating Moral Events:  
Characterization and Structure of Autobiographical Moral Memories

Thesis by

Jessica R. Escobedo

In Partial Fulfillment of the Requirements for the

degree of

Doctor of Philosophy

CALIFORNIA INSTITUTE OF TECHNOLOGY

Pasadena, California

2009

(Defended October 20, 2008)

© 2009

Jessica R. Escobedo

All Rights Reserved

## ACKNOWLEDGEMENTS

Above all, my thanks go to Ralph Adolphs, my mentor, in every sense of the word.

Neither the research contained in this volume nor its development into a dissertation could have happened without his guidance, support and patience. I am grateful beyond words.

My sister, Leah, is my hero. Only she and I know how many phone calls made this work possible.

My wonderful husband, Jesse, has carried all the burdens in my life, making it possible for me to devote myself completely to writing. He is my one true love; an amazing friend and partner and my life is filled with joy because of him.

My parents taught me that I could be anything I wanted in life. Their support, through 25 years of school, has taught me what love means.

My brother and the rest of my family provided unending support. Leigh and John have given me the gift of a second set of parents. I am so lucky.

Jessica Levine and Catherine Holcomb, who are now masters of Word formatting and Endnote. Thank you for your tremendous energy and help when I needed it most.

My friends have sustained me through the hard times and laughed with me during the good times. Sarina Mohanty, Becca Delahanty, Arthur Fitzmaurice, Aurelia Flaming, Meghana Bhatt and Eun Jung Choi, you are so special to me.

Aditya Rajagopal was there for help with math when I needed it most.

The Adolphs Lab members have taught me a tremendous amount about the intersection between science and life.

Dr. Drange and Dr. Tourani deserve tremendous thanks for helping me conquer the illness that would have kept me from finishing.

Out of Context , TACIT and Ricketts House brought richness to my Caltech experience.

All of my students: Nalini Colaco, Audrey Carstensen, Andrew Molina, Viviana Gradinaru, Jessica Stockburger, Lionel Jingles, Agnieszka Lesniak, Michelle Jiang. You have taught me that real learning comes through teaching someone else.

For technical assistance, thanks go to the programmers who helped make this possible and the transcribers who listened to all the memories: it couldn't have been done without you. Also, to Kodi Scheer and Amanda Hornaday and Remya Nair, for all of your help.

All the participants and all the raters. There would be no data without your contributions.

## ABSTRACT

Moral events and the actions, decisions and people they involve, are judged as right or wrong, and the moral responsibility associated with them generates further judgments, often legal in nature, of blame and punishment or praise. Not only do moral events and the normative judgments they presuppose define essential aspects of human nature, they are also ubiquitous at the level of society as well as the individual. Despite their importance, characterizing the sociological, psychological, and neurological features of moral events is in its infancy. Much of the recent research has focused on a priori philosophical frameworks and has used artificial events as probes, in part because collecting, characterizing and analyzing real-life moral events is a major undertaking. This dissertation attempts such an undertaking.

758 autobiographical memories of personal moral events were collected from a well-characterized and representative sample of 100 healthy Californian adults. Transcriptions of the events were further characterized and all data were entered into a large, searchable database. An initial set of results provides a detailed description of the participants and the memories of moral events they generated. This description showed that participants were highly representative of the general population of California; that the overall amount and patterns of moral events recollected was relatively universal and not influenced by gender, ethnicity, IQ, or personality; and that the moral events produced could generally be judged quite reliably both by the participants themselves as well as by independent raters.

The database was further analyzed with respect to three specific aims: (1) to study the semantic structure of real-life moral events; (2) to study the effects of focal

lesions to emotion-related brain regions on recollection of moral events; (3) to study the temporal distribution of autobiographical moral events. We found that real-life moral events have a hierarchical structure, with two broad categories of “good” and “bad”, and subordinate categories of “good”, “lying”, “stealing”, and “hurting another person”. These categories define the most common scripts encountered in real life that have strong moral value. In studying neurological patients with focal lesions to the ventromedial prefrontal cortex or the amygdala, we found no evidence for a notable skew in the moral events that were recollected, further evidence for the universality and robustness of such events in our autobiography. Finally, we found that positively valenced moral events were systematically recalled as being more recent in time than negatively valenced moral events, a temporal bias that was independent of absolute participant age. The methods used here, the database that was constructed, and the scientific questions that were analyzed constitute the first comprehensive investigation of a large number of real-life moral events and provide a rich resource for future studies.

## TABLE OF CONTENTS

ACKNOWLEDGEMENTS.....	iii
ABSTRACT.....	v
TABLE OF CONTENTS.....	vii
LIST OF FIGURES.....	viii
LIST OF TABLES.....	x
Chapter 1: Review of the Literature.....	1
Section 1: Moral Philosophy.....	2
Section 2: Moral Psychology.....	13
Section 3: Neuroscience of Moral Cognition.....	17
Section 4: Autobiographical Memory.....	46
Chapter 1: Figures and Tables.....	53
Chapter 2: Methods and Materials.....	58
Section 1: Participants.....	59
Section 2: Methods.....	62
Section 3: Cue Words.....	71
Chapter 2: Figures and Tables.....	79
Chapter 3: Moral Narrative Database.....	81
Section 1: Rationale for a Moral Narrative Database:.....	81
Section 2: Narrative Selection.....	83
Section 3: Methods.....	87
Section 4: Descriptive Database Statistics.....	91
Section 5: Characterization of Moral Events.....	93
Part 1: Development of the Rating Tool.....	93
Part 2: Raters.....	102
Section 6: Inter-rater reliability:.....	110
Section 7: Database Development.....	111
Section 8: Database Documentation: A Guide to the Tables in the Database.....	112
Chapter 3: Figures and Tables.....	137
Chapter 4: Results.....	142
Section 1: Characterization and Representativeness of the Participants.....	143
Part 1: Demographics.....	143
Part 2: Neuropsychology.....	151
Section 2: Characterization of the Moral Narratives.....	162
Part 1: Analysis by Cue Word.....	162
Part 2: Language Inquiry and Word Count Analysis.....	166
Part 3: Analysis by Rater Categories.....	169
Part 4: Analysis of Session Three Data.....	173
Section 3: Cluster Analysis Reveals Implicit Categories in Moral Narratives.....	178
Section 4: Using the Database for Further Research.....	184
Section 5: Effects of Brain Damage on Moral Narratives.....	189
Section 6: Temporal Biasing of Autobiographical Moral Events.....	194
Chapter 4: Figures and Tables.....	206
Appendices.....	278
References.....	315

## LIST OF FIGURES

Figure 1: Timeline of Writing about Morality	54
Figure 2: Kohlberg's Stages of Moral Development	55
Figure 3: Features of the Two Cognitive Processing Systems	56
Figure 4: Structure of Memory	57
Figure 5: Distribution of Number of Moral Memories per Participant	138
Figure 6: Distribution of Number of Words per Narrative	139
Figure 7: Table Structure of Major Database Divisions	140
Figure 8: Age Distribution	207
Figure 9: Handedness Distribution	208
Figure 10: Educational Distribution	209
Figure 11: Income Distribution	210
Figure 12: Language Distribution	212
Figure 13: Ethnicity Distribution	213
Figure 14: Marital Status Distribution	214
Figure 15: Distribution of Number of Children	215
Figure 16: Sexual Orientation Distribution	216
Figure 17: Sexual Preference Distribution	217
Figure 18: Religious Distribution	218
Figure 19: Degree of Religious Influence	219
Figure 20: Comparison of Religious Beliefs	220
Figure 21: Political Value Distribution	222
Figure 22: IQ Distribution	223
Figure 23: Model of Relationship Among Values	225
Figure 24: Distribution of Ethical Positions	230
Figure 25: EPQ Distribution of Relativism Scores	231
Figure 26: EPQ Distribution of Idealism Scores	232
Figure 27: Utilitarian Score	233
Figure 28: Correlation Matrix Organized by Participant Data	234
Figure 29: Correlation Matrix Organized by Narrative Data	235
Figure 30: Cue Word Frequency Distribution	237
Figure 31: Cue Frequency Distribution in Presentation Order	238
Figure 32: Gender Effects on Cue Word Frequencies	239
Figure 33: Effect of Age on Cue Word Frequencies	240
Figure 34: Frequency in Lexicon Does Not Determine Number of Elicited Memories	241
Figure 35: Distribution of Positive and Negative Words in Narratives	242
Figure 36: Narrative Categorization: What Did They Do?	243
Figure 37: Narrative Categorization: Why Did They Say They Did It?	244
Figure 38: Narrative Categorization: How Did They Feel?	245
Figure 39: Rater Judgments of Narrative Morality	247
Figure 40: Distribution of Change in Moral Ratings	248
Figure 41: Change in Moral Behavior Ratings Over Time	249
Figure 42: Dendrogram of Hierarchical Cluster Analysis of All Narratives	251
Figure 43: Plot of Scree Test	252



Figure 44: Factor Loadings for the 3 Bad Clusters	253
Figure 45: Varimax Rotation Good Cluster	254
Figure 46: Word Frequency Analysis Bad Factor 1	255
Figure 47: Word Frequency Analysis Bad Factor 2	256
Figure 48: Word Frequency Analysis Bad Factor 3	257
Figure 49: Word Frequency Analysis Good Narratives	258
Figure 50: Distribution of Selected Narratives from Each Bad Cluster	259
Figure 51: Select Query Structure for Gender Ratings	263
Figure 52: Number of Words per Narrative	267
Figure 53: Number of Moral Narratives Told	268
Figure 54: LIWC Analysis of Positive and Negative Emotion Words	269
Figure 55: Cue Word versus Mean Age of Memory	271
Figure 56: Valence of Cue Word Predicts Age of Memory	272
Figure 57: Valence of Rating Impacts Recency of Memory	273

## LIST OF TABLES

Table 1: Linguistic and Valence Properties of Cue Words	80
Table 2: Comparison of Database Narratives to Language and Inquiry Word Count Standards	141
Table 3: Occupations	211
Table 4: Religious Distribution in Southern California	221
Table 5: Value Definitions from the Schwartz Value Inventory	224
Table 6: Schwartz Value Inventory Results	226
Table 7: Attributes of Affect	227
Table 8: Comparison: Database Participant Affect Scores with Normative Values	228
Table 9: Forsyth's Taxonomy of Ethical Ideologies	229
Table 10: Location of Action in the Narratives	246
Table 11: Rating Categories that Contribute to Each Bad Factor	260
Table 12: Top 5 Eliciting Cue Words	261
Table 13: Outcome from Data Mining Example	264
Table 14: Background and Neuropsychological Data	266
Table 15: Demographics and Neuropsychological Profiles	274
Table 16: Cue Words Used in Recall Protocol	276

## *Chapter 1: Review of the Literature*

The events that constitute our autobiography largely define who we are individually and what we have experienced, making them of prime importance in developmental and social psychology. Closely intertwined with the events that actually transpired – and of necessity linked in the studies described in this dissertation – are our memories of them. The events we are able to recollect form the explicit part of our autobiographical memory and the part most amenable to experimental investigation, although they are a filtered version of the entirety of events we actually experienced.

This dissertation thus concerns two large topics: morality and memory, and investigates their intersection. This introductory chapter reviews some of the literature pertinent to each of these topics. I begin with a brief overview of philosophical views about morality, because this provides an important backdrop both for the way we as scientists may think about the data I will show, but also because the data themselves may well be influenced by philosophical, moral views that participants hold. This first section will also bring out a key point that motivated my study and that makes the data produced especially valuable for future experiments: there are several philosophical theories of morality (not to speak of psychological theories, of which there are also several). There is little agreement on how to structure moral events or how best to probe moral reasoning or judgment. One response to this situation is the data of this dissertation, which sought to characterize the moral events that people actually experience in everyday life and which now provides an organized database from which well-characterized, rated and veridical moral vignettes could be drawn as future experimental stimuli.

After reviewing some of the philosophical views on morality, I very briefly summarize some of the psychological views before going on to review in more detail current work in cognitive neuroscience that attempts to look at the neural substrates of moral processing. This most recent work in cognitive neuroscience also relates back to the philosophical and psychological theories that are presumably informed by the neuroscience data. I conclude with a quick review of autobiographical memory, a huge topic in its own right and not the primary focus of this dissertation.

## Section 1: Moral Philosophy

*We do not act rightly because we have virtue or excellence, but we rather have those because we have acted rightly. – Aristotle*

### **Moral Philosophy:**

Moral philosophy can be divided into three branches: metaethics, normative ethics and applied ethics<sup>1</sup>. Metaethics addresses broad philosophical questions about the role and meaning of ethical principles; normative ethics describes the framework of human moral standards; and applied ethics uses the concepts of metaethics and the framework of normative ethics to solve current social ethical dilemmas. While metaethics and applied ethics are interesting topics in moral philosophy, it is normative ethics that provides the intersection between moral philosophy and neuroscience that drives this work. (It also provides an intersection between moral philosophy and moral psychology,

---

<sup>1</sup> Fieser, J. (2006). "Ethics." from <http://www.iep.utm.edu/e/ethics.htm> .

as described by Anscombe<sup>2</sup>.) Thought experiments in moral philosophy, particularly ‘The Trolley Problem’ (see Appendix A), have inspired researchers across disciplines to design experiments to test moral judgment and the theories of normative ethics<sup>3</sup>.

### **Normative Ethics:**

Normative ethics seeks to describe which actions people should believe are right and wrong; that is, the norms for moral behavior. It is a prescriptive discipline – rather than describing those actions that people do believe are right or wrong, normative ethics is intended to develop a system that encompasses what people should embrace as moral standards. It is important to understand this distinction when approaching moral philosophy from an experimental perspective. The three major branches of normative ethics are not theories seeking to be proven (moral psychology, instead, has such theories); they are descriptions of what philosophers believe should be. This understanding helps to inform the logic behind the development of the Trolley Problem as a thought experiment, but also undermines its use as an experimental tool.

In seeking to describe what moral beliefs people ought to have, three branches (or theories) of normative ethics have emerged: virtue ethics, consequentialism and deontology. Each of these describes a system of moral conduct that highlights a different part of a moral action. Virtue ethics focuses on the moral character of the agent who performs the action. The morality of an action depends on whether an agent brings virtues like courage or compassion to the action. Consequentialism focuses on the results of an action. Under this theory, the most moral action is that which brings the best

---

<sup>2</sup> Anscombe, G. E. M. (1958). "Modern Moral Philosophy." *Philosophy* 33(124): 1.

<sup>3</sup> It is worth noting (as described in the next paragraph) that normative theories were not developed to be tested, but rather to prescribe what should be morally acceptable.

consequences. (What is ‘best’ is a question developed in detail by several consequentialist theories.) In consequentialism, the ends may justify the means of many actions. Deontology, however, focuses on the action itself<sup>4</sup>. Deontologists believe that actions carry with them a moral value, regardless of the agent who performs them or the results of the action. Under deontology, to be morally correct, some actions must always be performed and others must never be done. While the focus of each theory is quite different, the underlying desire of all normative ethics is to describe a system that uses a single set of principles to determine what morally ought to be done.

### **Virtue Ethics:**

Virtue ethics focuses around the central question, “How should I live?” or “What kind of person should I be?” Thus, rather than emphasizing total welfare as consequentialism does or the normative value of actions, virtue ethics focuses on the person. Although predated by similar writings of Chinese philosophers and Plato, Aristotle’s ‘Nicomachean Ethics’ is thought of as the founding work in virtue ethics. Aristotle believed that achieving a virtuous character should be the goal for any human’s life<sup>5</sup> and used this premise to develop the principles of virtue ethics: *arête* (virtue), *phronesis* (practical wisdom) and *eudaimonia* (usually translated as happiness or

---

<sup>4</sup> Deontology is from the greek word ‘deon’ which means duty.

<sup>5</sup> From Aristotle Nicomachean Ethics. H. Rackham. Book 1, Ch. 7: “If then the function of man is the active exercise of the soul's faculties in conformity with rational principle”... “and if we declare that the function of man is a certain form of life, and define that form of life as the exercise of the soul's faculties and activities in association with rational principle, and say that the function of a good man is to perform these activities well and rightly, and if a function is well performed when it is performed in accordance with its own proper excellence – from these premises it follows that the Good of man is the active exercise of his soul's faculties in conformity with excellence or virtue, or if there be several human excellences or virtues, in conformity with the best and most perfect among them. Moreover, to be happy takes a complete lifetime; for one swallow does not make spring, nor does one fine day; and similarly one day or a brief period of happiness does not make a man supremely blessed and happy.”

flourishing). The interplay of the first two – a virtuous character which is ruled by practical choices – leads to a eudaimon life; one that is fulfilling and good.

Personal character or disposition is paramount in virtue ethics. It is both the problem and the solution to the question, “How should I live?” To live a good life, one must have a virtuous character and to achieve this, one must practice virtues. Aristotle believed that virtuous behavior could be learned and, indeed, virtues must be repeated many times to become part of one’s disposition<sup>6</sup>. While some of the specifics of modern virtues have changed, the principle of striving for a virtuous character has remained. To have a virtuous character a person’s life must embody the virtue in question. For example, a courageous person not only is brave when faced with a dangerous circumstance, but is also one who teaches courage to others, who is disappointed when hearing about someone who is not courageous and who does not have difficulty with the idea of being brave in future situations. Thus, a courageous person cannot be judged by one action, but instead by the sum of all of her choices, beliefs and actions. In this way, a person’s virtuous character is a fixed part of who she is and one would expect to find this virtue displayed in all situations. However, humans are not perfect and this definition of a virtuous *person* necessarily suggests that this is an ideal to strive for and that most people will not live their lives with perfect virtue. However, since Aristotle believed that striving for virtue is the path to a flourishing life, it is enough, in virtue ethics, to try.

There are thus at least two differences between virtue ethics and the consequentialist and deontological views that are described in more detail below. First, in focusing on the person, it acknowledges real-life constraints and limitations. It is

---

<sup>6</sup> It is worth noting that habit alone, however, does not make one virtuous. A virtuous character must be developed through choice, not just habituation. Aristotle felt that becoming virtuous could take a lifetime of moral development.

focused on practice as opposed to a more abstract normative stance. Second, in focusing on effort and acquisition, it emphasizes learning and change, which are also not concepts that consequentialism or deontology have. Arguably, these features of virtue ethics make it the most relevant with respect to this dissertation, since the data here speak to moral actions and outcomes experienced by real people in the real world – often under constraints and extenuating circumstances and often with subsequent regret and change in the person.

While Aristotle encouraged the development of a virtuous character, the principle of practical (or moral) wisdom is necessary to ensure the balance needed to be virtuous in all situations. Aristotle advocated that virtue lies in the mean between two extremes<sup>7</sup>. Courage, for example, can become foolhardiness when faced with an unbeatable danger or fear when faced with an intimidating challenge. Phronesis encapsulates the idea that knowledge and choice are integral parts of virtue. A virtuous disposition must arise out of choice and the application of knowledge gained through experience. In order to be virtuous, one must understand a situation, be able to weigh the relative merits of competing interests and decide upon the path that is most virtuous. The recognition of those factors that are morally salient can be thought of as the key to practical wisdom. Understanding those things that ‘matter’ allows a virtuous person to do the ‘right’ thing in every situation, even when the nuances of the situation may change what the correct action is. Aristotle differentiates between natural virtue (as in kindness expressed by a child) and a virtuous character. A child, when doing the right thing, does not do so out of an understanding of possible choices, rather merely out of feeling. Aristotle felt

---

<sup>7</sup> From Aristotle *Nicomachean Ethics*. H. Rackham. Book 2, Chapter 6: And it is a mean state between two vices, one of excess and one of defect. Furthermore, it is a mean state in that whereas the vices either fall short of or exceed what is right in feelings and in actions, virtue ascertains and adopts the mean.



compelled to distinguish this specifically because doing the virtuous thing without knowing why cannot lead to eudaimonia. It is the combination of wisdom with a character that is attuned to being virtuous that leads to a life of eudaimonia.

Eudaimonia<sup>8</sup> is the root of most of the variance in modern virtue theories. Aristotle explains that eudaimonia is not merely the outcome of a virtuous life, but also the purpose for any human's life. This claim comes from his idea that each item in the world has one primary function. Using Aristotle's examples, a knife's function is to cut and a human's function is to reason. To this end, the purpose of a human life must be to reason well, thereby living a eudaimon life. This naturalistic teleological argument is part of the reason that virtue ethics fell out of favor during the 19<sup>th</sup> century. While modern virtue ethics has been split into several types<sup>9</sup>, nearly all agree that there is a link between virtue and eudaimonia<sup>10</sup>. However, the description and understanding of what sort of link it might be is the topic of a great deal of modern discussion. It is worth noting that modern virtue ethics encompasses more than just Aristotelian virtue ethics. Ethics of care and agent-based approaches to virtue ethics are interesting and emerging fields in virtue ethics.

### **Consequentialism:**

Consequentialism is based on the general principle that an action is morally correct when it leads to the best consequences. The details of what consequences are,

---

<sup>8</sup> Eudaimonia is most commonly translated as 'happiness' or 'flourishing'. Both translations have elements that capture the spirit of the word and elements that don't. For an explanation of these, see Hursthouse, R. (2007). "Virtue Ethics." from <http://plato.stanford.edu/entries/ethics-virtue/>. The literal translation is "having a good guardian spirit".

<sup>9</sup> As many as 75 varieties! Swanton, C. (2003). 75 Virtue Ethics. International Encyclopedia of the Social and Behavioral Sciences, Elsevier Science.

<sup>10</sup> A major exception to this is Michael Slote. Slote, M. (1992). From Morality to Virtue. New York, Oxford University Press.

how one determines what the ‘best’ is, and for whom the consequences should be considered have lead to a great number of different types of consequentialism. Rather than discuss these at length, this discussion will focus on the relevant historical elements of consequentialism.

The classic form of consequentialism, utilitarianism, was championed by John Stuart Mill<sup>11</sup> during the 19<sup>th</sup> century (as well as by Jeremy Bentham<sup>12</sup> and Henry Sidgwick during the 18<sup>th</sup> and 20<sup>th</sup> centuries, respectively). Utilitarianism is the strictest application of consequentialist thinking; defining all moral actions as a function of their utility<sup>13</sup>. Utilitarian philosophers before Mill had generally agreed that the utility of any action was defined by the amount of happiness that it produced. This conclusion was drawn from the idea that achievement of pleasure and avoidance of pain are the only true driving factors for human behavior and thus what is right and wrong must grow out of that drive. (Bentham espoused this belief so fully that he believed that the value of an action could be calculated by measuring the amount of pleasure that resulted, the length of time this pleasure lasted and so forth<sup>14</sup>.) This pleasure-driven line of thinking was often misrepresented by the public of Mill’s day as a purely hedonistic theory<sup>15</sup> despite Bentham’s ‘greatest happiness’ principle specifically taking into consideration the community at large. Mill set forth to reframe utilitarian thinking in his book,

---

<sup>11</sup> Mill, J. S. (1998). Utilitarianism New York, Oxford University Press.

<sup>12</sup> Bentham, J. (1961). An Introduction to the Principles of Morals and Legislation. Garden City, Doubleday.

<sup>13</sup> Utility in this sense can be thought of as the amount of satisfaction brought about by an action.

<sup>14</sup> Kemerling, G. (2002). "Utilitarianism." from <http://www.philosophypages.com/hy/5q.htm#util>.

<sup>15</sup> At least according to Mill, J. S. (1998). Utilitarianism New York, Oxford University Press. Ch. 2: “Yet the common herd, including the herd of writers, not only in newspapers and periodicals, but in books of weight and pretension, are perpetually falling into this shallow mistake. Having caught up the word utilitarian, while knowing nothing whatever about it but its sound, they habitually express by it the rejection, or the neglect, of pleasure in some of its forms; of beauty, of ornament, or of amusement. Nor is the term thus ignorantly misapplied solely in disparagement, but occasionally in compliment; as though it implied superiority to frivolity and the mere pleasures of the moment.”

*Utilitarianism*, as well as provide a counter-theory to moral intuition – the reigning theory of the day<sup>16</sup>.

Mill's definition of utilitarianism can be used as the framework to discuss all other consequentialist theories, based on their divergence from this basic statement.

“The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. By happiness is intended pleasure, and the absence of pain; by unhappiness, pain, and the privation of pleasure.”<sup>17</sup>

Mill specifically defines the measurement of consequences in terms of happiness – firmly placing utilitarianism on a track away from pure hedonism. He manages to do this without rejecting the idea that happiness may encompass pleasure by suggesting that happiness also involves the avoidance of pain. While Mill settled upon utility as the key to his doctrine, other philosophers began to argue that the satisfaction derived from an action might not be the only necessary measurement of consequences.

Many consequentialists agree that the best outcome can mean the most happiness, but perhaps it is not the only value that should be considered. One of the driving questions of modern consequentialism is: “What is the ‘best’ consequence?”

### **Deontology:**

---

<sup>16</sup> Again, according to Mill, J. S. (1998). *Utilitarianism* New York, Oxford University Press. Ch. 1:

“...recourse to the popular theory of a natural faculty, a sense or instinct, informing us of right and wrong.”

<sup>17</sup> Mill, J. S. (1998). *Utilitarianism* New York, Oxford University Press. Ch. 2.

The term ‘deontology’ is used to describe duty-based moral theories. First coined by C.D. Broad in 1930, the Greek words *deon*, meaning ‘obligation’ or ‘duty’, and *telos*, meaning ‘goal’, ‘end’ or ‘purpose’, were used to differentiate the two major ethical theories of his era<sup>18</sup>. (At the time virtue ethics had fallen out of favor, mainly due to difficulties with the naturalistic fallacy.) Deontology focuses on the obligation of an individual to adhere to universal moral rules, while teleology focuses on the outcome of an action to determine its moral value. Frequently, deontology and consequentialism are discussed in tandem as they are, in many ways, opposing theories. However, it is important to note that Broad’s distinction was slightly different from the typical contrast. Not all consequentialists are teleologists<sup>19</sup>. Teleologists, in Broad’s definition are fundamentally concerned only about the “rightness of actions in terms of the goodness of the state of affairs that occurs because of that action”. This means that regardless of the intentions of the person who carried out the action, something that accomplishes greater good in the world is right.

Several types of deontological theories have been posited but the one most commonly discussed is that of Immanuel Kant<sup>20</sup>. In 1785, Kant published his theory of duty-based ethics, developed from the position of calculated rationality. From this perspective, Kant found that only one concept, ‘a good will’, could be thought of as good without qualification. All other good things could be dependent on circumstances, but

---

<sup>18</sup> Broad, C. D. (1930). Five Types of Ethical Theory. New York, Harcourt, Brace and Co.: 277-278.

<sup>19</sup> Broad himself notes this, saying: ‘Utilitarianism, in some of its forms, would be an example of [teleological ethics]. But [Henry] Sidgwick, though a Utilitarian, definitely rejects the view that "right" means "conducive to good"’. Broad, C. D. (1930). Five Types of Ethical Theory. New York, Harcourt, Brace and Co. p.278.

<sup>20</sup> Notable others include Samuel Pufendorf’s 17<sup>th</sup> century theories about duties to God, oneself and others; Divine Command Theory, put forth by 18<sup>th</sup> century Calvinists and Rene Descartes; Rights Theory championed by John Locke (17<sup>th</sup> century) and Thomas Jefferson (18<sup>th</sup> century); and W.D. Ross’s (1930) *prima facie* duty theory, also known as pluralistic deontology.

since a good will is purely intentional, it cannot be impacted by external influences. Importantly, since a good will is encapsulated, even the outcome of an action guided by good will cannot affect it. This is a critical element of Kant's philosophy and it is the main distinction between deontology and consequentialism; deontology finds that the consequences cannot matter in a theory of morality while consequentialism holds that maximizing good consequences is the defining feature of a moral system.

Kant believes that a duty to principle (especially in the face of choices that might be personally preferable) defines the behavior of a rational, moral person. Additionally, any principle that defines a moral system must be abstract enough to be applied in all circumstances to all people equally; that is, it must have the property of universalizability<sup>21</sup>. Kant developed his 'categorical imperative' as distinct from 'hypothetical imperatives', which impel one to do something for a particular purpose. Morally, deontologists find this troubling as it requires the person to be duty bound to take an action only if they believe both in the goal and that the action taken will achieve the desired results. By comparison a categorical imperative requires (morally) only that the person perform the action. They are duty bound to this principle without reference to its possible outcomes. In designing his categorical imperative, Kant created three forms of a moral rule. Its first (and most commonly used) formulation is: "I ought never to act except in such a way that I could also will that my maxim should become a universal law."<sup>22,23</sup> For Kant, from this principle all moral behavior ought to follow.

---

<sup>21</sup> Other philosophers, like Richard Hare, also believe that universalizability is a defining feature of moral judgments. They must apply equally to all people. Hare, R. M. (1991). The Language of Morals, Clarendon.

<sup>22</sup> Kant, I. (1785). Grounding for the Metaphysics of Morals, Hackett. p.30.

<sup>23</sup> This formulation is usually called the Formula of the Universal Law of Nature. The other two formulations of the categorical imperative are the Humanity Formula: "Act in such a way that you treat humanity, whether in your own person or in the person of any other, always at the same time as an end and

Deontology's strength and weakness lies in its moral absolutes. If all actions have inherent rightness or wrongness and universal moral rules exist, what happens when these rules conflict? Benjamin Constant, a Swiss philosopher, challenged Kant about the universality of these absolutes by asking if it is acceptable to lie to a murderer who is asking for a victim's location. Kant replied, "Truthfulness in statements that cannot be avoided is the formal duty of man to everyone, however great the disadvantage that may arise therefrom for him or for any other."<sup>24</sup> While this may seem to merely refute Constant's question without further explanation, Kant's 'formal duty' asserts that if one person lies then all people must assume that all people lie all the time. Were that to be the case, no lies would be believed and lying to the murderer would not result in any help to the victim. More importantly, however, is that in an absolutist deontological system the outcome that may result from telling the truth does not affect the moral value of that choice<sup>25</sup>. It is upon this point that deontology and consequentialism have stood, diametrically opposed, for more than 200 years.

The parallel development of these two conflicting theories and the resurgence of virtue ethics has led to the development of three mature, moral philosophies. While these philosophies are designed to describe how we ought to behave most morally, analytic philosophers have found there is much to be learned about how these philosophies are applied in the real-world using thought experiments. The trolley problem (described in Appendix A), introduced by Philippa Foot and Judith Jarvis Thompson in the 1970s, was

---

never merely as a means to an end.”; and the Kingdom of Ends Formula: “Therefore, every rational being must so act as if he were through his maxim always a legislating member in the universal kingdom of ends.” Kant, I. (1993). Grounding for the Metaphysics of Morals, Hackett.

<sup>24</sup> “On a Supposed Right to Lie Because of Philanthropic Concerns” from Kant, I. (1993). Grounding for the metaphysics of morals; with, On a supposed right to lie because of philanthropic concerns (Third Edition). Indianapolis, Hackett .

<sup>25</sup> Not all deontologists are as absolutist as Kant. W.D. Ross is an example of one such philosopher.

devised to help tease apart exactly that<sup>26</sup>. Depending upon the dilemma and its permutation, people in the real world can be found to make either deontological or utilitarian judgments. There are individual differences, of course, but there is good general agreement among the testing population about most of these moral dilemmas<sup>27</sup>. The addition of these dilemmas to cognitive neuroscience has provided insight about the effects of brain damage on utilitarian judgments and the brain areas that are involved in these decisions. Perhaps most of all, however, they have illuminated how much more there is to learn about moral judgment. The data in this dissertation seek to provide new tools for philosophers to ask the questions that will best elucidate not just how we *do* behave, but how we *ought* to behave.

## Section 2: Moral Psychology

Thinking in moral psychology has been dominated by the work of Lawrence Kohlberg, a developmental psychologist whose work was defined by his description of the six stages of moral development<sup>28</sup>. Kohlberg believed that moral behavior was a function of moral reasoning and that the development of moral reasoning from childhood to adulthood could be described in stages. The idea of a stage theory of moral development was influenced by work from Jean Piaget, a Swiss philosopher who

---

<sup>26</sup> Foot, P. (1978). The Problem of Abortion and the Doctrine of the Double Effect in Virtues and Vices and Other Essays in Moral Philosophy. Berkeley and Los Angeles, University of California Press.; Thomson, J. J. (1986). Rights, Restitution, and Risk: Essays in Moral Theory. Cambridge, MA, Harvard University Press.

<sup>27</sup>Hauser, M. D., F. Cushman, et al. (2007). "A dissociation between moral judgments and justifications." Mind & Language **22**(1): 1-21.

<sup>28</sup> Kohlberg, L. (1958). The Development of Modes of Thinking and Choices in Years 10 to 16, University of Chicago. Ph.D. dissertation.

developed a theory of cognitive development<sup>29</sup>, explaining how children progress in cognitive development from stage to stage. Piaget's stages were defined as constructive; once a child has progressed into a new stage, they rarely lose the ability to use those cognitive abilities. Stages were also necessary – it was not possible to skip a stage because each was dependent on the cognitive growth during the previous stage. Kohlberg adopted the same rules for his stages, for reasons nearly identical to Piaget's. Kohlberg did acknowledge, however, that particularly in moral behavior, people do not necessarily function at the highest level of their development all the time.

Kohlberg's six stages of moral development are divided into three categories: pre-conventional (typically from birth to age 9), conventional (typically from age 9 to 20) and post conventional (typically after age 20; with most people never reaching the final stage). (See Figure 2 for a graphical depiction of these stages.) In the pre-conventional stages (1 and 2), morality is judged by the consequences of an action. Stage one is thought of as the obedience and punishment orientation and is typified by judging morality based on how much one will be punished for doing a particular action. This often leads to a belief that things that incur more punishment are more wrong. Stage two is primarily focused on self-interest. Reasoning in this stage is primarily involved with what the person wants and can be thought of as morally relativistic. Punishment is considered as a risk instead of a definition of morality. The conventional stages of morality (3 and 4) are driven by an individual's understanding of themselves as a part of a larger society. Stage 3 reasoning is focused on interpersonal relationships. The morality of an action is determined by its impact on those around us. Motivations for an action begin to play a role in determining how right or wrong an action is. People using this

---

<sup>29</sup>Piaget, J. (1977). The Essential Piaget. New York, Basic Books.



reasoning also typically believe that others will all feel the same way (hence, the title ‘conventional’)<sup>30</sup>. Stage 4 reasoning focuses on the role that moral behavior plays in maintaining a social order. It acknowledges an individual’s right to make bad choices, but concludes that if everyone did bad things, society would fall apart; thus one must make morally good choices to prevent this from happening. Some adults do not progress beyond this stage. In addition Kohlberg later proposed an intermediate stage here between stages 4 and 5. The post-conventional stages (5 and 6) of morality balance the rights of the individual with the rights of the society/world. Stage 5 reasoning focuses on the balance between society’s laws and individual happiness. Utilitarian thinking – the most happiness for the most people – is often associated with stage 5 and democratic government is proposed to arise out of this stage of moral development as well. People using stage 5 reasoning consider that rights and laws do not always coincide. Stage 6 reasoning is rarely achieved. Kohlberg had difficulty finding subjects who routinely employed this form of reasoning in real life and eventually came to consider it as largely theoretical. Reasoning in stage 6 is driven by a principle of universal ethical values. Stage 6 reasoning finds that all people require equal consideration and no law or rule should be implemented that disadvantages one person in favor of another. Kohlberg’s belief is that moral reasoning is justice driven and stage 6 is the pinnacle of justice-oriented reasoning.

To measure these stages, Kohlberg developed a series of moral dilemmas that were administered in a guided interview-style testing session<sup>31</sup>. The justification (not the

---

<sup>30</sup> It is also important to note that this is the stage that best meshes with Carol Gilligan’s care orientations.

<sup>31</sup> The classic example is that of “Heinz Steals the Drug”: “In Europe, a woman was near death from a special kind of cancer. There was one drug that the doctors thought might save her. It was a form of radium that a druggist in the same town had recently discovered. The drug was expensive to make, but the druggist was charging ten times what the drug cost him to make. He paid \$200 for the radium and charged \$2,000 for a small dose of the drug. The sick woman's husband, Heinz, went to everyone he knew to borrow the money, but he could only get together about \$ 1,000 which is half of what it cost. He told the druggist that

answers themselves) that subjects provided for their answers were used to determine an individual's moral stage. We considered administering an updated version<sup>32</sup> of the Kohlberg dilemmas to our participants but found that even the 'updated' version was too dated to seem relevant to our participants.

Kohlberg's work received criticism for being too focused on Western values and not adequately accounting for theories of morality espoused by other cultures<sup>33</sup>, but perhaps most famously for being sexist. Carol Gilligan, a colleague of Kohlberg's, took umbrage with Kohlberg's findings that women were generally less morally developed than men and began to question his fundamental assumption of justice as the determining orientation of morality. She felt that "... men *tend* to embrace an ethic of rights using quasi-legal terminology and impartial principles ... women *tend* to affirm an ethic of care that centers on responsiveness in an interconnected network of needs, care, and prevention of harm. Taking care of others is the core notion<sup>34</sup>." Gilligan developed an alternative stage theory based on an orientation of care, rather than one of justice<sup>35</sup>. The care stage theory of moral development has a pre-conventional stage with a focus on individual needs, a conventional stage that transitions from selfishness to responsibility to others, where goodness is defined by self-sacrifice and a post-conventional stage that

---

his wife was dying and asked him to sell it cheaper or let him pay later. But the druggist said: "No, I discovered the drug and I'm going to make money from it." So Heinz got desperate and broke into the man's store to steal the drug-for his wife. Should the husband have done that?" Kohlberg, L. (1963). "The development of children's orientations toward a moral order. I: Sequence in the development of human thought." *Vita Humana* 6: 11-33(a). p. 19.

<sup>32</sup> Hogan, R. and E. Dickstein (1972). "A measure of moral values." *Journal of Consulting and Clinical Psychology* 39(2): 210-214.

<sup>33</sup> Simpson, E. (1974). "Moral development research: A case study of scientific cultural bias." *Human Development* 17: 81-106.

<sup>34</sup> Beauchamp, T. L. and J. F. Childress (2001). *Principles of Biomedical Ethics, 5th edition*. New York, Oxford University Press. p. 371.

<sup>35</sup> Gilligan, C. (1977). "In a Different Voice: Women's Conceptions of Self and Morality." *Harvard Educational Review* 47(4): 1-37.

focuses on a principle of doing no harm (to self or others), which involves recognizing that one's own needs are as important as others.

The re-emergence of social intuitionism has moved current moral psychology away from the rationalism embraced by Kohlberg. Moral motivation<sup>36</sup>, moral hypocrisy and integrity<sup>37</sup>, moral intuition<sup>38</sup>, moral perspective taking (simulation)<sup>39</sup> and imitation<sup>40</sup>, and altruism are all topics being currently studied.

### Section 3: Neuroscience of Moral Cognition

The advent of recent technologies has allowed us to take the study of morality beyond the realm of philosophical thought experiments and into the analytical realm of the scientific method. Functional magnetic resonance imaging (fMRI) gives us a way to look specifically at which brain structures are involved in moral decision making.

Approaching the questions of moral cognition from a biological point of view, however, presents an interesting dilemma: Can the traditional philosophical methods be usefully applied to a biological inquiry of moral decision-making? Several neuroscientists have attempted to bridge this gap using variations of philosophical thought experiments<sup>41</sup>.

While this continuity of academic thought is admirable and has elicited some interesting

---

<sup>36</sup> Hardy, S. A. (2006). "Identity, Reasoning, and Emotion: An Empirical Comparison of Three Sources of Moral Motivation." *Motivation and Emotion* **30**: 207-215.

<sup>37</sup> Batson, C. D. and E. R. Thompson (2001). "Why don't moral people act morally? Motivational considerations." *Current directions in psychological science* **10**(2): 54-57.

<sup>38</sup> Knobe, J. (2005). "Theory of mind and moral cognition: exploring the connections." *Trends in Cognitive Sciences* **9**(8): 357.

<sup>39</sup> Batson, C. D., D. A. Lishner, et al. (2003). "... As you Would have Them Do Unto You: Does Imagining Yourself in the Other's Place Stimulate Moral Action?" *Personality & Social Psychology Bulletin* **29**(9): 1190.

<sup>40</sup> Meltzoff, A. N. and J. Decety (2003). "What imitation tells us about social cognition: a rapprochement between developmental psychology and cognitive neuroscience." *Philosophical Transactions of the Royal Society of London* **358**(1431): 491.

<sup>41</sup> See Appendix A for a discussion of the trolley problem.

results, it falls far short of the biological ideal of ecological validity. These philosophical problems do not mimic real life situations and in that way can never give us accurate information about what real people are doing when making real moral choices in their daily lives. They may give us insight into the more abstract functions of moral judgment; indeed, there is much of value that can be gained from these dilemmas. However, by their design, they cannot accurately assess the moral decision making that we encounter every day. In this section, the current neuroscience research about moral cognition will be discussed, as well as a brief overview of the research on moral intuition and moral reasoning and moral emotions.

To review the intersection of neuroscience, philosophical thought and moral cognition, it will be useful to progress from those researchers and ideas that fundamentally come from a philosophical school of thought to those that approach the problem from a biological perspective. In this way, we will cover the breadth of the field and examine the underlying biases that come from each way of examining moral problems.

Recent work by William Casebeer<sup>42</sup>, a moral philosopher, approaches neuroscience as a tool to help illuminate the debate in moral philosophy among the three leading theories in normative ethics: virtue ethics, deontology and utilitarianism<sup>43</sup>.

Casebeer's argument can be distilled as follows: Each of the three leading philosophical theories can be used to design a model for how the brain computes moral problems.

Utilitarianism, generally, seeks to maximize happiness. (For Casebeer, it is irrelevant to

---

<sup>42</sup> Casebeer, W. D. (2003). "Moral cognition and its neural constituents." *Nature Reviews Neuroscience* 4: 841-846., Casebeer, W. D. (2003). "The Neural Mechanisms of Moral Cognition: A Multiple-Aspect Approach to Moral Judgment and Decision-Making." *Biology and Philosophy* 18: 169-194.

<sup>43</sup> Casebeer chooses to discuss utilitarianism rather than the broader idea of consequentialism Casebeer, W. D. (2003). "Moral cognition and its neural constituents." *Nature Reviews Neuroscience* 4: 841-846.

the model whether this is personal happiness or happiness of the masses.) In order to achieve this goal, a brain must be capable of calculating which actions will eventually produce happiness. Deontology seeks to create categorical imperatives<sup>44</sup>. A brain that makes deontological judgments must facilitate the creation of universal rules and the application of those rules to the current moral situation. Virtue ethics seeks eudaimonia and a virtuous character. A brain that fulfills the needs of virtue ethics must be able to recognize virtuous behavior and happiness. Casebeer argues that, given the current neuroscience data, virtue ethics is the best theory for describing moral cognition. He does recognize that the goals of virtue theory are much more broadly described, but feels that the “give and take required by co-evolution of theories... would, hopefully, correct any such bias eventually.”<sup>45</sup>

While Casebeer’s idea that neuroscience may inform philosophy has strengths, he is strongly criticized by some<sup>46</sup> for failing to address a critical distinction between the two disciplines: their goals. Moral neuroscience seeks to understand the biological substrates of moral cognition and behavior while normative ethics seeks to describe a system of moral standards that people should uphold. The philosophical assertion that what *is* and what *ought to be* are distinct from one another was originally made by Hume<sup>47</sup>; the mistake in deriving one from the other is referred to as the ‘is/ought’ problem<sup>48</sup>. The is/ought problem does not necessarily prevent philosophers from using neuroscience to

---

<sup>44</sup> Kant, I. (1993). Grounding for the Metaphysics of Morals, Hackett. p. 30.

<sup>45</sup> Casebeer, W. D. (2003). "Moral cognition and its neural constituents." Nature Reviews Neuroscience 4: 841-846. p 842.

<sup>46</sup> Greene, J. D. Ibid. "From neural ‘is’ to moral ‘ought’: what are the moral implications of neuroscientific moral psychology?": 847-850.

<sup>47</sup> Hume, D. (2003). A Treatise of Human Nature, Project Gutenberg. Book III, Part I, Section I.

<sup>48</sup> The is/ought problem is often closely tied to the ‘naturalistic fallacy’; another problem that Casebeer’s critics have with his work. The ‘naturalistic fallacy’ is used to describe the claim that those things that are natural are good. It was first described by G.E. Moore. Moore, G. E. (1903). Principia Ethica, Cambridge University Press.

inform their theories, but it does require care when constructing those uses. In his critique of Casebeer, Joshua Greene points out that both Mill and Kant intended their theories to describe a more perfect system of moral thinking; a moral maturity for people to aspire to. To this end, neuroscience cannot and should not necessarily influence philosophical thought. What it can do, however, is to help better describe how humans generate their current moral judgments, which may help inform philosophical debates. Casebeer's claims seek to define which philosophical theories may be true, based on how people behave, not how they ought to behave. The use of neuroscience to investigate what people believe ought to be moral requires entirely different questions.

Casebeer's critics, especially Michael Ruse, also feel that his naturalized ethics<sup>49</sup> only "hack away" at the is/ought problem without addressing the larger naturalistic fallacy (i.e., ethical realism)<sup>50</sup>. Ruse believes instead in ethical non-realism – the idea that there are no real moral absolutes defined by our biology (or maybe even otherwise).

Other philosophers have had greater success in using moral neuroscience to investigate the merits of philosophical claims. Jesse Prinz has used findings from fMRI studies to investigate the debate between moral rationalism and moral sentimentalism/intuition<sup>51</sup>. His research has taken him so far as to write a compelling integration of appraisal theory and the somatic marker hypothesis. Walter Sinnott-Armstrong has used psychological research to develop a theory of affect heuristics – mental shortcuts for moral situations<sup>52</sup>. These findings will be discussed further in later

---

<sup>49</sup> Especially as described in Casebeer's recent book Naturalized Ethical Facts. Casebeer, W. D. (2003). Natural Ethical Facts: Evolution, Connectionism, and Moral Cognition. Cambridge, MA, The MIT Press.

<sup>50</sup> Ruse, M. (2004). "Book Review: Natural Ethical Facts: Evolution, Connectivism, and Moral Cognition by William D. Casebeer, Cambridge, MA: MIT Press, 2003." Evolutionary Psychology 2: 89-91.

<sup>51</sup> Prinz, J. (2006). "The Emotional Basis of Moral Judgments." Philosophical Explorations 9(1): 29-43.

<sup>52</sup> Sinnott-Armstrong, W. (2008). Framing moral intuitions. In Moral psychology, volume 2: the cognitive science of morality. W. Sinnott-Armstrong. Cambridge: MIT, Bradford Books: 47-76.

parts of this work. In order to better understand the context from which these and other moral researchers view moral cognition, it is important to understand the framework that has been developed. Specifically, one must understand the division of moral cognition into two component parts: moral intuition and moral reasoning.

### **Moral Intuition:**

The philosophical debate about the source of moral cognition is typified by the work of David Hume and Immanuel Kant. Kant believed so strongly in rationalism that he felt that behaving immorally could be nothing other than irrational. “The fundamental principle of morality — the categorical imperative — is none other than the law of an autonomous will. Thus, at the heart of Kant's moral philosophy is a conception of reason whose reach in practical affairs goes well beyond that of a Humean ‘slave’ to the passions<sup>53</sup>.” Hume, on the other hand, famously stated that “Reason is, and ought only to be, the slave of the passions, and can never pretend to any other office than to serve and obey them<sup>54</sup>.” This philosophical debate has informed thinking about moral cognition for hundreds of years. Work in psychology and recently in neuroscience has contributed empirical data to the question. As data have begun to accumulate about the neural mechanisms of moral cognition, theories have begun to arise to explain them. While variations abound, some consensus about the basic building blocks of moral cognition has begun to emerge. (1)Moral cognition appears to be a whole brain, distributed process,

---

<sup>53</sup> Johnson, R. (2008). Kant's Moral Philosophy. [The Stanford Encyclopedia of Philosophy](#). E. N. Zalta.

<sup>54</sup> Hume's most famous sentence occurs at Treatise, II, III, iii from Hume, D. (2003). *A Treatise of Human Nature*, Project Gutenberg.

rather than a ‘moral module’. Initial evidence from lesion patients like Phineas Gage<sup>55</sup> suggested the opposite – that patients with specific brain damage could selectively lose their ability to make moral choices in daily life. (These findings will be discussed in more detail later in this section.) However, recent work suggests that these patients had subtle underlying deficits in other cognitive domains as well. (2) Moral cognition involves both intuition and abstract reasoning and these may be dissociable processes.

Moral intuition is a fast, automatic evaluation that provides an immediate perception of right- or wrong-ness. Moral reasoning is a cognitively-intense, relatively slow process characterized by ‘thinking’ about the problem. (Details of these two systems can be seen in Figure 3.) John Allman, a biologist, has proposed that moral intuition arises in part from a unique class of bipolar cells found only in great apes and humans<sup>56</sup>; the von Economo neurons (VENs)<sup>57</sup>. VENs are located in the frontoinsula cortex (FI) and anterior cingulate cortex (ACC), regions responsible for interoception. These cells (especially because of their large size) may help to integrate ‘gut feelings’ (literally) with situational social information to produce moral intuitions. This is in keeping with Damasio’s somatic marker hypothesis (discussed later) and Haidt and Prinz’s models of embodied appraisals. In philosophy and psychology, rationalist models – which held that moral reasoning caused moral judgment – dominated academic thinking about moral

---

<sup>55</sup> Phineas Gage was a railroad worker who, in 1848, received a very traumatic brain injury when a tamping rod was blasted through his head. Remarkably, Gage was able to walk away from the accident, but a large part of his frontal lobe was destroyed. Although his intellectual abilities remained intact, Gage was a changed man, according to his family and friends. Always a responsible man before the accident, he became impatient, inconsiderate and quite profane in his language. Damasio, A. R. (1994). Descartes' Error: emotion, reason, and the human brain. New York, Grosset/Putnam.

<sup>56</sup> Allman, J. M., Watson, K.K., Tetreault, N.A., Hakeem, A.Y. (2005). "Intuition and autism: a possible role for Von Economo neurons." Trends in Cognitive Sciences **9**: 367–373.

Allman, J. M., A. Hakeem, et al. (2001). "Anterior cingulate cortex: the evolution of an interface between emotion and cognition." Annals of the New York Academy of Sciences **935**: 107–117.

<sup>57</sup> Allman, J. and J. Woodward (2007). "Moral intuition: Its neural substrates and normative significance." Journal of Physiology – Paris **101**: 179-202.



judgment until recently. Work by Jonathan Haidt proposed an alternative view: the social intuitionist model<sup>58,59</sup>. This model focuses on the intuitive feelings of right or wrong that emerge without conscious consideration when faced with a moral situation. Haidt stresses that both intuition and reasoning require cognition and that intuition should not be considered synonymous with emotional appraisal. Despite this distinction in the social intuitionist model, many think about intuitive processing as an external representation of emotional appraisals of the moral situation.

Philosophers, psychologists and neuroscientists investigating moral intuition often frame the idea of moral intuition in the following way: ‘When faced with a moral choice, if it feels bad, it is morally wrong’. There are a number of assumptions in this statement that may make it biologically or psychologically implausible. Moral non-realism questions whether there are any moral absolutes. Biology (especially in the field of vision research) makes us question whether conscious perceptions bear any similarity to the brain processes that lead to those perceptions. Despite many questions about the validity of this statement, it *seems* correct. It represents our conscious report about how we make some moral judgments – “they just feel wrong”<sup>60</sup>. Given the pervasiveness of this statement in considering moral cognition, many researchers have chosen to study moral intuition (and by extension moral reasoning) from this perspective.

---

<sup>58</sup> Haidt, J. (2001). "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment." *Psychological Reviews* **108**(4): 814-834.

<sup>59</sup> At nearly the same time, philosopher John Bolender published a similar model: a two-tiered cognitive architecture. Bolender’s work, however, did not receive nearly the acclaim that Haidt’s did. Bolender, J. (2001). "A Two-Tiered Cognitive Architecture For Moral Reasoning." *Biology and Philosophy* **16**: 339-356.

<sup>60</sup> It is important to note here that often people create reasons for their moral judgments that do not really fit the situation and when challenged on those beliefs revert to the above statement. The social intuitionist model considers these to be post-hoc reasoning. See Haidt, J. (2001). "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment." *Psychological Reviews* **108**(4): 814-834. p. 818.

Jonathan Baron and Walter Sinnott-Armstrong independently argue that moral intuitions can be best understood as heuristics – rules of thumb that help guide our daily behavior<sup>61,62</sup>. Baron suggests that these moral heuristics may be equivalent to the ‘naïve theories’<sup>63</sup> that people apply in other contexts. That is, moral heuristics are based on general understandings of the world around us, but are not necessarily reliable in all situations. ‘If it feels bad, it is morally wrong’ may not be correct in all situations, but it feels like it should be. Schnall, Haidt and Jordan<sup>64</sup> demonstrated that externally imposed feelings of disgust (arising from disgusting smells, sights or recollections) can influence the moral judgments that people make about unrelated moral scenarios. In this circumstance, ‘feeling bad’ was manipulated to make things ‘morally wrong’.

Sinnott-Armstrong claims that moral intuitions are an affect heuristic and that, as such, “work well in common circumstances, but also lead to systematic errors in unusual situations.”<sup>65</sup> Sinnott-Armstrong cites research that many heuristics employed by people on a daily basis (the ‘representativeness heuristic’<sup>66</sup>, the ‘recognition heuristic’<sup>67</sup>) can be unreliable. The implications of this unreliability are debated by psychologists.

---

<sup>61</sup> Baron, J. (1995). "A Psychological View of Moral Intuition." The Harvard Review of Philosophy Spring: 36-40.

<sup>62</sup> Sinnott-Armstrong, W. (2008). Framing moral intuitions. In Moral Psychology, Volume 2: The Cognitive Science of Morality. W. Sinnott-Armstrong. Cambridge: MIT, Bradford Books: 47–76.

<sup>63</sup> These ‘naïve theories’ are often demonstrated with respect to physics. Baron gives the following example: “In one study, students graduating from Harvard were asked why it was hot in the summer and cold in the winter. Many students said that the earth was closer to the sun in the summer. This is reasonable: if you move closer to a fire, you get warmer. If this were the explanation, though, it would be warmer in July than in January in the Southern Hemisphere too. People’s intuitions about science are often based on everyday experience, but this experience can mislead.” From Baron, J. (1995). "A Psychological View of Moral Intuition." The Harvard Review of Philosophy Spring: 36-40.

<sup>64</sup> Schnall, S. (2008). "Disgust as embodied moral judgment." Personality & Social Psychology Bulletin 34(8): 1096-1109.

<sup>65</sup> Seminar Materials: Sinnott-Armstrong, W. (2008). Moral Intuitions and Heuristics. William Bennett Munro Memorial Seminar Pasadena, CA, California Institute of Technology.

<sup>66</sup> Tversky, A. and D. Kahneman (1974). "Judgment under Uncertainty: Heuristics and Biases." Science 185(4157): 1124-1131.

<sup>67</sup> Goldstein, D. G. and G. Gigerenzer (2002). "Models of Ecological Rationality: The Recognition Heuristic." Psychological Review 109(1): 75-90.

Gigerenzer, for example, suggests that heuristics are an adaptive tool, allowing humans to develop with a strategy that serves us well most of the time, rather than one that fits the rules of formal logic<sup>68,69</sup>. This unreliability, however, troubles Sinnott-Armstrong when applied to moral situations. “[T]he view that moral intuitions result from heuristics raises doubts about whether and when we should trust moral intuitions. Just as non-moral heuristics lack accuracy in unusual situations, so do moral intuitions, if they are based on moral heuristics. It would be interesting and important (though challenging) to do the empirical work needed to determine which moral intuitions are reliable in which circumstances.”<sup>70</sup> While this dissertation does not address the veracity of moral intuition directly, the data presented within do provide the tools to do exactly the work Sinnott-Armstrong proposes. This is one great strength of this work: the data are available to address a wide range of research questions. In this case, real moral narratives are available to address the role of moral intuition in circumstances that participants have experienced themselves.

### **Moral Emotions and the Somatic Marker Hypothesis:**

Emotions are hypothesized to play a role in complex decision making and social behavior. The somatic marker hypothesis (SMH), postulated by Antonio Damasio<sup>71</sup>, posits that representations of emotions as bodily states stored in the orbitofrontal cortex are used during complex cognitive tasks to weight possible options. This biasing can be

---

<sup>68</sup> Gigerenzer, G. (2006). Bounded and Rational. Oxford, Blackwell.

<sup>69</sup> Michael DePaul feels similarly: “[T]hese days the majority of moral intuitionists hold that moral intuitions are only prima facie justified, that one should reflect on them and attempt to integrate them into a coherent moral system, and that they are to be rejected if this cannot be done. Even supposing moral intuitions are the products of the affect heuristic, it is not at all clear that they are so unreliable that they cannot play the role contemporary intuitionists assign to them.” DePaul, M. "Comments on Walter Sinnott-Armstrong." Retrieved September 16, 2008, from <http://www.slu.edu/colleges/AS/philos/PDF/depaul.pdf>.

<sup>70</sup> Sinnott-Armstrong, W. (2008). Framing moral intuitions. In Moral Psychology, Volume 2: The Cognitive Science of Morality. W. Sinnott-Armstrong. Cambridge: MIT, Bradford Books: 47–76.

<sup>71</sup> Damasio, A. R. (1994). Descartes' Error: Emotion, Reason, and the Human Brain. New York, Grosset/Putnam.

either unconscious or conscious, depending on the situation. Evidence for the SMH comes from research done with subjects who have damage to the ventromedial prefrontal cortex (VMPFC). In the Iowa Gambling Task, Bechara et al. found that subjects with VMPFC damage were unable to discriminate disadvantageous decks of cards, while normal subjects were unconsciously biased against these decks<sup>72</sup>.

Described narrowly, the somatic marker hypothesis suggests a specific link between emotion and cognition. Drawn more broadly, the SMH suggests a link between moral emotions and social behavior as well. Moral emotions can be distinguished from basic emotions by their social component. Embarrassment, for example, can only exist in a social context – you would never be embarrassed (a moral emotion) that you had an apple unless there was another person’s opinion involved. By comparison, you might be quite happy (a basic emotion) that you had an apple without any other input. Moral emotions are defined by this social input and can include: guilt, shame, pride, embarrassment, envy, schadenfreude<sup>73</sup>, empathy/sympathy, compassion, contempt and gratitude. In some situations, some of the basic emotions can also take on a moral character, such as disgust, awe and indignation or moral anger. There is evidence from primatology and evolutionary biology that moral emotions have evolved over time. Von Economo neurons in the frontoinsular cortex and anterior cingulate cortex of humans and great apes are thought to play a role in shame and guilt<sup>74</sup>. Work by Frans deWaal and

---

<sup>72</sup> Bechara, A., H. Damasio, et al. (2000). "Emotion, Decision Making and the Orbitofrontal Cortex." *Cerebral Cortex* **10**: 295-307.

<sup>73</sup> Schadenfreude is a German word which has no direct English translation. It is usually translated as taking pleasure in another person’s misfortune. (2006). Schadenfreude. *Dictionary.com*, Random House, Inc. Unabridged (v 1.1).

<sup>74</sup> Allman, J. and J. Woodward (2007). "Moral intuition: Its neural substrates and normative significance." *Journal of Physiology – Paris* **101**: 179-202.

others have shown that great apes demonstrate empathy, reciprocity and altruism<sup>75</sup>.

Joshua Greene suggests that the distinction between personal and impersonal moral violations is evolutionary in nature. He describes this as the ‘ME HURT YOU’ criterion. Personal violations involve an acting agent who is specifically inflicting damage (ME), a primitive type of harm (like “assault, not tax evasion”) (HURT) and a victim (YOU)<sup>76</sup>. Greene suggests that these are the kinds of violations that a chimpanzee could understand and might reflect the degree of sophistication with which human moral intuition still operates.

### **Moral Neuroscientists:**

Joshua Greene, a cognitive neuroscientist and philosopher, has made important progress in expanding the study of moral cognition using biological methods. In his seminal work, Greene chose to use fMRI and the trolley problem as a way to investigate the parts of the brain that make moral decisions<sup>77</sup>. Importantly, Greene was particularly interested in whether fMRI could be used to help elucidate some of the debate between moral reasoning and moral intuition. Greene designates these using the more philosophical terms ‘cognitive’ and ‘emotional’ judgment, which reflects the utilitarian/deontological divide that prompted this work.

In this initial study<sup>78</sup>, subjects read a variety of moral dilemmas while in the MRI and responded about the appropriateness of each moral choice. Greene theorized that

---

<sup>75</sup> de Waal, F. B. M., Ed. (2002). Tree of Origin: What Primate Behavior Can Tell Us About Human Social Evolution, Harvard University Press.

<sup>76</sup> Greene, J. D. and J. Haidt (2002). "How (and where) does moral judgment work?" Trends in Cognitive Sciences 6(12): 517-523.

<sup>77</sup> Greene, J. D., R. B. Sommerville, et al. (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." Science 293: 2105-2108.

<sup>78</sup> Greene, J. D., R. B. Sommerville, et al. (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." Science 293: 2105-2108.

moral choices that involved personal violations (like that of the ‘footbridge dilemma’) would selectively elicit brain activation in emotional processing areas, when compared to impersonal violations (like that of the trolley problem)<sup>79</sup>. As hypothesized, Greene found that judgments about personal violations selectively activated the medial frontal gyrus (Brodmann’s Areas 9 and 10), posterior cingulate gyrus (BA 31) and angular gyrus (BA 39), as compared to moral impersonal or non-moral judgments. In addition, Greene found that working memory structures which are known to be less active during emotional processing (right middle frontal gyrus and bilateral parietal lobe) were significantly less active during the personal dilemma judgments. In addition to these findings, they found that the reaction time when subjects judged personal violations as appropriate was longer than when they judged them as inappropriate.

In a subsequent paper, Greene and colleagues hypothesized that the differences in how long it took subjects to judge each moral dilemma not only depended on the subject’s final judgment but was also concomitant with the level of difficulty<sup>80</sup>. (As an example, the ‘crying baby’ problem was found to be much more difficult than the ‘infanticide’ problem<sup>81</sup>.) Dilemmas with long reaction times were presumed to create

---

<sup>79</sup> In subsequent work, Greene operationalizes the personal/impersonal distinction. To be personal, a moral violation must meet three requirements: 1. be “likely to cause serious bodily harm”; 2. “to a particular person”; 3. “in such a way that the harm does not result from the deflection of an existing threat onto a different party”. Greene, J. D. (2005). *Cognitive neuroscience and the structure of the moral mind. The Innate Mind: Structure and Contents*. S. Laurence, P. Carruthers and S. Stich. New York, Oxford University Press.. Recently, Greene has noted that he feels that the third criterion needs to be revised. Greene, J. D., S. A. Morelli, et al. (2008). "Cognitive load selectively interferes with utilitarian moral judgment." *Cognition* **107**: 1144-1154.

<sup>80</sup> Greene, J. D., L. E. Nystrom, et al. (2004). "The Neural Bases of Cognitive Conflict and Control in Moral Judgment." *Neuron* **44**: 389-400.

<sup>81</sup> The crying baby problem from the original study is slightly different from the one used in our study. The wording from the original study is as follows: “Enemy soldiers have taken over your village. They have orders to kill all remaining civilians. You and some of your townspeople have sought refuge in the cellar of a large house. Outside you hear the voices of soldiers who have come to search the house for valuables. Your baby begins to cry loudly. You cover his mouth to block the sound. If you remove your hand from his mouth his crying will summon the attention of the soldiers who will kill you, your child, and the others

more conflict between emotional processing (intuitive responses) and cognitive processing (reasoned responses). On an individual basis, subjects were tested in the MRI on putative ‘easy’ and ‘difficult’ dilemmas. When comparing the brain activity between difficult moral dilemmas and easy ones, difficult dilemmas showed increased activation bilaterally of the anterior DLPFC (BA 10/46) and inferior parietal lobes (BA 40/39), as expected for problems that required more cognitive control. This contrast also showed activation of the ACC (hypothesized to be involved in ‘conflict monitoring’) and posterior cingulate cortex (an activation that has not yet been explained). These findings suggest that some moral dilemmas do require more cognitive input (‘reasoning’) than others and there is a distinct pattern in brain activation that represents that cognitive input. To address the philosophical idea that the differences between the trolley problem and the footbridge problem are due to a struggle between emotional and cognitive processes, they compared subjects’ brain activations when they accepted a utilitarian decision as appropriate (a personal violation that helped more people) to those where they rejected the utilitarian decision (e.g., said it was inappropriate to push the fat man off the footbridge). As before, brain areas involved in cognitive control were more strongly activated during utilitarian than non-utilitarian decisions. Specifically, bilateral anterior DLPFC (BA 10) and right inferior parietal lobe (BA 40) were activated in this contrast. These findings, while significant, were not the only activations found with this

---

hiding out in the cellar. To save yourself and the others you must smother your child to death. Is it appropriate for you to smother your child in order to save yourself and the other townspeople?” The infanticide problem is as follows: “You are a fifteen-year-old girl who has become pregnant. By wearing loose clothing and deliberately putting on weight you have managed to keep your pregnancy a secret. One day, while at school, your water breaks. You run to the girls locker room and hide for several hours while you deliver the baby. You are sure that you are not prepared to care for this baby. You think to yourself that it would be such a relief to simply clean up the mess you’ve made in the locker room, wrap the baby in some towels, throw the baby in the dumpster behind the school, and act as if nothing had ever happened. Is it appropriate for you to throw your baby in the dumpster in order to move on with your life?”

comparison or in subsequent whole brain analyses, suggesting to this researcher that the utilitarian/non-utilitarian distinction may not be as clearly delineated in the brain as the previous finding suggests.

From this work, Greene began to develop a dual-process theory of moral decision making, which “associates utilitarian moral judgments (approving of harmful actions that maximize good consequences) with controlled cognitive processes and associates non-utilitarian moral judgment with automatic emotional responses.”<sup>82</sup> More simply, the footbridge dilemma is best understood as a conflict between two competing systems: a cognitive system and an emotional system. This theory suggests that a manipulation of the cognitive load would inhibit a subject’s ability to make utilitarian judgments. To test this, Greene and colleagues required subjects to respond to a standard battery of moral dilemmas while simultaneously completing a digit search task<sup>83</sup>. This digit search task significantly increased the cognitive load that subjects were under when responding to moral dilemmas. As hypothesized, increasing the cognitive difficulty of the task significantly increased the amount of time that it took subjects to make utilitarian judgments. The same effect was not found for non-utilitarian judgments<sup>84</sup>. Importantly, the cognitive load did not change the likelihood that subjects would make utilitarian judgments, only the length of time that it took them to make the decision. Greene suggests that this finding provides true evidence for post-hoc analysis of moral decisions.

---

<sup>82</sup> Greene, J. D., S. A. Morelli, et al. (2008). "Cognitive load selectively interferes with utilitarian moral judgment." *Cognition* **107**: 1144-1154.

<sup>83</sup> In this task, digits were streamed across the bottom of the screen during the entire task and subjects were instructed to hit a button each time the number 5 appeared. The subjects were told that they would be judged for accuracy and the rate of the digit stream increased partway through the task to account for subjects’ adjustment to the task.

<sup>84</sup> It is worth noting that these results were computed using only the ‘high-conflict’ dilemmas (as identified in Koenigs, et al.) rather than in all the personal moral dilemmas. Koenigs, M., L. Young, et al. (2007). "Damage to the prefrontal cortex increases utilitarian moral judgements." *Nature* **446**: 908–911.



The increase in reaction time without an impact on the subject's decision implies that subjects are making the decision and then thinking about it (longer under increased cognitive load) and then pressing the decision button. (This is in keeping with Haidt's social intuitionist model discussed above.) Greene's work establishes that dissociable processes are at work when we make moral decisions about these philosophical dilemmas. Further work, using the data in this dissertation may be able to assess the role that these processes play in every day decision making.

Marc Hauser, a psychologist and evolutionary biologist, believes that the automatic moral processing investigated by others as moral intuition arises from a deeper underlying neural process: a moral 'grammar'. In 1971, John Rawls suggested that Noam Chomsky's critical insight about linguistics could be applied to morality as well<sup>85</sup>. In the 1950s, Chomsky postulated that all language was based on an innate, universal grammar<sup>86</sup>. This hard-wired set of rules provides a flexible base from which many languages can evolve. An important element of universal grammar is that the rules of grammar are often inaccessible to the speakers of that language. (It is easy to know that a sentence is grammatically incorrect, but often difficult to explain why.) Rawls considered that moral rules might be similarly inaccessible, allowing people to know when an action was wrong but not necessarily be able to explain why. Since Rawls, only a few researchers have considered the idea of a universal moral grammar<sup>87</sup>. Hauser, however, finds that the linguistic analogy parallels strongly with his own insights from

---

<sup>85</sup> Rawls, J. (1971). *A Theory of Justice*. Cambridge, Harvard University Press.

<sup>86</sup> Chomsky, N. (1955). *The Logical Structure of Linguistic Theory*. Cambridge, MA, MIT.

<sup>87</sup> A notable exception to this is John Mikhail. Mikhail, J. (2007). "Universal moral grammar: theory, evidence and the future." *Trends in Cognitive Science* **11**(4): 144-152.

evolutionary biology and has recently written extensively on the idea of an innate moral faculty.

Hauser claims that there are two ways to interpret the linguistic analogy, a strong and a weak interpretation. The strong interpretation suggests that moral processing and language processing are quite similar processes. This includes suggesting that moral judgment has the same features that have been discovered about linguistics. The system must have dedicated machinery that is encapsulated from other systems; its principles must be unconscious and inaccessible; and the system must provide constraints on the possible moralities available. The weak interpretation suggests that there is no real parallel between the actual neural constructs of language and morality, but instead that looking at moral decision making differently may be as beneficial for its study as it was for language. Hauser says that the weak interpretation provides researchers with a new set of questions about morality; questions that address the underlying, unconscious behavior as well as the evolutionary and developmental aspects of moral judgment.

Hauser is particularly interested in the differentiation of moral decision making among cultures, as this may provide insight about the roots (i.e., 'phonemes') of a moral faculty. To investigate the variation among cultures, he developed the Moral Sense Test, an internet-based version of many moral dilemmas<sup>88</sup>. During the first year that the MST was online, data were collected from more than 30,000 subjects from 120 countries, distributed across age, gender, religion and education<sup>89</sup>. The MST is also designed to compare subjects' explicit justifications for their moral decisions with their choices on

---

<sup>88</sup> The Moral Sense Test can be found online at: Hauser, M. "The Moral Sense Test." 2008, from <http://moral.wjh.harvard.edu/index.html>.

<sup>89</sup> Hauser, M. D., F. Cushman, et al. (2007). "A dissociation between moral judgments and justifications." *Mind & Language* 22(1): 1-21. Quote from p. 18.

harm-based dilemmas. Subjects who visited the Moral Sense Test website<sup>90</sup> were asked to respond to 19 moral dilemmas involving harm to innocent people. As is typical, some of these dilemmas involved action and some omission of action. Differing from the dilemmas used by Greene, these dilemmas assign names to the subjects in each narrative (e.g., “Frank is standing on a footbridge above the railroad tracks ...”<sup>91</sup>). Subjects were asked whether it was permissible for ‘Frank’ to make the salient choice, rather than whether it was appropriate for the first-person subject to make the choice<sup>92</sup>. After responding to all of the moral dilemmas, subjects were provided with pairs of dilemmas where the subjects provided incongruent responses (for example, the trolley/footbridge dilemmas) and asked to provide short, written justifications for each of their decisions. Hauser refers to these pairs of dilemmas where subjects provide incongruous responses as manifestations of the ‘principle of double effect’<sup>93</sup>. Hauser found that the principle of double effect was demonstrated among all social, cultural and demographic groups that contained enough members to compute statistical significance, suggesting that the principle of double effect is common at a very basic level of moral decision making. In addition, a majority of subjects were not able to justify the differences in their decisions. Hauser finds that the inability to provide rational, explicit justifications supports the idea that at least some moral judgments are inaccessible to conscious reasoning.

---

<sup>90</sup> {Hauser, #95}

<sup>91</sup> Hauser, M. D., F. Cushman, et al. (2007). "A dissociation between moral judgments and justifications." *Mind & Language* 22(1): 1-21.

<sup>92</sup> While this may be an unimportant difference, I personally feel that the variations in the dilemmas used as ‘standards’ may be quite significant. The judgment one makes about another person may be quite different than the judgment one makes for herself. Is Frank more culpable for pushing someone than you would be yourself? Is the distinction between something being ‘permissible’ and ‘appropriate’ morally relevant? These are important, but unexplored questions.

<sup>93</sup> The principle (or doctrine) of double effect is a philosophical description of a case where an otherwise legitimate action will result in an additional, foreseen, negative outcome. E.g. Killing someone to prevent them from killing you involves protecting yourself (legitimate action) and killing someone (not permitted action). However, the unpermitted action is a necessary consequence of the legitimate action.

While Hauser provides a great deal of support for the strong interpretation of the linguistic analogy in his recent book, Moral Minds,<sup>94</sup> this position is not without its detractors. Philosopher Jesse Prinz argues that the linguistic analogy is simply a different route to return to ethical naturalism, with all of the pitfalls that philosophers have described before<sup>95</sup>. Prinz argues that all of the data that support the idea of innate moral processes can also be interpreted to support the idea of the specialized use of general cognitive systems for moral cognition. In addition, he finds that without deeper understanding, the presence of similar moral rules in all cultures to be insufficient to establish that those rules are innate<sup>96</sup>. Dupoux and Jacob<sup>97</sup>, on the other hand, find that there is no evidence to specifically dispute a nativist argument for moral cognition, but rather find weakness in Hauser's weak claim – that the linguistic analogy provides new and important questions for investigating moral behavior. Dupoux and Jacob find that there is not enough similarity between moral behavior and linguistic behavior to assert a moral grammar (e.g., moral competence lacks grammatical structure) and that using the linguistic analogy to drive investigations of moral cognition only confounds the research<sup>98</sup>.

While work by Greene and Hauser has progressed from theory to brain experiments, other researchers have chosen to begin with examination the brain to

---

<sup>94</sup> Hauser, M. D. (2006). Moral Minds: How Nature Designed a Universal Sense of Right and Wrong. New York, Harper Collins.

<sup>95</sup> Prinz, J. J. (2007). Is Morality Innate? Moral Psychology, vol. 1: Evolution of Morals. W. Sinnott-Armstrong. Cambridge, MA, MIT Press.

<sup>96</sup> As an example, Prinz points out that fire, clothing and marriage are universal or nearly universal cultural behaviors, but are not generally thought to be innate.

<sup>97</sup> Dupoux E, J. P. (2007). "Universal moral grammar: a critical appraisal." Trends in Cognitive Sciences **11**(9): 373-378.

<sup>98</sup> Dupoux, E. and P. Jacob (2008). "Response to Dwyer and Hauser: Sounding the retreat?" Trends in Cognitive Sciences **12**(1): 2-3.

uncover the features of moral cognition. Functional MRI has helped researchers investigate the brain networks that support moral cognition in normal subjects as well as in studying subjects with brain damage that show deficits in moral behaviors. Jorge Moll and his colleagues did a meta-analysis of fMRI and lesion data from studies of moral cognition. While the tasks varied greatly (from pictorial to written stimuli), the areas of activation were remarkably consistent across modalities and task requirements<sup>99</sup>. Orbital frontal cortex (OFC) (BA 10/11/25) is activated, particularly the medial OFC and damage to mOFC may cause deficiencies in pride, embarrassment and regret. Posterior superior temporal sulcus (STS) (BA 21/39) is thought to be a key area for social perception<sup>100</sup> and damage to this area may result in problems attributing intentionality (as in autism) which may reduce the experience of pride and embarrassment<sup>101</sup>. The anterior temporal lobes (BA20/21/38) are activated reliably and it is known that developmental damage to these areas can impair moral behavior. Other areas of activation include the insula, precuneus (BA 7/31), and anterior cingulate cortex (ACC) (BA 24, 32). ACC activation has been shown, however, to be correlated to the difficulty of the moral decision. The anterior PFC (including the frontal polar cortex, BA 9/10) is highly activated, however these findings were mediated by task somewhat. Anterior PFC is activated by moral judgments when compared to non-emotional factual judgments, but not when compared to a social-emotional condition. (In that case a more ventral region of the PFC was activated.) Moll

---

<sup>99</sup> Moll J., Z. R., de Oliveira-Souza R., Krueger F., Grafman J. (2005). "The neural basis of human moral cognition." *Nature Reviews Neuroscience* 6(10): 799–809.

<sup>100</sup> It is worth noting that the STS is thought to do more than just social perception. Heberlein, A. S. and R. Adolphs (2005). The anatomy of human social cognition. *The Cognitive Neuroscience of Social Behaviour*. A. E. Easton, Nathan. Philadelphia, Psychology Press: 157-194. "These three component processes – recognition of body form and biological motion, recognition of goal directedness, and mental state attribution – are all components of social cognition, and recent evidence indicates that distinct regions within posterior STS cortices subserved each of these three processes."

<sup>101</sup> Problems with the STS in autism are thought to be caused by difficulty with neuronal targeting rather than damage, per se.

suggests that there is emerging evidence for a “partially dissociable PFC-temporal- limbic” network that represents moral emotions like guilt and embarrassment.

In other work, Moll and colleagues have shown that moral emotions (evoked with visual stimuli) selectively activate the OFC, medial PFC and STS when compared to evoked basic emotions<sup>102</sup>. They found similar results when using a sentence verification task to show that the left medial OFC, temporal pole and STS show selective activation to moral judgments but not to unpleasant emotional (non-moral) judgments<sup>103</sup>. Recently, they have found that a fronto-mesolimbic network is activated by charitable donations, including the medial and lateral OFC. Most interestingly, they found that more anterior areas of the PFC were activated when subjects made altruistic choices instead of selfish ones<sup>104</sup>.

Moll proposes that understanding the neural architecture that underlies moral behavior presents evidence for a new model of moral cognition<sup>105</sup>. Specifically, “moral cognitive phenomena emerge from the integration of content- and context-dependent representations in cortical-limbic networks.”<sup>106</sup> This model asserts that event-feature-emotion complexes (EFECs) give rise to moral behaviors. Briefly, this framework incorporates the ability to represent large numbers of possible outcomes, evaluate them based on their emotional salience and their relevance to social features and make

---

<sup>102</sup> Moll, J., R. de Oliveira-Souza, et al. (2002). "The neural correlates of moral sensitivity: A functional magnetic resonance imaging investigation of basic and moral emotions." The Journal of Neuroscience **22**(7): 2730-2736.

<sup>103</sup> Moll, J., de Oliveira-Souza, Ricardo, Bramati, Ivanei E., Grafman, Jordan (2002). "Functional networks in emotional moral and nonmoral social judgments." NeuroImage **16**(3): 696-703.

<sup>104</sup> Moll, J. (2006). "Human fronto-mesolimbic networks guide decisions about charitable donation." Proceedings of the National Academy of Sciences of the United States of America **103**(42): 15623.

<sup>105</sup> Moll J., Z. R., de Oliveira-Souza R., Krueger F., Grafman J. (2005). "The neural basis of human moral cognition." Nature Reviews Neuroscience **6**(10): 799–809. Moll, J., de Oliveira-Souza, Ricardo, Eslinger, Paul J. (2003). "Morals and the human brain: a working model." Neuroreport **14**(3): 299-305.

<sup>106</sup> Moll J., Z. R., de Oliveira-Souza R., Krueger F., Grafman J. (2005). "The neural basis of human moral cognition." Nature Reviews Neuroscience **6**(10): 799–809. p. 804.

decisions that result in behavior that we perceive as moral. As an aside, it is interesting to consider the ramifications of this model for a naturalized moral position. Moll suggests that this neural framework arises out of neural structures that produce some proto-moral emotions in other primates (empathy, altruism). However, this model leaves open the possibility that while structures that produce these behaviors may be similar throughout the human population, that the behaviors that are produced by these EFECs could be dramatically different across the globe.

A number of researchers have shown that adult onset-damage to the PFC creates a selective deficit in actual moral behavior without impairing explicit moral knowledge<sup>107</sup>. Work by Anderson, et al. demonstrated that this dissociation did not exist in subjects whose brain damage was incurred in childhood<sup>108</sup>. These subjects had intact cognitive abilities but were unable to learn appropriate moral conventions (despite behavioral training); and showed abnormal skin conductance responses (SCRs) to the Iowa Gambling Task, suggesting an insensitivity to risky behavior. It is worth noting that this is distinctly different than the findings in psychopaths, who appear to behave abnormally from childhood, show abnormal SCRs, but have the ability to explicitly describe moral rules (but not distinguish them from conventional (non-moral rules)<sup>109</sup>.

---

<sup>107</sup> Damasio, A. R. (1994). Descartes' Error: Emotion, Reason, and the Human Brain. New York, Grosset/Putnam. Grafman, J. (1995). Similarities and distinctions among current models of prefrontal cortical functions. Structure and Functions of the Human Prefrontal Cortex. K. J. H. J. Grafman, & F. Boller. New York, New York Academy of Sciences: 337-368.

<sup>108</sup> Anderson, S. W., A. Bechara, et al. (1999). "Impairment of social and moral behavior related to early damage in human prefrontal cortex." Nature Neuroscience 2: 1032-1037.

<sup>109</sup> Psychopathy and its intersection with moral neuroscience has been extensively studied by scientists like Adrian Raine and James Blair, but a detailed discussion of that work is beyond the scope of this dissertation.

These findings prompted researchers in philosophy, psychology and neuroscience to consider their theories in an empirical light. Adina Roskies, a philosopher and neuroscientist, used data about adult-onset damage to the ventromedial prefrontal cortex (VMPFC) to challenge the philosophical theory of ‘motive internalism’, the idea that moral judgments and their motivating factors are fundamentally connected<sup>110</sup>. Shaun Nichols, an empirical philosopher, has used data about psychopaths to argue that moral rationalism is challenged by the findings that psychopathy is fundamentally an emotional disorder, not due to disordered reasoning<sup>111</sup>.

Michael Koenigs and colleagues used the standard Greene dilemmas to test the utilitarian judgments of subjects with adult-onset VMPFC damage<sup>112,113</sup>. The personal moral dilemmas were categorized as high conflict or low conflict based on the reaction times necessary for subjects to respond. (This yielded a similar set to the ones that Greene used.) As in Greene’s work, reaction time was used to measure the difficulty subjects were likely having when balancing strong emotional reactions with the more pragmatic (or utilitarian) choice to save more people. Subjects with damage to the VMPFC were much more likely to endorse a utilitarian choice (e.g., push the man off the footbridge) than the normal or brain damaged control subjects when faced with high conflict dilemmas. There was no difference from controls, however, on the low conflict dilemmas, suggesting that subjects with VMPFC damage do not have the same emotional

---

<sup>110</sup>Roskies, A. (2006). "Patients with ventromedial frontal damage have moral beliefs." Philosophical Psychology 19(5): 617-627. Cummiskey, D. (1996). Kantian Consequentialism, Oxford University Press.

<sup>111</sup>Nichols, S. (2002). "How Psychopaths Threaten Moral Rationalism, or Is it Irrational to Be Amoral?" The Monist 85: 285-304.

<sup>112</sup> Koenigs, M., L. Young, et al. (2007). "Damage to the prefrontal cortex increases utilitarian moral judgements." Nature 446: 908–911.

<sup>113</sup>Ciaramelli et al conducted a quite similar experiment with nearly identical findings. Ciaramelli, E., M. Muccioli, et al. (2007). "Selective deficit in personal moral judgment following damage to ventromedial prefrontal cortex." Social Cognitive and Affective Neuroscience 2(2): 84-92.



conflict that normal subjects do. The interpretation of these findings has been publicly debated by Greene and by Moll, et al.<sup>114</sup> Both groups feel that these data can be used to exclusively support their model of moral decision making. Koenigs, et al. have also publicly responded to philosophical arguments that their chosen dilemmas do not adequately reflect utilitarian choices<sup>115</sup>, by reanalyzing their data in terms of consequence, duty and self-interest. These analyses further supported the original finding<sup>116</sup>, showing that subjects with VMPFC damage differ only in scenarios that were designated ‘consequence vs. duty’. The debates that surround these research serve to highlight the interdisciplinary interest in moral cognition and the need for further research in this field.

A small number of neuroscientists have also begun using stimuli other than the trolley problem dilemmas to investigate moral judgment. Heekeren, et al. asked subjects to judge the correctness of sentences presented to them<sup>117</sup>. These sentences contained moral content (e.g., stealing a car) and grammatical errors. Before each trial, subjects were told to either judge the moral correctness or the grammatical correctness of the sentence. fMRI analysis showed that posterior STS and VMPFC were activated during the ethical dilemmas but not the grammatical ones. Borg, et al., investigated a range of philosophical theories (consequentialism, doctrine of doing and allowing, doctrine of

---

<sup>114</sup> Greene, J. D. (2007). "Why are VMPFC patients more utilitarian? A dual-process theory of moral judgment explains." *Trends in Cognitive Sciences* 11(8): 322-323. Moll, J., and de Oliveira-Souza, R. (2007). "Response to Greene: Moral sentiments and reason: friends or foes?" *Trends in Cognitive Sciences* 11(8): 323-324.

<sup>115</sup> Kahane, G. and N. Shackel (2008). "Do abnormal responses show utilitarian bias?" *Nature* 452(7185): E5-E5.

<sup>116</sup> Koenigs, M., L. Young, R. Adolphs, D. Tranel, F. Cushman, M. Hauser, A. Damasio Ibid. "Koenigs et al. *reply*." E5-E6.

<sup>117</sup> Heekeren, H. R., I. Wartenburger, et al. (2003). "An fMRI study of simple ethical decision-making." *Neuroreport* 14(9): 1215-1219.

double effect) using systematically varied, third-person, moral scenarios<sup>118</sup>. Subjects were asked about the moral correctness of each scenario and whether they would do the action in the scenario. fMRI showed that these various philosophical perspectives could be differentiated by the types of activation each elicited. Berthoz, et al., used first-person and third-person scenarios to examine whether intentionality and personal involvement interacted in moral judgment<sup>119</sup>. They found differential amygdala activation for situations where the subject imagined themselves to be personally involved in the transgression. Young, et al., investigated intentionality using scenarios developed by Knobe<sup>120</sup> with subjects with VMPFC damage<sup>121</sup>. These third-person scenarios have been used to show an asymmetry of assigning intention when hurting is compared to helping<sup>122</sup>. VMPFC damage does not affect the way that subjects respond to these dilemmas, suggesting that attribution of intention is not an emotional process<sup>123</sup>.

---

<sup>118</sup> Borg, J. S., C. Hynes, et al. (2006). "Consequences, Action, and Intention as Factors in Moral Judgments: An fMRI Investigation." *Journal of Cognitive Neuroscience* **18**(5): 803-817.

<sup>119</sup> Berthoz, S., J. Grezes, et al. (2006). "Affective response to one's own moral violations." *NeuroImage* **31**: 945-950.

<sup>120</sup> Knobe, J. (2003). "Intentional action and side effects in ordinary language." *Analysis* **63**: 190-193.

<sup>121</sup> Young, L., F. Cushman, et al. (2006). "Does Emotion Mediate the Effect of an Action's Moral Status on its Intentional Status? Neuropsychological Evidence." *Journal of Cognition and Culture* **6**: 291-304.

<sup>122</sup> An example of the Knobe scenarios:

*Harm vignette:*

The vice-president of a company went to the chairman of the board and said, "We are thinking of starting a new program. It will help us increase profits, and it will also **harm** the environment." The chairman of the board answered, "I don't care at all about **harming** the environment. I just want to make as much profit as I can. Let's start the new program." They started the new program. Sure enough, the environment was **harmed**. *Did the chairman intentionally **harm** the environment?*

*Help vignette:*

The vice-president of a company went to the chairman of the board and said, "We are thinking of starting a new program. It will help us increase profits, and it will also **help** the environment." The chairman of the board answered, "I don't care at all about **helping** the environment. I just want to make as much profit as I can. Let's start the new program." They started the new program. Sure enough, the environment was **helped**. *Did the chairman intentionally **help** the environment?*

<sup>123</sup> Of course, this assumes that emotional deficits cause the problems with moral judgment shown in other studies.

All of the stimuli used by these researchers have been designed to elicit results based on specific models. They vary in realism, style and perspective. One of the goals of this work is to provide a set of standardized stimuli from which many researchers can draw, independent from the models or ideas that they are interested in studying. For example, a researcher could choose a large set of stimuli with which to study intentionality from the database. These narratives could be modified in much the way Knobe did in order to create harm/help variations. The distinction lies in the way that the subjects perceive the narrative. Using real narrative provides a sense of understanding that is missing when reading a third-person narrative about a ‘chairman of the board’. The collaboration of these researchers over time will contribute even further to the characterization of these narratives. These narratives provide a uniform set of stimuli from which to draw, facilitating the study of moral cognition.

In biology, there is always an interplay between *in vivo* and *in vitro* studies. Scientists need the simplicity of experiments in a petri dish to understand what is happening in a live animal. However, the effects in an experimental system are never assumed to be the same as the effects in real life. Similarly, current moral neuroscientists feel that there is a need for both of these techniques for studying moral cognition. Moll highlights the need for realistic moral stimuli saying, “The making of moral judgements [sic] on extreme and unfamiliar situations, such as those posed by classic moral dilemmas, offers interesting ways to probe philosophical points of view, but can hardly be taken as a proxy for everyday moral reasoning”, while Hauser justifies using the trolley problem in moral neuroscience in a similar way:

“[B]y using artificial cases we can guarantee that subjects will have no familiarity with or personal attachment to the particular details of the case. ... [E]ach case can be modified in critical ways in order to isolate salient dimensions.

Consequently, the use of artificial moral dilemmas to explore our moral psychology is like the use of theoretical or statistical models with different parameters; parameters can be added or subtracted in order to determine which parameters contribute most significantly to the output.<sup>124</sup>»

This point is well-taken and there is a valid position for moral dilemmas in studying moral cognition. However, as in other biological systems, these findings do not tell us what happens in normal brains making everyday moral judgments, only in brains analyzing moral dilemmas.

The permutation of the trolley problem by ‘adding or subtracting parameters’ has taken these dilemmas far from the realm of believability. Foot’s original trolley problem involved a tram with a driver at the wheel<sup>125</sup>, a situation that is not so hard to imagine in San Francisco in the 1970s. The current version used by Greene, Hauser and others requires the subject to imagine him or herself as a bystander to a runaway trolley in a remote location with hikers on two parallel tracks and a switch nearby. These scenarios lack any element of realism. The presentation of only two possibilities and the way that the binary choice is worded, “Is it permissible?” or “Should the person switch the trolley?” further requires the subject to make a non-natural decision. Why can’t the bystander yell to the hikers on the track? How do we know that the hikers will be killed?

---

<sup>124</sup> Hauser, M. D., F. Cushman, et al. (2007). "A dissociation between moral judgments and justifications." *Mind & Language* **22**(1): 1-21.

<sup>125</sup> Foot, P. (1978). *The Problem of Abortion and the Doctrine of the Double Effect in Virtues and Vices and Other Essays in Moral Philosophy*. Berkeley and Los Angeles, University of California Press.

In Foot's original paper, she cites the possibility that the single person won't come to harm as a reason why we might find the switching action acceptable. "In real life it would hardly ever be certain that the man on the narrow track would be killed. Perhaps he might find a foothold on the side of the tunnel and cling on as the vehicle hurtled by. The driver of the tram does not then leap off and brain him with a crowbar."<sup>126</sup>

Realistic moral stimuli, like the ones included in this database, will allow researchers to investigate the nuances of moral judgments that people make in their daily lives. These insights, combined with the findings from more 'experimental systems' like the trolley problem may help to elucidate the mechanisms of moral judgment, the role of decision-making and the brain areas that subsume these processes.

### **Commentary on Moral Reasoning:**

In their paper, "How (and where) does moral judgment work?", Greene and Haidt make a strong case for moral intuition dominating moral reasoning in human behavior. They specifically cite evidence from evolutionary biology, primatology and psychology that suggests that the human ability to make moral decisions evolved from altruistic behaviors within groups, cheating and cooperative behaviors. These arguments are convincing and I believe that moral intuition likely drives much of human moral decision making and perception. However, in documenting the evidence for moral emotions as the primary player, they relegate moral reasoning to a position of 'has-been'; a post-hoc explanatory tool for why we make choices that are hidden from our conscious perception. I suspect that this downplays moral reasoning too much. I would suggest that moral

---

<sup>126</sup> Foot, P. (1978). The Problem of Abortion and the Doctrine of the Double Effect in Virtues and Vices and Other Essays in Moral Philosophy. Berkeley and Los Angeles, University of California Press. p. 23

reasoning evolved to help us with a problem that Greene and Haidt have also relegated to the back burner – the problem that Kohlberg defined as the delineation between people who achieve the conventional stage and those who arrive at post-conventional reasoning. More simply, this is the problem of perspective-taking. We encounter many situations that require us to weigh multiple perspectives when making a decision. An example:

A staunch liberal rises through the ranks of an industry and eventually is able to start a business of his own. Throughout his days as a worker, he was a union member and staunch supporter of union workers. His new business, however, is suffering from the heavy burden of union wages and may be forced to close if he does not choose to hire workers who are non-union (thus less expensive).

This business owner has several choices available to him and will make some decision about what to do. Moral intuition would focus our attention on the owner's possible choices and how he feels about them in determining his decision. Whatever his decision was, we would claim that his purported reasons for deciding were merely a post-hoc interpretation of his 'gut' feelings. And I do not dispute that this may be what the owner does when he makes his choice. However, I would argue that moral reasoning is available for doing something different. It is allowing the owner to weigh different 'gut' feelings. If he puts himself in the shoes of a business owner, he finds that he feels strongly that he must protect his company from harm. From this perspective, the union workers who want more money are greedy; bad people damaging his company. We know that he once was a union worker and may be able to put himself in the union worker's shoes. From that perspective, the union workers are not greedy; they're only trying to make enough money to feed their families. They need the money and they are good people, trying to earn a

living. (Additionally, there are probably many other perspectives – his stockholders, his family, etc.) The owner may use moral intuition to decide which of these perspectives is best, which ‘feels’ right, but I suggest that this may be an iterative process, rather than a singular one. The owner cannot merely ‘feel’ what the right decision is, he must ‘feel’ many things and then ‘feel’ the choice that he decides upon. One can argue that this is merely iterative moral intuition, but I think that as the choices become more difficult, as the perception of moral ‘rightness’ becomes more cloudy, conscious perspective-taking becomes necessary for people to parse out what exactly they ‘feel’, and that is moral ‘reasoning’.

There is some support for this idea in the current literature. Moll, et al.’s model of EFECs seems to incorporate this view of moral reasoning nicely. “Our view posits a central role for the human ability to represent and evaluate large sets of possible event outcomes, which are linked to motivational salience through cortical-limbic integration<sup>127</sup>.” The brain structures involved in Moll’s model seem to support moral reasoning acting as a mental simulation of moral intuition under various circumstances. Dupoux and Jacob suggest that “moral dilemmas are cases in which moral judgment is achieved by a process of adjudication between two (or more) conflicting intuitions, and that an agent’s explicit moral beliefs might contribute to this process<sup>128</sup>.” Experimental philosophers like Shaun Nichols feel that “judgments of whether an action is wrong, all things considered, implicate a complex set of psychological processes, including

---

<sup>127</sup> Moll J., Z. R., de Oliveira-Souza R., Krueger F., Grafman J. (2005). "The neural basis of human moral cognition." *Nature Reviews Neuroscience* 6(10): 799–809. p. 806.

<sup>128</sup> Dupoux, E. and P. Jacob (2008). "Response to Dwyer and Hauser: Sounding the retreat?" *Trends in Cognitive Sciences* 12(1): 2-3.

representations of rules, emotional responses, and assessments of costs and benefits<sup>129</sup>.”

More research may provide further data about the role of moral reasoning in moral cognition. The current trend, however, of dismissing it entirely because rationalism does not dominate moral cognition seems short-sighted, at least to this scientist.

#### Section 4: Autobiographical Memory

*“Life is that property of matter whereby it can remember. Matter which can remember is living; matter which cannot remember is dead. Life, then, is memory.”*

*-Samuel Butler, Life and Habit (1910)*

While the primary aim of my dissertation was to investigate moral events that people experience in everyday life, the method for obtaining descriptions of those moral events was, of necessity, through memory. Thus, the corpus of data represented in the database of moral events (Chapter 3) is filtered through the lens of autobiographical memory, which we review briefly in this section.

Textbooks and lectures on memory in its broadest sense typically subdivide it into two broad kinds: declarative and nondeclarative. These two kinds show different psychological properties, and are known to rely on different neural systems<sup>130</sup>. While declarative memory is relatively homogeneous and depends primarily on a single memory system, non-declarative memory is heterogeneous and lumps together forms of memory as diverse as Pavlovian fear conditioning (which depends on the amygdala),

<sup>129</sup> Nichols, S. and R. Mallon (2006). "Moral dilemmas and moral rules." *Cognition* **100**(3): 530-542.

<sup>130</sup>Squire, L. R. (1986). "Mechanisms of memory." *Science* **232**: 1612-1619.; Squire, L. R. and S. M. Zola (1996). "Structure and function of declarative and nondeclarative memory systems." *PNAS* **93**: 13515-13522.



eyeblick conditioning (which depends on the cerebellum), skill learning (which depends on the basal ganglia), and forms of motor and perceptual learning (which can depend on specific sensory cortices). Non-declarative memory is essentially defined by exclusion: all of those forms of memory that are not declarative. (Figure 4 shows the organization of memory.) Since this dissertation only concerns declarative memory, this is briefly discussed next.

Declarative memory is that form of memory that we can recollect, talk about, think about, and experience in consciousness. However, since its initial psychological definition, there have been several other views that try to get at the essence of declarative memory. One currently popular view is that declarative memory is distinguished by the large number of relationships between stored items. Thus, “declarative memory” is sometimes rephrased as “relational memory”, and has been studied in animals as well as in humans as the ability to encode relationships, which can be temporal, spatial, or semantic<sup>131</sup>. Spatial memory is one aspect of declarative memory that has been studied in great detail in rodents, for instance in their ability to remember where a hidden platform is located in a water tank in relation to the configuration of a number of spatial cues.

Declarative memory in humans is usually divided into two further categories. One is semantic memory, the other is episodic memory. Semantic memory is memory for facts, such as knowing that Canada is located north of the USA or that Paris is the capital of France. Episodic memory, the form that is relevant to this dissertation, is memory for a particular, dated, personal, episode of experience. Thus, if you were able to recollect a specific experience (say, sitting in class in grade school on one sunny morning) where

---

<sup>131</sup> Eichenbaum, H., T. Otto, et al. (1994). "Two functional components of the hippocampal memory system." *Behavioral & Brain Sciences* 17: 449-518.

you learned that Paris is the capital of France, this would count as episodic memory (technically, a form of episodic memory that is called “source memory”, since it is memory for the source of a semantic memory).

Episodic memory is the form of memory that is most prominent in our autobiography (although autobiographical memory can also include semantic components, such as knowing your name and age, for example). The psychologist Endel Tulving<sup>132</sup> first elaborated on the nature of episodic memory, and stressed one of its essential features: the recollection of episodic memories is always accompanied by the conscious re-experience of the episode, often complete with strong personal emotions. Tulving called this aspect of episodic memory “autonoetic memory” to emphasize its personal nature. Episodic memories include all specific memories that we can consciously bring to mind as particular events when we recollect them, but it is important to emphasize several criteria that have become standard in memory research for ensuring that the memory is really episodic. These are criteria I also used in order to select memories in the studies described in this dissertation.

- (a) It must be a specific event. Generic memories that abstract over several events are not episodic memory. Typically, a subject should be able to date the memory, or to locate it within a temporal context (what came before, and what came afterwards).
- (b) It must be a personal memory. Memories for things that happened to other people but were not personally witnessed are not episodic.

---

<sup>132</sup> Tulving, E. (1972). Episodic and semantic memory. Organization of Memory. New York, Academic Press: 381-403.

- (c) It must be a memory that can be recollected in vivid conscious detail. Typically, the recollection of episodic memories involves sensory experiences in the same modalities as those in which the experienced memory initially occurs. For instance, remembering the explosion of the Challenger space shuttle typically involves a visual re-experience of seeing the explosion on TV and may include auditory experiences of the voice of the newscaster, etc.
- (d) Episodic memories are usually emotionally laden, often highly so. Strongly emotional personal memories are a particular class of episodic memory that has been studied in great detail. These are the kinds of memories most relevant to this dissertation.

While there is essentially unanimous agreement that nondeclarative memory types are found in most other animals, and while there is agreement that declarative memory is found in many higher animals, there is heated debate about the phylogenetic continuity of episodic memory. Because episodic memory, at least according to Tulving's view, essentially involves a specific link to space, time, and person, and because it involves conscious experience, it has been difficult to demonstrate in animals. There is some evidence that other animals, notably corvid birds, exhibit the specificity in space, time, and person, but it is unclear how to assess the ingredient of conscious re-experience.

Declarative memory is known to depend on the medial temporal lobe, notably the hippocampus and adjacent perirhinal and parahippocampal cortices. It also involves other structures of the brain, such as the basal forebrain, the anterior thalamus, and the mamillary bodies. The best studied neuronal circuit that is an essential component of

declarative memory is long-term potentiation in the hippocampus, a mechanism that has been investigated in molecular detail and that is known to involve a particular form of Hebbian learning mechanism involving a specific subtype of the glutamate receptor. However, the neural basis of episodic memory more specifically has been much more elusive. Functional imaging studies in humans have suggested that the frontal lobes are particularly important, although all of the structures more generally involved in declarative memory, including the hippocampus, are also essential.

Human lesion studies have clearly and famously tied declarative memory to the medial temporal lobe. Following the description of Milner and Scoville's famous patient H.M. (who had bilateral medial temporal lobe lesions and a severe amnesia)<sup>133</sup>, subsequent studies have argued that damage to specific sectors of the hippocampal formation is sufficient to prevent the encoding of new declarative memories<sup>134</sup>. The kind of amnesia produced by medial temporal lobe lesions has two important features. First, it is mainly anterograde, that is, an inability to encode new memories. Second, there is a lesser impairment in the retrograde direction (the time prior to the onset of the lesion), but this is graded in time such that memories that were encoded close in time to the point of the lesion are the most damaged, whereas remote memories that were acquired long before the lesion are intact. This feature supports the hypothesis that the hippocampus is essential for encoding new declarative memories, but that over time the declarative memories become independent of the hippocampus, a process called "consolidation". There is vigorous debate about the role of the hippocampal formation in episodic

---

<sup>133</sup> Scoville, W. and B. Milner (1957). "Loss of recent memory after hippocampal lesions." Journal of Neurology, Neurosurgery and Psychiatry **20**: 11-21.

<sup>134</sup> Zola-Morgan, S., L. R. Squire, et al. (1986). "Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus." The Journal of Neuroscience **6**(10): 2950-2967.

declarative memories, with some arguments that the hippocampus may be necessary at all times for the encoding, consolidation and possibly even recollection of episodic events.

Finally, it is worth noting that emotional episodic memories, as already hinted above, are a category of memory of particular interest to psychologists and have been studied in great detail. Part of the reason is that they in fact constitute the most salient part of our autobiography. Another part is that it is known that emotional memories are encoded preferentially over less emotional memories, with the result that we remember best those events in our lives that were the most emotionally arousing. And the topic is, of course, of interest also from the perspective of traumatic memories and eyewitness testimony.

The neural systems involved in emotional memory have been studied by a number of people, perhaps most notably James McGaugh at U.C. Irvine<sup>135</sup>. Studies in both animals and humans have argued that emotional memory depends on modulation of hippocampal-dependent memory by other structures, such as the amygdala. In humans, lesions of the amygdala impair the acquisition of emotional episodic memories, and in functional imaging studies, investigators have found reliable activation of the amygdala when people encode emotional, but not neutral, events into declarative memory<sup>136</sup>.

The type of memory involved in producing the collection of moral events described in this dissertation is a form of emotional, episodic memory. It is autobiographical, personal, dated, and specific, and as far as the selection of moral memories is concerned, it is typically emotional. It would thus be expected to feature some of the attributes of

---

<sup>135</sup>McGaugh, J. L. (1990). "Significance and remembrance: the role of neuromodulatory systems." *Psychological Science* **1**: 15-25. McGaugh, J. L. (2000). "Memory-- a century of consolidation." *Science* **287**: 248-251. McGaugh, J. L. (2002). "Memory consolidation and the amygdala: a systems perspective." *TINS* **25**: 456-461. McGaugh, J. L. (2004). "The amygdala modulates the consolidation of memories of emotionally arousing experiences." *Annual Review of Neuroscience* **27**: 1-28.


<sup>136</sup> Cahill, L., R. J. Haier, et al. (1996). "Amygdala activity at encoding correlated with long-term, free recall of emotional information." *Proceedings of the National Academy of Sciences* **93**: 8016-8021.

emotional memories in general – the memories should be vivid, highly memorable, and able to re-evoked the emotions with which they were associated at encoding. All of these features no doubt contributed to the rich and detailed data that was obtained on our task. Some additional features of emotional memory contributed to the hypotheses that we tested on our data, described in Chapter 4 – for instance, we expected them to depend on the prefrontal cortex and we expected their recollection to show a particular temporal distribution that might highlight adolescence/early adulthood, when the highest density of emotional events is normally encountered in life.

*Chapter 1: Figures and Tables*

“A man's moral worth is not measured by what his religious beliefs are but rather by what emotional impulses he has received from Nature during his lifetime.”  
-- Albert Einstein

Figure 1: Timeline of Writing about Morality



427 BCE	Plato born in Athens
384 BCE	Aristotle born in Northern Greece
1751	David Hume publishes "An Enquiry Concerning the Principles of Morals"
1785	Immanuel Kant publishes "Grounding for the Metaphysics of Morals"
	Jeremy Bentham publishes "An Introduction to the Principles of Morals and Legislation"
1789	
1861	John Stuart Mill publishes "Utilitarianism"
1930	C.D. Broad coins the term 'deontology'
1932	Jean Piaget publishes "The Moral Judgment of the Child"
	Lawrence Kohlberg publishes "The Development Modes of Thinking and Choices in Years 10 to 16", his PhD work on moral development
1958	
1974	John Rawls makes the linguistic analogy in "A Theory of Justice"
	Carol Gilligan publishes "In a Different Voice: Women's Conceptions of Self and Morality"
1977	
	Philippa Foot publishes the trolley problem in "Virtues and Vices and Other Essays in Moral Philosophy"
1978	
	Jonathan Haidt publishes "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Psychology"
2001	
	Joshua Greene publishes "An fMRI Investigation of Emotional Engagement in Moral Judgment"
2001	
2003	Marc Hauser puts "The Moral Sense Test" online
Present	Publications in moral neuroscience become much more frequent



**Figure 2: Kohlberg's Stages of Moral Development**

	Approximate Age Range	Stage	Sub-Stage	View of Person	Social Perspective Level
6	Age 20+, maybe never	Postconventional	Personal moral standards	Sees how human fallibility and frailty are impacted by communication	Mutual respect as a universal principle
5			Agreed upon rights	Recognizes that contracts will allow persons to increase welfare of both	Contractual perspective
4	Age 9 to 20	Conventional	Duty & Guilty	Able to see abstract normative systems	Social systems perspective
3			Gain approval and avoid disapproval	Recognize good and bad intentions	Social relationships perspective
2	Birth to 9	Preconventional	Gain Reward	Sees that a) others have goals and preferences, b) wither conform to or deviate from norms	Intrumental egoism
1			Avoid punishment	No VOP: only self & norm are recognized	Blind egoism

This figure was compiled from two existing figures: one found in (Walker, 1989) and one found online (Huff).

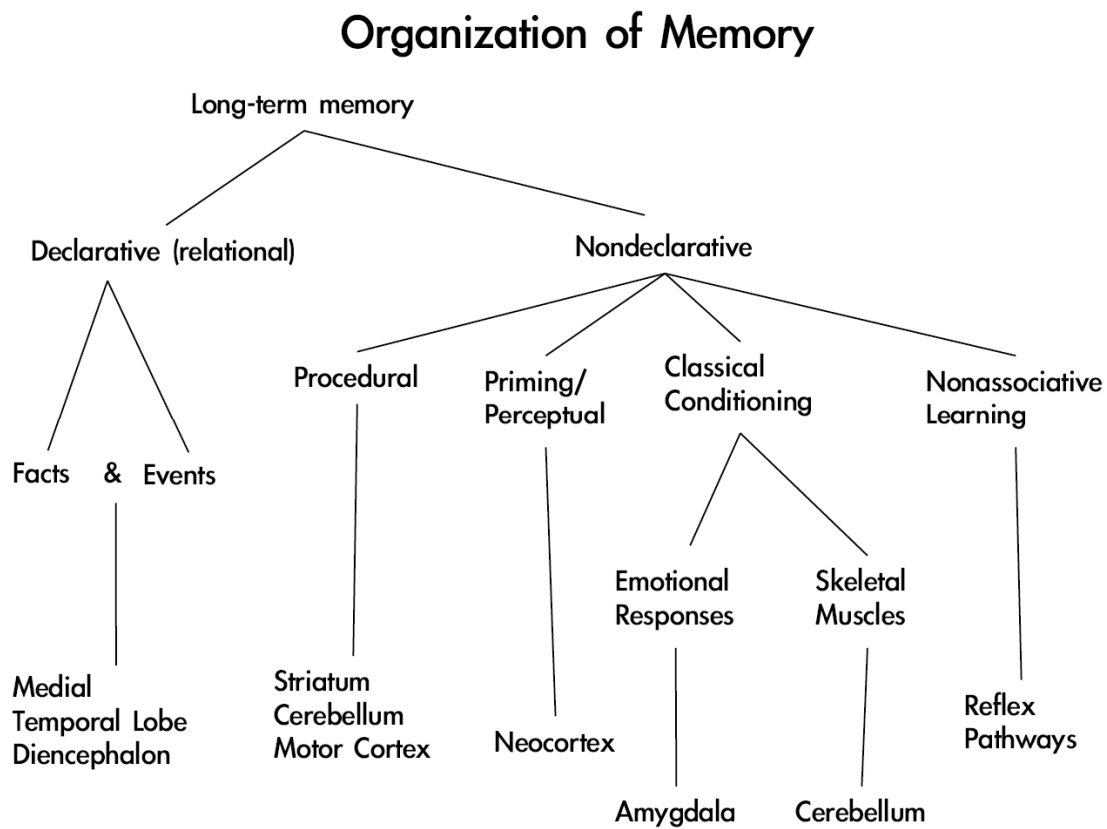
**Figure 3: Features of the Two Cognitive Processing Systems**

*General Features of the Two Systems*

The intuitive system	The reasoning system
Fast and effortless Process is unintentional and runs automatically Process is inaccessible; only results enter awareness Does not demand attentional resources Parallel distributed processing Pattern matching; thought is metaphorical, holistic Common to all mammals Context dependent Platform dependent (depends on the brain and body that houses it)	Slow and effortful Process is intentional and controllable Process is consciously accessible and viewable Demands attentional resources, which are limited Serial processing Symbol manipulation; thought is truth preserving, analytical Unique to humans over age 2 and perhaps some language-trained apes Context independent Platform independent (the process can be transported to any rule following organism or machine)

This figure is reproduced from (Haidt, 2001).

Figure 4: Structure of Memory



*Chapter 2: Methods and Materials*

This dissertation had two main aims. First, I wanted to characterize in detail representative moral events, as remembered by typical people. This meant recruiting a representative sample of participants in the first place, then probing their recollections of real-life moral events in as unbiased a manner as possible. To this end, I recruited participants with diverse gender, ethnicity and socioeconomic status and queried their memories with an evenly distributed set of cue words. These data were entered into a comprehensive database for the second main aim: to investigate their semantic structure, temporal distribution and association with demographic or psychological factors. This second aim tested several specific hypotheses about the nature of moral events that people would have encountered in real life. In addition to these two main aims, we also carried out an independent study in neurological patients in order to gain some preliminary insight into the possible role that certain brain structures (notably the ventromedial prefrontal cortex) might play in recollecting moral events. Finally, by setting up such a database containing the transcripts of real-life moral events together with their eliciting cues and the characteristics of the people who generated them, these narratives can be made available as stimuli for future experiments and the database for further data mining.

To address all these aims, we first needed to develop a database of real-life moral memories that were well-characterized in all respects. We recruited participants to participate in a cued memory recall protocol, designed to elicit morally laden memories. These were subsequently selected, transcribed, entered into a database described in

Chapter 3, and analyzed to yield results to several experimental questions described in Chapter 4.

### Section 1: Participants

Participants contacted us in response to an advertisement, were screened over the phone and in the lab, provided informed consent once they met our inclusionary/exclusionary criteria, then filled out background demographic and psychological information, before proceeding to the experimental task.

#### **Recruitment:**

Healthy adults (final n =100) were recruited at Caltech (Pasadena, CA) using an online job posting website, Craig's List<sup>137</sup>. Participants were paid \$15/hour for their time throughout the experiment and the recruitment advertisement stated that participants could make \$100 or more. (Most participants made slightly more than \$100.) This sample was self-selecting and all had access to computers, but given the wide availability of computers in local libraries and job resource centers, we do not believe that we skewed the distribution by using a computerized recruiting technique. Prospective participants were initially screened over the telephone and additionally screened in the lab. The following criteria were used to determine our experimental pool.

Inclusionary criteria: All participants were between 40 and 60 years old and were screened for coherence and fluency in their English language skills, although they were

---

<sup>137</sup> Craig's List Los Angeles can be found at: <http://www.losangeles.craigslist.com>.

not required to be native English speakers. All must have been residents of Southern California for the past 15 years or more. In addition, all participants were asked to provide their occupation and permanent address in order to be included in the study.

Exclusionary criteria: We excluded participants if they were currently under the care of a neurologist or had any history of head injury or seizures, tumor or brain surgery. Participants were also excluded if they had ever been diagnosed with any major psychiatric illness or were currently taking any medications normally prescribed for psychiatric conditions (such as antidepressants). Participants who were habitually unemployed or homeless were excluded from the study, as were those with a history of schizophrenia or major depressive or bipolar disorder.

These inclusionary and exclusionary criteria provided us with a diverse sample of participants fairly representative of the population of California. Our sample was limited in that it only included English-speaking, healthy adults.

### **Demographic Information:**

Demographic information was collected about each participant through a series of pencil and paper questionnaires. The population demographics of this group matched the California census averages on IQ (mean = 110), gender distribution and ethnicity. The mean income range was \$15,000 - \$30,000 per year and participants' ages were evenly distributed within the range (mean = 48.9). We included the expected proportion of right and left handed people and participants varied in occupation, sexual orientation, political, religious and cultural backgrounds. Overall, we believe our experimental group was a

representative sampling of the general, healthy, English-speaking population in Southern California within the given age range. Chapter 4 presents detailed demographic and psychological characterization of the participants.

**General Assessment:**

We administered a battery of background information and neuropsychological tests to gather further information about our participants. (These are described in detail later in this chapter.) This information was necessary both to characterize and examine the characteristics of the people who generated the moral memories for our database and for providing a further filter for exclusion.

Further Exclusionary Criteria: Participants were excluded from the study after the first testing session if they did not meet certain minimum standards. Any participant whose IQ was measured in the lab to be below 80 was excluded. A short, very easy measure of general memory was given to each participant (including questions about the identity of the current president and capital of the state they were born in). All participants passed this memory measure, but any participant who hadn't would have been excluded. Participants who were non-fluent English speakers or who spoke with an accent or speech impediment that made it difficult to understand them were excluded. (Since all data were spoken and needed to be transcribed, this was a necessity of the study.) Participants who displayed emotional or erratic behavior to the experimenter who was conducting the IQ test were excluded. Participants who took too long (greater than 2 hours) to complete the preliminary questionnaires were not invited to complete the rest of

the study. (Most participants took about 1 hour to complete this task.) Participants were excluded from continuing the study if they provided inappropriate, incomplete or incomprehensible memories during the cued recall protocol. 11 participants were excluded after Session 1 due to these criteria; 7 participants were excluded after Session 2. While several of these criteria were judgment calls on the part of the experimenters, all exclusionary criteria were made blind to the content of the moral memories a participant produced.

Participants answered a variety of questions about their cultural background, religious and political viewpoints. To assess cultural background, a paper and pencil questionnaire was administered. Each participant was asked to describe the cultural background in which they grew up. If the participants felt that their current culture was different than the one that they grew up in, they were asked to describe their current culture as well. Similar questions were asked about past and current religious affiliations. Participants wrote, on average, between one and three sentences for each of these questions.

## Section 2: Methods

### **Initial Recruitment and Screening:**

Following telephone recruitment and screening, participants were scheduled for a four hour testing session to complete the initial assessments and the recall protocol. All participants were tested individually, in closed testing rooms, where they interacted only



with the two primary experimenters<sup>138</sup> and the testing computer. There was no telephone or general computer access available. Throughout the testing session, participants were encouraged to take breaks as needed and were provided with snacks and drinks. Tissues were also available for emotional participants. For the researchers' convenience, the first testing session was divided into two parts, participant assessment (Session One) and cued memory recall (Session Two).

### **Session One Protocol:**

During Session One testing, the participants completed a group of assessments. After having the study explained and providing informed consent, participants began with a pencil and paper background questionnaire which included general demographic information and open-ended questions about religion and culture. Subsequent to completion of the background questionnaire, the Wechsler Abbreviated Scale of Intelligence (WASI)<sup>139</sup> was administered by one of the primary experimenters, to test the participant's IQ. The WASI is typically used for measuring IQ when the amount of time available is short. Our administration of the WASI took approximately 20 minutes and provided us with a measure of full scale IQ (FSIQ) only.

Following the WASI, participants completed the battery of neuropsychological assessments (described below), in a computerized format. They were encouraged to take breaks between each task or as needed. Instructions were provided on the computer screen before each task began. Detailed descriptions of each of these tasks can be found

---

<sup>138</sup> The two primary experimenters were Jessica Escobedo, graduate student, and Jessica Stockburger, Summer Undergraduate Research Fellow.

<sup>139</sup> Wechsler, D. (1999). Wechsler Abbreviated Scale of Intelligence. San Antonio, Psychological Corporation.

in Chapter 4, Section 1. The task order was fixed and tasks were presented in the following order:

1. NEO Personality Inventory<sup>140</sup>. The NEO Five Factor Inventory (NEO-FFI) was used to evaluate participants' personality profiles. These 5 factors describe a participant's overall approach to situations and problem solving. This measure took approximately 20 minutes to complete.

2. Schwartz Human Value Inventory (SVI)<sup>141,142</sup>. This measure generates 10 dimensions that measure cultural values and can be compared cross-culturally using the data from the European Social Survey. This task took approximately 15 minutes.

3. Political Values Questionnaire. The political values questionnaire consisted of scenarios based on current political issues (including abortion, pollution and taxes). Participants were given three viewpoints about each issue (one conservative, one intermediate, one liberal) and were asked to choose the viewpoint that best fit their own. The task included 6 questions and took approximately 5 minutes to complete.

4. Positive and Negative Affect Schedule (PANAS)<sup>143</sup>. The PANAS was used to measure a participant's general mood and affect. We intended to use this measure to exclude participants who scored exceptionally high on negative affect but did not encounter any participants who needed to be excluded on this criterion. (This is likely due to the prior screenings to exclude participants with psychiatric disorders.) The test takes approximately 10 minutes to administer.

---

<sup>140</sup> Costa, P. T. and R. R. McCrae (1992). *NEO PI-R Professional Manual*. Lutz, FL.

<sup>141</sup> Schwartz, S. H. (1992). "Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries." *Advances in Experimental Social Psychology* **25**: 1-65.

<sup>142</sup> This version of the SVI was taken from (2007). "European Social Survey." from <http://www.europeansocialsurvey.org/>.

<sup>143</sup> Watson D, C. L., Tellegen A. (1988). "Development and validation of brief measures of positive and negative affect: the PANAS scales." *Journal of Personality and Social Psychology* **54**(6): 1063-1070.

5. Ethical Position Questionnaire (EPQ)<sup>144</sup>. The EPQ was used to measure participants' degree of idealism and relativism. The newer version of the EPQ, using a 5 point scale instead of the original 9 point scale, was administered for this study. Testing took approximately 5 minutes.

6. Religious Value Questionnaire: The religious values questionnaire was adapted from a test available on the internet called the Belief-O-Matic<sup>145</sup>, which measures the degree to which a person's beliefs fit each major religion. This task uses 6 questions to evaluate participants' religious values. It took approximately 15 minutes to complete. This task data is available, but scoring proved problematic and these data are not included in the database.

7. Bem Sex Role Inventory (BEM)<sup>146</sup>. The BEM is designed to measure degree of sexual androgyny or differentiation in participants. We initially felt that gender might play a role in moral recollection and wanted to include a measure that allowed us to more finely discriminate participants. We were not satisfied with the results from this measure and chose not to include them in the database. This task took approximately 10 minutes to complete.

The computer assessments took most participants approximately one hour and a half to complete. Participants who had low IQ (<80), difficulty using the computer or difficulty completing the assessments in a reasonable amount of time were excluded at

---

<sup>144</sup> Forsyth, D. (1980). "A Taxonomy of Ethical Ideologies." Journal of Personality and Social Psychology **39**(1): 175-184.

<sup>145</sup> (2000-2008). "Belief-O-Matic - A personality quiz about your religious and spiritual beliefs." from [www.beliefnet.com/story/76/story\\_7665\\_1.html](http://www.beliefnet.com/story/76/story_7665_1.html).

<sup>146</sup> Bem, S. L. (1981). Bem Sex-Role Inventory. Redwood City, CA, Mind Garden; Consulting Psychologists Press, Inc.

this stage. The latter was generally, but not always, due to poor literacy skills. These background and neuropsychological tasks were administered for two reasons. The first was to characterize our participants as broadly as possible. In some cases, this characterization provided us with covariates for our analyses. For example, we expected that the NEO personality types might interact with moral recollection. In other cases, these data provided us with criteria for matching participants in our analyses. It was important to measure participants' IQ for this reason (as well as for screening purposes). The second reason these tasks were administered was to provide a wide range of measures for other researchers who might use these data. We considered that researchers from disciplines as diverse as anthropology and philosophy might be interested in these data as stimuli and included measures that we thought were applicable to these disciplines (e.g., cultural measures, moral dilemmas).

**Session Two Protocol:**

After completion of the Session One tasks, participants were encouraged to take a break and then the cued memory recall task was explained. To elicit autobiographical memories of morally laden events, we designed a cued memory recall protocol. (The complete task can be found in Appendix B.) Participants received both verbal and written instructions about the task and were reassured at this point that their recollections would be confidential. Participants were set up alone in a closed, comfortable room. They spoke into a microphone and their recollections were digitally recorded in two formats: to

the computer hard drive and to an external recording device<sup>147</sup>. These recordings were later transcribed by research assistants.

The computer-administered task began by asking participants to recall a memorable life experience. Participants were instructed to try to make each memory approximately 3 minutes in length and a countdown timer was provided for their reference. The timer did not require that the participants stop at the end of three minutes and continuing to the next memory was controlled by the participant. After completing their first recollection, participants were asked to recall an additional 4 memorable life experiences. These recollections served to familiarize the participants with the equipment and the timing and provided them the opportunity to request help from the researchers before beginning the experimental portion of the task.

After participants completed the memorable experiences portion of the task, they began the cued recall portion. It is important to note that at no point were participants informed that this task was specifically about *moral* memories. The computer presented a cue word in the form of a sentence (e.g., “Please talk about a time when you did something that made you feel *embarrassed*.”), and participants were asked to relate a memory that matched the cue, in the same way that they had done in the previous section. Participants continued this way through the first set of cue words, the emotion cues, and then through the second and third sets of cues, the action words and the superlative cues. (The details of the cue words used in this task can be seen in Section 3 of this chapter.) After each section, participants were encouraged to take a break. Many participants

---

<sup>147</sup> This backup system proved invaluable since some participants seemed to have an unerring ability to turn off the computer recording.

became emotional during the telling of their memories. Participants spent approximately four hours (including breaks) completing the Session Two task.

### **Section Two Exit Questionnaire:**

After completing Session Two, participants were asked to complete an exit questionnaire. This questionnaire was designed to elicit information about the amount of time pressure that participants felt during the recall protocol, whether the Session Two task affected their mood/emotions and whether participants thought that it would affect their general thoughts in the near future. Participants generally endorsed ratings on their exit questionnaire that indicated a high level of emotional involvement (mean = 7.16  $\pm$  1.89)<sup>148</sup>. In addition, participants generally did not feel very pressured by the time limit and the amount of pressure a participant reported did not change the length of their narratives ( $R^2 = 0.02$ ,  $p = 0.61$ ).

### **Data Transcription and Processing:**

The data from Session Two were manually transcribed from the original recordings by twelve different research assistants who were blind to the identity of the participants. The transcribers were instructed to replace any names and locations in the recollection. Transcribers made every effort to replace names with other names that retained the age, regional location and ethnicity of the name being replaced. Locations were similarly substituted with other similar locations (e.g., replacing large cities with other large cities in the region). After reading each participant's corpus, the memories

---

<sup>148</sup> These participants were responding to the question: "How emotional did you feel during this study?" using a Likert scale that ranged from 1 = not at all to 10 = very.

which met our criteria for ‘moral’ were selected; the remainder of this dissertation discusses only data pertaining to that subset of ‘moral’ memories. (Chapter 3 contains a detailed explanation of the criteria used for moral memory selection.) These memories were edited for grammar, coherence and readability, but vernacular and agrammatical figures of speech were maintained as much as possible, to allow the memories to retain the participant’s natural speech patterns. The average participant generated 6 moral memories from the 35 cues provided<sup>149</sup>.

### **Session Three Protocol:**

After the memories were transcribed, edited and moral memories were selected, participants were scheduled for a second four hour testing session, Session Three. Typically, 1 to 3 months elapsed between testing sessions. This length of time was built into the testing procedure to ensure that participants were not still feeling emotional from the initial recollection. During Session Three, participants were provided with transcripts of their moral memories and asked to provide ratings about their own memories.

As in Session Two, participants were provided with a confidential and comfortable environment and all instructions were given on the computer. They were given a folder with transcripts of their moral memories, but again they were not told that this research was specifically about moral behavior. They were shown that the names and locations in their transcripts had been changed, further reinforcing the confidentiality of this process. Participants were instructed to read each transcript closely and make any

---

<sup>149</sup> This is the mode; the mean number of memories generated was 7.66.

corrections necessary. Typically, corrections were due to transcriber difficulties, not to the major content of the memories.

After correcting each memory, participants were asked to answer the computer-administered questionnaire provided. (The complete testing protocol can be seen in Appendix C.) This 52 item questionnaire was intended to elicit information about several factors: general background information about the action in the memory, the emotions the participant felt about their behavior in the memory and moral judgments about their behavior in the memory. The moral judgments were subdivided; we asked participants to rate and explain their behavior at the time the memory occurred, their perceptions of their actions now as they look back on the memory, their beliefs about how others would view their actions and general questions about how ethical they believed themselves to be. Finally, we asked participants to provide any further justification of their actions that they felt were warranted. These questions were repeated for each moral memory that the participant had generated. The number of memories varied widely among participants (between 3 and 17 memories), so the length of time on this task varied as well. Generally, participants took between 2 and 4 hours to complete this task.

### **Moral Dilemmas:**

After completing the Session Three questionnaires for each memory, the participants also completed a 33 question, computer-administered version of the standard moral philosophy problems. (The complete set of these dilemmas can be seen in Appendix D). These dilemmas were provided by Greene<sup>150</sup> and have been discussed in

---

<sup>150</sup> Greene, J. D., R. B. Sommerville, et al. (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." *Science* **293**: 2105-2108.



Chapter 1. These questions are similar to (and included) the trolley problem. These questions are forced-choice, yes or no decisions, based on short paragraphs explaining each scenario. This took most participants about 20 minutes to complete.

### **Session Three Exit Questionnaire:**

After completing the moral dilemmas, participants were asked to complete a short questionnaire about their sexual preferences, hormone and steroid use and menstrual cycle. These data are not included in the database.

### Section 3: Cue Words

#### **Cue Word Selection:**

To begin selecting cue words that might elicit moral recollections, we generated 88 potential cues during a brainstorming session among several researchers. These words encompassed a wide variety of words related to morality. After examining these words, it became clear that they fell into two major categories: words that described moral feelings and words that described moral actions. Since both feelings and actions are important triggers for recollection, the decision was made to split the cued recall task into two sections. (This was later expanded to three sections with the addition of superlative cues.)

In order to better select among the potential cue words, the MRC Psycholinguistic Database<sup>151</sup> was used to provide initial examination of their linguistic properties. The

---

<sup>151</sup>These data were collected using the Wilson, M. (1987). "MRC Psycholinguistic Database: Machine Usable Dictionary. Version 2.00." from [http://www.psy.uwa.edu.au/mrcdatabase/uwa\\_mrc.htm](http://www.psy.uwa.edu.au/mrcdatabase/uwa_mrc.htm). Version 1 of the MRC Database is published as: Coltheart, M. (1981). "The MRC Psycholinguistic Database." *Quarterly Journal of Experimental Psychology* **33A**: 497-505. and as Coltheart, M. (1981a).

MRC database contains 150,837 words and provides information about 26 different linguistic properties, including the Kucera-Francis word frequencies<sup>152</sup> and measures of concreteness, familiarity, imageability and meaningfulness<sup>153</sup>. (Table 1 shows the final cue words and these measures.) Concreteness, familiarity, imageability and meaningfulness were used to eliminate words that were obvious outliers. It is important to note that the size of the databases from which each of these measures is taken varies and thus many words do not have values for all of these measurements. Words were not eliminated if measurements of them were not available; however, priority was given to words where data were available. While words that were obvious outliers were eliminated, a few words with low written frequencies (from the Kucera-Francis normative values)<sup>154</sup> were ultimately added back into the list. These words included:

---

"MRC Psycholinguistic Database User Manual: Version 1." from  
[http://www.psych.rl.ac.uk/User\\_Manual\\_v1\\_0.html](http://www.psych.rl.ac.uk/User_Manual_v1_0.html).

<sup>152</sup> Kucera-Francis word frequencies are taken from: Kucera, H. and W. N. Francis (1967). Computational Analysis of Present-Day American English. Providence, Brown University Press.

<sup>153</sup> Details for the concreteness, familiarity, imaginability and meaningfulness measures are as follows, taken from the MRC website referenced above.

Concreteness: "is derived from a merging of the Pavio, Colerado, and Gilhooly-Logie norms: details of merging are given in Appendix 2 of the MRC Psycholinguistic Database User Manual (Coltheart, 1981a). Concreteness values are integer, in the range 100 to 700 (min 158; max 670; mean 438; s.d. 120)."

Imageability: "is derived from merging the three sets of norms referred to above, and having values in the range 100 to 700 (min 129; max 669; mean 450; s.d. 108)."

Familiarity refers to "printed familiarity. The values were derived from merging three sets of familiarity norms: Pavio (unpublished), Toglia and Battig (1978) and Gilhooly and Logie (1980). The method by which these three sets of norms were merged is described in detail in Appendix 2 of the MRC Psycholinguistic Database User Manual (Coltheart, 1981). This method may not meet with everyone's approval. FAM values lie in the range 100 to 700 with the maximum entry of 657, a mean of 488 and a standard deviation of 99: note that they are integer values (in the original norms the equivalent range was 1.00 to 7.00)."

Meaningfulness ratings are taken from "the Toglia and Battig (1978) [paper], multiplied by 100 to produce a range from 100 to 700 (min 127; max 667; mean 415; s.d. 78)."

Pavio, A., Yuille, J.C. and Madigan, S.A. (1968). *Concreteness, imagery and meaningfulness values for 925 words*. Journal of Experimental Psychology Monograph Supplement, 76 (3, part 2). The MRC website notes that in addition to the data in this paper, further unpublished work was included in the MRC database.

Toglia, M.P. and Battig, W.R. (1978). *Handbook of Semantic Word Norms*. New York: Erlbaum.

Gilhooly, K.J. and Logie, R.H. (1980). *Age of acquisition, imagery, concreteness, familiarity and ambiguity measures for 1944 words*. Behaviour Research Methods and Instrumentation, 12, 395-427.

<sup>154</sup> Kucera, H. and W. N. Francis (1967). Computational Analysis of Present-Day American English. Providence, Brown University Press.

‘bittersweet’ – a word that was needed to balance the emotional valences of the emotion cues; and ‘unfaithful’ which was included in the action cues because we felt that it was an important concept for moral behaviors. The final set of cue words was also checked for outliers in word length (mean =  $7.95 \pm 2.37$ , range = 4-13 letters) and written frequency (mean =  $21.35 \pm 23.14$ , range = 1 -71).

In addition to the linguistic properties, we wanted to balance the cue words’ emotional impact. Using the Affective Norms for English Words (ANEW)<sup>155</sup>, we examined the potential cue words’ affect. ANEW measures words on three dimensions: valence, arousal and dominance. Each word is rated on a 9 point scale from most unhappy/calm/controlled by others (1) to most happy/excited/in control (9). Slightly less than half (47.8%) of the potential cue words were rated in the ANEW set, which was not enough to choose words based on their ANEW valence, but available valence ratings correlated with ratings given to the cue words by independent raters (described in Chapter 4). Available valences were also used to help organize the final emotional cue words based on how positive or negative they were. The final set contains three positive cues, four intermediate cues and three cues intended to evoke negative moral memories. The ANEW was most important in normalization of the arousal of the cue word set. It is often difficult to match the arousal felt on pleasant things with the arousal felt by unpleasant things<sup>156</sup> so it was especially important to keep the arousal of the cues as similar as possible. With only half of the cues’ arousal values available, the absolute arousal value of each set can not be calculated, but the available values all fall within one

---

<sup>155</sup> Bradley, M. M. and P. J. Lang (1999). Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings. Technical report C-1. Gainesville, FL, The Center for Research in Psychophysiology, University of Florida.

<sup>156</sup> Lewis, P. A., H. D. Critchley, et al. (2007). "Neural Correlates of Processing Valence and Arousal in Affective Words." *Cerebral Cortex* 17(3): 742–748.

standard deviation of the mean arousal score. (This can be seen in the final column of Table 1.)

### **Development of Cue Word Testing Procedure:**

The chosen cue words were separated in two categories: emotions (10 cues) and actions (10 cues). A third category was added to help participants recollect the very best and worst moral events in their lives. This ‘superlatives’ category contained 6 cues. Four control words (2 emotions, 2 actions) were added to the sets to use for comparison. Control words were common words that fit well into the category (i.e., were not obvious as controls) but that were not expected to elicit any moral memories<sup>157</sup>. Memories elicited by these cue words could be used for comparison to those that elicited moral memories, with respect to basic variables such as the average number of words contained in a memory or its recency. The complete list of cues in their presentation order can be seen in **Appendix E**.

The testing protocol was designed with 5 open recollections that were not biased by any experience with cue words, a procedure common in many studies of autobiographical memory. The question: “Please tell us about a very memorable thing that you have done” was used to help participants adjust to the testing setup, timing and equipment. While these cues were not specifically intended to elicit moral memories, occasionally (n=20), participants recalled moral memories that met our criteria during this phase. These memories are included in the database under the cue ‘memorable’.

---

<sup>157</sup> In fact, this is the case for 3 of the 4 control words as can be seen in Chapter 4, Figure 30.

The first testing category used emotional cue words to help the participant recollect experiences that felt morally laden. Each cue was presented in the form: “Please talk about a time when you did something that made you feel \_\_\_\_\_.” The emotional cues were chosen to span the valence spectrum, with three positive cues – proud, compassionate and virtuous; four intermediate cues – responsible, relieved, bittersweet and doubtful; and three negative cues – regretful, embarrassed and guilty.

The second testing category used cue words that described morally laden actions like hurting someone or being honest. These cues filled in the sentence: “Please talk about a time when you \_\_\_\_\_.” While the action cues were not as specifically valenced as the emotion cues, concepts were chosen to elicit both good and bad past moral experiences (which was corroborated by the available valences from the ANEW).

The final category, superlatives, was used to elicit the most meaningful moral experiences in the participants’ lives. For these cues, participants were asked to “Please talk about \_\_\_\_\_.” This sentence was filled in with three of the best things the participant has ever done (e.g., the thing you are most proud of having done) and three of the worst things (e.g., the thing that you would most like to change if you could go back and do it differently).

During the testing session, these cues were presented in a randomized but fixed<sup>158</sup> order within each category, as shown in Appendix 1. The cues were randomized within each category to determine the final order of presentation, but each group was kept together. Thus, the participants recalled all of their emotional, moral memories first, then their moral behaviors and finally their moral superlatives. A fixed order was chosen to

---

<sup>158</sup> All randomization was done using the random number generator at Random.org. Haahr, M. (1998-2008). "random.org." from <http://www.random.org/>.

reduce the variation caused by cuing from previous recollections. Although this increases the likelihood that participants might all tell slightly less happy memories because the cue 'happy' followed the negative cue 'regretful', it eliminated the need for corrections due to between participants cue order. Typically, one might use a randomized presentation to avoid this, but we suspected that some cues might have considerably stronger effects than others and our participant sample size was modest relative to the number of cue words used, so keeping them in the same order was a more effective way of avoiding between participant order effects. For example, participants often thought of more than one memory when the cue 'cheated' was given. These memories often percolated down through the next cues to the cue 'sneaky', where participants would tell another memory about cheating. We were particularly interested in the effects that demographics and other personality factors might have on the memories more than we were interested in the priming effects of particular cues. By presenting them in a fixed order, all participants generated memories that would have been participant to the same order effects, allowing us to disregard order effects in our analyses. However, when examining these data for the types of moral memories people generally recall, it should be borne in mind that the presentation order was fixed and likely had some (unknown) effect.

**Validity Testing:**

Our participants were provided with a private environment and assurances of anonymity before beginning the cued recall protocol. They were encouraged to be as honest as possible with their recollections. However, we felt it was still prudent to assess the validity of these memories. In some research protocols, family members or friends

can be polled to measure the accuracy of memories. In this case, preserving the anonymity of the participants prevented us from using that technique. Instead, we developed a two part process that allowed participants to be contacted by a research assistant who did not know who they were.

Two years after testing, we attempted to contact all of the original 100 participants by telephone<sup>159</sup>. 44 were still available for this validity testing. These participants were told that we were doing a follow-up to the original experiment about honesty and asked whether they told the truth during the experiment. All 44 reported that they did tell the truth, although 2 participants said that they may have elaborated their memories somewhat. Several participants expressed surprise that anyone might have not told the truth, especially in a scientific experiment. Other responses included “I had no reason to lie”, “I was exceptionally honest” and “I was gruesomely, painfully honest and telling those memories still haunts me to this day.” Not one of the 44 participants asked gave any response that would have led us to question their truthfulness in recounting the original memories.

In addition to collecting these self-reports of honesty, participants were asked if they would be willing to speak to an assistant about the memories they had produced. They were provided with a participant ID and a separate phone number to call if they were interested in participating. This second layer of testing protected the participants’ anonymity, since the research assistant only had access to the participant’s ID, not their name or phone number. 11 participants participated in this second testing session<sup>160</sup>. The

---

<sup>159</sup> It is our hope that the large time gap between the initial testing and this calling would encourage participants to give us honest answers.

<sup>160</sup> A twelfth subject called in to participate but stated that although he “was totally honest but the memories are too painful to recall.”

research assistant prompted the participant to re-tell a memory by providing a short sentence about the original memory. For example, “You told a story about returning a lost wallet. Can you tell me about that?” Each participant was asked to recall three memories and the assistant rated each recollection as a 0 (no recollection), 1 (vague recollection, not detailed or descriptive), or 2 (very descriptive, clearly recalls the situation). All 11 participants received a score of 2 on all memories. While there is an inherent bias in this second part of the validity testing (i.e., perhaps only honest participants chose to participate), we feel that these results combined with surveying nearly half of the original participants in the first part of this testing gives us confidence in the validity of the memories included in the database. The second part, in particular, also suggests that the memories were reliable in the details.



*Chapter 2: Figures and Tables*

“We have, in fact, two kinds of morality side by side; one which we preach but do not practice, and another which we practice but seldom preach.”  
-- Bertrand Russell

Table 1: Linguistic and Valence Properties of Cue Words

Cue words	Kucera-Francis Written Frequency	Concreteness	Familiarity	Imageability	Meaningfulness	Valence [Mean (SD)]	Arousal [Mean(SD)]
<b>Emotion Cues</b>							
Proud	50	327	570	434	446	8.03 (1.56)	5.56 (3.01)
Compassionate	2						
Virtuous*	6					6.22 (2.06)	4.52 (2.52)
Responsible	71		588	348			
Relieved	13						
Bittersweet	1						
Doubtful	22						
Regretful*	9	260	529	359		2.28 (1.42)	5.74 (2.32)
Embarrassed						3.03 (1.85)	5.87 (2.55)
Guilty*	29	299	559	381	421	2.63 (1.98)	6.04 (2.76)
<b>Action Cues</b>							
Honest	47		607	366		7.7 (1.43)	5.32 (1.92)
Tempted*	2	283	495	437	443		
Qualms							
Reckless	9		517	395			
Sneaky	2						
Hurtful*	37	368	579	465	530	1.9 (1.26)	5.85 (2.49)
Cheated*	3	329	549	457	446		
Lied*	59	357	548	385	447	2.79 (1.92)	5.96 (2.63)
Unfaithful	1					2.05 (1.55)	6.2 (2.7)
<b>Control Cues</b>							
Happy	98	355	621	511	568	8.21 (1.82)	6.49 (2.77)
Tired	48		630	419			
Exercised	58					7.13 (1.58)	6.84 (2.06)
Funny	41		617	468			

Table 1: It is important to note that the cue ‘took something that didn’t belong to you’ is left out of this table. The word ‘steal’ was used for comparison purposes during the initial process of choosing cue words. The superlative cues are also not included because they were concepts rather than single words. \*Cue words with this notation did not have measures of concreteness, familiarity, imageability, meaningfulness and/or valence. These data are taken from closely related words – typically from the present tense of each verb (i.e., ‘cheat’ for the cue ‘cheated’). This was done to increase the amount of comparison that could be done among cue words.

### *Chapter 3: Moral Narrative Database*

We constructed a database of 758 transcribed personal, episodic, moral memories using multiple steps. Raw transcripts of spoken memories were selected by two experimenters, edited using in several careful steps detailed below and finally given back to the participants who originally produced them to check for accuracy. Editing was done with particular attention to retaining the unique vernacular and slang of each participant. The final memory transcripts were stored, together with background information about the participants and ratings of the memories, in a searchable database. Also entered into the database were ratings of the moral memories obtained from an independent set of raters. Since one aim of this study was to construct a detailed database that could be utilized by other researchers, we included considerably more variables in this database than were analyzed within the scope of this dissertation.

#### Section 1: Rationale for a Moral Narrative Database:

Three underlying rationales drove the creation of a moral narrative database: The first was to obtain an inventory of real-life autobiographical moral memories that could be compared with more theoretical accounts of moral events and behaviors. The second was to obtain a rich characterization of the causes, intentions and justifications that accompany those moral events. This will help explore the semantic structure of moral events in everyday life. The third rationale was to allow the use of these data as stimuli about which moral judgments can be made. The construction of this database from a unique body of moral recollections generates a truly novel stimulus set for the understanding of moral events, behaviors and decision making.

The purpose for collecting true moral memories from a relatively typical and representative group of participants was to create a corpus of moral autobiographical events experienced in real life. These memories provide insight into the types of moral challenges people experience during their daily lives as well as data about how they face these challenges. This body of data contains information about both good and bad choices; ones that the participants felt were right and ones that were wrong. Using the cued recall protocol, 3,300 memories were elicited from these participants about all aspects of their lives. This chapter describes the methods by which nearly 800 moral recollections were selected from this group and carefully edited into a database of moral narratives.

A database of autobiographical moral narratives also provides a rich dataset to use in characterizing the types of moral events experienced by normal adults. Using a group of volunteer raters, the moral narrative database was sorted and rated into its many component parts. These elements of moral experiences can be used to understand the behaviors underlying normal moral decision making and the milieu in which these decisions are made.

Finally, the creation of a well-characterized database of moral narratives provides a unique set of stimuli for a myriad of research purposes. A searchable database was designed to allow researchers from many fields to use these data to explore questions of moral behavior. Future research, from cultural anthropology to functional MRI, can be facilitated by the existence of a consistent, excellent stimulus set for testing participants. When creating a database with this intent, three characteristics were required. The data must have widespread usability, the database must be highly characterized and it must be

easily searchable. These features were built into the design of the database and implemented throughout, from the selection of moral memories to the construction of a Microsoft Access database, searchable with SQL. The database has also been converted to Filemaker to allow easier access for Macintosh users.

## Section 2: Narrative Selection

During the cued recall testing session (Session Two), participants generated memories for each of 35 cued recall questions. While the cues were intended to elicit as many moral memories as possible, only a subset of the memories generated by each participant was expected to be moral. A set of criteria were developed to select which memories would become narratives in the database. These criteria were made intentionally broad after considering many different definitions of morality. By including narratives that fit many different schools of thought, the database retains the greatest possible usability for the greatest number of researchers. The criteria were applied in the order presented below to progressively eliminate memories at each level. The criteria used to select each narrative were:

### 1. Is the memory episodic and personal?

The first criterion defines features that are critical to describing a ‘narrative’. These memories are told in the first-person about a specific event or events. This criterion is necessary to eliminate memories where the participant simply spoke about the cue word in general. For example, a participant given the cue, ‘Please talk about a time when you lied’ might have discussed the act of lying itself, whether telling white lies was bad

or how they never liked lying to friends. While many of these memories might involve moral beliefs, they are not specific enough to lend themselves well to research stimuli and therefore were removed. We followed the convention in the memory literature here to choose only episodic memories – that is, memories about specific, personal events to which a date could be assigned, in principle.

## 2. Does the memory involve the speaker making a decision or choice?

The second criterion addresses the idea that a moral action requires some specific choice about the situation, whether that choice is an action or an omission of action<sup>161</sup>. Requiring that each narrative involves some sort of decision making process, selects for memories where the action was intentional. For example, a memory about a time when a participant hit someone with a baseball by accident would not meet this criterion, even though it meets the first criterion, while a memory about deciding to throw a ball at someone to exact revenge would meet this criterion. This criterion further ensured that our moral memories reflected moral judgments, choices or outcomes related to the participant in a direct way and it ties in with conceptions of moral responsibility that typically require an agent's choice or decision in order to assign blame or praise.

## 3. Does the memory involve harm or prevention of harm?

-or-

Does the memory involve feelings of right or wrong?

---

<sup>161</sup> Action versus omission of an action is often discussed in deontology. Alexander, L. and M. Moore. (Fall 2008 Edition). "Deontological Ethics." The Stanford Encyclopedia of Philosophy 2008, from <http://plato.stanford.edu/archives/fall2008/entries/ethics-deontological/>.

These criteria constitute of the core definition of the sense of “moral” used in this project. For an action, behavior or choice to count as moral, it must involve harm (or the prevention of harm) or feelings of right or wrong. While by itself this criterion would yield a definition of morality that would be very broad and overinclusive, together with the other criteria above it yields a concept of “moral” that is consistent with both intuitive and technical senses of the term<sup>162</sup>. Keeping our concept of “moral” relatively broad is an intentional strength of the database. Rather than choosing a particular definition of ‘moral’ from one of the many schools of thought, this definition was created to include as many different definitions of the idea of morality as possible. In order to accomplish this, we focused on two features. First, the feature of harm: The distinction between harm and the prevention of harm is a critical feature of the consequentialist debate in moral philosophy<sup>163</sup>. Consequentialism claims that doing harm to someone is no worse than allowing harm to come to them while deontology stands firmly against this claim. By including memories with either harmful actions or omissions that resulted in harm, both these points of view were included. Second, the feature of ‘right or wrong’: Using the concept of right or wrong as a defining feature of a moral memory was a simple way to encapsulate a number of other theories of morality. Almost any lay definition of moral generates the words ‘right’ and ‘wrong’; Merriam-Webster’s defines moral as “of or relating to principles of right and wrong in behavior.”<sup>164</sup> Virtue ethics (a branch of moral

---

<sup>162</sup> Borg, J. S., C. Hynes, et al. (2006). "Consequences, Action, and Intention as Factors in Moral Judgments: An fMRI Investigation." *Journal of Cognitive Neuroscience* **18**(5): 803-817.

<sup>163</sup> Hooker, B. (Fall 2008 Edition). "Rule Consequentialism " *The Stanford Encyclopedia of Philosophy*, from <http://plato.stanford.edu/archives/fall2008/entries/consequentialism-rule/>.

<sup>164</sup> Definition taken from the Dictionary, M.-W. O. "moral." from <http://www.merriam-webster.com/dictionary/moral..> This is the first of three definitions given by this source for the word ‘moral’.

philosophy discussed in Chapter 1) and a number of moral psychology theories can also be included under this umbrella theme.

4. Are the feelings of right or wrong in this memory based on religious interpretations of morality? If so, discard unless the participant specifically discussed their personal beliefs.

Religion and morality are often intertwined in Western belief systems but within this database they were defined as separate ideations. The final criterion in selecting memories was designed to assess whether a religious bias had influenced the inclusion of this memory. Ideas of what is moral vary among religions and since the participants providing these memories were a normally distributed sample, their religious beliefs varied as well. (See Chapter 4, Section 1 for details about participants' religious beliefs.) In addition, the database was developed to be used for a wide range of scientific pursuits and therefore it seemed prudent to eliminate narratives which depended on their moral status entirely by invoking particular beliefs from specific religions. However, moral behavior can certainly be influenced by religious beliefs and some participants were specific about their choices being both moral and religious. These memories were included in the database. For example, most memories where participants had abortions were excluded from the database because the decision of whether this was a case of 'right or wrong' was based in religious beliefs. However, in circumstances where the participant explicitly discussed basing her personal decision about having an abortion on her religious beliefs, the memory was included. This criterion was applied stringently to only include memories where the participant was detailed and specific about his or her beliefs, so that it was clear that an action was 'right' or 'wrong' from the personal



perspective of the participant, rather than only from the inferred perspective of any particular religion.

### Section 3: Methods

To select the memories on the basis of the criteria enumerated above, two independent readers were used. The readers received transcriptions of each participant's cued recall testing session (Session Two). These transcriptions typically included 30 – 35 distinct memories. (On average, participants skipped 1.7 cues during the cued recall protocol.) The first reader read the entire transcription and chose all the memories that she felt met all the criteria. This reader applied a liberal interpretation of these criteria, to select all memories that might be appropriate for the database. The second reader only read the memories selected by the first reader. We found that this allowed her to be more objective than the first reader about the moral content of the memories. The second reader applied a more conservative standard in selecting the final memories. In every case, the memory needed to meet all the criteria listed above, except in cases where more than one episode had been related during one memory. These were moved forward to the editing process where they were split up into narratives that contained only a single episode. The first and second readers were always the same and always read the transcriptions in the same order to impose as much similarity on the selection process as possible. The first reader selected 788 memories to pass along to the second reader, the second selected 758 final memories for the database. Both readers were blind with respect to information about the participants' background information (e.g., gender,

ethnicity), although some of this information could be gleaned from the text of the narratives themselves.

While most of the criteria were straightforward, the readers received further specific instructions to ensure consistency with our aims. The first reader was instructed to be cognizant of cases where the participant talked about the same episode during different cued memories and include all of these possible narratives in the group provided to the second reader. In some cases, these memories were joined during the editing process to provide the most coherent narrative. When applying the question, “Does the memory involve feelings of right or wrong?” to each memory, the reader was instructed to interpret this from two perspectives, her own and that of the narrator of the memory. The reason for this was to include memories where others (in this case the reader) felt the participant had clearly done something wrong but the participant did not state that they felt that way or vice versa. When evaluating memories in terms of the fourth criterion (the underlying religious bias), readers were asked to deeply consider their own religious beliefs and frequently check that they were not applying them to the memory selection process. While it is impossible to say that no religious bias was used when selecting the memories, the readers made every effort to avoid this. It is also important to note that we analyzed the demographic and neuropsychological profiles of the participants after selection to ensure that there was no systematic bias in the selection process.

### **Creating Narratives out of Memories:**

After selecting memories that met all the criteria, the memories were edited into narratives by the second reader. This process began with splitting memories that

contained multiple episodes and joining memories that were told during multiple cues. Splitting memories occurred when participants told more than one story during the 3 minute recording session for each cue. Generally, these were simply recollections of two separate incidents where the participant felt the cue pertained. For example, given the cue “guilty”, Subject 1313 told three stories, one about her husband, one about an accident and one about unintentional stealing. In this case, two of the three memories met our criteria for inclusion in the database. These two memories were separated and given separate story numbers and ratings. Both memories were noted as having the same eliciting cue word. The same was done for all other similar cases (n=3)<sup>165</sup>.

In some cases, the final edited version of a memory came from two cue words. This resulted from situations where a participant talked about the same situation during two separate cue words. Typically, participants would say something like, “Now, I was telling you before, when I ...” and then proceed to give more detail about the situation. In most circumstances, only one version of this memory was used. In some cases, however, the participant gave different, pertinent details in each retelling. In order to have the most accurate and complete memories in the database, these memories were combined. As such, these memories have two (n=13) or three (n=3) eliciting cue words. This adds a small degree of complication to some of the data entry and analysis. For the purposes of the database, only one eliciting cue word was entered. Since choosing which cue word contributed more to the final memory is quite subjective, the cue word entered was chosen by random number generation<sup>166</sup>. However, the additional data was retained for

---

<sup>165</sup> There were 4 total cases of this. Subject 1094 on two different cues, Subject 1313 and Subject 1317.

<sup>166</sup> Random number generation was done with the generator at Random.org. Haahr, M. (1998-2008). "random.org." from <http://www.random.org/>.

use in future analyses, as needed. After joining or splitting, these memories were added to the collection for editing.

Each memory was edited for length, readability and coherence. Editing for length included eliminating memories that were too short and too long. Length was a particularly subjective measure, but the second reader made every effort to include all memories that told a specific story, no matter how short or long, in order to prevent introducing a length bias into the database. There are narratives as short as 36 words and as long as 837 that were included. Only short memories that did not tell a detailed enough story to become a useable stimulus in an experiment were purged<sup>167</sup>. Similarly, memories that were quite long were included as intact as possible. However, in some of these cases (approximately 10 – 15 memories), the memory was embedded in unnecessary reminiscences and these were removed while retaining the moral memory within. Editing for readability was done as sparsely as possible to retain the participant's personal vernacular and word choice. This was very significant in retaining the uniqueness of each narrative. The individual voice in each narrative makes them quite different from the stimuli used by other researchers to study moral behavior. Typically, editing for readability included adding or correcting punctuation<sup>168</sup>, reorganizing sentences into clearer paragraphs and correcting pronouns and tense changes that disrupted comprehension of the narrative. Tangents made by the participant during the

---

<sup>167</sup> Here is an example of a very short narrative that was included in the database: Memory 5235. "I remember when public transportation buses started charging a dollar. My friends and I would tear the dollar into 4 pieces and fold it up so it looked like a dollar. I would use a dollar for four rides."

<sup>168</sup> The transcribers did their best to represent the participants' pauses with punctuation. These often were revised during the editing process since the pauses of the original narrator didn't contribute to the flow of the narrative.

retelling were rearranged into more logical flow or eliminated when unrelated to the moral episode.

Finally, the rough draft of each narrative was reread and compared against the criteria that shaped the database. During this process, the narratives were selectively edited for content. In some cases, the entire memory was included in the narrative, even though only part of the memory was specific to a particular episode. In those cases, the rest was typically included because the participant was elaborating on their beliefs that shaped the actions or choices in the narrative. In other cases, extraneous recollections were removed from the narratives to pare them down to be better stimuli. Throughout the editing process as little was changed from the original memory as possible in order to retain the true character of each participant's recollection in the final narrative.

In addition to this editing, each participant read their own final narratives during Session Three (as described in Chapter 2). They were provided with red pens and allowed to edit anything in the narrative that was incorrect. These corrections were made to the final narratives. Most of these corrections were trivial, but a few participants found significant transcription mistakes (usually involving the wrong pronoun) that would otherwise have gone uncorrected. Overall, the editing techniques used were as consistent, unobtrusive and accurate as possible which produced concise and unique narratives for the database.

#### Section 4: Descriptive Database Statistics

The database provides rapid access to almost any information one would like to know about the contents. A few statistics that describe the overall scope of the database

are listed here. The database contains 758 moral narratives, 382 told by women and 376 told by men. On average, each participant told 7.66 moral memories; the maximum number of moral memories told by a participant was 17 and the minimum was 3. (Figure 5 shows this distribution.) The mean word count was 217.6 words per narrative; the range was from 36 – 837 words. The word count distribution for the entire database can be seen in Figure 6. The database contains narratives that were elicited by 28 of the 31 unique cues. The three cues that did not elicit any moral narratives were control words that were not intended to elicit moral memories. One of our four control cues, *happy*, elicited 7 of the narratives included in the database, suggesting that it was not truly a control.

The Language and Inquiry Word Count 2007 analysis program was used to evaluate the content of the narratives, as well as providing a measure of how normal the narratives are compared to other sample narratives. These data are described in Table 2. The narratives in the database are compared to three LIWC control samples; the first from participants who wrote on “deeply emotional topics”, the second from participants who wrote on “relatively trivial topics, such as plans for the day” and the third from people who were “tape-recorded while engaging in conversations with others.”<sup>169</sup> While none of these categories individually capture the contents of the narratives, a similar structure can be found in the intersection of the three. As shown in the table, the narratives are quite comparable to the control data as well as the control memories collected from our own participants.

---

<sup>169</sup> Pennebaker, J. W., Chung, Cindy K., Ireland, Molly, Gonzales, Amy, & Booth, Roger J. (2007). The development and psychometric properties of LIWC2007. Austin, TX, LIWC.net.

## Section 5: Characterization of Moral Events

The compilation of these moral narratives into a searchable database was a natural extension of this work. The narratives could be sorted by cue word, word count and participant characteristics, but although the narratives were rich with detail, there was no data available to use to sort the narratives by content. In order to create a truly searchable database, the narratives needed to be extensively characterized. In addition to facilitating the creation of a searchable database, characterization of the narratives provided an opportunity to examine the structure of the narratives more closely. Detailed characterization allowed investigation of the types of moral judgment included in the database and the temporal dynamics of autobiographical moral recall.

A rating system was devised to collect detailed characterizations of each narrative. Each narrative would be read and characterized (rated) by a number of volunteer ‘raters’. A rating tool was created to ensure uniform characterization of each narrative. The completed rating tool can be found in Appendix G, in the form of the worksheet these raters completed. Raters received a set of instructions with the rating sheet to explain the details of each category. These instructions are detailed later in the text.

### **Part 1: Development of the Rating Tool**

A subsample of 100 narratives was chosen randomly from the database<sup>170</sup> to use during development of the rating tool. The development of the rating tool was done by

---

<sup>170</sup> All randomizations done in this work used the random number generator found at Random.org. Haahr, M. (1998-2008). "random.org." from <http://www.random.org/>.

two researchers<sup>171</sup> in tandem to ensure completeness. After reading the subsample, a standard journalistic style of “Who, What When, Where, Why and How?” emerged as a clear way to categorize the narratives. These categories provided broad coverage of the elements of each narrative and organized them in a manner with which raters would be familiar. Creating a tool that raters would be able to easily use was critical to this process because the ease of use of the tool would directly affect the quality of the ratings acquired. The six major categories (below) used in the final rating tool evolved from this initial journalistic style inquiry. Throughout the discussion of each category below, the term ‘narrator’ is used to refer to the participant who told the original memory narrative.

To develop the items that comprise each category of the rating tool, the following technique was used throughout. Simultaneously, each of the two readers carefully read a narrative from the subsample. After reading, each reader generated as many ‘actions in the narrative’ as possible. These ‘actions’ included actual actions, reasons given by the narrator for their behavior and feelings the narrator explicitly talked about. (You may see here how these evolved into the categories in the rating tool.) The readers also listed all the locations in each narrative. Upon completing this process with all of the narratives in the subsample, the items were refined and categorized appropriately. The final rating tool was edited into a worksheet for the raters to mark their answers.

This worksheet, referred to as ‘the rating sheet’ throughout this discussion, was comprised of two parts: the rating tool (Part One) and an additional group of rating questions intended to address research hypotheses (Part Two). A pilot set of ratings was collected on the rating sheet using narratives that were not in the 100 story subsample, using lab members as the raters. These pilot data were used to further refine a number of

---

<sup>171</sup> Myself and J. Stockburger.



the items in Part One, primarily to clarify the wording. Several new items that were generated by multiple pilot raters were also added to the list. Room for write-in choices was also added to the three large categories.

### **Designing the Rating Tool (Part One of the Rating Sheet):**

The Who? category was developed to characterize the reader's perception of who the narrator was. Specifically, it was used to ascertain the perceived gender of the narrator. It is critical to understand that this 'gender' as found in the database table 'Rating Data' is the rater's perception, not the actual gender of the participant who told the story. This rating was augmented by a question about how certain the rater was of their decision about gender of the narrator. This question is represented on the rating sheet with the choice 'Guess?' These two questions were designed specifically for two purposes: to identify narratives where the gender was unknown and to ascertain how accurate the raters were in guessing the narrator's gender. In both cases, these findings were intended to facilitate research about gender and moral judgment being conducted using the database narratives. (The findings from some of this research can be found in Chapter 4, Section 4.) To ascertain characteristics about the narrator besides gender that might be relevant to the narrative, the tool includes a question about the age of the narrator. We hypothesized that within the same category of moral behavior (i.e., taking something that doesn't belong to you), the age of the narrator might play a critical role in generating a moral judgment about the actions in the narrative. Given this hypothesis, it was important to include a characterization of the age of the narrator. Similarly, characteristics like race and religion might be of particular importance in some narratives.

A write-in space was included to allow raters to note race and religion, if relevant, and to write in other group affiliations that were particularly important to the content of the narrative.

The Where? category was developed to assess the types of milieus in which the narratives took place. The purpose for this category grew out of the intuition that the moral rules people apply to their decision making may change depending on specific circumstances. For example, a business person might obscure the truth (or outright lie) frequently at work but would never lie at home or to family. To identify these specific locations and situations, the readers listed every relevant scene of the action in each narrative. These scenes included non-location context like family and animals because there was a general feeling from the readers that these were also important to the narratives in the subsample. During the course of data collection a number of other scenes were written in. This was expected; while the subsample could be expected to cover almost all of the actions in the database, it could not possibly contain all of the locations encountered by the participants in their lifetimes.

The three large categories in the rating sheet cover three major features of the narratives: What did the narrator do in the narrative?; Why did the narrator say they did the actions in the narrative?; How did the narrator feel about their actions? The purpose of the What, Why and How? categories was twofold. First, detailing the actions and behaviors in the database provides insight into the range of moral behaviors in a normal sample of participants. Developing the items in these categories using the readers' compilation technique described above initially produced a list of the activities and feelings of those participants. The data collected using this category, however, were

necessary to understand the relative frequencies of these activities. Secondly, the ratings in these categories are integral to the development of the database as a tool to sort the narratives. The ability to select narratives based specifically on what actions were (or were not) performed in the story is much more powerful than selecting narratives based on the relatively general cue words that elicited them.

To further define the items in each category, a scale from 0 to 3 was assigned to each item. An item was to be given a score of three when it was an excellent example of the action described in the item. A score of zero was given when the action was not included in the narrative at all. This allows the narratives to be further subdivided based on the specific content of the item. It is important to note that acts as well as omissions are considered to be part of any given item. A time when the narrator chooses not to hit someone is as much a part of the item ‘physical harm to someone’ as a story where someone chooses to start a fight. Within each category the items were grouped based on the readers’ general sense of how similar the items were. The decision to group the items rather than randomize them within each category was made to increase the reliability the ratings. Given the number of items in each category, organizing them (in any fashion) seemed likely to facilitate the accuracy of the raters. Analysis of inter-rater reliability can be found later in this chapter.

The category Why Did They Say They Did It? was specifically designed with considerations from moral philosophy in mind. The items in the category were gleaned in the same way as the other categories, from the reading of the subsample. These items were then evaluated with philosophical discussions about moral decision making in mind. As discussed in Chapter 1, consequentialism and deontology are two competing schools

of philosophical thought about moral behavior. Two of the items listed by the readers dovetailed nicely with these two schools of thought and the philosophical positions were noted with the item. Thus, ‘Doing lesser harm to prevent greater harm’ had the label ‘consequentialist’ added to it; and the ‘Means do NOT justify the ends’ had the label ‘deontological’ added to it<sup>172</sup>.

The question “Did the narrator rationalize his or her behavior?” was added to the rating tool after consideration of the questions answered by the participants in Session Three. In Session Three the participants were asked to read their own memories and answer questions about them. (For more details, see Chapter 2.) In particular, the participants were allowed to provide a justification for their actions. After reading the subsample, it became clear that many participants provided justifications or rationalizations for their actions during the narrative itself. This category was developed to identify narratives with strong rationalizations already included, so that they could be compared to narratives with rationalizations added after the initial recall<sup>173</sup>. The identification of rationalizations is also a powerful sorting choice to add to the database for other researchers.

### **Part Two of the Rating Sheet:**

The collection of rating data for sorting and analysis purposes (as described above) provided an opportunity to address other research questions that arose during the collection of the moral narratives with a greater number of independent raters than would

---

<sup>172</sup> It is worth noting that some raters did not like the fact that this item was reversed from the way it is often given in common parlance: ‘the ends do not justify the means’. We chose to keep the item as the readers had originally listed it.

<sup>173</sup> As it turns out, this will need to be a future direction for the research.

otherwise be available. Specifically, Part Two of the rating sheet was designed to: examine the accuracy of the cue words as descriptors of the narratives; collect pilot data about the moral correctness of the actions in each narrative; and to investigate Richard Lazarus' theory of core relational themes.

The 'Cue Word' section of the rating sheet was used to address two issues with the moral narratives. The first is the accuracy with which the cue word can be applied to the narrative. During narrative selection the readers noticed that there were circumstances where, when responding to an eliciting cue in the cued recall protocol, participants gave memories that were not actually related to the cue word. This presents a problem for using the eliciting cue word as a category for sorting the database. Quantifying these occurrences using the rating data provides better data for the database as well as generating information that may help to refine the cued recall protocol. While accuracy was a small issue that was addressed when designing the rating sheet, a much larger one was the co-occurrence of multiple cue concepts within a single narrative. For example, a narrative elicited by the cue 'unfaithful' might also include feelings of guilt and regret<sup>174</sup>. While it is important to know which cue elicited the narrative, it seems equally useful to know which other cue word elements exist in the text. To collect these data, the rating sheet was designed with a list of all the cue words<sup>175</sup> and the instructions to circle all the items that applied.

---

<sup>174</sup> Also embarrassment, relief, bittersweetness, lying, having qualms, etc.

<sup>175</sup> The cue 'a very memorable thing that you have done' was left out of this list because it was confusing in pilot studies.

The 'Judgment' question was designed to collect preliminary<sup>176</sup> data about the moral valence of each narrative, using a Likert scale from -5 to +5. The usefulness for a moral valence rating in the database is illustrated by other research tools like the International Affective Picture System, which allows researchers to choose visual stimuli based on widely accepted normative ratings of their valence. Normative ratings of moral valence will allow researchers to consistently select narratives based on a standard rating rather than on personal belief or individual pilot data. Data about how our raters rated these narratives would facilitate these research questions and therefore was added to the rating sheet. This question also provides a way of sorting the data using these approximate ratings of valence for other research questions using the narratives. Given the relatively small number of raters, we did not make any effort to examine how normal these ratings are, however they may help researchers select narratives on which they can generate their own normative ratings<sup>177</sup>. In addition, we are developing a website that will allow us to collect data from the general population through the internet in the future. This methodology has been used by Marc Hauser<sup>178</sup> to generate thousands of ratings for moral dilemmas and we hope will one day produce normative ratings for this database.

The final section of the rating sheet was designed based on Richard Lazarus' theory of emotions as adaptive tools. Lazarus claims that emotions are generated out of an appraisal process and that each emotion can be defined by the action that provokes it. Lazarus defines these actions quite specifically (as detailed on the rating sheet). Given that the rating tool was intended to elicit both the actions in each narrative as well as the

---

<sup>176</sup> It is important to note at the outset that the raters used to collect these data were not a statistically normal sample. This is discussed in detail below (in the Raters section) but it bears brief mention here as well.

<sup>177</sup> It is our hope that the database project will remain active and collect these data as they are produced.

<sup>178</sup> Hauser, M. "The Moral Sense Test." 2008, from <http://moral.wjh.harvard.edu/index.html>.

feelings experienced while engaged in those actions, the narratives provided a rich opportunity to explore Lazarus' appraisal structure in relation to real-life moral events. The wide range of actions contained in the database, the inclusion of emotions in every narrative and the normal distribution of the narrators who generated these memories made this a uniquely appropriate dataset with which to examine Lazarus' claims. Furthermore, psychologists working in appraisal theory might wish to select stimuli (narratives) based on particular appraisal categories as described by Lazarus.

### **Similarity Ratings:**

The rating tool was designed to sort the narratives into categories generated by researchers reading the 100 story subsample. These categories were based on the knowledge, intuition and experience of those researchers. While this categorization system is an excellent way of classifying the narratives, it is inherently biased by the researchers' perceptions. We were curious about other possible underlying categories, those that might be inherent in the data but not accessible (or at least not obvious) to a researcher categorizing the narratives. A similarity rating scale was designed to generate data that would allow us to find other categories in the data. Its purpose was merely to generate a rating of how similar any two given narratives were judged to be, using a Likert scale from -2 to +2, without providing the raters with any specific dimension or category with respect to which to make this similarity judgment. With enough unique pairwise comparisons of narratives, these data could allow us to create a multidimensional space that contained the complete contents of the database, and to examine this space for particular category structures (using tools such as principal

components analysis, clustering, or multidimensional scaling). In order to facilitate the rating process, these data were collected on a separate worksheet, not on the rating sheet.

## **Part 2: Raters**

Raters were initially recruited from a group of local colleagues who were interested in the database narratives, but it became clear that the sheer number of ratings needed was going to require more volunteers to rate the entire database. Alumni from the California Institute of Technology were approached through the alumni email list as a volunteer pool. This pool was ideal for a rater group because, as a whole, they were interested in the science being done at their alma mater, they were all college graduates, and they were easily accessible by email. Raters, both alumni and colleagues, were approached with the understanding that all of their ratings would be anonymous. No demographic information was collected about any raters. The only piece of information that is consistent about all raters is that they have (or are in the process of achieving) at least a bachelor's degree. This was important in the choice of raters for two reasons: 1. Raters all share a minimum level of reading comprehension. Since participants in the database have a normal distribution of IQs, the narratives in the database vary in reading level. Given their educational background, all raters needed to be capable of easily reading the narratives with the highest level of vocabulary. 2. Raters were able to follow the long set of directions that accompanied their rating packet. Although the rating process was relatively simple, the instructions that went along with the rating were detailed, in order to ensure consistent rating. Setting a minimum level of education for the



raters made certain that all raters would be able to complete the rating process accurately. (For the complete set of rating instructions and the rating sheet, see **Appendix H** and **Appendix G**.)

While, for the purpose of obtaining consistent if preliminary ratings, we had chosen a relatively homogeneous and educated group of raters, the raters nonetheless varied somewhat in age, gender, and education. Although there is no demographic information, it can also be assumed that they varied to some extent in ethnicity and religion. In total, 55 raters<sup>179</sup> were used to generate the 2169 ratings collected. It is important to reiterate, however, that there is no way of knowing how ‘normal’ a group the raters are. The rating data was intended for use as a categorization tool, not to generate normative data of any kind. The data about “right or wrong” generated by the raters should never be used as normative or standardized ratings. In any study where a subset of the database is used, researchers should be sure to generate their own normal ratings. It is hoped that over time, as various researchers generate normal rating data for particular narratives, the database will begin to have a pool of normal ratings which can be used for those purposes.

### **Rater Tasks:**

After being recruited, raters received an email explaining the rating project. This email informed raters that the content of the narratives would range from very positive to very negative. Additionally, it assured raters that they were under no obligation to rate any narratives that made them uncomfortable and they could stop rating at any point in the process.

---

<sup>179</sup> 74 raters were recruited. Of those, 55 completed at least one rating.

Raters were given the option of doing their ratings on paper and mailing them back to the laboratory or completing a spreadsheet with the ratings and emailing them back. This system allowed us to recruit the greatest number of raters since volunteers had different preferences about the style of rating that would be most comfortable for them. (Several raters lived in other countries and felt more comfortable emailing their ratings, for example.) Each rater began with 30 narratives to rate, however the number of ratings completed by each rater varied from 2 to 194 ratings.

Each rater received a link to an individual website after initial recruitment. The website included a downloadable packet of 30 narratives, a set of instructions about how to rate each narrative and the rating sheet and similarity worksheet. Each website also included contact information. Raters were encouraged to email questions about the rating instructions before beginning and during the rating process as they came up. Approximately 10% of raters sent in questions; these were clarified by the experimenters in all cases.

### **The Rating Instructions:**

The rating sheet contained 9 distinct categories, as described above. These were: Who?, Where?, What Did They Do?, Why Did They Say They Did It?, How Did They Feel?, Did They Rationalize?, Cue Word Choice, Judgment and Core Relational Themes. (The “they” in all of these categories refers to the narrator of the narrative being read.) For each category, the rating instructions were quite detailed, in order to obtain the greatest reliability from the ratings. In the ‘Who?’ category, raters were first asked to guess the gender of the narrator of the story. This was the only rating category where

raters were supposed to guess at anything. In a previous study, we had hypothesized that many of the narratives could be gender-interchangeable. This guess was necessary to substantiate that hypothesis. (Results of this and other gender related research can be found in Chapter 4, Section 4.) Raters were also asked to report whether they were certain about their choice of gender or uncertain. Second in the ‘Who?’ category, raters were asked to designate the age of the narrator, only if it was important to the story. Specifically, raters were told, “... being an adult is not relevant to the story”. However, many raters appeared to be slightly overzealous in this categorization and frequently chose ‘adult’ when it may not have been relevant to the narrative. Lastly, raters were asked to identify any race, religion or other “category” which might be relevant to the story. Other categories that raters identified included Little League, single parent households and homosexuality. The goal of this category was to provide additional insight about other organizations, groups or affiliations that might be an interesting way to categorize the narratives.

The ‘Where?’ category was subtitled ‘Scene of the Action’ and was intended to gather information about locations where the narrative took place. We hypothesized that the participants might have behaved differently in different parts of their lives. For example, there were several narratives where participants talked about doing things of questionable morality while on vacation; things they would not have done at home. We were interested in whether this pattern could be seen more generally in the database as a whole. This category included items such as ‘family’ and ‘animals’ as well, under the assumption that these might also affect how a person chooses to apply their moral

standards. Raters were encouraged to write in additional scenes and their write-ins included scenes as varied as: Church, Rehab, Online and Funeral.

The three major categories, ‘What Did They Do?’, ‘Why Did They Say They Did It?’ and ‘How Did They Feel?’ were all rated the same way. The instructions given asked the raters to “...put each story in EVERY category that it belongs in.” They were asked to rate each item in these categories on a scale from zero to three. (For a complete list of the items in each category, see Appendix G.) Zero was to be used when the actions in the narrative did not apply to the category at all. Three was to be used when the actions in the narrative were a great exemplar of the category. Raters were told to consider that a story about a child taking a piece of candy and a professional jewel thief who robs a museum are both stories where someone stole something and should both be categorized as ‘stealing/taking someone else’s money’ but that they might not earn the same degree of rating. Raters were specifically not given any actual instructions on how to rate these two examples, to prevent biasing the rating process. In addition to rating stories where the narrator did the action in the category, raters were also instructed to look for omissions of the action. In the case where the narrator talks about deciding not to shoplift, for example, that story would be categorized as an ‘omission of stealing’. Omissions were designated with negative signs, to distinguish them from actions. This is critical to remember when using the rating data in the database. Negative signs are simply a convenient code for ‘omission’. They are not indicative of a negative action. The raters were not asked to rate on a Likert scale from -3 to +3. (They were not, in fact, asked to rate on a Likert scale at all.) Omissions are a useful part of the category for any kind of action since they provide an additional aspect to moral decision making – that of making a concrete decision not to

act. The morality of making a decision to act or omit an action has been discussed by moral philosophers for many years. For research purposes, all stories that are rated higher than one are considered members of that category. Further subdividing them based on their score between one and three (or negative one and negative three) should be used on a case-by-case basis.

For their last zero to three judgment, raters were asked to answer the question, “Did they rationalize an action or behavior?” Again, “they” refers to the narrator in each narrative. In response to preliminary rater questions about how to apply a scale to a yes or no question, the instruction sheet clarified this question as: “How strongly does rationalization play a role in this story?”

The category ‘Cue Word’ was used to generate a more complete list of the cues that could have prompted each narrative. Raters understood that each narrative they were reading had been produced using a cued recall protocol, using a single cue word. During the initial editing process of turning spoken memories into narratives, we had discovered that many of the narratives contained elements of several cues. For example, a story about cheating on a spouse, cued by ‘cheating’ might also have been cued by the cues ‘guilty’, ‘regretful’ and ‘worst thing you ever did’. Additionally, the priming effect that has been discussed in Chapter 2 meant that some cue words overlapped with each other in the participants’ minds. A good example of this would be the cues ‘cheating’ and ‘unfaithful’. Often participants told a story about cheating on a spouse when given the cue ‘cheating’, but then were stymied when the cue ‘unfaithful’ came up and instead told a story about cheating on a test. To characterize these effects, the raters were asked to select every word in the cue word list that could have elicited the narrative they just read.

The raters were also instructed not to “over-think this category; just do it quickly and move on.” This prevented raters from adding too much extraneous information from their own experience to the choices they made.

The category ‘Judgment’ is perhaps the most interesting of the categorizations that the raters completed. However, as has been discussed above, the rater judgments in this category cannot be used as ‘normal’ ratings for these narratives. They are merely for initial investigation purposes. Raters were given a visual representation of a Likert scale from -5 to +5 and asked to provide their own judgment of the following question: “How right or wrong do YOU think the action in the story is?” Negative five was designated “completely wrong”, positive five was designated “completely right” and zero was specifically labeled ‘gray area’. In the instructions, the raters were further told that zero meant “completely morally ambiguous” or “equally right and wrong”. They were also told to write the words ‘not moral’ in the rating space if they believe that there was no moral component to the narrative. These specific instructions about zero were intended to prevent a discontinuity at that point on the scale and to re-emphasize the continuous nature of a Likert scale to the raters. The raters were also told in the instructions that this rating was unique because it was purely their own subjective opinion about the actions in the narrative. They did not need to take into account the beliefs of the narrator, only their own beliefs. This was important, for example, when considering a narrative that involved an action like taking a family member off life support.

The final category that raters completed was ‘Core Relational Themes’. This section was based on Richard Lazarus’ concept of the core relational themes of emotion. Briefly, Lazarus claims that each core emotion is generated from a specific action. For

example, “A demeaning offense against me and mine” would cause one to feel angry. He has a core relational theme for each of fifteen emotions. The raters did not receive any background about Lazarus’ theory because it was felt that they could choose the actions in their narrative without further explanation. Instead, they were asked to choose all the themes that applied to their narrative “based on the description of the theme, NOT the emotion word.” They were also told to leave this blank if none of the themes fit their narrative. Analyzing these data is outside the scope of this current project; however, the data can be found in the database for other research projects.

**Similarity Rating Instructions:**

The similarity rating sheet was explained to the raters with a separate set of instructions. These instructions explained that they were to generate a rating on a Likert scale from -2 to +2 of how similar two narratives were. Specifically, they were told that “similarity ratings [may] help identify “categories” that might not be intuitively obvious”. They were told to complete these pairwise ratings after completing one rating sheet and reading the next narrative, before filling out its rating sheet. In this way, they were best able to compare their intuitive sense of how similar the two narratives were. To emphasize the organic nature of this rating, raters were told to “do [this rating] by feel”. This part of the rating task generated many questions, both during pilot ratings and from the alumni raters over email. Many raters were concerned about what ‘similar’ meant in this context. Rather than giving examples, raters were simply encouraged to go with their instincts. They were also told that they didn’t need to have a reason for why the two narratives were similar, just a sense of how similar they were. In this way, we hoped to

generate the most objective similarity ratings with the least researcher bias. In order to generate as many unique pairwise comparisons between narratives as possible, the raters were encouraged to read the narratives in any order they chose. Each packet with a particular set of 30 narratives sent out was only sent to a maximum of two raters and since most raters printed out their narratives, their similarity ratings were not likely to be done in the same order.

**Returning Rating Data:**

When raters completed the rating process, they emailed and indicated whether they would like a pre-paid mailing envelope mailed to them through the US Postal Service so that they could return their rating sheets cost-free. Most raters opted for this choice, but a few emailed their ratings in instead. After the ratings were received, they were either manually entered into the database (if they were received on paper) or digitally transferred into the database (if they were received as spreadsheets).

**Section 6: Inter-rater reliability:**

To assess the reliability of ratings among raters, we evaluated ratings from 111 narratives (all the narratives with exactly 5 raters, approximately 15% of the database). Since ratings were inclusive and different raters may have had different tolerances for inclusion, we calculated the number of narratives where at least 4 of the raters agreed on at least one category as our minimum standard for rater agreement. Under these conditions we found that 105 out of 111 narratives showed rater agreement. (Using a more stringent standard where all 5 raters must agree found 76 of 111 narratives had



agreement.) By comparison, when the ratings were shuffled (within raters but among narratives), only 26 of 111 narratives showed rater agreement. These findings give us confidence that there is consistency among raters and the ratings can be considered reliable.

## Section 7: Database Development

After some consideration about the requirements of the database, Microsoft Access was chosen as the platform for this database. Several items were influential in choosing Access. 1. It is readily available to most researchers because it is included in the Microsoft Office Professional Suite. Also, the fact that it is a Microsoft product suggests that it is likely to be available and supported for the foreseeable computing future. 2. Access databases are relatively easy to implement. This was critical for creating a well-organized database with features such as accurate primary keys. 3. Access databases can be searched with SQL, a common computer language used for databases. Although Access is somewhat difficult to use on a Macintosh computer, the SQL compatibility means that researchers using only Macs can still use the database. 4. Access databases can be placed on websites to allow remote users the ability to view the data. This substantially improves the ease with which the database can be administrated. In particular, we hope that other researchers will contribute ratings to the database and allow it to grow with time. In order to add new data while preserving the integrity of the original data, it is important to maintain a 'primary' database where all researchers can come to view data that has been newly added. 5. Lastly, Access is straightforward to use.

This provides ease of use for future users as well as a minimum of support that needs to be maintained for the database. The database has also been converted to Filemaker, a database program that can be used on Macintosh computers. These file formats allow users to access the data in the two most commonly used database programs.

## Section 8: Database Documentation: A Guide to the Tables in the Database

All tables are organized using the field ‘Subject ID’ or ‘Story ID’ as the primary key. This allows for straightforward linking between the demographic data and the rating data. These descriptions are provided in alphabetical order, for ease of reference. For a better understanding of the structure of the tables, refer to Figure 7.

Commonly recurring fields:

**Subject ID** – All subjects were given a four digit code between 1000 and 2000 upon entry into the study. These are unique identifiers.

**Story ID** – All narratives were given a four digit code between 5000 and 6000. These are unique identifiers. To ensure that the narrative identifiers could not be used to identify subjects, a table of random numbers was used to select the Story ID.

**Rater ID** – All raters were given a two letter, two digit code of the form: (RA or RR)(01-54). Each Rater ID is a unique identifier. The distinction between RA and RR was for internal purposes only and should be disregarded for research purposes. There is no difference between the two groups.

In the text describing each table, ‘subjects’ refers to the 100 participants who told the narratives that comprise the database. ‘Raters’ refers to the 55 anonymous people who read and rated each story based on the information it contained.

**Biographical Subject Data:**

This table contains all of the demographic data collected about the subjects during their initial intake into the study. The headings in this table are:

Subject ID – This field is the primary key in this table.

Year of Birth – All subjects were between the ages of 40 and 60 when the study was conducted, so all years of birth are between 1943 and 1965.

Gender – Subjects were only given a choice between Male and Female. All subjects chose to identify themselves by these labels.

Marital Status – Subjects wrote in their marital status. From these write-ins, a list of possible marital statuses was generated: Single, Married, Divorced, Widowed, Separated and Partnered.

Children – Subjects provided a numerical answer to how many children they had.

Ethnicity – Subjects chose from a list of ethnicity labels: American Indian, Asian/Pacific Islander, Black, Caucasian, Hispanic, Multiracial, Other. The label Hispanic was included separately from Caucasian to allow demographic distinction between these two populations. This was especially important in this study because the population of Southern California has a large Hispanic population.

Education – Subjects chose the highest degree of education completed from the list of: Elementary School, Middle School, Some High School, High School, Some College, Bachelor's Degree, Master's Degree, Graduate Degree.

Income – Subjects chose from a list of possible income ranges. Data in this column refer to the following key: 1- \$0 – \$15,000; 2- \$15,000 - \$30,000; 3- \$30,000 - \$60,000; 4- \$60,000 - \$90,000; 5- \$90,000 – 100,000; 6- \$100,000 - \$200,000; 7- Over \$200,000.

Occupation – Subjects provided a written answer about their occupation. Data in this column are transcribed exactly. Some cells include multiple occupations; some include prior occupations (before retirement or unemployment).

Handedness – Subjects chose from the list: Right Handed, Left Handed, Ambidextrous. Subjects also completed a modified Edinburgh Handedness Inventory and their answers were checked against these choices for accuracy.

Other Language – Subjects provided a written answer about any other languages they spoke fluently.

Religion Then – Subjects were given room to write a short explanation of their religious background. An independent rater read these and extracted the religion or religions that the subject grew up with. Some cells may include multiple religions.

Religion Now – Subjects were given room to write a short explanation of their current religious beliefs and practices. As with previous column, a rater extracted the current religion that subjects practiced. Cells are non-uniform as the subjects' own descriptions were used wherever possible.

Degree of Religious Influence – After reading the subjects' descriptions of their childhood and current religions, two independent raters gave each subject a subjective score on degree of religious influence in their life. A 0 was scored if subjects did not have any religious influence in their lives. A 1 was scored if subjects: were raised with a religion and did not currently practice a religion; were occasionally involved in some

religious activity; or appeared to have some religious influence that was not strongly significant in their lives. A 2 was scored if subjects identified themselves as practicing members of a religion. Raters did not discriminate between established religious organizations and atypical forms of religious belief. All data in this column are whole numbers; any disagreement between the two raters was resolved by a third rater.

Sexual Orientation – Subjects chose from the following list: Heterosexual, Homosexual, Bisexual, Transgender, Polyamorous, None. There is one cell that contains a write-in answer. A small number (n=6) of subjects chose not to answer this question. A blank cell is distinct from the answer None.

#### **Cue Word Ratings:**

This table contains the data collected from raters about cue words. Raters were given a table of all cue words and asked to select every word that might have elicited the story they just read. In this table, all three fields combined provide the primary key. All fields are duplicated, but the combination is unique since Story ID and Rater ID are each unique identifiers. The headings in this table are: Story ID, Rater ID and Cue Word.

Cue Word – Data in this field were chosen from the table ‘LookupCues’ and include the 31 possible choices given to the raters.

#### **Eliciting Cue Word:**

This table contains the cue word that elicited each narrative. For more detail about the cue words and cued recall protocol, see Methods and Materials. The headings in this table are Story ID and Cue Word.

Story ID – This field is the primary key in this table.

Cue Word – Data in this field were chosen from the table ‘LookupCues’. They are the 31 unique cues given to subjects during the cued recall protocol.

**Eliciting Cue Word (Multiple Cues):**

This table contains data about the narratives that were elicited by multiple cue words. In some circumstances, subjects told the same story more than once during the cued recall protocol. In most cases, only one of these narratives was selected for the database, but in a few (n=16) cases the stories were combined to provide the best narrative. (This is explained more fully elsewhere.) In these cases, multiple cues (2 or 3) elicited a single story. The headings in this table are Story ID and Cue Word. The primary key is the combination of both fields.

Cue Word – Data in this field were chosen from the table ‘LookupCues’.

**LIWC Data:**

This table contains the results from a complete Language Inquiry and Word Count 2007 analysis of each narrative in the database. (For more information about LIWC analysis, see Chapter 4. For more information about LIWC, see the LIWC2007 Language Manual<sup>180</sup>.) The headings for each field were not altered from the LIWC output to allow easier comparison between these data and other LIWC analyses. For easier comprehension in this database, each field is briefly described below. For examples of each field, see the LIWC2007 Language Manual, Table 1. The headings in this table are

---

<sup>180</sup>Pennebaker, J. W., C. K. Chung, et al. (2007). The development and psychometric properties of LIWC2007. Austin, TX, LIWC.net.

divided into categories below based on the hierarchical structure of the LIWC dictionary.

Be sure to understand the hierarchy before using these data.

Story ID – This field is the primary key in this table.

WC – Word Count. Note that each narrative includes the words “Memory 5xxx” in the text that was analyzed, so this word count is two words longer than the actual length of the narrative.

WPS – Number of words per sentence.

Sixltr – Number of words with greater than six letters.

Dic – Percentage of words in the text that appeared in the LIWC dictionary.

All data in the fields below are given as percentages of the total narrative. As an example, the heading ‘pronoun’ calculates the number of pronouns in the text and presents it as a percentage of the total length of the narrative.

funct – Function words. Included in this hierarchy are the headings:

pronoun – Total pronouns, which includes:

ppron – Personal pronouns, which includes:

i – 1st person singular; we – 1st person plural; you – 2nd person; shehe – 3rd person singular; they – 3rd person plural

ipron – Impersonal pronouns;

article – Articles;

verb – Common verbs. Some verbs are included in total function words and some are not.

For details be sure to read the LIWC manual. The category verb includes:

auxverb – Auxiliary verbs; past – Past tense verbs; present – Present tense verbs; future – Future tense verbs;

adverb – Adverbs; preps – Prepositions; conj – Conjunctions; negate – Negations; quant – Quantifiers; number – Numbers;

swear – Swear words

Psychological Processes: The words in this category are all related to psychological behaviors and include:

social – Social processes, which includes:

family – Family; friend – Friends; humans – Humans;

affect – Affective processes; which includes:

posemo – Positive emotion;

negemo – Negative emotion, which includes:

anx – Anxiety; anger – Anger; sad – Sadness;

cogmech – Cognitive processes, which includes:

insight – Insight; cause – Causation; discrep – Discrepancy; tentat – Tentative; certain – Certainty; inhib – Inhibition; incl – Inclusive; excl – Exclusive;

percept – Perceptual processes, which includes:

see – See; hear – Hear; feel – Feel;

bio – Biological processes, which includes:

body – Body; health – Health; sexual – Sexual; ingest – Ingestion;

relative – Relativity, which includes:

motion – Motion; space – Space; time – Time;



Personal Concerns: The words in this category are all related to personal behavior and include:

work – Work; achieve – Achievement; leisure – Leisure; home – Home; money – Money; relig – Religion; death – Death;

Spoken Categories: The words in this category have to do with word usage that is unique to spoken words (as compared to written word) and include:

assent – Assent words, like “OK”; nonfl – Nonfluencies, like “err” ;  
filler – Filler words, like “I mean”;

Punctuation Marks: This field calculates the total percentage of the text that is comprised of each type of punctuation. These types of punctuation include:

Period – Periods; Comma – Commas; Colon – Colons; SemiC – Semicolons;

QMark – Question marks; Exclam – Exclamation points; Dash – Dashes;

Quote – Quotes; Apostro – Apostrophes; Parenth – Parentheses; OtherP – Other punctuation marks;

AllPct – Total of all the punctuation marks in the text.

### **LookupCues:**

This is a lookup table used to generate identical lists of cue words throughout the database. Lookup tables are used to eliminate data errors based on spelling or capitalization inconsistencies. Throughout this work, there are references to the 35 cues in the cued recall protocol. There are only 31 choices in this table because the cue “Talk about a memorable experience” was repeated 5 times for the subjects, bringing the total number of unique cue words to 31. The heading in this table is Cue Word.

Cue Word – Cue words were presented to subjects in complete sentences in the form, “Talk about a time when you did something that made you feel \_\_\_\_\_” or “Talk about a time when you \_\_\_\_\_”. A complete list of these cues can be seen in Table 1. Most of the words used in this lookup table are exactly the cue words used. However, some of the cues used longer prompts and were abbreviated for length in this table. Those cues are listed with each heading below:

Best Thing – Please talk about the best thing you ever did.

Memorable – Please tell us about a very memorable thing that you have done.

Most Afraid – Please talk about the thing that you are most afraid that other people will find out that you did.

Most Like Others to Know – Please talk about the thing that you would most like other people to know that you did.

Most Like to Change – Please talk about the think that you would most like to change if you could go back and do it differently.

Most Proud – Please talk about the thing that you are most proud of having done.

Worst Thing – Please talk about the worst thing you ever did.

### **LookupEmotions:**

This is a lookup table for the table ‘Session Three: Emotion Data’. It contains a list of the 12 emotion choices subjects were given to choose from when rating their own narratives during Session Three. There were no write-in options given in this choice.

**LookupScenes:**

This is a lookup table that contains a list of all of the ‘scenes’ chosen by raters. For each narrative, raters were asked to choose the location (or scene) of the story from the following choices: Job/Business Situation; Family/Close Friend; Abroad/Vacation; Home/Neighborhood; School; Animal; Party/Club; Car. Raters also wrote in additional locations and the lookup table was used to codify those choices. This prevented items like ‘clinic’ and ‘doctor’s office’ from appearing separately, rather than together.

**Neuropsychological Subject Data:**

This table contains all of the neuropsychological data collected from subjects during their initial intake into the study. For more complete descriptions of each of these tests, see Chapter 4, Section 1. The headings in this table are:

Subject ID – This field is the primary key in this table.

IQ – Full-scale intelligence quotient (FSIQ) was tested in the laboratory using the 2-subtest Wechsler Abbreviated Scale of Intelligence.

NEO-FFI Scores – The NEO personality inventory generates five scores based on five personality factors: Neuroticism, Extroversion, Openness, Agreeableness and Conscientiousness. Scores are reported as T scores. These headings include: N NEO T Score, E NEO T Score, O NEO T Score, A NEO T Score and C NEO T Score.

PolV – Subjects completed a questionnaire about their political values, which generated a score between 1 and 3. A low score (1) indicates more conservative values, an intermediate score (2) indicates middle of the road values and a high score (3) indicates more liberal values.

PANAS Scores – The Positive and Negative Affect Schedule generates two scores to assess general mood, the positive affect score (PA) and the negative affect score (NA). PA and NA are orthogonal, not opposite dimensions. (For more discussion of this, see Results.) The PANAS headings include: PA PANAS and NA PANAS.

Ethical Position Questionnaire (EPQ) Scores – The Ethical Position Questionnaire is used to measure ethical ideologies. It creates two subscores, relativism and idealism, as well as an overall score. Although the overall score can be used, the two subscores (particularly the idealism score) are more commonly reported in the literature, so they have been included in the database. The headings in the table are: Idealism EPQ, Relativism EPQ, Total EPQ.

Weighted Utilitarian Score – Subjects completed a 30 question version of philosophical moral dilemmas, including the standard trolley problem. A score of each subject's utilitarianism was generated from the results of these dilemmas. These scores were weighted based on the severity of the dilemma. Lower scores indicate subjects who are less utilitarian in their decision making, higher scores indicate subjects who are more utilitarian.

Schwartz Value Inventory (SVI) Scores – The Schwartz Value Inventory is used to measure differences in human value structures. It generates 10 scores based on 10 core human values. For greater discussion of each of these headings, see Chapter 4, Section 1 and Table 6. The SVI headings include: Conformity SVI, Tradition SVI, Benevolence SVI, Universalism SVI, Self-Direction SVI, Simulation SVI, Hedonism SVI, Achievement SVI, Power SVI, Security SVI.

**Rating Data:**

This table contains all of the data collected from the raters about narratives. The rating process and information about the raters can be found earlier in this chapter. Raters were asked to make decisions about 3 major categories: What Did They Do?; Why Did They Do It?; and How Did They Feel?. In the description of the fields below, ratings will be grouped by category. Within these categories fall the majority of the fields in this table. For each item, raters gave a rating between 0 and 3 based on how well the narrative they were reading typified the action. A negative sign was used to designate narratives where the person specifically chose not to perform the item.

The headings in this table are:

Story ID – This field is combined with Rater ID to create the primary key for this table.

Rater ID – As above.

Gender? – This field is the rater’s guess about the gender of the author. It is NOT the actual gender of the author. The question marks in headings throughout this table are intended to designate the raters’ opinions or feelings, not those of the subjects.

Gender Guess? – Raters were asked if they were certain of their choice of gender for each narrative. If raters felt that they were guessing, the answer is yes. If they were certain of the gender of the person in the narrative (the narrator), the answer is no. It is worth nothing that although raters believed they were certain, there are many cases in which they were still wrong. (This is discussed in Chapter 4, Section 4.)

Age – Raters chose the age of the narrator based on the information in the narrative, only if they were relevant to the narrative. Specifically, the raters were told that being an adult

was typically not relevant. The choices given to choose from are: Child, Preteen, Adolescent, Adult.

Other Feature? – This field is unique in its open-endedness. Raters were instructed to note any group affiliations that might be relevant to the story. This included ethnicity and religion in the instructions to the raters. The items that raters generated also included things as diverse as scouting groups, single parenting and more.

What Did They Do? Headings – The headings in this category include: Helping someone, Protecting someone, Taking care of someone, Being kind to someone, Self-sacrifice, Cheating in school/work, Cheating in business, Sneaking into a movie theatre, Falsifying a document, Tricking someone else, Physically unfaithful, Mentally unfaithful, Neglecting someone, Being hurtful to someone, Being unkind to someone, Lying, Telling a white lie, Stealing/Spending someone's money, Stealing an idea/plagiarizing, Shoplifting, Returning lost money/wallet/things, Bank error/too much change/undercharged, Enlisting someone's help to do the wrong thing, Enlisting someone's help to do the right thing, Fighting for someone else's rights, Taking advantage of someone, Breaking a rule/law, Vandalism/Damaged property

Other What? – In this field, raters were able to write in any additional actions that they felt were missing from the What Did They Do? headings.

Why Did They Do It? Headings – The headings in this category include: Taking the high road, Doing the wrong thing, Doing the right thing, Making extra effort to do the right thing, Making extra effort to do the wrong thing, Doing lesser harm to prevent greater harm, Means do NOT justify the ends, Fear of retribution/karma, Cheating 'the system'/'the man', Righting a wrong, Peer pressure/social pressure, Altruism, Teaching a

lesson/revenge, Self-preservation, Getting out of a tight spot, Selfishness, Taking a chance, Because of money.

Other Why? – This field is the same as Other What?. The raters generated additional reasons that they felt the narrator had acted the way that he or she did. The raters were not limited to reasons specifically stated in the narrative; they could add their own intuition about the narrator’s reasons in this category as well.

How Did They Feel? Headings – The headings in this category include: Regret about doing the right thing, Regret about doing the wrong thing, Regret about behavior, Knowledge of personal moral weakness, Knowledge of personal moral strength, Uncertainty about moral correctness of action, Not taking responsibility for actions, Taking responsibility for actions, Learning a lesson, Putting self in someone else’s shoes, Shame/ashamed, Powerful.

Other How? – This field is the same as the other Other headings. Raters were allowed to write in their any rationale that they felt was not included above.

Did They Rationalize? – Raters judged the degree to which the narrator rationalized his or her actions. This judgment was also made on a scale from 0 to 3.

Rating – After reading each narrative, raters gave their personal judgment of how right or wrong the actions of the narrator were. This was done with a Likert scale from -5 to +5. Raters were instructed to use ‘0’ to mean “neither right nor wrong” and to write in “not moral” if they believed the story was not moral at all. These specific instructions were intended to prevent discontinuity in the scale at the zero point. In the database itself, the words “not moral” were coded as the code 999. When raters left the rating judgment

blank, it was coded as 998. These codes were used since the field uses a ‘number’ data type.

### **Scene Ratings:**

This table contains the data raters generated about the locations in which the narratives took place. Often, the narratives took place in more than one location and raters chose all of the appropriate locations. Most of the other rater data is included in the table ‘Rating Data’. However, in order to create a primary key in the ‘Rating Data’ table, it was necessary to display each Story ID and Rater ID pair only once in the table. Since the same rater selected multiple scenes for the same story, it precluded this data being included directly in the table. It is available, instead, as a subdatasheet drop down menu in ‘Rating Data’. This maximizes the amount of information that can be contained in that table. The headings in this table are: Story ID, Rater ID and Scene. The primary key is the intersection of all three fields.

Scene – This field uses the lookup table ‘LookupScenes’, which includes all 30 possible locations generated by the raters.

### **Session Three Data:**

This table contains the data acquired during Session Three, the testing session where subjects returned to the lab and read the narratives that had been transcribed from their memories. For more information about this testing session, see Methods and Materials. The headings in this table are quite abbreviated; the complete questions asked



for each item can be found below as well as in the description that accompanies the design view of the table. The headings in this table are:

Story ID – This field is the primary key for this table.

Rating the Narrative Headings – For each of the following headings, the subjects were asked to answer the question of the general form, “how \_\_\_ was this experience?”, using a Likert scale from 1 to 10, where 1 was “not \_\_\_ at all” and 10 was “extremely \_\_\_”.

For exact wording for each question, see Appendix C. The headings in this category include: Pleasant, Vivid, Intense, Emotional, Significant, Confidential, Regret<sup>181</sup>, Guilt, Embarrassment, Pride<sup>182</sup>.

At the Time Headings – The second category of questions asked of the subjects was about their feelings at the time that the narrative occurred. Most of these questions involved Likert scale decisions. The endpoints for these scales are listed with the heading. The headings in this category include:

Then RoW – At the time, how wrong or right did your actions in this episode seem to you? Very Wrong = 1; Very Right = 10.

Then Options – At the time, did you feel like you had options about what you could do in the situation? Could Do Nothing Else = 1; Could Just As Easily Have Done Many Other Things = 10.

Then Choices – To the best of your recollection, approximately how many distinctly different things could you have done at the time?

---

<sup>181</sup> The wording for regret was “do not regret at all” to “regret a great deal”.

<sup>182</sup> The question asked for this heading was “how proud are you now about your behavior during this experience?” The heading ‘Pride’ was used to distinguish this question from others in Session Three about pride.

Morally Better Could – At the time, did you think there was something morally better that you could have done?

Morally Better Should – At the time, did you think there was something morally better that you should have done?

Content with Actions – Immediately after you made the decision and did these actions, how content were you with your actions? Completely Satisfied = 1; Completely Discontent = 10.

How Difficult – At the time, how hard was it, morally, for you to do what you did? Specifically, subjects were instructed to think about their psychological struggle with making the decision in the narrative. Very Easy To Do = 1; Agonizing and Difficult = 10.

Second Guess Self – At the time, did you second-guess your decision?

Frequency of Recollection – To the best of your recollection, how many times in the year following this episode did you think about this decision?

Wish Done Differently – To the best of your recollection, how many times in the year following this episode did you wish that you'd done something different?

Current Feeling Headings – The third category of questions elicited the differences between then and now. The same questions as the previous category were asked. The headings in this category include:

Now RoW – Reading this today, how wrong or right do your past actions in this episode seem to you? Very Wrong = 1; Very Right = 10.

Now Options – Reading this today, do you feel like you had options about what you could have done in the situation? Could Do Nothing Else = 1; Could Just As Easily Have Done Many Other Things = 10.

Now Other Choices – Thinking about the situation today, approximately how many distinctly different things could you have done at the time?

Now Morally Better Could – Today, do you think there was something morally better that you could have done?

Now Morally Better Should – Today, do you think there was something morally better that you should have done?

Now Content with Actions – Today, how content are you with your past actions?

Completely Satisfied = 1; Completely Discontent = 10.

Now How Difficult – If you were in the same situation again, how hard would it be, morally, for you to decide what to do? Very Easy To Do = 1; Agonizing and Difficult = 10.

Still Think About It – Do you still think about whether you made the right decision in this situation?

Now Frequency of Recollection – How many times in this past year have you thought about the decision you made in this episode? Subjects were instructed not to count times that they had thought about the episode since beginning the study.

Now Wish Done Differently – When you think about this episode today do you wish that you'd done something different? This field uses the data type 'memo' to contain the complete text of the written-in answers given by the subjects.

Taking an Outside Perspective Headings – This category asked subjects to imagine how others reading these narrative would view them. Specifically, these imaginary other people were described as strangers to the subject and "other normal people who might be in a research study". The headings in this category include:

Others RoW – If other people read this situation, how wrong or right would the actions in this episode seem to them? Very Wrong = 1; Very Right = 10.

Others Options – If other people read this situation, would they feel like you had options about what you could do in the situation? Could Do Nothing Else = 1; Could Just As Easily Have Done Many Other Things = 10.

Others Morally Better Could – Would other people think there was something morally better that you could have done?

Others Morally Better Should – Would other people think there was something morally better that you should have done?

Others How Difficult – If other people reading this story were rating on a scale of 1 to 10, how hard would they say that it was, morally, for you to decide what to do? Very Easy To Do = 1; Agonizing and Difficult = 10.

Others Right Choice – Do you think other people reading this would say that you made the right decision in this situation?

Highly Ethical Other Do Same – Do you think a very strictly ethical person (for example Gandhi) would do the same thing?

Highly Ethical Other RoW – How right or wrong would a strictly ethical person rate your decision? Extremely Wrong = 1; Extremely Right = 10;

Personal Reflection Headings – The final category included additional questions that reflect on the subject's personal state of mind during the narrative.

Consider Judgments – Did you think about what other people would say at the time that this happened?

Consider Opinions – Did you think about other people’s opinions before you made the decision?

Whose Opinions – Whose opinions did you consider? This field uses the data type ‘memo’ to contain the complete text of the written-in answers given by the subjects.

Thought About Self – How much did you think about yourself (your own needs and wants) when you made your decision? Did Not Think of Self At All = 1; Thought Entirely About Self = 10.

Decision Based on Self – How much did you base your decision on your own needs and wants? Based Decision Entirely on Other Factors = 1; Based Decision Entirely on Self = 10.

Self Ethical Rating – Rate yourself as an ethical person. Not Ethical At All = 1; Highly Ethical = 10.

Ethical Definition – Please define ethical. Subjects were asked this question after reading each narrative. Some answered differently each time, some referred to their first answer.

To determine the order in which these answers were given, this table can be queried using the ‘Subject ID to Story ID’ table.

Reflect General Behavior – How well does this story reflect your general behavior?

Never Act Like This = 1; Always Act Like This = 10.

Reflect On Self – How well does this story reflect on you? Reflects Very Badly On Me = 1; Reflects Very Well On Me = 10.

### **Session Three: Emotion Data:**

This table contains the emotion choices subjects made after reading their own narratives. During testing session three, subjects were asked to select from a given list the emotions that “best describe how this experience made you feel”. They were instructed to choose up three emotions from the list of twelve. In addition, they were told to use the order of their choice to designate the relative degree among the emotions (i.e., the emotion that best described how they felt was to be chosen as Emotion 1). This table is intended to be used as a subdatasheet of the table ‘Session Three Data’. The headings in this table are:

Story ID – This field is the primary key for this table.

Emotion 1, Emotion 2, Emotion 3 – The data for these fields are drawn from the lookup table ‘LookupEmotions’ to generate the list of the given emotions. No write-in emotions were allowed. Subjects were allowed to choose no emotions if they did not apply and were given an upper bound of 3 choices. The order of these headings contains information about the degree that each emotion was elicited by the narrative.

### **Session Three: Temporal Data:**

This table contains several representations of the age of each narrative. During testing session three, subjects were asked to report when the actions in the narrative took place for each narrative. If they were unable to give a date, they were asked to give their “best guess of their age when this event occurred”. This table can be used as a subdatasheet for the table ‘Session Three Data’ or it can stand alone. It may be especially useful when combined with the ‘Biographical Subject Data’ table about the age of each subject. The headings in this table are:

Story ID – This field is the primary key for this table.

Year of Story – This field has been generated from the original text given by the subject.

In some cases, these dates have been estimated based on the text. The original text can be seen in the field Original Text of Date Recollection. The code 9999 has been used to indicate narratives that the subject said were recurring or ongoing. This distinguishes them from stories where the subject was unable to provide a date.

Age of Memory – This field is described by the equation:  $\text{Age of Memory} = 2004 - \text{Year of Story}$ . It is included purely for convenience of the researcher. Session Three was completed for all subjects in 2004. As above, the code 9999 was used to replace information for narratives that were recurrent.

Original Text of Date Recollection – This field contains the original text typed by each subject. It has not been edited in any way. It is included in this table for researchers who would like to use the original text to discriminate smaller temporal units than years.

### **Session Two Exit Questionnaire:**

This table contains the data gathered from subjects after completion of Session Two. An exit questionnaire was administered to collect data on subjects' general feelings about the process of generating autobiographical memories and the effect that imposing a time limit had on the cued recall protocol. In all applicable cases below, the subjects were given a Likert scale from 1 to 10 to use when answering these questions. 'Not at all' was represented by 1 and 'very much' was represented by 10. The headings in this table are:

Subject ID – This field is the primary key for this table.

Degree of Emotion – Subjects answered the question “How emotional did you feel during this study?”

Degree of Difficulty – Subjects answered the question “How difficult was this task?”

Reflection – Subjects answered the question “How much will you think about these memories in the next week?”

Reflection Comparison – Subjects answered the question “How much do you think you would have thought about these memories in the next week if you hadn’t participated in the study?”. This field is intended to be a comparison with the previous field. The goal as an exit question was to elicit some information about the affect of the autobiographical recall on the subject’s overall behavior.

Suggested Time Limit – Subjects were asked to fill in a blank about how long they were supposed to spend recording each memory. The correct answer was 3 minutes. However, the question may have been ambiguous since several subjects responded with ‘4 hours’ which was the length of time the entire testing session was supposed to take. This field also gives insight into subjects’ general comprehension abilities.

Adherence – Subjects answered the question “How closely did you try to adhere to the time limit?” using a Likert scale from 1 to 5.

Pressured – Subjects answered the question “How pressured did you feel by the time limit?” using a Likert scale from 1 to 5.

Comments – Subjects were given an open-ended space for comments on the testing session. This has been included in the database because they may provide insight for other researchers about the state of mind of the subjects after the cued recall protocol.



### **Session Two Exit Questionnaire: Emotion:**

This is a table that describes how subjects were feeling after completing Session Two. Data in this table was part of a larger exit questionnaire; however since subjects were able to choose more than one answer to this question, the data are better presented in a separate table. This table is most useful as a subdatasheet of the table ‘Session Two Exit Questionnaire’. The headings in this table are: Subject ID and Emotion. The intersection of these fields is the primary key.

Emotion – Subjects responded to the question “How do you feel now that you have completed this task?” by choosing all of the ways they were feeling from the following list: Happy, Satisfied, Pensive, Tired, Relieved, Wrung Out. Subjects were also allowed to write down any other emotions they had. Written answers were taken verbatim from subject responses, resulting in some inconsistent emotions.

### **Subject ID to Story ID:**

This table links Subject ID to Story ID. This generates a relationship between the two major divisions in this database – data that are linked to subjects and data that are linked to stories. For a graphical representation of this see Figure 7. The headings in this table are:

Subject ID

Memory # – Each subject produced multiple memories. These memories were numbered sequentially, so this field makes it possible to order the memories of each subject in the order that they were generated.

Story ID – This field is the primary key for this table.

**How the Database Can Be Used:**

The breadth of the database makes a complete description of its contents overly repetitive. However, a concise presentation of some of the data can be instructive in demonstrating how the database can be used to address research questions. This can be found in Chapter 4, Section 4. The example presented will address the presumed gender of the narrators based on the content of the narrative and compare that to the actual gender of the subject who generated the memory. In addition, the queries that were necessary to collect these data are included in the database to help others become familiar with Access and queries in general.

*Chapter 3: Figures and Tables*

“Men are more moral than they think and far more immoral than they can imagine.”  
-- Sigmund Freud

Figure 5: Distribution of Number of Moral Memories per Participant




Figure 5: The mean number of moral memories elicited during the cued recall protocol was 7.66; however, the mode shows that it was most common to generate 6 moral memories.

Figure 6: Distribution of Number of Words per Narrative

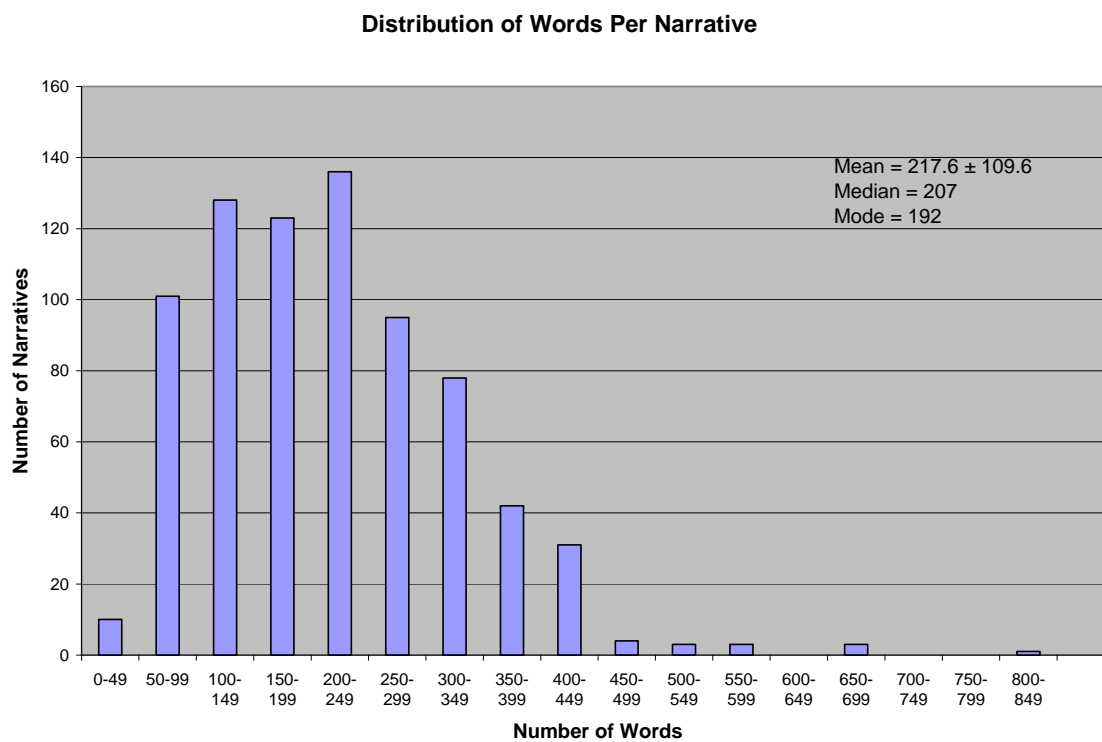


Figure 6: The number of cue words in each narrative was calculated using the Linguistic Inquiry and Word Count 2007 text analysis tool.

Figure 7: Table Structure of Major Database Divisions

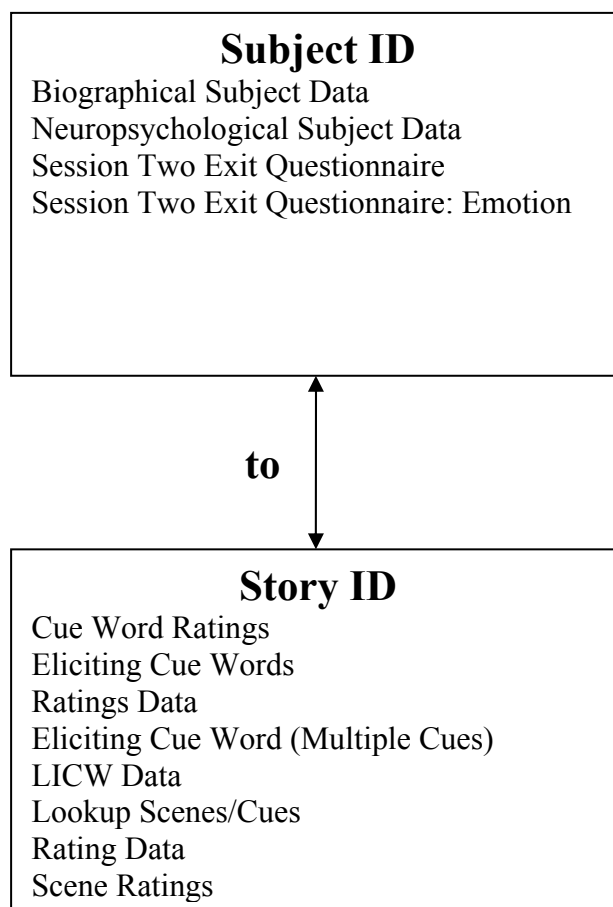


Figure 7: The database structure reflects the primary dichotomy among the data: participant oriented (Subject ID) or narrative oriented (Story ID). All of the primary tables can be classified into one of these categories, as shown in the table above. The table Subject ID to Story ID is the primary linking table.

Table 2: Comparison of Database Narratives to Language and Inquiry Word Count Standards

	Database Narratives	LIWC: Emotional writing	LIWC: Control writing	LIWC: Talking	Database Control Narratives
Words Per Sentence (mean)	16.08 ± 3.69	19.56	19.84	25.87	20.18 ± 0.95
Words in LIWC Dictionary	93.9%	93.42%	88.55%	91.49%	92.73%
Words > 6 letters	11.9%	13.27%	13.87%	9.43%	13.04%
Total function words	65.71%	63.93%	57.53%	60.48%	63.76%

Table 2: Narratives from our database are quite typical when compared to the standards found in the Language and Inquiry Word Count manual as well as control narratives from our own research. The three LIWC categories being compared to are (from left to right): participants who wrote on “deeply emotional topics”; participants who wrote on “relatively trivial topics, such as plans for the day”; and people who were “tape-recorded while engaging in conversations with others”. The percentages given for ‘words with more than six letters’ and ‘total function words’ are the percent of the total text that was devoted to this category. The fourth comparison is to the non-moral narratives that were elicited by our control cues: *funny*, *tired* and *exercised*. It should be noted that the control narratives were not edited, unlike the database narratives, which may account for the difference in length between the two groups.

### *Chapter 4: Results*

The collection and characterization of autobiographical moral memories generated considerable data about the participants who generated the memories and the narratives prepared from them. This chapter seeks to examine some of these data in detail and cover others more broadly. As discussed in Chapter 2, we had various filters in place for screening data for their validity. For example, memories generated by participants whose cognitive condition was dubious on other criteria were omitted, as were memories that were difficult to understand. Valid data comprised all those entered into the database described in Chapter 3. More detailed descriptions of specific analyses and the statistical tests used to test their reliability are given in the following chapter. We followed general recommendations for data analysis<sup>183</sup>. All data were initially plotted and examined visually. When outliers were clearly apparent, they were removed. Any parametric statistical analyses were carried out only on data whose distribution did not deviate too much from normality or whose distribution was transformed to be normal. Some of our data analysis was entirely exploratory whereas other components arose from focused hypotheses. These cases are discussed in their respective sections below.

The detailed collection of data with consideration for other researchers' possible interests and the in-depth characterization of the moral narratives has yielded this extensive collection of data. These findings comprise only a portion of the possible research that could be completed from the rich dataset collected from this research. This chapter begins by analyzing the participants in detail to examine both their demographic

---

<sup>183</sup> Motulsky, H. (1995). *Intuitive Biostatistics*. Oxford, Oxford University Press.



and neuropsychological representativeness. We find that the sample population used for collecting these moral memories is normal on every measure used to assess them. The second section seeks to characterize the moral narratives contained in the database. The narratives are analyzed with respect to the cue word used to elicit them to show that the variation in narratives arises from the moral content of the narratives and not the demographics or personalities of the participants. The subsequent sections analyze specific hypotheses about personal moral autobiographies. Section 3 examines the implicit structure of moral memories and we show that negative memories can be differentiated into three groups: lying, stealing and hurting another person. Section 4 demonstrates how a hypothesis-driven analysis about gender can be augmented by the database. We find that raters of narratives are only moderately good at being able to distinguish the gender of the narrator in each narrative. Section 5 examines the effect of brain damage to emotion-related areas on the ability to generate moral autobiographical memories. Finally, we examine the effect of temporal biasing on moral memory recollection and find that positively valenced moral memories are recollected more recently in a personal autobiography than negatively valenced moral memories.

## Section 1: Characterization and Representativeness of the Participants

### **Part 1: Demographics**

General intuition might suggest that a wide variety of demographic factors may influence the moral stories generated by our participants. One might expect that participants growing up in different cultures, religions or socioeconomic classes have different formative experiences. Gender, IQ and education seem likely, at least

intuitively, to play a role in the moral choices that participants make. However, after collecting nearly 800 stories from 100 participants, we find just the opposite. In general, people from all backgrounds and walks of life seem to have experienced similar moral dilemmas throughout their lives and recall similar moral memories.

There are two key questions about the demographics of our participant sample. First, are they a representative sampling of the population? To begin assessing the representativeness of our sample, we must begin by defining the population they describe. We chose to compare them to the population from which they were drawn: the general population of southern California. In the following section, we provide a detailed characterization of the participant group. Briefly, participants were specifically matched to California state averages on ethnicity and gender. They fall within the normal range for IQ; they are both right and left handed; and their religions, secondary languages, socio-economic statuses and political beliefs are diverse. This characterization of the participants provides the basis for our second question: How might individual differences along demographic or neuropsychological dimensions influence the moral memories that participants generated? For instance, do men and women generate different kinds of moral memories? Does age or ethnicity impact these memories? These types of exploratory questions are used to probe this very large database and the interesting findings are treated in detail in the subsequent sections.

### **Gender, Age and Handedness:**

Our participants were diverse and representative of the general population on all of the demographic measures that we collected. The gender ratio of the participants in the

database was quite close to the state average (53% female, 47% male and 50.2% female, 49.8% male, respectively). We did not have any transgendered participants. All of our participants were between the ages of 40 and 60 years old. Participants were normally distributed (mean =  $48.9 \pm 5.9$  years) within this range (see Figure 8). 83% of our participants were right-handed. The distribution of the others can be seen in Figure 9. Although the proportion (and definition) of handedness in the general population is disputed<sup>184</sup>, most sources suggest that approximately 90% of people are right handed.

### **Education:**

Participants' education (see Figure 10) ranged from only elementary education (1 participant) to advanced graduate degrees, including medical and law school. Most participants (90%) had attended at least some college, which does make their general education level slightly higher than the local population. 56.7% of the general Californian population has attended "some college or more"<sup>185</sup>. The elevated educational status may be accounted for by the techniques used to recruit participants. All participants were recruited through computer advertising (or word of mouth) and nearly all of our participants were comfortable using computers. We do not think this difference in education level is significant to the recall of autobiographical memories.

### **Income:**

The median income range for these participants was \$30,000 - \$60,000 per year. This is in keeping with local census findings. Per capita income for Californians was

---

<sup>184</sup> Oldfield, R. C. (1971). "The assessment and analysis of handedness: The Edinburgh Inventory." *Neuropsychologia* **9**: 97-113.

<sup>185</sup> Bauman, K. J. and N. L. Graf (2003). Educational Attainment: 2000, US Census Bureau.

\$22,711; median household income was \$47,493<sup>186</sup>. Participants were asked to report individual incomes, not household incomes, during our demographic collection; however it is difficult to ascertain how accurately they followed this instruction. A significant number of our participants were not (or no longer) married, so the majority of these data are likely to be individual incomes. The distribution of incomes in our sample was somewhat skewed toward lower income ranges (Figure 11). This is likely due, in part, to participant recruitment. Our participants were paid \$15/hour for their time, without reimbursement for travel. Although we would like to believe that at least some of our participants participated because of a genuine interest in research, we suspect that many of them were interested in participating to make extra money. Given this, we expected to (and did) recruit more participants with lower incomes. Despite this, 16% of our participants made more than \$60,000, supporting the claim that this is a representative sample of wage-earners in Southern California.

### **Occupations and Foreign Language Skills:**

Our participant pool also included nearly 50 different occupations, spanning a wide range of both white and blue collar jobs, including: teachers, financial analysts, musicians, hair dressers, a farm grower, a photographer and a substance abuse counselor. For a complete list, see Table 3. Participants were also culturally diverse. All participants spoke English as their primary language (by adolescence), but many (34%) were raised bi- or trilingual (see Figure 12). This exact statistic is not available for the Californian

---

<sup>186</sup> Income data is from Welniak, E. and K. Posey (2005). Household Income: 1999, US Census Bureau. Our participants were tested in 2004. Data from US Census 2005 are not yet completely available at the time of this writing.

population at large, but it is in keeping with the available statistic – 39.5% of Californians speak a language other than English at home.

**Ethnicity:**

Choosing ethnically diverse participants was of primary importance to creating a database that was representative. The ethnicity distribution of participants in our study closely matches the California state averages (see Figure 13.) Given the high proportion of Hispanic citizens (of several races) in Southern California, we chose to match our demographics to ethnicity instead of race alone. Normative data are taken from the US Census 2000<sup>187</sup>. The participant pool slightly over-represents Black and Hispanic participants (by 4.6% and 3.7%, respectively) and slightly under-represents Caucasian participants (by 6.5%).

**Marital Status and Children:**

The marital statuses of our participants were varied (see Figure 14). The number of married and partnered participants in the database was smaller than the average population<sup>188</sup> (27% versus 52.3%) and the number of divorced participants was higher than average (24% versus 9.5%). Although the proportions of the marital statuses of our participants were not particularly close to the state demographics, we feel that the sample is still a good representation of the general population. All marital status categories (single, married, partnered, divorced, separated and widowed) were represented. The variance in these data is likely attributable to the size of the sample. 53% of our

---

<sup>187</sup> Data taken from "US Census Bureau American FactFinder", from <http://factfinder.census.gov/>. All data is from US Census 2000.

<sup>188</sup> Kreider, R. M. and T. Simmons (2003). Marital Status: 2000, US Census Bureau.

participants had at least one child (see Figure 15). (One participant had 7 children.) This number is in line with the state averages; 56.8% of families in California have children<sup>189</sup> and the average family size is 3.43 people.

### **Sexual Orientation:**

Our participant pool included people of most sexual orientations. The population frequency data for sexual orientation are widely disputed in current literature. A review of the literature in *The New England Journal of Medicine* suggests that between 2 and 6 percent of the adult population are exclusively homosexual or bisexual<sup>190</sup>. Many sources suggest that these reported numbers may be low, due to the social stigma associated with identifying as homosexual or bisexual. Our database includes participants who reported their sexual orientation as homosexual (4%) and bisexual (3%) as well as several participants who chose not to respond (6%). (See Figure 16.) One of the findings of the well-known Kinsey study<sup>191</sup> was that sexuality cannot always be defined by a single parameter, like “sexual orientation”. We also collected attractiveness ratings and degree of sexual desire on a Likert scale (1 – 10). These findings were in concert with the sexual orientation rating (see Figure 17). Despite the controversy about the exact frequency of homosexuality in the population, the frequency of various sexual orientations in our database seems to be in line with most findings, suggesting that our participants are representative on this demographic dimension.

---

<sup>189</sup> Census data is complicated here because not all of our participants would qualify as “families” by census data. Nonetheless, these numbers are probably representative.

<sup>190</sup> Friedman, R. C. and J. I. Downey (1994). "Homosexuality." *New England Journal of Medicine* **331**(14): 923-930.

<sup>191</sup> Gebhard, P. H. and A. B. Johnson (1998 ). *The Kinsey Data: Marginal Tabulations of 1938-1963 Interviews Conducted by the Institute for Sex Research* Bloomington, IN, Indiana University Press.

### **Religious Views:**

Participants were also religiously diverse. Since religion and morality are often entangled, we felt it was particularly important to collect data about participants' religious beliefs and practices. Data were collected on religious upbringing, current religious practices and degree of participation in religion. Figure 18 shows a trend towards religious exploration in our participants during their adult lives. A number of participants (19%) participated in multiple religions or non-conventional spiritual practices. To gauge the effect of religion on our participants' lives (Figure 19), they were asked to answer a free-response question about their religious upbringing and their current religious practices. These responses (usually 2-3 sentences for each question) were scored by one reader, to give each participant a 'degree of religious influence' rating. 92% of our participants had some degree of religious influence in their lives, either during their childhood or through their current religious practices.

Our participants were generally well-matched to state averages in the distribution of religious beliefs (Figure 20). Overall, Catholics seem to be under-represented, but we believe that some of this discrepancy is subsumed by the 'Other Christian' group. Participants were performing a free-response task when asked about religion and we believe that some of them may simply have written 'Christian', without specifying 'Catholic'. While reliable data about the proportion of Californians who practice many of these religions are not always available, the Center for Religious and Civic Culture<sup>192</sup> in

---

<sup>192</sup> "USC Center for Religious and Civic Culture." from <http://www.usc.edu/schools/college/crcc/demographics/lacounty.html>.

conjunction with the Glenmary Research Center<sup>193</sup> provides us with some statistics we can use to compare our sample to the general population (Table 4). In our sampling, we did not have any participants who were Muslim, although 1% of the population in California report being Muslim. We did, however, recruit several Buddhists (3%) and it appears that there are nearly 3 times as many Buddhist temples in Los Angeles County as Muslim mosques. The numbers of Muslims mosques and Hindu temples are fairly equal and we did not manage to recruit any participants from either religion, suggesting that our sampling of the population is relatively uniform.

### **Political Views:**

Lastly, our participants were diverse in their political views. Participants' political value systems were assessed using a questionnaire; those methods are described in Chapter 2. Figure 21 shows that 50% of our participants were best affiliated with liberal values, 40% were moderate and 10% espoused conservative beliefs. The average rating for our overall participant pool was slightly skewed to the left (mean =  $2.4 \pm 0.67$  on scale of 1 (Conservative) to 3 (Liberal)). In 2004, 43.2% of voters in California were registered as Democrats and 35.7 % were registered as Republicans<sup>194</sup>. Since the questionnaire used to test political values included questions that are divisive for some Republicans (like abortion), we suspect that some of our 'moderate' participants are likely to be affiliated with the Republican Party.

---

<sup>193</sup> (2000). Religious Congregations and Membership in the United States. Nashville, TN, Glenmary Research Center. Description of this group and these statistics:  
<http://www.usc.edu/schools/college/crcc/demographics/>

<sup>194</sup> "Report of Registration: January 2, 2004." from [http://www.ss.ca.gov/elections/ror/regstats\\_01-02-04.pdf](http://www.ss.ca.gov/elections/ror/regstats_01-02-04.pdf).



One caveat worth noting is that our participants were selected to have a somewhat uniform cultural context. They were all residents of Southern California who had lived in the area for at least 15 years. The intent of this restriction, originally, was to ensure that our database of memories had some uniformity in terms of social mores, societal rules and general tenets of behavior. Given the diversity demonstrated in the demographics, we feel that even given this restriction, we have constructed a broad, representative sample of Southern Californian society.

## **Part 2: Neuropsychology**

Characterization of our participants' demographic profiles provides a detailed picture about where our participants come from and the groups they belong to, but it tells us little about the individual differences among the members of these groups. One might expect that personality has a strong effect on moral decisions. People who worry more might make less risky choices; people who are happier might recall more positive memories. A battery of neuropsychological tests was administered to our participants to characterize them as well individually as we did demographically. Interestingly, the findings are the same with individual measures as they are with demographics – all kinds of people, nervous, extroverted, happy and sad, recall similar types of moral memories.

Our participants were also assessed on a wide variety of neuropsychological measures for three purposes: 1. To ensure that they did not differ from the general population in some way that was not encapsulated by demographics. Their IQs, personality types, social values and affect were measured, as well as their positions on an ethical diagnostic and performance on standard moral philosophy problems. 2. To

provide covariates for moral memories. Large group correlations allowed us to compare ratings from the memories with the neuropsychological features of the participants<sup>195</sup> and assess whether particular individual traits were correlated with specific types of memories. 3. To provide background data for future studies. As with the demographics, detailed information about the participants who provided these narratives will allow other researchers to select the best stimuli for their specific purposes.

### **Wechsler Abbreviated Scale of Intelligence:**

Each participant completed the Wechsler Abbreviated Scale of Intelligence (WASI) to measure their general intelligence. The Wechsler Adult Intelligence Scale-III (WAIS-III) is a widely used and highly validated measure which allows us to compare all participants in a standardized way (using IQ). In order to collect the most information in the shortest testing period, the 2-subtest version of the WASI was used to assess full-scale IQ (FSIQ), rather than the WAIS<sup>196</sup>. The 2-subtest version uses the Vocabulary subtest to assess an individual's verbal knowledge as well as "cognitive abilities, such as memory, learning ability, and concept and language development". It uses the Matrix Reasoning subtest to assess "nonverbal fluid reasoning and general intellectual ability"<sup>197</sup>.

The distribution of our participants' IQ can be seen in Figure 22. The mean IQ for our group was  $110 \pm 13.4$ , which is not significantly different from the published normative values for IQ ( $t[98] = 0.54$ ,  $p = 0.60$ ). IQ is used as a covariate throughout this

---

<sup>195</sup> While correcting for multiple comparisons, of course.

<sup>196</sup> Wechsler, D. (1999). Wechsler Abbreviated Scale of Intelligence. San Antonio, Psychological Corporation.

<sup>197</sup> Wechsler, D. (1999). Wechsler Abbreviated Scale of Intelligence. San Antonio, Psychological Corporation. p.4.

work to assess the effect that intelligence has on production of moral memories. As discussed throughout, we have found no impact of IQ on any feature of the narratives.

### **NEO Personality Inventory:**

The NEO-FFI personality inventory was used to evaluate participants' general personality types. Research using the NEO has shown that five major factors can be used to describe an individual's "emotional, interpersonal, experiential, attitudinal and motivational styles"<sup>198</sup>. The five factors are: Neuroticism, Extroversion, Openness, Agreeableness and Conscientiousness. Each is a measure of the range of normal personalities, so high or low scores should not be thought of as indicative of psychological problems. A more complete description of each personality type has been included in

---

<sup>198</sup> McCrae, P. T. C. a. R. R. (1992). NEO PI-R Professional Manual. Lutz, FL.p.14.

## Appendix F.

The NEO addresses the need for both diversity and representativeness in our participant pool by allowing us to assess the variety of personality types in our sample. Our participants do not differ from the NEO-FFI standards ( $t[98] = 1.38, p = 0.167$  (N);  $t[98] = 1.73, p = 0.083$  (E);  $t[98] = 1.30, p = 0.193$  (O);  $t[98] = 0.56, p = 0.564$  (A);  $t[98] = 1.01, p = 0.314$  (C)) on any of the five factors. (p values are uncorrected for multiple comparisons.)

The NEO personality factors lend themselves easily to hypotheses about moral decision making. The comparison between personality types and cue word is examined in Section 2, Part 1. In addition, personality types were used in a variety of correlations to consider the impact of individual differences on the narratives.

### **Schwartz Human Value Inventory:**

The Schwartz Value Inventory (SVI) was used to collect extensive data about the participants' cultural values. Since many people believe that culture affects morals, it was important to collect as much information about our participants' cultural values as possible. The SVI was developed for two purposes: to allow the cross-comparison of cultural values<sup>199</sup> (specifically those values shared by individuals in any given country) and to acquire the data necessary to develop a "theory of the basic content and structure of human values"<sup>200</sup>. In its two developmental studies, the SVI was used to survey nearly

---

<sup>199</sup> This is a goal shared by Milton Rokeach (discussed in: Rokeach, M. (1973). The Nature of Human Values. New York, Free Press.).

<sup>200</sup> Schwartz, S. H. (1992). "Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries." Advances in Experimental Social Psychology **25**: 1-65. and Schwartz, S. (1994). "Are There Universal Aspects in the Structure and Contents of Human Values?" Journal of Social Issues **50**(4): 19-45.

60,000 participants from more than 50 countries about their values (referred to during testing as the ‘guiding principles in their lives’). From this sample, 10 core values emerged: Power, Achievement, Hedonism, Stimulation, Self-direction, Universalism, Benevolence, Tradition, Conformity and Security. (Definitions for these values can be found in Table 6.) These 10 values can be arranged in a circular pattern to represent their relative positions on a continuum with each other<sup>201</sup>, shown in Figure 23. Schwartz describes these as motivational values, with two focal dimensions: openness to change/conservation and self-enhancement/self-transcendence.

Our subjects completed the full SVI. Though there are major data collection projects underway using the SVI, most notably the European Social Survey<sup>202</sup>, there are not yet “standards” for citizens of the United States or Californians on the SVI. Given the diversity of American culture, it seems unlikely that the entire country would show a uniform pattern of value loadings, in any case. Despite the lack of normative data, we can make some interpretations about the SVI data collected. As shown in Table 6, the three most highly rated values for our subjects were Self-Direction<sup>203</sup> (independent thought and action), Universalism (understanding, appreciation, tolerance and protection for the welfare of all people and for nature) and Benevolence (preservation and enhancement of the welfare of people with whom one is in frequent personal contact). Schwartz identifies five values as ones that deal with “the needs of individuals as biological organisms”. In light of the stereotype of Americans as individualists, it is interesting to note that four of

---

<sup>201</sup> For detailed definitions, see Schwartz, S. H. (1992). "Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries." *Advances in Experimental Social Psychology* **25**: 1-65. p.5-12. or Table 6.

<sup>202</sup> (2007). "European Social Survey." from <http://www.europeansocialsurvey.org/>.

<sup>203</sup> The definitions in parentheses are taken from Schwartz, S. (1994). "Are There Universal Aspects in the Structure and Contents of Human Values?" *Journal of Social Issues* **50**(4): 19-45. p. 22.

these five ‘organism’ values appear as the top four values for our subjects. The lowest rated values were Conformity (restraint of action, inclinations and impulses likely to upset or harm others and violate social expectations or norms) and Power (social status and prestige, control or dominance over people and resources). It is also interesting to note that Power, a value that would be stereotypically attributed to Americans is the lowest rated value on our survey. It may be that subjects in our study were self-selecting on this dimension to some extent. There are three factors that may have influenced this. Our subjects were available during working hours on weekdays, interested in participating in research and/or in need of income enough to respond to our advertising. These factors may have contributed to recruiting participants who were not as ‘power-driven’ as the average American stereotype. It is also possible, of course, that the stereotype is merely incorrect and Americans are not as driven by power as one might believe.

Although the SVI data may not provide a great deal of insight about the participant population as a whole, the value dimensions of the SVI may, at a later date, be very informative in selecting stories from subjects with particular profiles. One can imagine comparing the moral choices of subjects with very different Hedonism or Conformity ratings, for example.

### **Positive and Negative Affect Schedule:**

Each participant completed the Positive and Negative Affect Schedule (PANAS) to assess their general mood. The PANAS provides a valid and reliable measure for assessing participant affect. It is based on a two-factor model of affect which finds that

‘positive affect’ and ‘negative affect’ are the two primary dimensions of mood. Despite the intuitive sense that these factors might be opposites, they are, in fact, orthogonal. Positive Affect (PA) “reflects the extent to which a person feels enthusiastic, active and alert.” Negative Affect (NA) “is a general dimension of subjective distress... [and] aversive mood states, including anger, contempt, disgust, fear and nervousness.”<sup>204</sup> (Descriptions of each affect profile can be found in Table 7). The PANAS is stable over a 2 month time period and highly internally consistent, providing an excellent assessment of the participant’s current mood state.

Although we want the database to reflect the full range of the general population, there is research suggesting that emotion affects recollection. Cutler, Larsen and Bunce<sup>205</sup> found that people who display repressive coping techniques show inhibited recall for unpleasant memories. van den Berg, et al. were able to demonstrate that affective focus can modulate subjects’ recall (for affective versus cognitive stimuli)<sup>206</sup>. Most importantly, McFarland and Buehler<sup>207</sup> found that subjects who ruminate on their feelings while recalling autobiographical memories will generate memories with similar moods. Thus, subjects who have strong negative feelings (especially those who are depressed) may recall different memories than those subjects who are feeling more positive. Given these findings, subjects were assessed using the PANAS to ensure that their general affect was within normal ranges. Since our subjects had previously been screened for psychiatric

---

<sup>204</sup> Watson D, C. L., Tellegen A. (1988). "Development and validation of brief measures of positive and negative affect: the PANAS scales." *Journal of Personality and Social Psychology* **54**(6): 1063-1070.

<sup>205</sup> Cutler, S. E., Larsen, R. J., Bunce, S. C. (1996). "Repressive Coping Style and the Experience and Recall of Emotion: A Naturalistic Study of Daily Affect." *Ibid.* **64**(2): 379-405.

<sup>206</sup> van den Berg, H., Manstead, A.S.R., van der Pligt, J., Wigboldus, D.H.J. (2006). "The impact of affective and cognitive focus on attitude formation." *Journal of Experimental Social Psychology* **42**: 373-379.

<sup>207</sup> McFarland, C. B., Buehler, R. (1998). "The Impact of Negative Affect on Autobiographical Memory: The Role of Self-Focused Attention to Moods." *Journal of Personality and Social Psychology* **75**(6): 1424-1440.

conditions and medications, we did not encounter any subjects who needed to be eliminated from the study due to PANAS scores<sup>208</sup>.

Using the normative data found in Watson, Clark and Tellegen<sup>209</sup>, we found that our subjects did not differ significantly from normative data on either the PA ( $t[98] = 0.40, p = 0.69$ ) or the NA scale ( $t[98] = 0.14, p = 0.89$ ). (See Table 8 for complete data.) Overall, the distribution of affect scores was quite broad (one participant had a positive affect of 49 out of 50).

### **Ethical Position Questionnaire:**

The Ethical Position Questionnaire<sup>210</sup> (EPQ) was used to assess subjects' ethical ideologies. The EPQ measures subjects' idealism and relativism and classifies subjects as belonging to one of four ethical ideologies: situationists, subjectivists, absolutists or exceptionists. For descriptions of each ideology, see Table 9. This measure was developed to explain individual differences in moral judgments in subjects who seem to be otherwise similar. (This work stems from the writings of F.C. Sharp<sup>211</sup>, a psychologist who worked at the turn of the 20th century.) Forsyth notes that a person's ethical ideology likely influences their moral judgments however, "the relationship between

---

<sup>208</sup> We feel strongly that the PANAS is an excellent measure to include in this study, especially with respect to future users choosing narratives based on PA or NA score. In retrospect, however, given the frequency in depression in the general population, it would have been prudent to administer the Beck Depression Inventory (BDI) as well, to ensure further depression screening. BDI reference: Beck, A. T. (1972). *Depression: Causes and Treatment*. Philadelphia, University of Pennsylvania Press.

<sup>209</sup> Watson D., Clark, L., Tellegen A. (1988). "Development and validation of brief measures of positive and negative affect: the PANAS scales." *Journal of Personality and Social Psychology* **54**(6): 1063-1070.

<sup>210</sup> Forsyth, D. (1980). "A Taxonomy of Ethical Ideologies." *Journal of Personality and Social Psychology* **39**(1): 175-184.

<sup>211</sup> Sharp, F. C. (1898). "An Objective study of some moral judgments." *American Journal of Psychology* **9**: 198-234.



ideology and behavior is more tenuous.” This test was included in the battery for two reasons. First, we were interested in the interplay between ideology and the types of stories our subjects related. Second, Forsyth showed that the EPQ is not correlated with Kohlberg’s developmental stages. Since most adults fall into the same developmental stage using Kohlberg’s models, it was our hope that it might help distinguish the moral positions of our subjects in a way that moral development models could not. While the EPQ allows us to segregate our subjects to some degree, it did not solve the latter problem. Like Kohlberg’s stages, most of our adult subjects (79%) espoused the beliefs of one ideological category: situationism (see Figure 24). Forsyth defines situationists as those with high relativism and high idealism. They tend to “distrust absolute moral principles and argue instead that each situation must be examined individually.”<sup>212</sup>

There are a few explanations for why our subjects might have been grouped primarily into one category. The ethical position score can be broken down into its two component parts: the idealism score and the relativism score. While the relativism score is normally distributed (see Figure 25), as would be expected, the idealism scores were skewed to the high end of the scoring range (see Figure 26). Fortunately, this does not necessarily suggest that our subjects are not ‘normal’. Forsyth has recently found an increase in idealism scores across all groups, as well as increases in idealism with age. He suggests that this may be a combination of “developmental tendency (people get more idealistic as they age), [and/or] widespread social shift (people are getting more idealistic, overall).”<sup>213</sup> Additionally, the 5 point scale used in our study may compress scores more than the 9 point scale used in older studies. Regardless of the reason, unusually high

---

<sup>212</sup> Forsyth, D. (1980). "A Taxonomy of Ethical Ideologies." *Journal of Personality and Social Psychology* 39(1): 175-184.

<sup>213</sup> Personal correspondence with Forsyth, D. (November 2006).

idealism scores overbalance the number of situationists and absolutists, as our data show. After some investigation, we would suggest that our subjects are “normal”, but the test is not well validated for different age groups. Forsyth tested the EPQ on college age students (mean = 21 years old  $\pm$  3.9 years). Three recent studies using the EPQ with adults all found that most subjects were situationists, followed by absolutists. (These are the two categories we would expect to be overweighted if the idealism scale is skewed to the left.) By comparison, studies using college students seem to have normally distributed idealism scores<sup>214</sup>.

These findings suggest that the EPQ data are most useful as the individual Idealism and Realism scores, rather than the overall Ethical Position. These measures are excellent covariates for looking at individual differences among moral decision makers or choosing narratives that are likely to have more idealistic styles. Future work may also include determining the extent to which Forsyth’s intuition (that moral ideology and moral behavior are not always congruent) is true, using EPQ scores and narratives from the database.

### **Moral Dilemmas:**

As discussed in Chapter 1 and Appendix A, the moral philosophy problem referred to as ‘the trolley problem’ and its application in moral neuroscience is a major driving force of this work. To that end, we felt that it was important to collect data from our subjects on a set of standard moral dilemmas. (The full text of these dilemmas can be

---

<sup>214</sup>Wuensch, K. L., & Poterat, G. M. (1998). "Evaluating the morality of animal research: Effects of ethical ideology, gender, and purpose." *Journal of Social Behavior and Personality* **13**: 139-150.

found in Appendix D.) 33 moral dilemmas<sup>215</sup> were used to assess non-moral decision making (n = 10), non-personal moral decision making (n = 10) and personal moral decision making (n = 13), including the trolley and footbridge problems and a ‘Sophie’s choice’ dilemma.

The comparison of these moral dilemma data with the narratives contained in the database will be a rich, scientific endeavor; however, it is outside of the scope of this work. Nonetheless, quantifying the data for inclusion in the database is important for future stimuli selection as well as for using as covariates in our analyses. 19 of the 33 dilemmas were chosen for developing a ‘utilitarian’ score<sup>216</sup>. Each dilemma was scored based on a weighted measure of the degree of utilitarian decision making that it elicited and each participant received a total utilitarian score between 0 and 42 (most utilitarian). Our subjects’ utilitarian distribution (see Figure 27) was somewhat left-skewed (toward the less utilitarian end of the scale) (skewness = 0.72). The mean utilitarian score was  $13.2 \pm 7.44$  (range = 2 – 33), on a scale with a maximum score of 42. Given the uniqueness of this scoring system, it is unknown how typical this may be of the general population.

Another score which may be worth quantifying in the future is the emotional component versus the cognitive component of these decisions. Greene reported differential brain activations when subjects were considering emotionally incongruent responses<sup>217</sup>. Our data do not include a timing component however, it may be possible to generally categorize them based on typical response patterns (i.e., subjects who say that it

---

<sup>215</sup> Dilemmas were provided by J. Greene as published in: Greene, J. D., Sommerville, R. B., et al. (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." *Science* **293**: 2105-2108.

<sup>216</sup> This methodology was developed by J. Woodward.

<sup>217</sup> Greene, J. D., Sommerville, R. B., et al. (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." *Science* **293**: 2105-2108.

is acceptable to push the man off the footbridge might be given some sort of cognitive score increase). The value of these data in selecting moral narratives for new stimuli is likely limited (and thus not included in the database) but could be interesting when investigating the similarities (or differences) between subjects' responses to moral dilemmas and their moral behaviors and decisions documented in the narratives.

### **Correlations between Demographics and Neuropsychological Measures:**

To examine the breadth of the data in an exploratory manner and assess the possible correlations among all the demographic and neuropsychological measures, we calculated Pearson correlations for all of the background measures. Figure 28 and Figure 29 show these matrices. There are some weak correlations that appear in the matrix, for example, NEO N score appears to be correlated with age ( $R^2 = 0.24$ ,  $p < 0.01$ ) and Utilitarian score (derived from the standard philosophy dilemmas) appears to be correlated with Relativism score (from the EPQ) ( $R^2 = 0.25$ ,  $p < 0.01$ ). However, after correcting for multiple comparisons (the Bonferroni correction<sup>218</sup>), none of these correlations remain significant. This lack of significance suggests that the demographics have no effect on production or recollection of moral memories. (This can be seen in Figure 29.) It suggests that the moral events encountered, at least in this sample, are fairly universal. We provide these figures for qualitative interest; perhaps they will spark interest in future, more directed hypotheses.

## Section 2: Characterization of the Moral Narratives

---

<sup>218</sup>Abdi, H. (2007). Bonferroni and Sidak corrections for multiple comparisons. Encyclopedia of Measurement and Statistics. N. J. Salkind. Thousand Oaks, CA, Sage.

## Part 1: Analysis by Cue Word

Fundamentally, this database is comprised of moral memories elicited by a cued recall protocol. It is necessary then, to begin an examination of the database with an examination of the cue words used to elicit these moral memories. We will demonstrate throughout this dissertation that the variance in moral memories is dependent on the life experiences of the story tellers themselves and not their demographic differences. The cue words used in the cued recall protocol were selected from common English words to elicit as many moral memories as possible. In this section, we will examine the degree of success that we achieved in eliciting moral memories with these cue words. We will also break down the demographics and one personality measure by cue word in order to demonstrate that the cue words did not, in and of themselves, affect the memories generated.

### Cue Word Distribution:

The cue words chosen were quite effective: 23% (n = 758) of the memories recalled were moral. All of the cue words elicited at least 1 moral memory and 3 of the 4 control cues elicited no moral memories. (The cue 'happy' elicited 7 moral memories, suggesting that it was not truly a control.) Some cue words, however, were more effective than others. Figure 30 shows the frequency distribution for the moral narratives in the database. The cues that elicited the most moral narratives are: *took something that didn't belong to you* (n = 78), *cheated* (n = 69), *lied* (n = 60), *were sneaky* (n = 56) and *guilty* (n = 52). The narratives elicited by these cues comprise 42% of the narratives in the database. The cues that elicited the fewest moral memories are: *embarrassed* (n = 1), *best*

*thing you ever did* (n = 4), *bittersweet* (n = 4), *thing you'd most like others to know that you did* (n = 5) and *proud* (n = 5). These narratives make up only 2.5% of the database.

While the cues were almost all good elicitors of moral memories, some features of the cues distinguished them. The action cues were significantly more effective elicitors of memory than the emotion cues ( $t[22] = 2.22, p = 0.0369$ ), a finding that seems intuitive. Concrete cues – like the verb phrases used for the action cues – evoke particular concrete events of the past. Emotional cues are more abstract, requiring the participant to re-experience the feeling (at least to some degree) in order to summon up an action or choice that caused the emotion. The cues were presented in groups based on cue type<sup>219</sup> and the order of presentation within each category<sup>220</sup> did not seem to affect the number of memories elicited (see Figure 31). The valence of the cue was significant in determining the number of moral memories elicited. Subjects also recalled more moral memories when cued with negative cues (e.g., *unfaithful*, *worst thing you ever did*) than they did when cued with positive cues (e.g., *compassionate*, *thing you are most proud of having done*) ( $t[17] = 2.27, p = 0.0363$ ). This finding prompted greater examination of the role that valence plays in moral memory recollection, which is covered in detail in Chapter 4, Section 5.

### **Demographic Impact:**

The first section of this chapter describes the demographic features of the participant pool in detail and emphasizes the normality of our subjects. These subjects are quite representative of the population at large. This representativeness, however, tells us

---

<sup>219</sup> See Chapter 2 for more detail.

<sup>220</sup> The cues were randomized within each section before being presented in a fixed presentation to all participants. The reasons for this are discussed in Chapter 2.

little about whether demographic measures can affect the way subjects perform in cued recall protocol tasks. Upon examination, we found that demographic measures were not influential on cue word usage. Men and women produced very similar numbers of moral memories (men's total  $n = 376$ ; women's total  $n = 382$ ) and also showed similar patterns of distribution of cue words (see Figure 32). The youngest (40-45 years old) and oldest (55-60 years old) subjects respond similarly to the cue words, showing little difference in their distribution patterns (see Figure 33) when normalized for the number of subjects in each group.

### **Effect of NEO Personality Types:**

To demonstrate that the cue words are robust across neuropsychological groups in their effectiveness to elicit memories, we have analyzed the effect of NEO personality type on cue word frequency. Using a median split (and eliminating subjects with a median score), we divided subjects into groups of 'High' and 'Low' scorers for each NEO factor (N, E, O, A, C)<sup>221</sup> and generated cue word frequency plots for each group. A Spearman's rank correlation coefficient was calculated to determine whether high or low score on any of the factors significantly impacted the frequency with which cue words elicited memories. All five factors showed strong correlation between the high and the low scoring groups, signifying that personality factor does not affect cue word usage. (Correlations are as follows (given as Spearman  $R^2$  values): Neuroticism = 0.940, Extraversion = 0.931, Openness = 0.925, Agreeableness = 0.934, Conscientiousness =

---

<sup>221</sup> For more detail about NEO types, see Appendix F.

0.938.) Cross correlations among the NEO factors (both high and low across all five factors) are all also highly correlated ( $R^2 \geq 0.88$ ).

### **Kucera-Francis Word Frequencies:**

Finally, in examining the frequency distribution of cue words, it is important to look back to the selection criteria themselves. The final selection of cue words was based on a number of criteria (described in Chapter 2) including the balance of positive, neutral and negative concepts, concreteness, imageability, meaningfulness and arousal. In addition, preference was given to cues with greater lexical frequency (as measured by the Kucera-Francis frequencies)<sup>222</sup>. The final distribution of cue words, however, includes words with frequencies ranging from 1 to 98, providing the opportunity to examine whether words that are more common in the English language elicit more moral memories than those that appear less frequently. However, when these data are graphed Figure 35 demonstrates clearly that there is no relationship between the frequency of a word's appearance in the lexicon and the number of memories that that cue elicited.

## **Part 2: Language Inquiry and Word Count Analysis**

The Language Inquiry and Word Count text analysis program was designed to “provide an efficient and effective method for studying the various emotional, cognitive, and structural components present in individuals’ verbal and written speech samples.”<sup>223</sup>

This program allows us to analyze the contents of each narrative in the database using an

---

<sup>222</sup> The use of the Kucera-Francis values is discussed in Chapter 2. Kucera, H. and W. N. Francis (1967). Computational Analysis of Present-Day American English. Providence, Brown University Press.

<sup>223</sup> Pennebaker, J. W., Chung, Cindy K., Ireland, Molly, Gonzales, Amy, & Booth, Roger J. (2007). The development and psychometric properties of LIWC2007. Austin, TX, LIWC.net. p. 7.



automated process. While automated text analysis of this sort cannot interpret complex moral narratives, it does provide an excellent way to quantify the overall structure of the narratives for comparison purposes. The LIWC analysis also allows us to compare our narratives to other samples in order to determine how typical they are. These data have been shown in Table 2 and demonstrate that the narratives are quite typical in their number of words per sentence, percentage of words found in the LIWC dictionary, number of words with greater than six letters and their overall grammatical content (as measured by function words).

LIWC2007 calculates the percentage of total text that is devoted to each of 80 variables. These categories include grammatical structures (e.g., nouns, verbs), personal motivators (e.g., work, religion) and psychological constructs (e.g., affective, social and cognitive processes). The affective processes variables are particularly useful for our analyses because they allow us to compare the percentages of positive and negative emotion words in each narrative. The dictionary of words that connote positive emotions contains 406 words (e.g., love, sweet) and the dictionary of words that connote negative emotions contains 499 words (e.g., hurt, nasty). The dictionary of 'negative words' can be further broken down into anxiety, anger and sadness related words but since 'positive words' is not further subdivided, we have chosen not to use these divisions for our analyses.

#### **LIWC Variable: Positive and Negative Emotion Words:**

Using the LIWC text analysis, we examined the percentage of emotion words (both positive and negative) found in each narrative. The distribution of these data can be

seen in Figure 35. The average narrative in the database contained  $2.7 \pm 1.85\%$  positive words and  $2.0 \pm 1.61\%$  negative words, a significant difference ( $t[1514] = 7.10, p < 0.0001$ ). This is a particularly surprising finding since there are more negative<sup>224</sup> narratives in the database.

Additionally, Figure 35 shows that 7.9% ( $n = 60$ ) of the narratives in the database contained no positive emotion words and 11.2% ( $n = 85$ ) of the narratives contained no negative emotion words. 10 of these narratives overlap, hence they are entirely devoid of emotion words as found in the LIWC dictionaries<sup>225</sup>.

When analyzing the data, we noticed that the narratives with the highest percentage of emotion words (both positive and negative) were particularly short. (The narrative with the highest percentage of positive emotion words (15.22%) was 92 words long and the narrative with the highest percentage of negative emotion words (11.59%) was 138 words long. The average word count for narratives was 217.6 words.) Examination of these two narratives found that each participant repeated the cue word (or versions of the cue word) frequently<sup>226</sup>. We hypothesized that this might have happened whenever a participant had difficulty recalling a memory that fit the cue word, accounting for both the high emotional word count and the length of the narrative. In order to make sure that short narratives were not being disproportionately represented in the word count analyses, we inspected the data further. The word count of the 100 stories with the highest percentage of positive emotional words is not significantly different ( $t[198] =$

---

<sup>224</sup> Negative narratives, in this case, are narratives that were cued by negative cue words.

<sup>225</sup> It is important to point out that this is an important caveat about analyzing data based on LIWC tallies. The program is only as accurate as its dictionaries are complete. The number of positive and negative emotion words is large (905 words combined) but it cannot contain all emotion words. More importantly, it cannot understand the nuance that may be imparted by the context of the narrative. However, it is this difficulty that is weighed against the difficulty of acquiring readers who can analyze each narrative individually.

<sup>226</sup> These are memories 5588 and 5441.

1.82,  $p = 0.07$ ) from the word count of the 100 stories with the lowest percentage of positive emotional words (including those stories with no positive emotional words).

### **Demographic Impact:**

Similarly to the cue word analyses done in the previous section, we were interested in the influence that gender and neuropsychological measures might have on the number of emotional words in a narrative. While there is a wealth of literature<sup>227</sup> to lead us to expect that men and women use emotion words with different frequencies, they are not statistically different for either positive emotional words ( $t[117] = 1.19, p = 0.24$ ) or negative emotional words ( $t[168] = 0.15, p = 0.88$ ). (These comparisons and those following were done by selecting the narratives with no emotional words ( $n = 60$  for positive and  $n = 85$  for negative) and comparing them with the same number of narratives with the highest percentage of emotional words in each group. The outliers discussed above were included in these analyses.) This finding is quite advantageous for using the narratives as new stimuli. Since there are no underlying gender cues in the emotional word usage, studies involving gender can use the narratives interchangeably. For an example of this, see Section 4, later in this chapter.

Research by McFarland and Buehler has found that when subjects involved in autobiographical recall are ‘ruminating’ on their own mood, they tend to recall memories with similar emotions. However, if instead the subjects are ‘reflecting’ on their own mood, they tend to recall memories with dissimilar emotions<sup>228</sup>. These data suggest that

---

<sup>227</sup> As one example: Brody, L. R. Hall, J.A. (2000). Gender, Emotion and Expression. *Gender Handbook of Emotions*. M. H.-J. Lewis, J.M. New York, Guilford Press: 338-347.

<sup>228</sup> McFarland, C. Buehler, R. (1998). "The Impact of Negative Affect on Autobiographical Memory: The Role of Self-Focused Attention to Moods." *Journal of Personality and Social Psychology* 75(6): 1424-1440.

subjects' mood may affect the emotional content of the narratives. The Positive and Negative Affect Schedule measures the general mood of a participant using two orthogonal scales: Positive Affect (PA) and Negative Affect (NA)<sup>229</sup>. Together, these scales provide a general picture of a participant's mood. Comparison of PA and NA scores for the subjects who told the narratives with the greatest percentage of emotional words with those for the subjects whose narratives contained no emotional words showed no significant difference for either positive emotional words (for PA,  $t[117] = 0.14$ ,  $p = 0.89$ ; for NA,  $t[117] = 0.42$ ,  $p = 0.68$ ) or for negative emotional words (for PA,  $t[168] = 0.36$ ,  $p = 0.72$ ; for NA,  $t[168] = 1.81$ ,  $p = 0.07$ ).

### **Part 3: Analysis by Rater Categories**

The automated analyses of the LIWC program provide tools for examining the narratives' content, but only human readers can provide insight about the details of the narratives themselves. The descriptive figures in this section give a window into the content of the narratives as only could be described by readers. These readers, a group of 55 anonymous raters, were recruited to categorize the narratives. (The details of this process are given in Chapter 3.) The ratings collected describe the actions in the narrative, the rationales given for various moral behaviors, and the emotions that these scenarios elicited in the participant. The rating sheet was designed to collect these data using three major categories: 'What did they do?', 'Why did they say they did it?' and 'How did they feel?' Using traditional journalistic style questions made it easy for the

---

<sup>229</sup> For more details see Chapter 4, Section 1.

raters to read each narrative in the context intended. Graphing the resulting data also makes it easy to get a sense of the overall contents of the database.

Figure 36 shows the actions that appear in the narratives. Most prominently, there are 533 ratings about narratives where the narrator lied and 507 ratings where the participant was helping someone. The raters were asked to use a scale from 1-3 when categorizing the ratings to designate how well the narrative exemplified the category. They were also asked to note stories in which the participant specifically chose to do the opposite of the category. For example, if the narrator decided not to cheat on a business deal, the rater would give the category 'cheating in business' an opposite rating of 1-3. This allows future users of the database to more easily select narratives where the narrator chose not to make a bad moral decision<sup>230</sup>. The rating categories are shown as a stacked bar graph (Figure 36 - Figure 38) to display the contribution of each to the total number of ratings in that category. The reasons narrators most commonly gave for their actions can be seen in Figure 37. They were motivated by selfishness (n = 702 ratings) and doing the wrong (n = 646) and right thing (n = 596). Finally, raters were asked to categorize how the narrators felt about their actions (shown in Figure 38), where we find that subjects mostly felt regret about their behavior (n = 893) and doing the wrong thing (n = 835). These emotions are not surprising since all of the narratives concern morally laden events. Along with categorizing the feelings of each narrator, the raters were asked to rate how much rationalization was given for the narrator's actions. In 1565 of the 2169 ratings (72.2%), the narrators provided some sort of justification for their behavior.

In designing the rating sheet, a subsample of 100 narratives was used to generate the categories included. During that process, the hypothesis arose that subjects might

---

<sup>230</sup> It also allows users to choose narratives in which the subject specifically chose *to* do the wrong thing.

change their moral behavior based on the location of their actions<sup>231</sup>. Generally, it seemed that subjects had different moral standards for different situations (work versus home, situations where children or animals were involved, etc.). In addition, knowing about where the narratives took place (as well as salient characters in those locations) provides another way to ‘see’ into the entire database at once. A group of 8 ‘scenes of action’ were developed from the subsample and were included on the rating sheet. Table 10 displays the results of those ratings. Only 11.6% (n = 251) of the ratings had no scene chosen while 37.5% (n = 831) of the ratings found that the narrative involved a family member or close friend. Not surprisingly, 29.2% of ratings also found that the narrative took place at home or in the narrator’s neighborhood. Raters were encouraged to write in other scenes that seemed important to the narrative. These scenes are enlightening in themselves, providing insight about other places where important moral decisions occur including places such as jail, hospital, rehab and on the internet.

After categorizing each narrative, raters were asked to give their own opinion<sup>232</sup> of how right or wrong the action in the narrative was (using a scale from -5 (wrong) to +5 (right)). Raters were told to use ‘0’ to represent a morally ‘gray area’ and to leave the rating blank or note if they felt that the narrative did not involve a moral situation. Only 3.3% of the ratings (n = 72) were thought by some raters not to be moral<sup>233</sup>. Figure 39 shows two different representations of the distributions of the raters’ judgments. The first

---

<sup>231</sup> One example of this is Memory 5096, a narrative where the subject decides while on vacation abroad to have sex on the beach with her boyfriend even though it is illegal. She regrets her decision in part because she knew she wouldn’t have to live with the consequences but her foreign boyfriend did.

<sup>232</sup> As was discussed in Chapter 3, this is NOT an unbiased sample and these data about the morality of the ratings should only be used as a way to separate the data. Any ratings for comparison to other groups should be done with an unbiased sample group.

<sup>233</sup> It is important to remember that some memories were chosen for inclusion in the database even though neither researcher reading the memories felt it was moral. This occurred when the subject specifically said that they felt the action was right or wrong. The database criteria were intentionally broad enough to allow the inclusion of these narratives.

(represented by the blue bars) is a distribution that shows the number of ratings that received the score shown on the x-axis. Thus, 130 ratings (5.9%) received a score of -5 from any rater. The red bars represent the number of narratives with a mean score as shown on the x-axis. There are 7 narratives that received a mean score of -5. This means that every rater believed that these narratives<sup>234</sup> were as morally wrong as the scale allowed. Together these distributions show us that individual raters rate similarly to the group, although the extremes (-5 and +5) are not used quite as reliably, which is to be expected with Likert scale ratings.

Any generalized description of these 758 moral narratives is complicated by the individuality and complexity of each narrative (as well as the number of narratives). However, using the categorizations generated by the raters provides an excellent overview of the contents and breadth of the narratives.

#### **Part 4: Analysis of Session Three Data**

When participants returned to the lab for Session Three, they were asked to respond to 54 rating questions about each of their own moral narratives. (The methods for this are described in Chapter 2; the questionnaire administered is found in Appendix C). These ratings allowed us to investigate a number of open questions about moral perceptions of self, including looking at accuracy of self-perception and the effects of demographic and neuropsychological variables. These results are most easily interpreted

---

<sup>234</sup> These narratives include instances of shoplifting, cheating in a relationship and stealing.

when they are framed by the questions that initiated them. To that end, these questions will be used as headings throughout this section.

Among the ratings collected in Session Three were three questions in which participants rated (on a scale from 1 to 10) the morality of their own narratives. They were asked to reflect on the morality of their behavior in the narrative at two time points: first, thinking back to their feelings at the time and second, looking back on the situation now (i.e., on the day of testing)<sup>235</sup>. For the third rating, the participants were asked to judge how others reading the narrative would rate the behavior in the narrative. These three ratings allow us to investigate whether participants' ratings of themselves change over time and how critical participants are of their own behavior. By comparing these ratings to those generated by the raters who read these narratives we can also examine whether the participants' own ratings are similar to outsiders' ratings as well as how accurate participants are in assessing how others judge them<sup>236</sup>.

### **Does the length of time since the memory occurred affect how moral the participant feels the action was?**

The saying that 'time heals all wounds' alludes to the idea that over time painful memories can become less painful. Chapter 4, Section 6 discusses the effects of the age

---

235 The specific questions given were: "At the time, how wrong or right did your actions in this episode seem to you? Use a scale from 1 to 10. If, at the time you did it, you felt this was a very wrong thing to do, you would give it a 1. If, at the time you did it, you felt this was a very right thing to do, you would give it a 10." and "Reading this today, how wrong or right do your past actions in this episode seem to you? Use a scale from 1 to 10. If now you feel this was a very wrong thing to do, you would give it a 1. If now you feel this was a very right thing to do, you would give it a 10."

236 It is very important to note that these data can only be considered preliminary findings. As was discussed in Chapter 3, the raters were chosen to be a homogenous group (with respect to education). All other demographic information about the raters is unknown. Thus, these data are biased at least by education. However, the raters are still diverse in the ratings they gave and we feel confident using them for these preliminary analyses.



of a memory on its recollection. Knowing that a memory's remoteness can affect its recall prompted us to explore whether the age of a memory can change how the participant feels about their actions, as well. As shown in Figure 40, 45% of participants did not change their feelings about how right or wrong their actions were between the time of the narrative and the time of the rating. Those who did change their feelings did so in a normally distributed way across the entire range of possible changes. The mean change was to lower the rating by 0.5 points, but the variance in this was quite large ( $SD = \pm 2.93$  points). When we calculate the difference between the ratings of how participants felt at the time of the event and how participants felt at the time of the rating, averaged over 5 year periods, we can see that the difference in ratings changes very little for different time frames (Figure 41). This leads us to conclude that time for reflection affects participants' moral ratings about their actions very little.

### **Do subjects judge themselves differently than they think others will?**

#### **How accurate are subjects at predicting what others think of them?**

Finding out whether we rate ourselves differently than others rate us in any situation is fundamental to many social behaviors<sup>237</sup>. We were able to compare our subjects' own ratings of their narratives with ratings from other people, as well as with their own perceived ratings of others' ratings. We find that subjects' ratings of themselves are highly correlated with the rating they think others would give them ( $R^2 = 0.64$ ,  $p < 0.001$ ) as well as with the actual ratings that raters give them ( $R^2 = 0.63$ ,  $p <$

---

<sup>237</sup> Gossip is an excellent example of this.

0.001). The ratings they think others would give them and the ratings that they actually receive are also highly correlated ( $R^2 = 0.57$ ,  $p < 0.001$ ).

While these ratings are all highly correlated, the correlation does not reveal the constant difference in rating scores. We find that subjects rate themselves as significantly more moral than they are rated by others, rating themselves an average of  $0.78 \pm 1.47$  points higher on a scale from one to ten ( $t[757] = 7.16$ ,  $p < 0.0001$ ). (Similarly, subjects rate themselves significantly higher than they think others will rate them (mean =  $0.77 \pm 1.25$  points,  $t[757] = 6.51$ ,  $p < 0.0001$ .) Interestingly, despite rating themselves as more moral, subjects are very accurate at predicting the rating that others will give them. There is no significant difference between the prediction the subjects gave and the actual moral score given by the raters (mean =  $0.07 \pm 1.51$  points,  $t[757] = 0.47$ ,  $p = 0.64$ ).

These findings suggest that, in general, subjects have an excellent sense of their ‘moral standing’ in the world but still feel that they are more moral than that standing would suggest. There are, however, demographic and neuropsychological factors that can impact how morally a participant perceives him or herself.

### **Does one gender rate themselves as more moral than the other?**

When asked to rate how ethical they were on a scale from one to ten, all subjects rated themselves as quite ethical (mean =  $8.4 \pm 1.65$  points for women, mean =  $7.7 \pm 1.70$  points for men). However, the women’s self-ratings were significantly higher than the men’s ( $t[745] = 6.18$ ,  $p < 0.0001$ ). Ratings of the narratives by outside raters found no significant difference between the behaviors of men and women ( $t[756] = 1.87$ ,  $p = 0.06$ ), which is to be expected since the raters had no specific knowledge about the gender of

the narrator (only that that they assumed from the narrative) and, as reported in Section 5, raters were not particularly good at intuiting the gender of the rater from the narrative.

### **Are certain personality types more likely to be self-critical?**

One might expect that personality would affect how a person views their morality, especially if, as shown above, we are able to construct two views of morality – one that judges how we feel about ourselves and one that judges how the world views us. It appears that three personality factors can affect how critical a person is about their own moral position. Subjects who score high on the Neuroticism factor on the NEO-FFI personality measure are significantly more critical of their own morality than subjects who score low on this dimension ( $t[654] = 4.31, p < 0.0001$ ), rating themselves 0.6 points lower on a 10 point scale. One can imagine that subjects who worry more might judge themselves more critically than subjects who are less prone to worry. However, these subjects are still able to make accurate assessments of how others will actually rate them. (Subjects who score high on N and subjects who score low on N are not significantly different when comparing their own judgments of how others will rate them to actual ratings by others;  $t[84] = 0.39, p = 0.70$ ). This suggests that they are able to create two different assessments of their moral value, a personal (internal) one and an objective (external) one. Interestingly, subjects who score high on Openness and Agreeableness also are more critical of themselves than their low scoring counterparts ( $t[622] = 2.85, p < 0.01$  for O;  $t[631] = 5.18, p < 0.0001$  for A), rating themselves 0.4 points lower (for O) and 0.7 points lower (for A). While subjects with high O scores retain the ability to accurately judge how others will rate their behavior ( $t[92] = 0.18, p = 0.86$  for O),

subjects with high A scores do have difficulty with this assessment when compared to their low A peers ( $t[85] = 2.58, p = 0.01$ ). While the interpretation for this finding is less straightforward than for the N dimension, it may be that subjects who are more open and more agreeable are more empathetic and therefore less critical of other people, rating themselves lower by comparison to the rest of the social group. Subjects who are less open and less agreeable may be more judgmental of others and therefore rate themselves higher by comparison. More research would be needed to investigate the ramifications of personality difference on self ratings; including other measures of self-worth could expand this finding significantly.

### **Does IQ affect a participant's personal moral rating?**

Since we have shown that subjects are capable of maintaining two separate constructs about moral value, it seems important to measure whether a participant's intelligence has any impact on their ability to generate these two values. A linear regression shows that IQ has no significant effect on a participant's personal moral rating ( $R^2 = 0.03, p = 0.29$ ). Similarly, we find that IQ also does not affect a participant's ability to predict others' ratings of his or her own behavior ( $R^2 = 0.00004, p = 0.99$ ). Together, the findings demonstrate that IQ does not affect a participant's ability to generate separate moral ratings for their own perception of moral value and the perception of others about them.

### Section 3: Cluster Analysis Reveals Implicit Categories in Moral Narratives

As described in Chapter 3, the narratives contained in the moral narrative database were characterized in great detail by the raters, using categories generated from

the narratives themselves. In addition to these categories generated by researchers, we were interested in whether there might be an implicit category structure to these moral narratives. One might suspect that there are features that are fundamentally different between stories about positive events in a person's life and stories about negative events, for example. We hypothesized that the moral narratives could be separated into these and other less intuitive categories based on their content using mathematical clustering techniques<sup>238</sup>.

Cluster analysis is a technique that can use a large number of different algorithms to collect data into meaningful groups. In a two-dimensional group of data, clusters can be physically proximate in space – a parameter which is easily graphed to see the grouping; with a multi-dimensional data set, as in the case of these narratives with many ratings, it is difficult to meaningfully display or interpret the groupings of these items. Cluster analyses are designed to address this problem. Central to the goals of all cluster analyses is the notion of degree of similarity (or dissimilarity) between the individual items being clustered.

To begin, we used hierarchical clustering to generate a dendrogram (tree-diagram) of all of the narratives. All calculations were done using STATISTICA<sup>239</sup>. Hierarchical clustering builds a hierarchy of data clusters, usually displayed as a tree, by successively changing restrictions on what comprises a cluster<sup>240</sup>. To begin, each narrative is treated as a single leaf on a large tree; each item is its own cluster. We used agglomerative

---

<sup>238</sup> The work in this section was done in conjunction with a Agnieszka Leśniak, Jagiellonian University, Cracow, Poland.

<sup>239</sup> STATISTICA. Tulsa, OK, Stat Soft, Inc at <http://www.statsoft.com>.

<sup>240</sup> (2007). *Electronic Textbook: Cluster Analysis*. Tulsa, OK, StatSoft. Online at: <http://www.statsoft.com/textbook/stathome.html>.

methods<sup>241</sup> to repeatedly lower the threshold on how large a Euclidean distance<sup>242</sup> was necessary to comprise a cluster. Each successive change adds more narratives to a cluster, until all of the narratives are joined in a single cluster. As more items are added to a cluster, it becomes necessary to determine which items within the cluster are used to calculate the distance for the cluster. While there are several possible methods, we chose to use the complete linkage (farthest neighbor) rule. This method is used when the data are expected to cluster naturally, an assumption that is substantiated by our initial hypotheses. In this way, the dendrogram shown in Figure 42 was generated<sup>243</sup>.

Figure 42 shows two major clusters that remain quite distinct until the final linkage (distance = 13.3). The smaller of these two clusters (on the right side) contains 174 narratives. An initial examination of this cluster found that the ratings ‘helping someone’ and ‘protecting someone’ were particularly important in generating this cluster. By comparison, the larger cluster contains 575 narratives where ‘lying’ and ‘regret about doing the wrong thing’ were important ratings. Reading a subsample of the stories in each group confirms this finding: the smaller cluster contains stories about doing something good, while the larger cluster contains stories about doing something bad. (These will be referred to as the ‘Good cluster and the ‘Bad cluster’ throughout.) This finding substantiates our hypothesis that narratives about positive events are measurably different from narratives about negative events in a person’s life.

---

<sup>241</sup> These are often referred to as ‘bottom-up’ methods. Divisive (‘top-down’) methods can be used for hierarchical clustering as well. However, since only agglomerative methods were used here, they are the only ones discussed in the text.

<sup>242</sup> Euclidean distances are measurements of the geometrical distance between data points in multidimensional space.

<sup>243</sup> These data represent 749 of the 758 narratives in the database. At the time of these analyses, 9 narratives did not have sufficient data to be included.

The dendrogram also appears to show several smaller clusters within the Bad cluster. (These can be seen most clearly as the Euclidean distance approaches 10.) However, without as clear a division as there is between the Good and Bad clusters, it is difficult to know which clusters are significant. In order to more closely examine the importance of these clusters, we performed a factor analysis on each cluster. Factor analysis is a statistical technique used to explain variability among observed random variables in terms of fewer unobserved random variables called ‘factors’. It assumes that all the rating data on different attributes can be reduced down to a few important dimensions. There are 2 ways of evaluating the results of a factor analysis. The first is the Kaiser criterion, which states that the main factors are those whose minimal eigenvalue is greater than 1<sup>244</sup>. More simply, this criterion requires that each factor extracts at least as much data as any one of the original variables has alone. Using this method, we found 59 factors for the Bad cluster. The other method is the Scree test, which uses a graphical plot of the number of eigenvalues to determine the point at which most of the variance in the data has been explained. (See Figure 43.) Using this, we found 3 factors for the Bad cluster. In both cases, we found only 1 factor for the Good cluster; these data will be presented in comparison to the Bad cluster findings in the following figures. While both the Kaiser criterion and the Scree test are widely used, it is known that the Kaiser criterion has a tendency to retain too many factors under some circumstances<sup>245</sup>. Given

---

<sup>244</sup> (2007). Electronic Textbook: Cluster Analysis. Tulsa, OK, StatSoft. Online at: <http://www.statsoft.com/textbook/stathome.html>.

<sup>245</sup> It is also known that the Scree test sometimes retains too few factors. (2007). Electronic Textbook: Principle Components and Factor Analysis. Tulsa, OK, StatSoft. Online at: <http://www.statsoft.com/textbook/stathome.html>.

our findings, we chose to use the number of factors selected by the Scree test ( $n = 3$ ) in order to best understand what distinguishes these clusters<sup>246</sup>.

To obtain a clearer picture of the relationship among these three factors, we applied a varimax rotation. This technique allows us to compare the factors and find the orientation that best distinguishes each group. The factors in the Bad cluster are orthogonal to each other (as seen in Figure 44a-c, as well as in three-dimensions in Figure 44d.). Therefore, each factor describes a distinct group. This supports the claim that each of these factors can be associated with a unique cluster from the dendrogram. These data can be examined in contrast with the varimax rotation of the data from the Good cluster (Figure 45). The homogenous cloud corresponds with the single main factor found for this cluster.

One of the challenges with factor analysis is that after compressing all of the variables into a few factors it can be difficult to identify what each factor describes. Despite this difficulty, identifying the factors can be important to interpreting the data. In fact, in this case, it is the very reason for these analyses. Each of these factors represents a non-intuitive category that can segregate different types of negative narratives. Knowing what each factor represents will allow us to name these non-intuitive categories.

To begin analyzing what each factor represented, we generated word frequency clouds<sup>247</sup> for each of the three factors (

Figure 46 – Figure 48). To generate these figures, a large subset of the narratives from each Bad factor was chosen ( $n = 128$  from Factor 1,  $n = 121$  from Factor 2,  $n = 74$  from

---

<sup>246</sup> It is worth noting that if you examine the dendrogram in Figure 42 carefully, there are about 59 clusters at a Euclidean distance of about 4. While this is an interesting test of the two criteria, the sheer number of factors given by the Kaiser criterion would make the results uninterpretable.

<sup>247</sup> These word clouds were generated using the software from Feinberg, J. (2008). "Wordle." from <http://wordle.net/>. Sincere thanks to Jonathan Feinberg for this application.



Factor 3). Subsets were chosen as shown in Figure 50, by selecting narratives from the outer portion of each cluster. (The clusters were chosen to include roughly two-thirds of the total narratives, using multiple 3-dimensional views to ensure clear separation between clusters.) Choosing from the outside of each cluster allowed us to avoid the overlapping origin for all three clusters, as well as select a group of narratives that were distinct. All of the narratives found in the Good factor (n = 174) were included in these analyses.

Using the software at Wordle.net, the text from all the narratives in each group was counted, generating a word frequency analysis for each group. These data are shown in Figure 46 - Figure 49 as word ‘clouds’ – images that display the frequency of each word in the total text analysis as a function of the word’s size. Bad factor 3 (Figure 48) shows the power that this kind of visual representation can have in interpreting the data. A qualitative analysis (even a quick glance) shows that the most important words in Factor 3 are: took, something, money one, back, really, take. This is substantiated by the word frequencies shown in the table on Figure 48. The most frequent words<sup>248</sup> used in narratives that are found in Factor 3 are words about taking things and money. It seems that within the overarching category of ‘Bad’, narratives about stealing exist as a distinct group. Similarly, we find that Factor 1 describes narratives about lying and Factor 2 describes narratives about being hurtful to another person. Factor 2 is the most difficult to interpret using only the most frequent words – however, visual analysis makes the topic

---

<sup>248</sup> It is important to note that the software removes the most common words in the English language before creating the word clouds. In addition, the following words were removed from all of the word clouds: *memory, just, really, like, time*. These were removed because they seemed to be unrelated to the actual content of the memories, but rather were common to the speech patterns of the participants in all of the word clouds.

quite clear (as does looking further down the frequency chart). This may be because a more varied vocabulary is used in these narratives or simply because the situations were more heterogeneous. This heterogeneity can also be seen in Figure 49, the Good cluster word cloud. Careful visual inspection makes the positivity of the content clear, but the most frequent words: ‘money’, ‘know’, ‘felt’ and ‘people’ do not paint as clear a picture as they might if the narratives were more similar.

To further substantiate the concepts we have been able to associate with the three Bad factors, we also examined two other measures that describe the contents of these clusters. First, we examined the rating data from which these clusters were derived. Table 11 shows the rating categories with the greatest influence on each factor. This is shown as the sum of all of the rating values in that category. The concepts that we have associated with each of the three factors can be seen reiterated in the ten most influential ratings: lying for Factor 1; regret about behavior for Factor 2; and stealing/spending someone else’s money for Factor 3. Second, we looked at the cue words that elicited the narratives in each cluster. A table of the 5 most common eliciting cues for each cluster can be seen in Table 12; data for the Good narratives is included as well. The most common eliciting cue for Factor 1 is *lied*, for Factor 2 is *was hurtful to someone* and for Factor 3 is *took something*. Similarly, the most common cue for the Good narratives is *compassionate*.

These data establish that most memories can be categorized in one of two ways – good or bad. The group of good narratives that we analyzed is very homogenous and cannot be further divided (as seen in Figure 49). Bad memories, however, can be subdivided into three main and independent factors. Using three separate data sets, we

have found that lying, hurting someone else and stealing are distinct types of behavior when recalled as moral memories.

#### Section 4: Using the Database for Further Research

As discussed in Chapter 3, the narratives collected in the database can be used both as data about normal moral behavior and as stimuli for future studies. The first use of these data as stimuli was for a research project investigating the role of gender in moral judgments<sup>249</sup>. A number of psychologists, most notably Gilligan, have investigated the gender differences in the moral development of adolescents and young adults<sup>250, 251</sup>. Recently, Overman, et al. found that contemplation of Greene's standard trolley problem dilemmas could affect women's performance on the Iowa Gambling Task<sup>252</sup>. While these findings are striking, the same problems exist with the dilemmas being used by these researchers as with researchers studying non-gender related moral judgment – the stimuli are not based on real-life moral challenges. To address this concern and to investigate gender bias in moral judgment, narratives were chosen from the database to use as

---

<sup>249</sup> This work was conducted by Jessica Stockburger as a Summer Undergraduate Research Fellowship project.

<sup>250</sup> Gilligan, C. (1982). In a different voice: Psychological theory and women's development. Cambridge, MA, Harvard University Press.

Muuss, R. E. (1988). "Carol Gilligan's theory of sex differences in the development of moral reasoning during adolescence." Adolescence **23**: 229-243.

Bjorklund, F. (2003). "Differences in the justification of choices in moral dilemmas: Effects of gender, time pressure and dilemma seriousness." Scandinavian Journal of Psychology **44**: 459-466.

<sup>251</sup> For other work about gender and morality, see: Bjorklund, F. (2003). "Differences in the justification of choices in moral dilemmas: Effects of gender, time pressure and dilemma seriousness." Scandinavian Journal of Psychology **44**: 459-466.; Aldrich, D. and R. Kage (2003). "Mars and Venus at Twilight: A Critical Investigation of Moralism, Age Effects, and Sex Differences." Political Psychology **24**(1): 23-40. ; and Agerstrom, J. and K. Moller (2006). "Moral Reasoning: The Influence of Affective Personality, Dilemma Content and Gender." Social Behavior and Personality **34**(10): 1259-1276.

<sup>252</sup> Overman, W., L. Graham, et al. (2006). "Contemplation of Moral Dilemmas Eliminates Sex Differences on the Iowa Gambling Task." Behavioral Neuroscience **120**(4): 817-825.

stimuli. These narratives were modified to examine whether people rate moral decisions less harshly when they share the gender of the narrator.

### **Choosing Stimuli:**

31 potential stimuli were selected from the moral memory database using the following criteria: (a) Is the memory gender charged? (b) Did this memory describe a “romantic” relationship? (c) Can the gender in this story be switched while maintaining the integrity of the memory? A pilot study was done to ensure that gender-reversed memories were believable. From these data, 20 narratives were selected where the gender of the narrator could be switched.

During the creation of these stimuli, a question arose about whether a narrative could be ‘gender neutral’ and whether there were narratives in the database that were ambiguous about the gender of the narrator. To address these questions, the rating tool was designed to include a question about the gender of the narrator. The raters were asked to identify the gender of the narrator and if they were unsure about the gender, to indicate that this gender choice was a guess<sup>253</sup>.

### **An Example of Mining the Database: Are There Gender Neutral Narratives?**

To identify candidate narratives, a ‘select’ query was designed in Access entitled, “How Good are Rater’s Gender Guesses?” Select queries are intended to combine and

---

<sup>253</sup> It is important to note that this was the only question on the rating sheet where raters were asked to guess. Specifically they were told not to guess in any other circumstances.

sort data from various tables to create a datasheet that contains only the necessary information. This query combines gender data about the subjects (from the table ‘Biographical Subject Data’) with gender guesses generated by the raters (from the table ‘Rating Data’). To combine the data from these tables, the table ‘Subject ID to Story ID’ is used to describe the relationship between the narratives (using the field ‘Story ID’) and the subjects (‘Subject ID’). Three fields are used to generate the initial datasheet: Story ID; Actual Gender – the gender of the subject<sup>254</sup>; and Gender? – the guess at the narrator’s gender made by the rater. Two other fields are included for further analyses: Gender Guess? – a true/false field based on whether the rater was certain about their guess at the narrator’s gender; and Sexual Orientation – the self-identified sexual orientation of the subject. The design view of this table can be seen in Figure 51.

To selectively choose ratings where the rater was incorrect about the gender of the narrator, the criteria ‘Female’ and ‘Male’ can be inserted as shown in Figure 51 in the ‘Criteria’ rows. Access implies the operator ‘AND’ across rows and the operator ‘OR’ down columns unless specified otherwise. The criteria in row one require that only data where the actual gender of the participant was ‘Female’ AND the rater guess was ‘Male’ be included in the table. Row two requires the opposite. These two rows are joined using the OR statement: “Row One OR Row Two must be true”, which generates a final datasheet that includes all cases where the rater’s guess about the narrator’s gender was incorrect. It is worth noting at this point that these ‘errors’ should not be considered a reflection of the rater’s reliability. They were specifically instructed to make a choice

---

<sup>254</sup> For clarification, ‘rater’ refers to the anonymous volunteers who rated the narratives, ‘narrator’ refers to the person who is telling the narrative and ‘subject’ refers to the person who initially recalled the memory that was turned into the narrative. In all cases here, ‘narrator’ is synonymous with ‘subject’. However, the people reading the narratives (the raters) did not know that narrator and subject were synonymous, so narrator is used here.

even if they were unsure and these data should be viewed as misperceptions about the narrative in general, rather than errors by the raters.

Of the 2169 ratings in the database, 509 ratings are incorrect about the gender of the narrator. Those errors were split nearly evenly; 281 (55.2%) of those errors believed that the narrator was male, when in fact the narrative was told by a woman and 228 of the errors were in the opposite direction. Interestingly, the raters were relatively good at knowing when they were unsure. Only 113 (22.2%) of those errors were made when the raters thought they were sure of the gender of the narrator. To generate data about the raters' guessing the criteria 'False' or 'True'<sup>255</sup> can be added to the 'Gender Guess?' column. These criteria must be added to each row in order to select the criteria in each row's logical statement.

The number of ratings in the database is nearly three times the number of narratives and knowing the percentage of correct guesses from each of these categories is informative. The same query can be used to generate the number of narratives, rather than the number of ratings by setting the 'unique values' property to 'Yes'<sup>256</sup>. These 509 mistaken ratings originate from 340 distinct narratives, suggesting that nearly half (44.9%) of the narratives in the database contain some sort of inconclusive gender information. Preliminary investigation suggested that some of these narratives were simply misclassified based on the sexual orientation of the narrator. The self-identified sexual orientation of the participant<sup>257</sup> is included in the query design as the field 'Sexual

---

<sup>255</sup> False indicates that the rater was not guessing, true indicates that the rater was guessing.

<sup>256</sup> This is done under the Properties menu which can be found by right-clicking on any empty space in the design view.

<sup>257</sup> These data were acquired at the end of Session Three from each subject.

Orientation'. Any of the possible selections for this field<sup>258</sup> can be used as criteria to limit the datasheet, but as before must be added to each row in order to generate an accurate datasheet. To acquire a conservative estimate of the narratives which are gender indeterminate, narratives told by subjects who were not identified as heterosexual<sup>259</sup> in the database were excluded from further analyses. This decision was based on the assumption that raters who misclassified these narratives did not feel that the gender of the narrator was unclear, rather that they fell victim to their own perceptual biases that all narrators were heterosexual. This removed 132 ratings of 110 unique narratives. This sorting almost certainly removes some narratives which are gender neutral from the dataset. Without any information about the raters' biases, however, this method generates the clearest list of narratives which may be gender neutral or have gender ambiguity.

After the all the filters were applied, 230 narratives were identified where at least one rater made a mistake about the gender of the narrator. In order to more clearly identify which narratives may contain interesting data, Table 13 was constructed. This table compares ratings where all the raters were wrong about the gender with ones where some raters were right and some were wrong, as well as looking at the guessing behaviors of both of these groups. (An alternative version of the table comparing narratives rather than ratings is included as well. However, it is important to note that this table is slightly more difficult to interpret as there are narratives in which all the raters were wrong but some of them guessed and some of them didn't.) After dividing the ratings into the four quadrants depicted in Table 13, two quadrants emerge as the most interesting. The top left quadrant shows that there are 19 ratings where the raters were sure of their gender

---

<sup>258</sup> Heterosexual, Homosexual, Bisexual, Polyamorous, Transgender, None.

<sup>259</sup> This excludes all participants who did not identify as any orientation as well as those who identified as homosexual, bisexual, polyamorous, transgender and none. 'None' is different from non-identifying.

choice and they were all wrong. These narratives contain strong but misleading gender cues, fooling all of the raters into choosing the wrong gender. For example, narrative 5356 refers to a situation in high school where a girl had a crush on the narrator; a strong gender cue if one makes the assumption of heterosexuality. Despite this gender cue, the narrator is female. These narratives are interesting especially because they highlight the assumption of heterosexuality in most people who read the narratives. Narratives like these might be useful when studying gender to examine the biases of a participant pool.

The second interesting quadrant in Table 13 is the bottom right quadrant, where 438 ratings are generated by raters who were mixed in both their accuracy and their certainty. The 165 narratives from these ratings show true gender ambiguity – overall, raters were unable to ascertain the true gender of the narrator from the narrative, whether or not they thought they were certain of their choice. These narratives provide a large pool from which to choose potentially neutral stimuli for gender experiments.

## Section 5: Effects of Brain Damage on Moral Narratives

Recent work by Koenigs et al.<sup>260</sup> among others has highlighted the role of the ventromedial prefrontal cortex (VMPFC) in utilitarian judgments on trolley problem tasks. Individuals with damage to the VMPFC are also known to have social impairments<sup>261</sup> and difficulty in real life situations<sup>262</sup> but retain intact explicit knowledge

---

<sup>260</sup> Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., & Damasio, A. (2007). "Damage to the prefrontal cortex increases utilitarian moral judgements." *Nature* **446**: 908–911.

<sup>261</sup> Mah, L. W. Y., Arnold, M.C., Grafman, J. (2005). "Deficits in Social Knowledge Following Damage to Ventromedial Prefrontal Cortex." *The Journal of Neuropsychiatry and Clinical Neurosciences* **17**: 66–74.

<sup>262</sup> Damasio, A. R. (1994). *Descartes' Error: emotion, reason, and the human brain*. New York, Grosset/Putnam.



of social mores and rules. Subjects with damage to the prefrontal cortex early on in life show particularly severe impairments and are unable to learn appropriate social rules or behaviors<sup>263</sup>. These findings suggest that the prefrontal cortex is a necessary structure for the ability to learn social rules and may be required for effective social decision making. However, since the PFC is not sufficient for remembering explicitly learned social rules, other structures must be involved in the creation of an internal set of moral rules. Our cued recall protocol requires the recollection of moral behaviors and decision making in a participant's own life. We hypothesize that damage to the VMPFC will decrease a participant's ability to reconstruct these personal, moral memories and may impact the number and kinds of justifications provided for his or her actions. By comparison, we expect subjects with other types of brain damage (amnesia, temporal lobe damage) should not be impaired in their ability to recall moral memories or the content therein<sup>264</sup>. To address these questions, we tested collected moral memories from 7 subjects with frontal lobe damage (2 developmental cases, 5 adult onset) as well as 14 subjects with other types of brain damage.

### **Subjects and Testing Procedures:**

21 subjects with brain damage (bilateral amygdala damage (n = 1), amnesia (n = 5), frontal lobe damage (n = 7), temporal lobe damage (n = 8)) were recruited from the Patient Registry of the Division of Cognitive Neuroscience at the University of Iowa. The general location of the damage is shown in Table 14 for each patient group, as well as

---

<sup>263</sup>Anderson, S. W., Bechara, A., Damasio, H., Tranel, D., & Damasio, A.R. (1999). "Impairment of social and moral behavior related to early damage in human prefrontal cortex." *Nature Neuroscience* **2**: 1032-1037.

<sup>264</sup> Obviously, we expect that the timeline for participants with memory loss will be altered compared to normal participants, but the number and content of memories should remain similar.

background data about each participant. Six normal subjects were recruited from the local Iowan community to serve as controls. Normal controls were particularly essential during this testing because a modification to the typical cued recall protocol was necessary. Unlike the cued recall protocol described in Chapter 2, these subjects were tested by an examiner in the room. This change was necessary because a number of the subjects (particularly those with amnesia) were not able to complete the task alone or on a computer. Instead, the examiner was given a worksheet with each of the cued recall questions and the same instructions used in the typical testing. (These instructions can be seen in Appendix B.) Each question was read aloud and the participant's answer was audio recorded for later transcription, while the examiner wrote a short note about the story on the worksheet.

A number of differences appeared to arise from these testing variations, even when comparing normal subjects. Most importantly, subjects gave much shorter answers (mean = 116.4 words per narrative from Iowa normals versus 217.6 words per narrative from the database normals;  $t[779] = 4.40$ ,  $p < 0.0001$ ; see Figure 52). Specifically, we feel that this difference results from the social difficulty of speaking at another person for 3 minutes, rather than at a computer screen. Secondly, the subjects skipped many more memories when speaking to the examiner than when speaking to the computer (mean = 4.8 cues skipped by Iowa normals versus mean = 1.7 cues skipped by the database normals;  $t[103] = 3.12$ ,  $p = 0.002$ ). Since some of the cues were intended to elicit very difficult memories (e.g. Please talk about the thing that you are most afraid that other people will find out that you did), it is likely that these cues were difficult for subjects to talk about with a stranger. While these differences were significant, we expected that the

differences between normal subjects and those with frontal lobe damage would be greater than the differences caused by testing. This did not turn out to be the case. As shown below, the difference between the normal subjects in the database and the normal subjects in Iowa is much greater than the differences among the subjects with brain damage.

### **Results:**

Our a priori hypothesis was that the number of moral memories generated by subjects with frontal lobe damage would be significantly different from that of brain damaged controls, although the total number of memories generated would remain constant. As shown in Figure 53, the number of moral memories generated by subjects with frontal lobe damage (mean = 3.8) is nearly exactly the same as the number generated by brain damaged controls<sup>265</sup> (mean = 3.38) and the number generated by the normal controls from Iowa (mean = 3.83). Since there was variance between the total number of memories generated by the database controls and the Iowa normals, the percentage of moral memories generated was also calculated and found to be similarly invariant among the frontal lobe damaged subjects, the brain damaged controls and the Iowa normals. The number of cue words skipped by subjects with frontal lobe damage was also not significantly different.

We measured the number of words per sentence in each group to ascertain whether there was something substantively different about the speech patterns of the subjects which might affect the overall narratives. However, we found all groups used approximately the same mean number of words per sentence (16.08 for database normals, 16.46 for Iowa normals, 15.28 for brain damage controls and 17.36 for frontal lobe

---

<sup>265</sup> These are participants who have had temporal lobectomies.

damaged subjects). As a preliminary examination of the content of the narratives, we did a Language and Inquiry Word Count (LIWC) analysis in the same way that it was done on the database narratives (described in Chapter 3). As shown in Figure 54, the number of positive emotion words in each sample was not different among any of the groups. The number of negative emotion words was also invariant.

The narratives were also examined qualitatively during the editing process. (As with the other narratives, these were edited for readability while retaining the individual character of each narrative.) The most obvious differences were in length and in the coherence of the narratives, which varied by group. Samples of the edited narratives from each group are included in Appendix I. The final, edited narratives are quite readable and could be used as stimuli; however, it was much more difficult to the narratives from most of the brain damaged groups than it was to edit the narratives from the normals. One salient feature was that the memories were often disordered. They did not begin at the beginning of a story, carry through the content and then end at an appropriate finishing point. While this finding is merely qualitative, it may be worth examining further if more data are collected.

While our hypotheses were not sustained by the data, the number of subjects in each sample group was quite small. The sample sizes were small due to the rarity of these subjects; however, without enough statistical power it is difficult to interpret the results. These studies seem to demonstrate a need for greater numbers of subjects, tested over a greater period of time and with many more normal controls. Given the overwhelming differences between the database normals and the Iowa normals, any further studies with these subjects would need to have their own substantive control group to compare to.

## Section 6: Temporal Biasing of Autobiographical Moral Events<sup>266</sup>

### **Abstract**

Our autobiographical self depends on the differential recollection of our personal past, which depends on both memory remoteness and emotional saliency. While both time and emotionality are known to exert strong effects on memory, little is known about how highly emotional, personal memories are distributed through time. To investigate this issue in detail, we collected a novel database of 758 autobiographical narratives for personal moral events from 100 exceptionally well-characterized healthy adults. Negatively valenced memories were significantly more remote than positively valenced memories, both as measured by the valence of the cue word that evoked the memory as well as by the content of the memory itself. The effect was independent of chronological age, ethnicity, gender, or personality, arguing for a universal emotional bias in how we construct our moral autobiography.

### **Introduction**

An essential part of how we think of people is normative: some are good, others bad, some should be praised, others punished. Moral judgment pervades not only how we think of others, but also how we view ourselves and it appears plausible that a large proportion of the memories that matter the most to us personally are morally laden.

---

<sup>266</sup> This section is in preparation for submission under the title: “Becoming a better person: temporal remoteness biases autobiographical memories for moral events”.

Curiously, despite great interest both in moral cognition<sup>267</sup> and in autobiographical memory<sup>268</sup>, memories for moral events have received scant investigation. Part of the reason for this neglect is no doubt the effort required to collect such memories in the first place; another may be the presumption that moral memories are no different from other memories.

Yet research on emotional memory has shown us that emotional memories are special. They are typically more vivid and recollected more easily<sup>269</sup>, even though they are not necessarily more accurate in their details<sup>270</sup>. Studies of so-called “flashbulb memories” have probed the strong effects of emotion on memory consolidation to events such as the 9/11 attacks on New York<sup>271</sup>, the explosion of the Challenger space shuttle<sup>272</sup>, or the Kobe earthquake in Japan<sup>273</sup>. These findings mirror laboratory experiments in both

---

<sup>267</sup> Moll J., Z. R., de Oliveira-Souza R., Krueger F., Grafman J. (2005). "The neural basis of human moral cognition." *Nature Reviews Neuroscience* 6(10): 799–809.; Greene, J. D. and J. Haidt (2002). "How (and where) does moral judgment work?" *Trends in Cognitive Sciences* 6(12): 517-523.; Greene, J. D., L. E. Nystrom, et al. (2004). "The Neural Bases of Cognitive Conflict and Control in Moral Judgment." *Neuron* 44: 389-400.; Haidt, J. (2001). "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment." *Psychological Reviews* 108(4): 814-834.; Haidt, J. (2007). "The new synthesis in moral psychology." *Science* 316: 998-1002.

<sup>268</sup> Berntsen, D. and D. C. Rubin (2002). "Emotionally charged autobiographical memories across the life span: the recall of happy, sad, traumatic, and involuntary memories." *Psychology and Aging* 17: 636-652.; Rubin, D. C. and M. D. Schulkind (1997). "The distribution of autobiographical memories across the lifespan." *Memory and Cognition* 25: 859-866.; Schacter, D. L. (1996). *Searching for Memory: The Brain, The Mind, And The Past*. New York, Basic Books.

<sup>269</sup> Burke, A., F. Heuer, et al. (1992). "Remembering emotional events." *Memory and Cognition* 20: 277-290.

<sup>270</sup> Christianson, S.-A. and E. F. Loftus (1991). "Remembering emotional events: the fate of detailed information." *Cognition and Emotion* 5: 81-108.; Heuer, F. and D. Reisberg (1990). "Vivid memories of emotional events: the accuracy of remembered minutiae." *Memory and Cognition* 18: 496-506.; Winograd, E. and U. Neisser (1992). *Affect and Accuracy in Recall: Studies of "Flashbulb" Memories*. Cambridge, MA, Harvard University Press.

<sup>271</sup> Phelps, E. A. (2006). "Emotion and cognition: insights from studies of the human amygdala." *Annual Review of Psychology* 57: 27-53.

<sup>272</sup> Neisser, U. and N. Harsch (1992). Phantom flashbulbs: false recollections of hearing the news about Challenger. *Affect and Accuracy in Recall: Studies of "Flashbulb" Memories*. E. W. U. Neisser. Cambridge, MA, Cambridge University Press: 9-31.

<sup>273</sup> Mori, E., M. Ikeda, et al. (1999). "Amygdalar volume and emotional memory in Alzheimer's disease." *American Journal of Psychiatry* 156: 216-222.

humans and rats that demonstrate that emotional arousal enhances memory encoding and consolidation via both neurological and endocrinological mechanisms<sup>274</sup>.

Given these findings, and given that moral events are typically judged to be highly emotional and involve some of the same brain structures that are also involved in the emotional modulation of memory<sup>275</sup>, one would predict that moral memories should also be special. Like emotional memories more generally, one would expect moral memories to predominate in our autobiography and to be recollected with especially vivid detail and accompanying reliving of some of the emotion. On the other hand, just as emotionally traumatic memories can be suppressed<sup>276</sup>, highly emotional moral memories, presumably especially negatively valenced ones, would be expected to be recalled less frequently.

All of these considerations argue for an important contribution to our autobiographical selves made by memories of morally laden personal events. We wanted to investigate two key questions: (1) are moral events recollected differentially through time (are moral memories, on average, more remote or more recent?) (2) are positive moral events recollected at different times than negative moral events?

## Method

---

<sup>274</sup> Cahill, L., R. J. Haier, et al. (1996). "Amygdala activity at encoding correlated with long-term, free recall of emotional information." *Proceedings of the National Academy of Sciences* **93**: 8016-8021.; Phelps, E. A. (2006). "Emotion and cognition: insights from studies of the human amygdala." *Annual Review of Psychology* **57**: 27-53.; McGaugh, J. L. (2004). "The amygdala modulates the consolidation of memories of emotionally arousing experiences." *Annual Review of Neuroscience* **27**: 1-28.

<sup>275</sup> Greene, J. D., R. B. Sommerville, et al. (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." *Science* **293**: 2105-2108.; Moll, J., R. de Oliveira-Souza, et al. (2002). "The neural correlates of moral sensitivity: A functional magnetic resonance imaging investigation of basic and moral emotions." *The Journal of Neuroscience* **22**(7): 2730-2736.; Moll, J., R. Zahn, et al. (2005). "The neural basis of human moral cognition." *Nature Reviews Neuroscience* **6**(10): 799-809.

<sup>276</sup> Schacter, D. L. and S. D. Slotnick (2004). "The cognitive neuroscience of memory distortion." *Neuron* **44**: 149-160.; Loftus, E. F. (1993). "The reality of repressed memories." *Journal of the American Psychological Association* **48**: 518-537.

### *Participants*

We tested 100 healthy adults recruited from the Los Angeles community, one at a time, over a period of approximately one year. We selected participants between the ages of 40-60, with equal gender distribution, and with a racial composition matching that of the state of California. All subjects were screened to be neurologically and psychiatrically healthy and all were given an extensive battery of background tests to assess IQ, mood, personality, socioeconomic status, as well as political and religious affiliation (Table 15). All participants gave informed consent and received written assurance of anonymity regarding their data.

### *Procedure*

Participants came in for 2 separate days of testing, typically spaced several weeks apart. In the first visit, we collected the background data given in Table 15 and eliminated subjects who were not normal, healthy individuals. After screening, we collected the memory data in a second session conducted on the same day. During their second visit, we collected self-ratings about the memories which had moral components. Participants read transcripts of their own memories and dated them. All participants also filled out an exit questionnaire about their experience during the experiment.

For the first visit, participants were administered a number of background questionnaires; here we report data collected from the Wechsler Abbreviated Scale of Intelligence (WASI)<sup>277</sup>, NEO Five-Factor Inventory (NEO-FFI)<sup>278</sup>, Positive and Negative Affect Schedule (PANAS)<sup>279</sup> and the Ethical Position Questionnaire (EPQ)<sup>280</sup>.

---

<sup>277</sup>Wechsler, D. (1999). Wechsler Abbreviated Scale of Intelligence. San Antonio, Psychological Corporation.



During the testing session, participants were seated in a room by themselves in front of a computer that recorded their spoken memories for later transcription. We used a common cue-elicitation protocol<sup>281</sup>. Instructions and cue words appeared on the computer monitor. We chose 30 cue words to prompt moral recollection, spanning both positive and negative affect (see Table 16). The cues were one of three types: emotion words, descriptions of moral behaviors and prompts about the most superlative moments in the participant's life. The cues were chosen using the MRC Psycholinguistic Database<sup>282</sup> to select words based on Kucera-Francis word frequencies<sup>283</sup> and measures of concreteness, familiarity, imageability and meaningfulness, when available. Each cue word appeared on the screen in a specific question ("Please talk about a time when you did something that made you feel GUILTY"), followed by a 3-minute period during which the participant was instructed to recount a specific, personal, autobiographical event related to that question. Participants' narratives were recorded onto the hard drive and transcribed. Participants filled out an exit questionnaire about their experience at the end of the session.

---

<sup>278</sup>Costa, P. T. and R. R. McCrae (1992). NEO PI-R Professional Manual. Lutz, FL.

<sup>279</sup>Watson, D., L. Clark, et al. (1988). "Development and validation of brief measures of positive and negative affect: the PANAS scales." Journal of Personality and Social Psychology **54**(6): 1063-1070

<sup>280</sup>Forsyth, D. (1980). "A Taxonomy of Ethical Ideologies." Journal of Personality and Social Psychology **39**(1): 175-184.

<sup>281</sup>Crovitz, H. and H. Schiffman (1974). "Frequency of episodic memories as a function of their age." Bulletin of the Psychonomic Society **4**: 517-518.; Rubin, D. C. and M. D. Schulkind (1997). "The distribution of autobiographical memories across the lifespan." Memory and Cognition **25**: 859-866.; Buchanan, T. W., D. Tranel, et al. (2005). "Emotional autobiographical memories in amnesic patients with medial temporal lobe damage." The Journal of Neuroscience **25**: 3151-3160.

<sup>282</sup>Coltheart, M. (1981). "The MRC Psycholinguistic Database." Quarterly Journal of Experimental Psychology **33A**: 497-505.

<sup>283</sup>Kucera, H. and W. N. Francis (1967). Computational Analysis of Present-Day American English. Providence, Brown University Press.

For the final visit, participants were given selected transcripts of the memories they had narrated on the first visit and were asked to rate them on several dimensions and date them. These transcripts were selected based on the moral content of the memories.

Finally, we contacted 44 of the 100 participants two years after data collection over the phone and discussed with them the validity of their data. Participants were explicitly asked if they had produced only completely truthful narratives, to the best of their knowledge. Every single participant of the 44 contacted endorsed this. We further had an independent and unfamiliar investigator contact 11 of the participants by phone a second time and probe specific memories for test-retest reliability. All 11 of these participants were able to provide descriptive, accurate retellings of their original memories.

### *Stimulus Selection*

From a total of 3300 memories produced by the 100 subjects, we selected 758 as autobiographical moral memories. Our criteria for this selection were that the memories needed to be: 1. episodic and personal, 2. moral and 3. involve a decision or choice. Events that were not recounted in the first person were excluded, as were events that were generic or overly vague. The moral status of the memory was judged by two independent raters. The two raters worked in series to determine the inclusion of the memories in the moral group. The first rater read all of the moral memories and used a more liberal interpretation of the criteria (selecting 788 memories) while the second rater read only those provided by the first and applied a more strict interpretation of the criteria to the memories. This process ensured that the memories were evaluated individually and not in

the context of the participant's life overall. Minimal editing was also performed on the memories to clarify the narrative of the memory while retaining the unique vernacular of each participant.

### *Memory Ratings*

All participants rated transcripts of their own memories during their second visit, as described above. Memory transcripts were also rated by an independent group of 55 adult raters, none of whom had participated in the main experiment. Raters judged the memories on a large number of attributes; here we include only ratings of six (paired) features. The raters categorized each memory into many categories about the actions in the story, the reasons for those actions and the feelings the subjects had about their actions. Among these categories, the raters judged whether each memory involved helping someone, hurting someone, doing the right thing, doing the wrong thing, feelings of personal moral strength and/or feelings of personal moral weakness.

### *Cue Word Ratings*

The cue words themselves were also rated by yet a third independent group of 4 raters, using a Likert scale from -5 to +5. All four raters showed very high agreement on rating the cue words as well as agreement with the Affective Norms for English Words<sup>284</sup> (see Table 16).

## **Results**

---

<sup>284</sup> Bradley, M. M. and P. J. Lang (1999). Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings. Technical report C-1. Gainesville, FL, The Center for Research in Psychophysiology, University of Florida.

As Table 15 bears out, all our participants were healthy and representative adults and our sample had a good distribution of ages between 40-60, of gender, and of ethnicity. Participants all treated the experiment quite seriously; several were in tears during their second session (where they recounted the memories) and generally endorsed ratings on their exit questionnaire that indicated a high level of emotional involvement (mean =  $7.16 \pm 1.89$ )<sup>285</sup>. Our follow-up phone calls assured us that the memories were not confabulated, since all of the 44 subjects we contacted strongly indicated they had provided honest accounts and since the 11 subjects whose memories we probed for reliability were all able to produce narratives that matched their original one in overall content. Our independent investigator rated every narrative from every participant with the highest rating score, indicating that each participant was able to recall specific and complete details without significant prompting.

### ***Memory Remoteness influences Valence of Cue Words***

We first examined the mean age of all moral memories elicited by each of the 26 cue words. (Four cues are left out of this analysis because they yielded too few memories. These are: three control cues ‘tired’, ‘exercised’ and ‘funny’ and the cue ‘embarrassed’.) Simply rank-ordering cue words according to the mean age of the memories (Figure 55) suggested a pattern whereby more remote memories were associated with more negative cue words. Specifically of note are the memories elicited by the cue ‘memorable’, which was itself not valenced. Subsequently, we grouped these most memorable events into those rated positive versus those rated negative, using ratings from three additional

---

<sup>285</sup> These participants were responding to the question: “How emotional did you feel during this study?” using a Likert scale that ranged from 1 = not at all to 10 = very.

independent raters. The mean age of the positive memorable events falls nearly 8 years earlier than the negative memorable events. However, given the small number of items in each group ( $n = 8$ ), this difference did not achieve significance.

Plotting memory remoteness against the rated valence of the cue word (Figure 56), yielded a significant positive regression ( $R(26) = 0.74$ ,  $p = 0.004$ ), bearing out the qualitative finding above. To test whether these effects might be due to the chronological age of our participants, rather than the remoteness of their memories, we carried out the same analyses using absolute chronological age at the time of the memory. We found that there was little difference in the distributions between absolute age of the participant at the time of the memory and the age of memory.

### *Memory remoteness influences valence of the memory*

We next examined the recollected narratives themselves. We began with an automated text analysis using the Linguistic Inquiry and Word Count database<sup>286</sup>, which counts the number of words in a sequence of text in terms of their frequency in various categories. We examined positive and negative emotion categories. There was no significant effect of the word count of emotionally valenced words as a function of memory remoteness, a null finding that is perhaps not surprising given that the emotional semantic content of the narrative is likely only weakly related to mere frequencies of emotion words within it.

As a second probe of the valence of the memories, we used each participant's own ratings of his/her memories. For each memory, the participant provided a rating of

---

<sup>286</sup> Pennebaker, J. W., C. K. Chung, et al. (2007). The development and psychometric properties of LIWC2007. Austin, TX, LIWC.net.

how well the memory reflected on him or her. Here we found a weak but significant effect. While participant's own ratings of valence correlated only weakly with temporal remoteness of the memory, a contrast between the two most positively rated versus 2 most negatively rated memories for each participant did reveal a significant time difference ( $t[134] = 3.22, p = 0.0016$ ).

We found the strongest effect when the memories were categorized by the 55 independent raters. Here we focused on pairs of categories that had clear opposites: moral weakness versus strength, doing the right thing versus the wrong thing, and hurting someone versus helping someone (Figure 57). Each of these pairs showed a significant difference with respect to the mean age of the memories ( $t[972] = 5.44, p < 0.001$ ;  $t[1192] = 8.68, p < 0.001$ ; and  $t[958] = 7.95, p < 0.001$ , respectively). When the three negative categories were contrasted with the three positive categories, a significant difference was also found ( $t[4] = 9.21, p < 0.001$ ).

### *Possible covariates*

To probe for possible factors that could influence the above effects, we examined gender, ethnicity, IQ, religion, political affiliation, and personality. We computed a Pearson product-moment correlation among these variables and found no significant effect.

### **Discussion**

We collected and analyzed an exceptionally rich set of data to investigate the temporal distribution of moral autobiographical memories. We found that memories of

more positive moral events were, on average, more recent than memories of more negative moral events. The effect was not attributable to chronological age as such, nor modulated by gender, IQ, ethnicity, personality or other factors. The findings suggest that there is a universal bias in how we construct our autobiographical past, with a tendency to remember the most recent events as those that are also the most morally positive.

An important aspect of the effect we found is that it was not strongly driven by participant's own subjective ratings of their memories, but rather by the objective ratings of the cue words that elicited the memories, as well as by categorization from independent raters. There are several plausible explanations for this. First, people may have relatively little insight into the emotional valence of the memory. Second, subjective ratings of experienced emotional valence may diverge substantially from objective third-party ratings. Third, and perhaps most intriguing, is the possibility that people distort the emotional valence of their memories when asked to provide explicit ratings, perhaps especially for memories of events for which they were responsible. This possibility could also result in people's unawareness of the memory bias we found, resulting in a systematic but implicit emotional bias on how we think of ourselves as moral agents.

It is surprising that there were no effects due to gender, personality or any of the other factors we examined, since one might have expected there to be such a relationship. In particular, we had expected associations between personality traits and the magnitude of our memory bias.

Our database of moral memories remains a valuable resource for probing additional questions and for generating targeted narratives that could themselves be used as stimuli, for instance, in functional imaging studies. An important future direction for research will be to investigate how specific the bias we found is to moral memories as such or whether it pertains to highly emotional personal memories more generally.



*Chapter 4: Figures and Tables*  
Section 1

“The depth and strength of a human character are defined by its moral reserves. People reveal themselves completely only when they are thrown out of the customary conditions of their life, for only then do they have to fall back on their reserves.”

-- Leonardo da Vinci

Figure 8: Age Distribution




Figure 8: The participants' ages were distributed between 40 and 60 years old (mean =  $48.9 \pm 5.9$  years).

Figure 9: Handedness Distribution

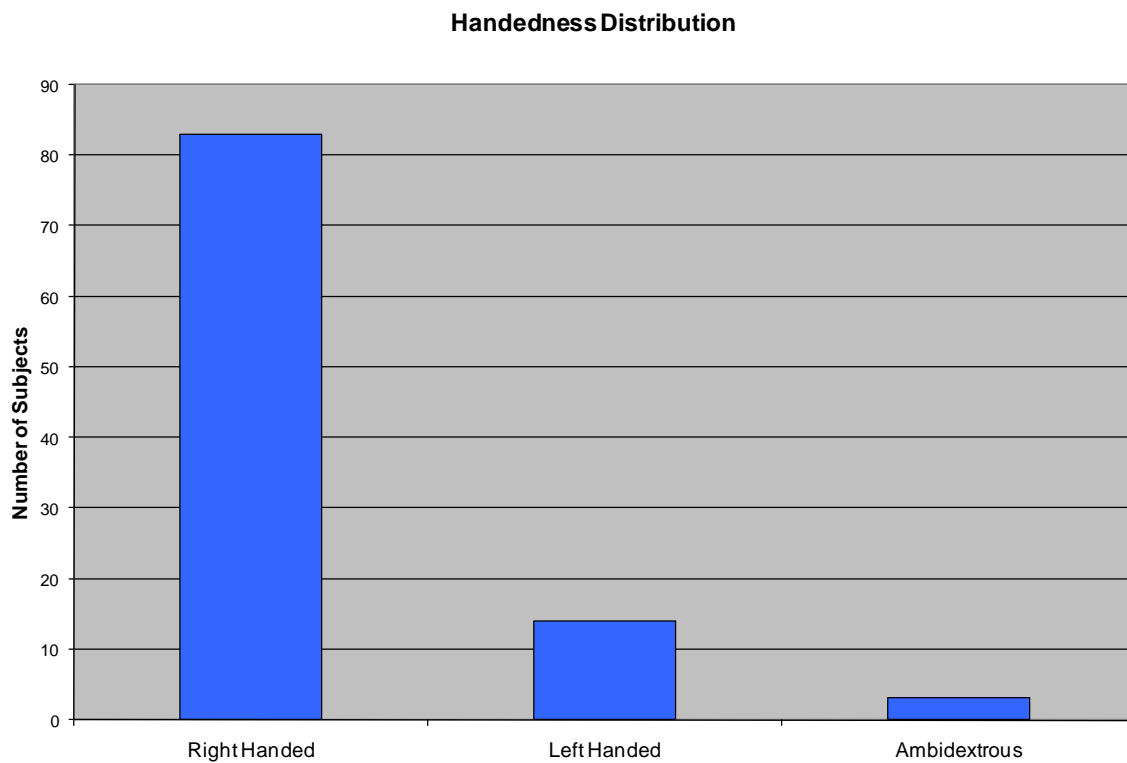


Figure 9: 83% of our participants were right handed. This is in keeping with published estimates of handedness in the population.

Figure 10: Educational Distribution

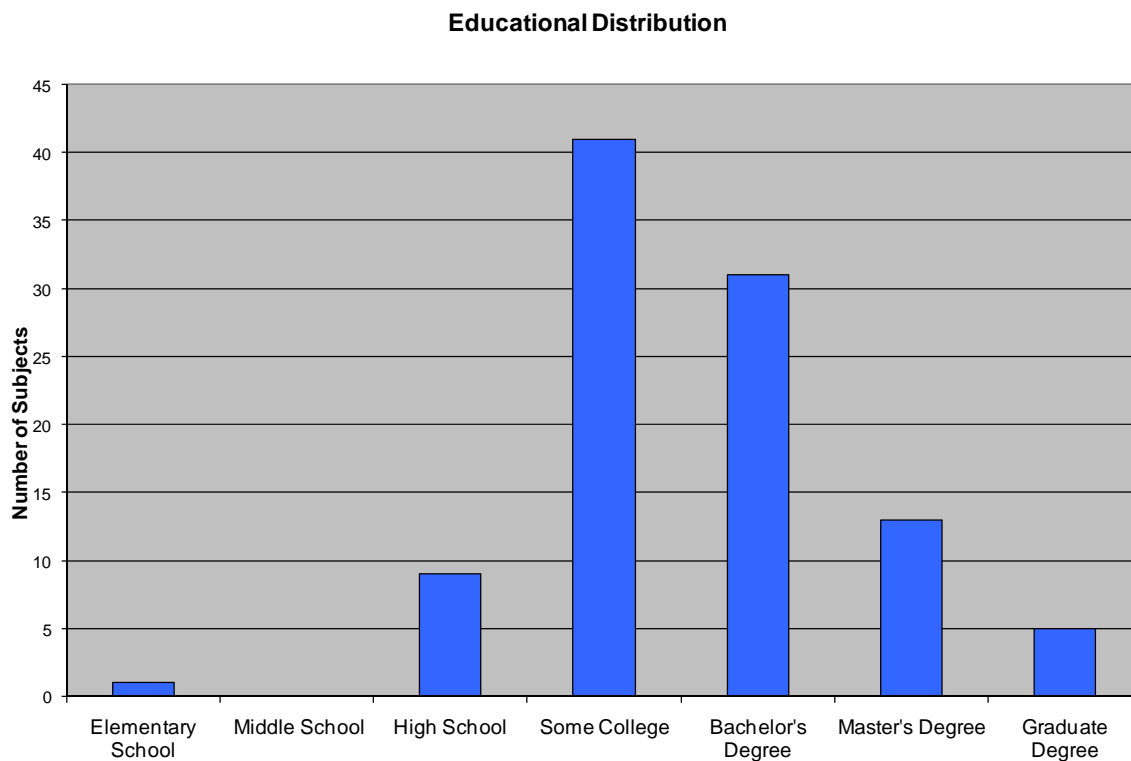


Figure 10: Participants' education ranged from elementary education to advanced graduate degrees, including medical and law school. Most participants (90%) had attended at least some college.

Figure 11: Income Distribution




Figure 11: Most participants' individual income was between \$15,000 and \$30,000, in keeping with published census figures that find the average Californian makes \$22,711.

Table 3: Occupations

Occupation	Number of Subjects	Occupation	Number of Subjects
Administrative Assistant	12	Property Supervisor	1
Writer	9	Product Planner	1
Teacher	6	Producer	1
Sales	5	Photographer	1
Actor	5	Office Manager	1
Musician	4	Mediator	1
Financial Services	4	Massage Therapist	1
Artist	4	Lawyer	1
None	3	Law Clerk	1
Currently Unemployed	3	Interior Design	1
Self Employed	2	Hatter	1
Retired	2	Hairdresser	1
Homemaker	2	Farm Grower	1
Esthetician	2	Engineer	1
Entertainment Industry	2	DSP Engineer	1
Data Collector	2	Drug and Alcohol Counselor	1
Caterer	2	Distributor - Health Product	1
Accountant	2	Designer	1
Warehouse Man (Forklift Operator)	1	Customer Service Representative	1
Training and Development Tech	1	Custodian	1
Tour Guide	1	Consultant	1
Student	1	Computer Specialist	1
Story Analyst	1	Chef	1
Security	1		

Table 3: Our participants' occupations were quite varied. These data are particularly important in establishing that our online recruiting methods were effective in reaching a wide variety of people.

Figure 12: Language Distribution

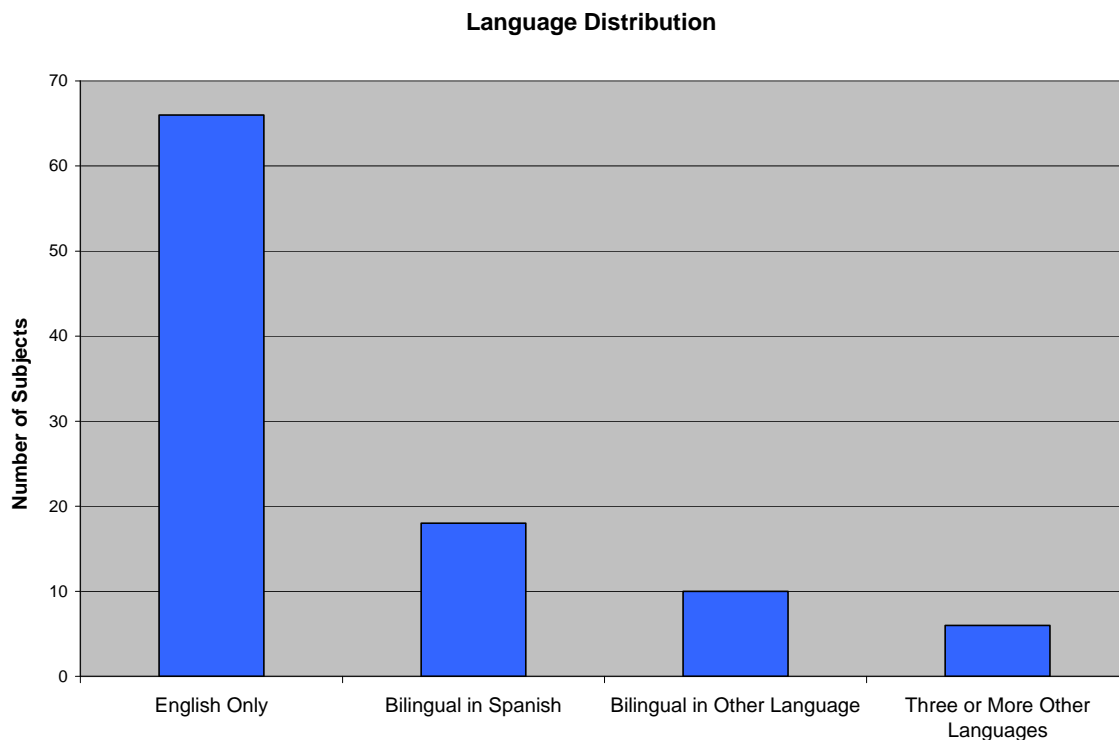


Figure 12: The majority of our participants ( $n = 66$ ) spoke English as their only language, while some were bilingual ( $n = 18$  in Spanish,  $n = 10$  in other languages) or spoke three or more languages ( $n = 6$ ). It is important to note that all participants were fluent in English by adolescence. The languages spoken by our participants were diverse, including: Armenian, Chinese (Mandarin), Czech, Filipino, French, German, Hindi, Hungarian, Japanese, Korean, Romanian, Russian, Sign Language, Slovak, Spanish and Tagalog.

Figure 13: Ethnicity Distribution




Figure 13: The ethnic diversity of our participant population was carefully monitored during recruitment to match the US Census findings for California. As described in Chapter 4, the inclusion of Hispanic as an ethnicity option for participants was critical given our location in Southern California.



Figure 14: Marital Status Distribution

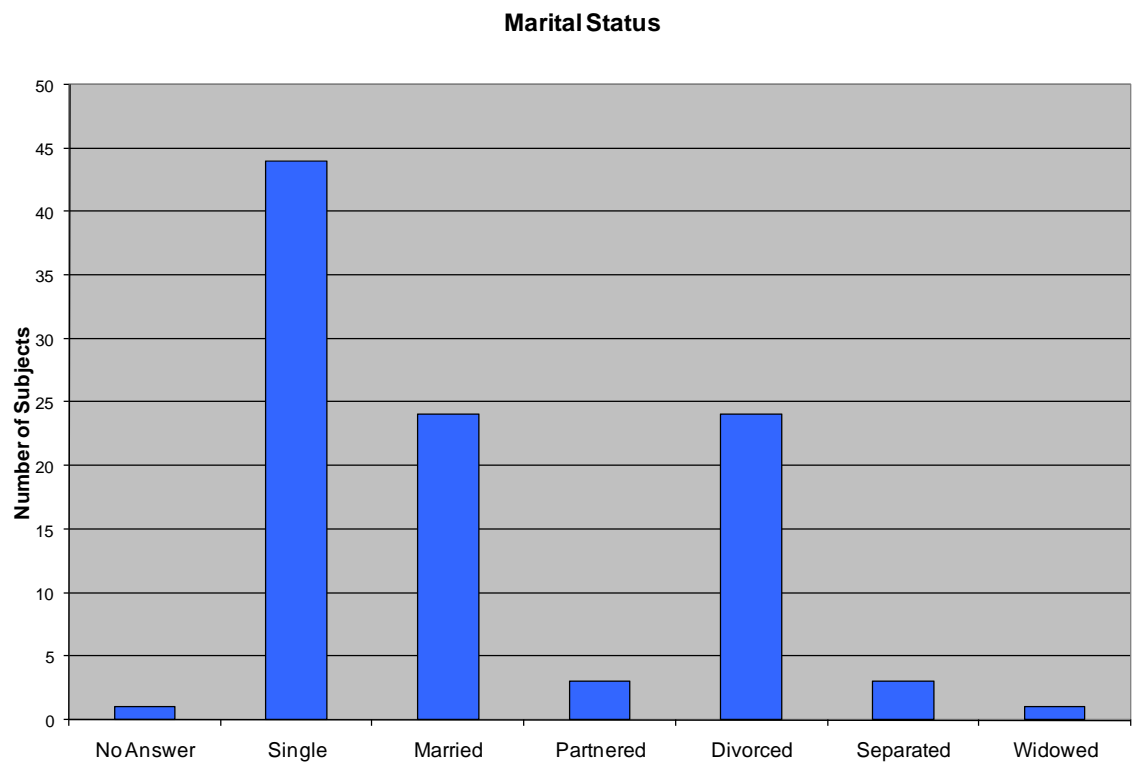


Figure 14: Our participants were distributed among all types of relationship statuses.

Figure 15: Distribution of Number of Children




Figure 15: 53% of our participants had at least one child, in keeping with US Census findings.

Figure 16: Sexual Orientation Distribution

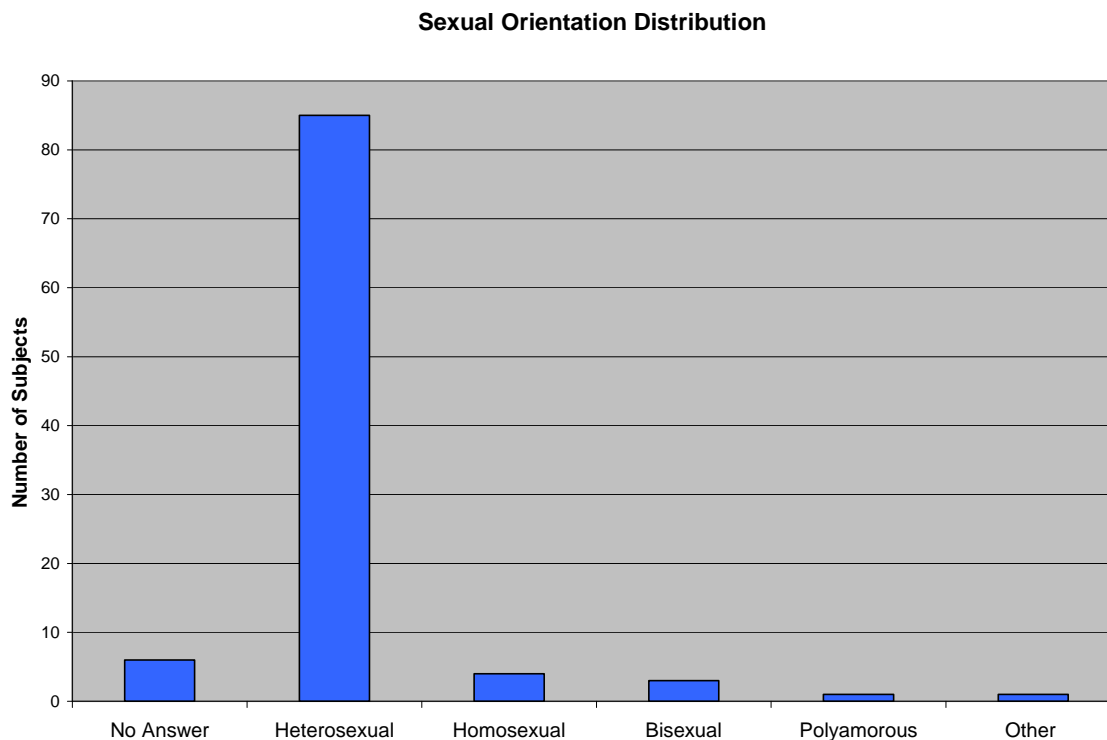


Figure 16: Most of our participants ( $n = 85$ ) chose heterosexual as their sexual orientation. The choice 'transgendered' was removed as an option after completing the study as it is not descriptive of a particular sexual orientation.

Figure 17: Sexual Preference Distribution

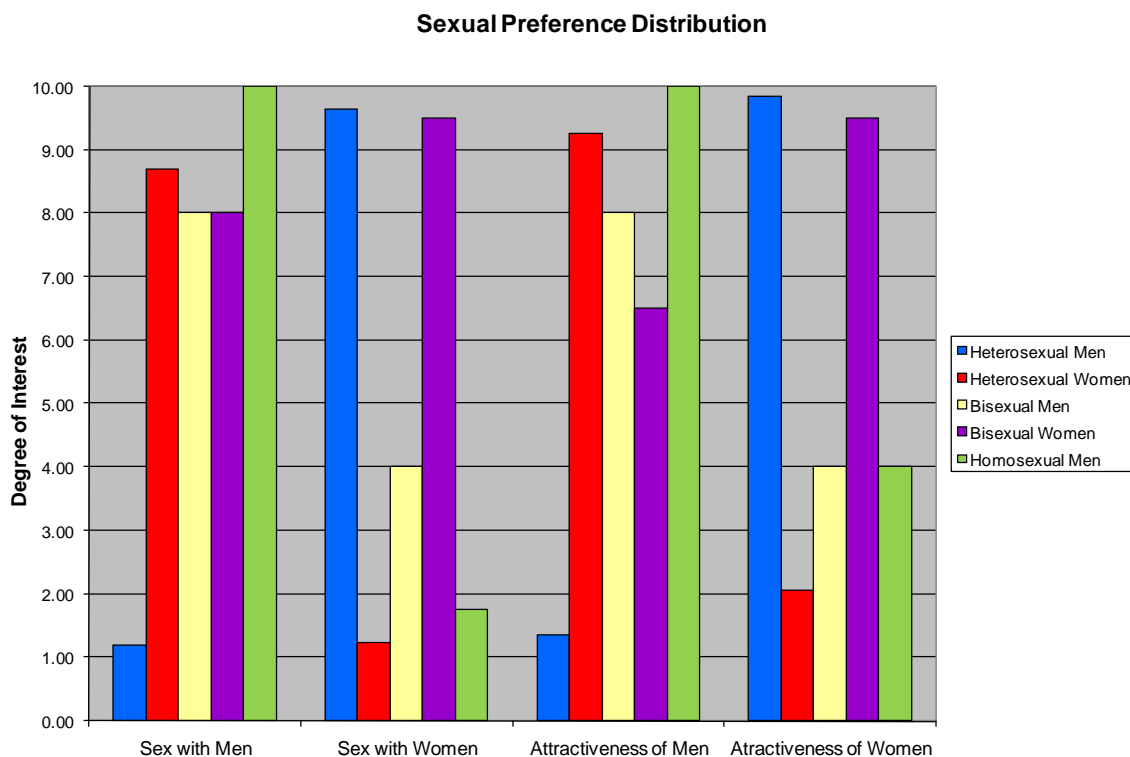


Figure 17: Our participants answered a number of questions about their sexual orientation to establish the degree to which they were interested in having sex with each gender as well as the degree of attractiveness they felt for each gender. These findings were generally in keeping with the participants' stated sexual orientation.

Figure 18: Religious Distribution




Figure 18: Our participants were asked to provide a short written response about their current religion and the religion with which they were raised. 25% of our participants practiced no religion at the time of testing.

Figure 19: Degree of Religious Influence

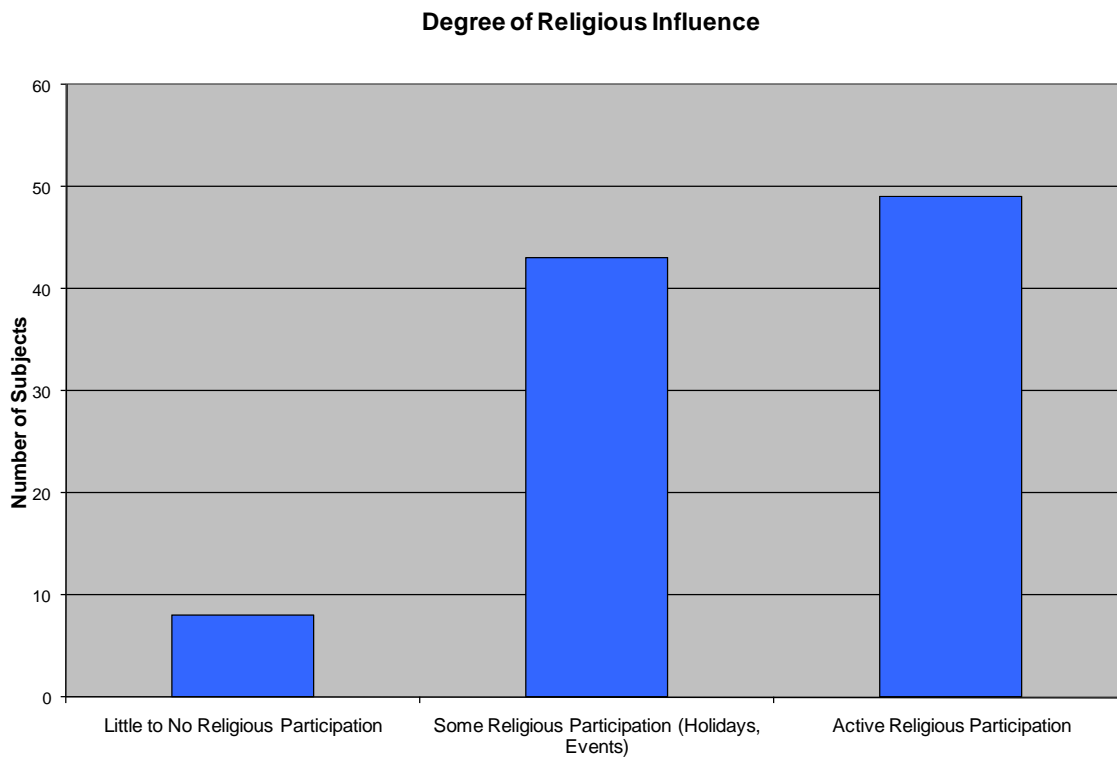


Figure 19: Using written paragraphs provided by the participants, a research assistant rated the religious influence of each participant on a scale from 1 to 3.

Figure 20: Comparison of Religious Beliefs

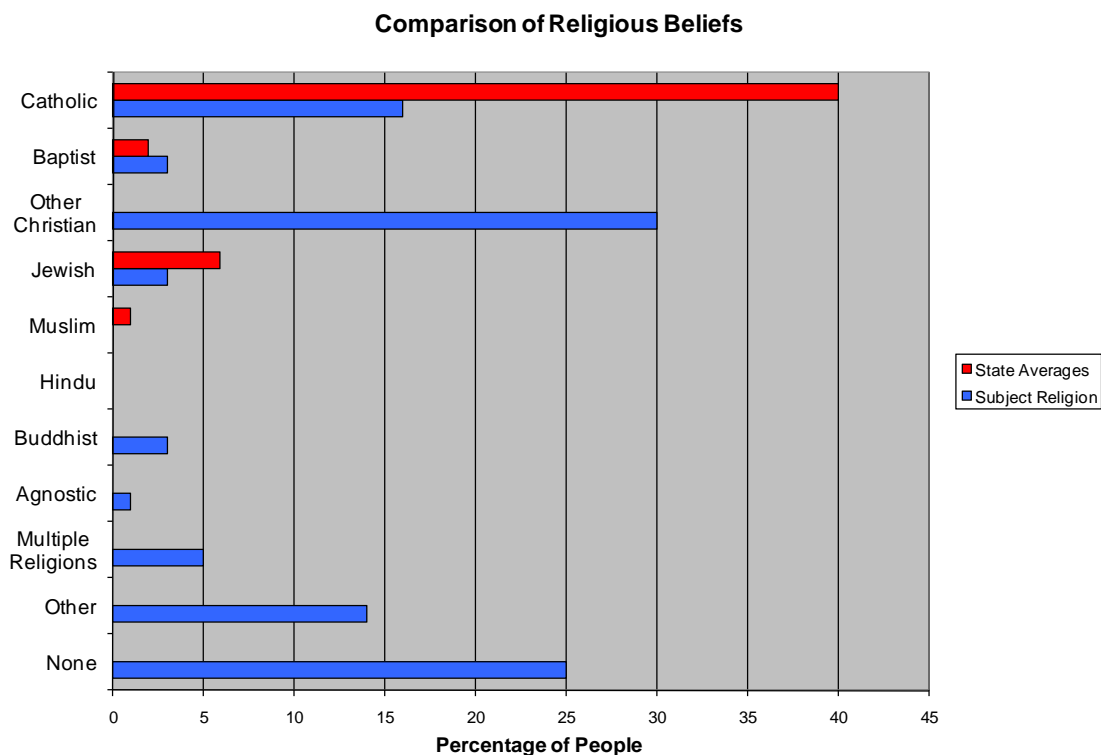


Figure 20: Using data from a myriad of sources, we were able to compare our participants' religious practices with the Los Angeles community. Our participants were generally well-matched to state averages. Overall, Catholics seem to be under-represented, but we believe that some these participants may be categorized in the less specific 'Other Christian' group.

Table 4: Religious Distribution in Southern California

Religious Group	Number of temples, churches, mosques or synagogues	Number of adherents	% of total population	% of recruited subjects
Roman Catholic	278	3,806,377	40	16
Jewish	202	564,700	5.9	3
Southern Baptist	312	111,634	1.2	3 (all Baptists combined)
American Baptist	211	73,217	0.8	--
Muslim	48	92,919	1	0
Hindu	37	NA*	NA*	0
Buddhist	145	NA*	NA*	3

Table 4: A detailed breakdown of religious centers in the Los Angeles area shows the comparison of our participant population to the general population.



Figure 21: Political Value Distribution

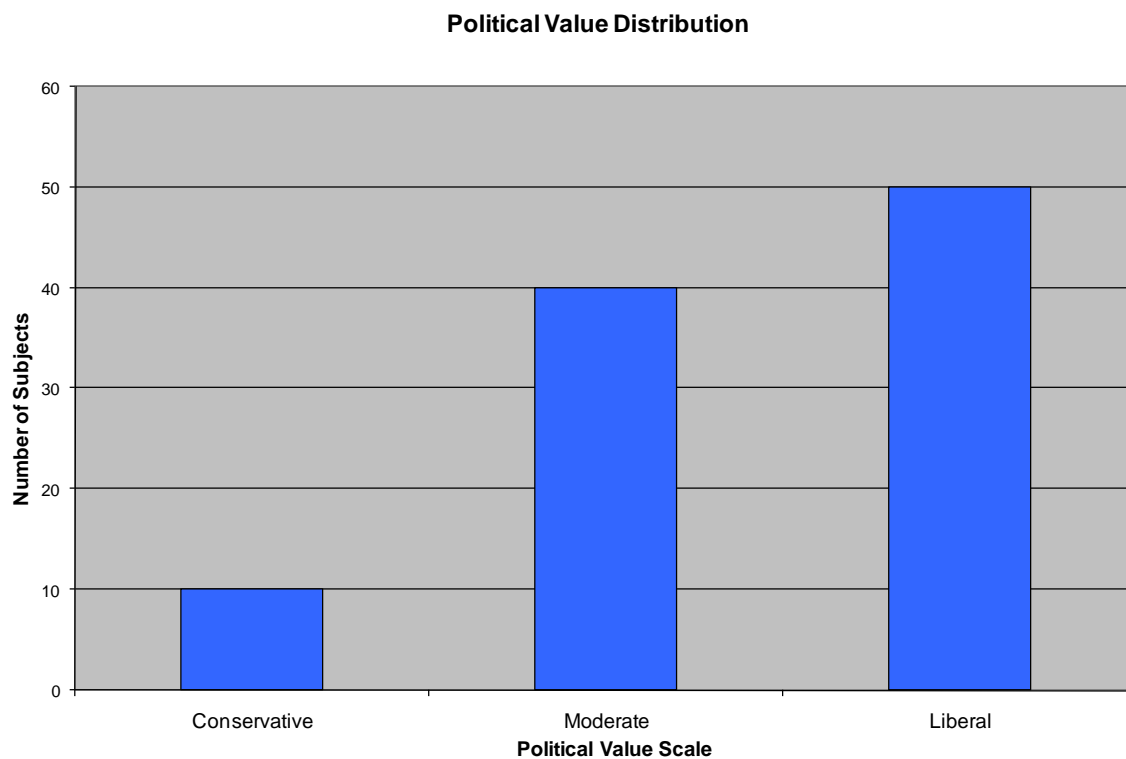


Figure 21: Our participants' political viewpoints were determined using a questionnaire that required them to choose among three possible perspectives on current political issues.

Figure 22: IQ Distribution




Figure 22: Participant's full scale IQ (FSIQ) was measured using the 2-subtest version of the Wechsler Abbreviated Scale of Intelligence (WASI-III). (Wechsler, 1999).

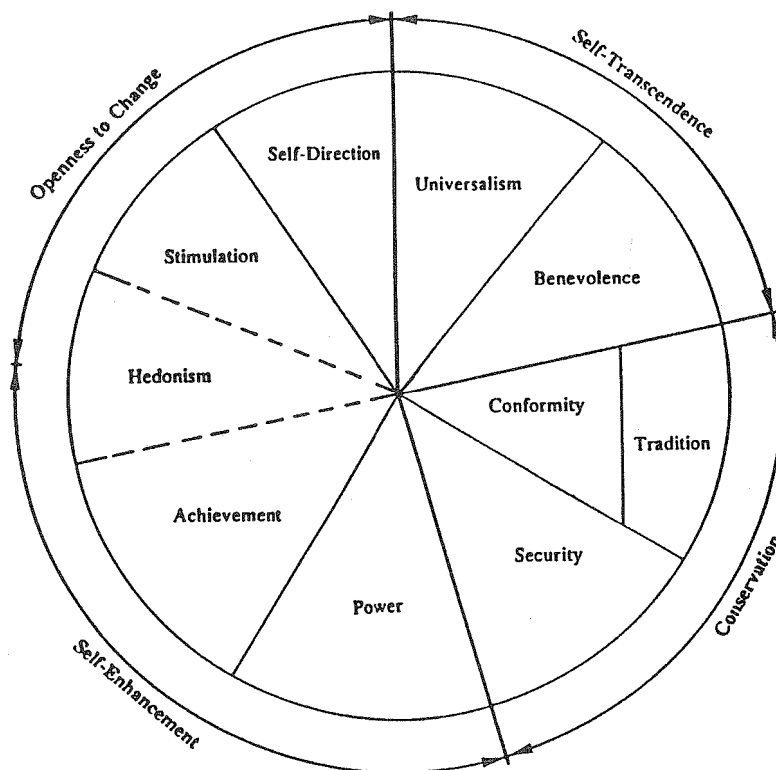
Table 5: Value Definitions from the Schwartz Value Inventory

Definition	Exemplary Values	Sources
<b>Power:</b> Social status and prestige, control or dominance over people and resources	Social power authority, wealth	Interaction Group
<b>Achievement:</b> Personal success through demonstrating competence according to social standards.	Successful capable ambitious	Interaction Group
<b>Hedonism:</b> Pleasure and sensuous gratification for oneself.	Pleasure Enjoying life	Organism
<b>Stimulation:</b> Excitement, novelty, and challenge in life.	Daring, varied life, exciting life	Organism
<b>Self-direction:</b> Independent thought and action—choosing, creating, exploring.	Creativity, curious Freedom	Organism Interaction
<b>Universalism:</b> Understanding, appreciation, tolerance, and protection for the welfare of <i>all</i> people and for nature.	Broad-minded, social justice, equality Protecting the environment	Group* Organism
<b>Benevolence:</b> Preservation and enhancement of the welfare of people with whom one is in frequent personal contact.	Helpful Honest Forgiving	Organism Interaction Group
<b>Tradition:</b> Respect, commitment, and acceptance of the customs and ideas that traditional culture or religion provide.	Humble, devout Accepting my portion in life	Group
<b>Conformity:</b> Restraint of actions, inclinations, and impulses likely to upset or harm others and violate social expectations or norms.	Politeness, obedient Honoring parents and elders	Interaction Group
<b>Security:</b> Safety, harmony, and stability of society, of relationships, and of self.	National security Social order, clean	Organism Interaction Group

*Note about Sources.* Organism: universal needs of individuals as biological organisms; Interaction: universal requisites of coordinated social interaction; Group: universal requirements for smooth functioning and survival of groups. \*Emerges when people come into contact with those outside the extended primary group, recognize intergroup interdependence, and become aware of the scarcity of natural resources.

Table 5: This figure is reproduced from (Schwartz, 1994).

Figure 23: Model of Relationship Among Values



**Fig. 1.** Theoretical model of relations among motivational types of values, higher order value types, and bipolar value dimensions (adapted, with permission, from Schwartz, 1992).

Figure 23: This figure is reproduced from (Schwartz, 1994).

Table 6: Schwartz Value Inventory Results

Value	Mean $\pm$ SD
Self-Direction	1.9 $\pm$ 0.89
Universalism	2.0 $\pm$ 0.81
Benevolence	2.0 $\pm$ 0.89
Achievement	2.3 $\pm$ 0.98
Simulation	2.3 $\pm$ 0.95
Security	2.5 $\pm$ 1.11
Hedonism	2.5 $\pm$ 0.88
Tradition	2.7 $\pm$ 0.88
Conformity	2.9 $\pm$ 0.96
Power	3.1 $\pm$ 0.83

Table 6: The SVI questions are phrased as sentences in the first person. Participants read each statement and then rate their answer based on how much the statement is 'like them'. The scale for these data is: 1 – very much like me; 2 – like me; 3 – somewhat like me; 4 – not like me; 5 – not like me at all.

Table 7: Attributes of Affect

	Positive Affect	Negative Affect
High Score	High Energy, Full Concentration, Pleasurable engagement	Aversive mood states (anger, guilt, fear)
Low Score	Sadness and lethargy	Calmness and Serenity

Table 7: The Positive and Negative Affect Schedule (PANAS) measures a person's affect using two orthogonal scales, positive and negative. This table details what scores on each of these scales suggest about a person's mood and behavior. (Watson, Clark et al., 1988).

Table 8: Comparison: Database Participant Affect Scores with Normative Values

	Mean $\pm$ SD	Range
Positive Affect Score: Database	33.4 $\pm$ 7.37	33 [16-49]
Positive Affect: Norm	29.7 $\pm$ 7.9	
Negative Affect Score: Database	13.6 $\pm$ 6.57	31 [10-41]
Negative Affect: Norm	14.8 $\pm$ 5.4	

Table 8: The scale maximum for the PANAS is 50 points. Database participants were not different from normative values: PA ( $p = 0.69$ ), NA ( $p = 0.89$ ).

Table 9: Forsyth's Taxonomy of Ethical Ideologies

	<b>Relativism</b>	
<b>Idealism</b>	High	Low
High	<b>Situationists</b> Rejects moral rules; advocates individualistic analysis of each act in each situation; relativistic.	<b>Absolutists</b> Assumes that the best possible outcome can always be achieved by following universal moral rules.
Low	<b>Subjectivists</b> Appraisals based on personal values and perspective rather than universal moral principles; relativistic.	<b>Exceptionists</b> Moral absolutes guide judgments but pragmatically open to exceptions to these standards; utilitarian.

Table 9: The two scales derived from the Ethical Position Questionnaire are Relativism and Idealism. The combination of scores on each of these measures defines 4 different types of ethical ideologies. This figure is reproduced from (Forsyth, 1980).



Figure 24: Distribution of Ethical Positions

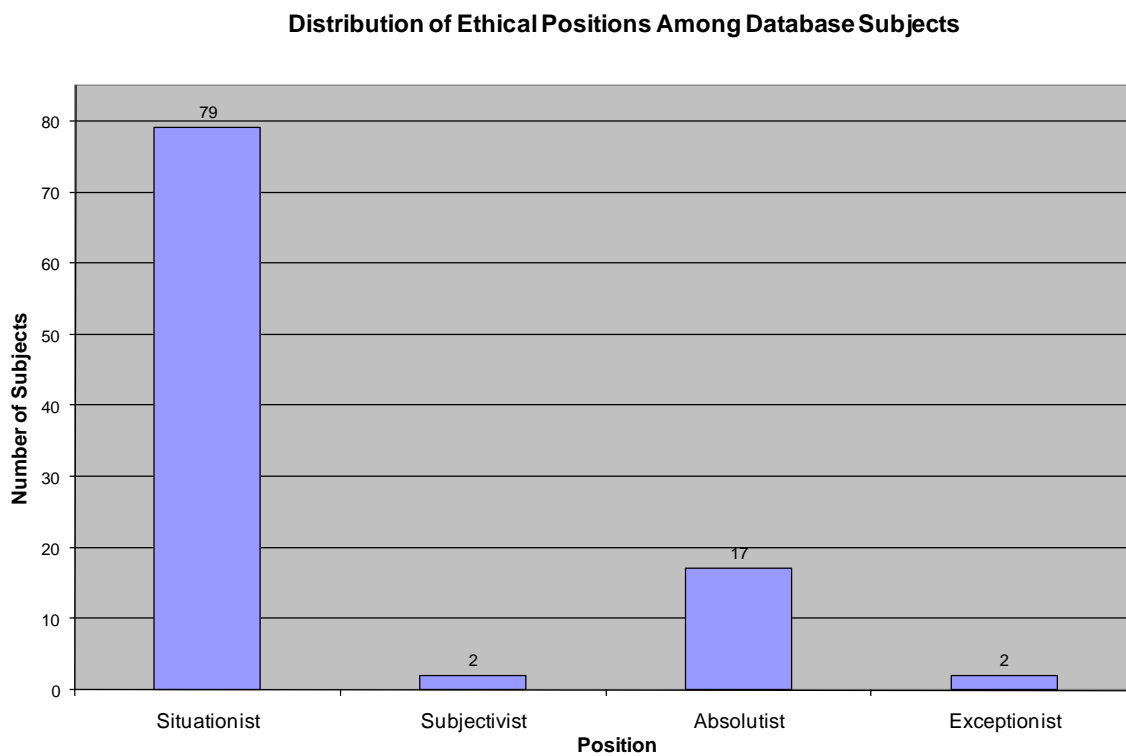


Figure 24: As measured by the EPQ, the majority of our participants were ‘Situationists’ – high scorers on both Idealism and Relativism. This may be caused by an increase in Idealism with age; see Chapter 4.

Figure 25: EPQ Distribution of Relativism Scores




Figure 25: The distribution of participants scores on Relativism(measured by the EPQ) was fairly normally distributed.

Figure 26: EPQ Distribution of Idealism Scores

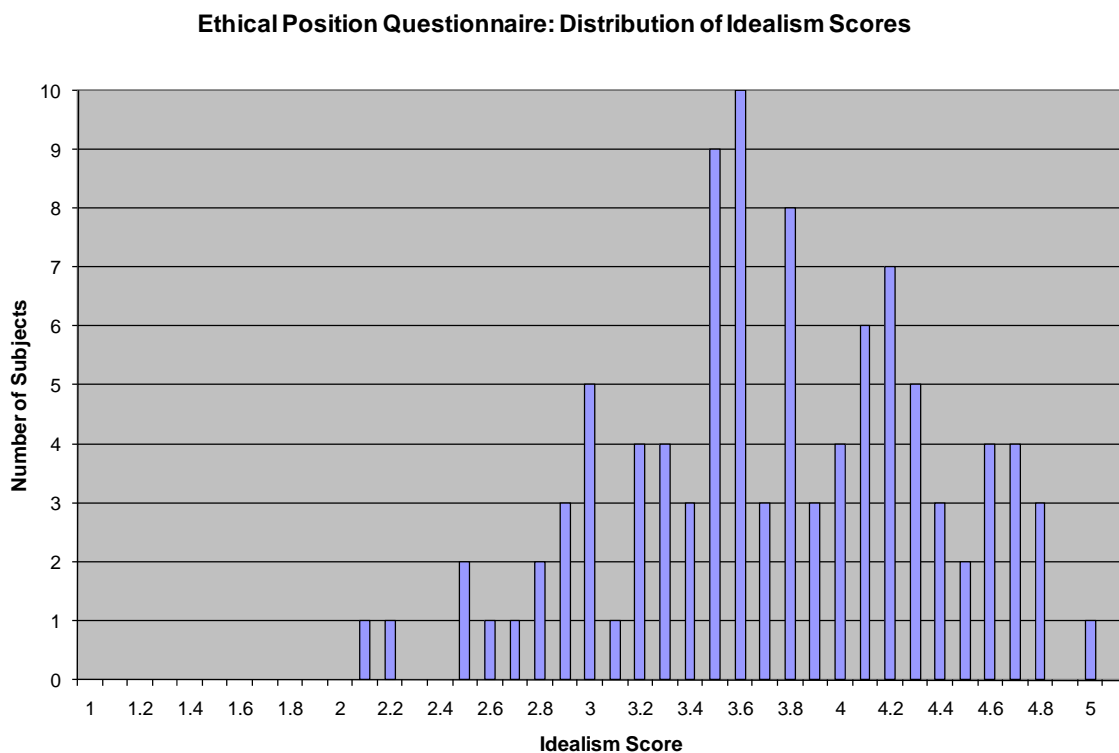


Figure 26: The distribution of Idealism scores (measured by the EPQ) was skewed. This may be due to an increase in general Idealism scores with age.

Figure 27: Utilitarian Score




Figure 27: Participants responded to a set of 33 standard moral dilemmas (see Appendix D). The choices on those dilemmas were used to generate a Utilitarian score – a measurement of the degree to which they espoused utilitarian choices on the dilemmas.

Figure 28: Correlation Matrix Organized by Participant Data

Correlations  
 Marked correlations are significant at  $p < .05000$   
 N=93 (Casewise deletion of missing data)

Variable	Year of Birth	N NEO T Score	E NEO T Score	O NEO T Score	A NEO T Score	C NEO T Score	PoIV	PA PANAS	NA PANAS	Idealism EPQ	Relativism EPQ	Total EPQ	Weighted Utilitarian Score
Year of Birth	1.00	0.24	0.03	0.12	0.19	-0.06	-0.01	-0.04	-0.12	-0.17	-0.07	0.04	0.06
N NEO T Score	0.24	1.00	0.18	0.19	0.49	0.08	-0.05	-0.16	0.07	0.05	0.12	-0.18	0.15
E NEO T Score	0.03	0.18	1.00	0.38	0.17	0.23	0.03	0.06	-0.06	0.02	0.01	-0.08	0.04
O NEO T Score	0.12	0.19	0.38	1.00	0.34	0.10	-0.08	0.07	-0.09	-0.12	0.00	-0.10	0.31
A NEO T Score	0.19	0.49	0.17	0.34	1.00	0.22	0.12	-0.12	0.02	0.01	0.15	-0.21	0.26
C NEO T Score	-0.06	0.08	0.23	0.10	0.22	1.00	0.06	0.22	0.07	-0.00	0.02	-0.02	-0.03
PoIV	-0.01	-0.05	0.03	-0.08	0.12	0.06	1.00	0.06	-0.09	0.05	0.25	-0.25	0.04
PA PANAS	-0.04	-0.16	0.06	0.07	-0.12	0.22	0.06	1.00	0.22	0.31	0.15	-0.09	-0.07
NA PANAS	-0.12	0.07	-0.06	-0.09	0.02	0.07	-0.09	0.22	1.00	0.16	0.19	-0.06	-0.13
Idealism EPQ	-0.17	0.05	0.02	-0.12	0.01	-0.00	0.05	0.31	0.16	1.00	0.17	-0.23	-0.08
Relativism EPQ	-0.07	0.12	0.01	0.00	0.15	0.02	0.25	0.15	0.19	0.17	1.00	-0.78	0.25
Total EPQ	0.04	-0.18	-0.08	-0.10	-0.21	-0.02	-0.25	-0.09	-0.06	-0.23	-0.78	1.00	-0.10
Weighted Utilitarian Score	0.06	0.15	0.04	0.31	0.26	-0.03	0.04	-0.07	-0.13	-0.08	0.25	-0.10	1.00

Figure 28: Correlation coefficients ( $R^2$  values) are given in this matrix comparing background and neuropsychological measures, organized by participant. Schwartz Value Inventory (SVI) data were included in this matrix as well, but were not displayed graphically.

Figure 29: Correlation Matrix Organized by Narrative Data

Correlations  
Marked correlations are significant at  $p < .05000$   
N=701

Variable	Year of Birth	Age of Memory	WC	Sixtr	IQ	N NEO T Score	E NEO T Score	O NEO T Score	A NEO T Score	C NEO T Score	PotV	PA PANAS	NA PANAS	Idealism EPQ	Relativism EPQ	Total EPQ	Weighted Utilitarian Score
mean rating	-0.09	-0.03	0.11	0.03	-0.01	-0.11	-0.02	-0.04	-0.09	-0.01	-0.00	0.01	0.05	0.02	-0.07	0.03	-0.13
Year of Birth	1.00	-0.02	0.03	-0.00	0.04	0.21	0.02	0.10	0.21	-0.06	-0.02	-0.03	-0.14	-0.12	-0.03	-0.01	0.08
Age of Memory	-0.02	1.00	-0.02	0.03	0.04	-0.08	-0.02	-0.03	-0.02	0.01	0.03	0.03	0.02	-0.01	-0.07	0.09	0.03
WC	0.03	-0.02	1.00	0.04	0.16	0.07	0.16	0.06	0.09	0.05	0.08	-0.05	-0.11	-0.01	-0.02	-0.01	-0.07
Sixtr	-0.00	0.03	0.04	1.00	0.16	-0.00	-0.02	-0.06	-0.03	0.02	-0.01	-0.07	-0.02	-0.09	-0.05	0.04	0.03
IQ	0.04	0.04	0.16	0.16	1.00	0.06	-0.15	-0.03	0.05	-0.08	0.10	-0.44	-0.22	-0.48	-0.13	0.19	0.13
N NEO T Score	0.21	-0.08	0.07	-0.00	0.06	1.00	0.18	0.22	0.52	0.07	-0.07	-0.23	0.09	0.06	0.15	-0.21	0.15
E NEO T Score	0.02	-0.02	0.16	-0.02	-0.15	0.18	1.00	0.40	0.20	0.23	0.02	0.10	0.01	0.09	0.08	-0.15	0.06
O NEO T Score	0.10	-0.03	0.06	-0.06	-0.03	0.22	0.40	1.00	0.39	0.13	-0.11	0.08	-0.05	-0.11	0.02	-0.12	0.26
A NEO T Score	0.21	-0.02	0.09	-0.03	0.05	0.52	0.20	0.39	1.00	0.23	0.11	-0.10	0.03	0.07	0.21	-0.27	0.26
C NEO T Score	-0.06	0.01	0.05	0.02	-0.08	0.07	0.23	0.13	0.23	1.00	0.01	0.23	0.11	0.05	0.01	-0.04	-0.07
PotV	-0.02	0.03	0.08	-0.01	0.10	-0.07	0.02	-0.11	0.11	0.01	1.00	0.02	-0.07	0.06	0.22	-0.20	0.06
PA PANAS	-0.03	0.03	-0.05	-0.07	-0.44	-0.23	0.10	0.08	-0.10	0.23	0.02	1.00	0.24	0.32	0.14	-0.09	-0.09
NA PANAS	-0.14	0.02	-0.11	-0.02	-0.22	0.09	0.01	-0.05	0.03	0.11	-0.07	0.24	1.00	0.16	0.19	-0.08	-0.13
Idealism EPQ	-0.12	-0.01	-0.01	-0.09	-0.48	0.06	0.09	-0.11	0.07	0.05	0.06	0.32	0.16	1.00	0.16	-0.24	-0.09
Relativism EPQ	-0.03	-0.07	-0.02	-0.05	-0.13	0.15	0.08	0.02	0.21	0.01	0.22	0.14	0.19	0.16	1.00	-0.79	0.26
Total EPQ	-0.01	0.09	-0.01	0.04	0.19	-0.21	-0.15	-0.12	-0.27	-0.04	-0.20	-0.09	-0.08	-0.24	-0.79	1.00	-0.10
Weighted Utilitarian Score	0.08	0.03	-0.07	0.03	0.13	0.15	0.06	0.26	0.26	-0.07	0.06	-0.09	-0.13	-0.09	0.26	-0.10	1.00

Figure 29: Correlation coefficients ( $R^2$  values) are given in this matrix comparing background and neuropsychological measures, organized by narrative. Schwartz Value Inventory (SVI) data were included in this matrix as well, but were not displayed graphically.

*Chapter 4: Figures and Tables*  
**Section 2**

“Virtue is a state of war, and to live in it we have always to combat with ourselves.”  
-- Jean-Jacques Rousseau

Figure 30: Cue Word Frequency Distribution

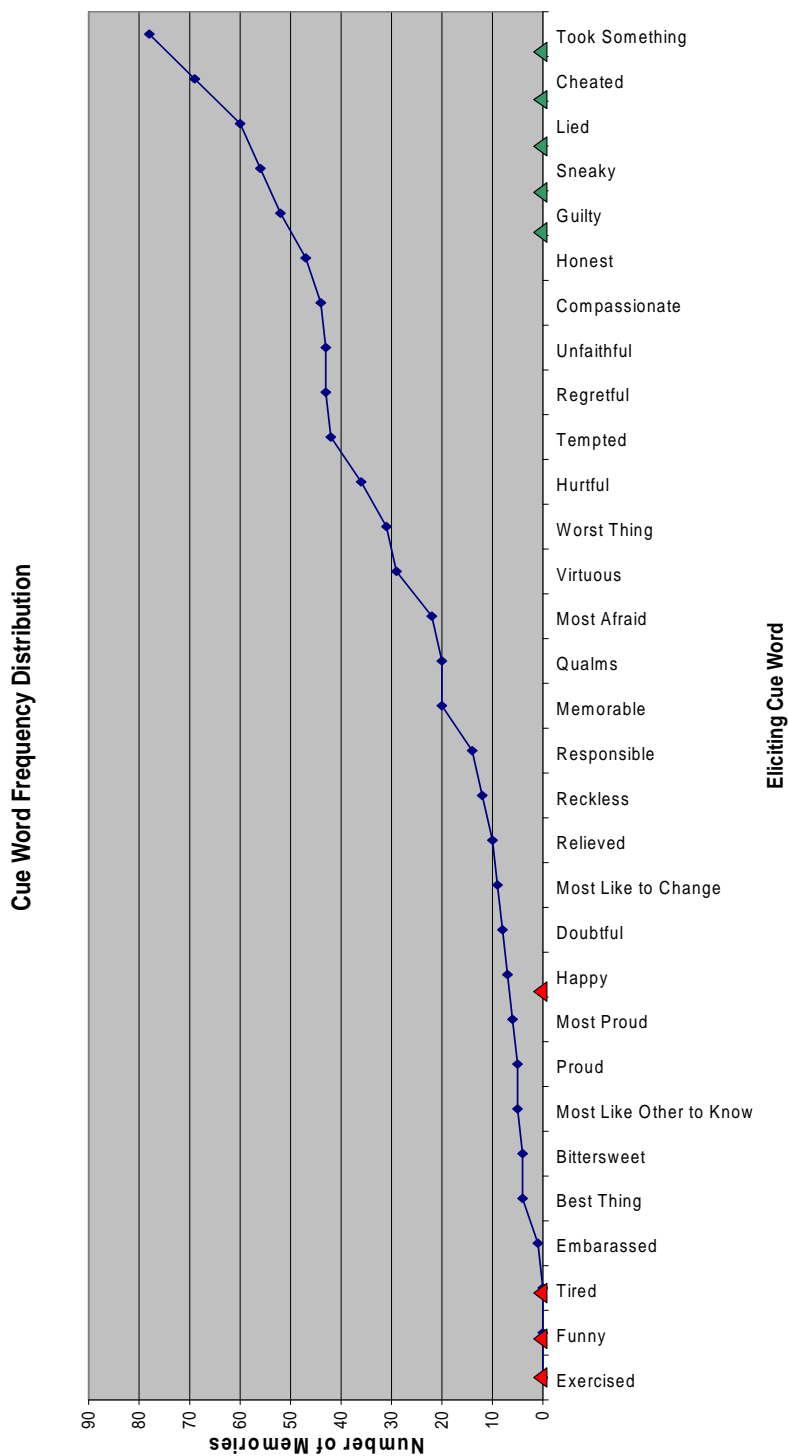


Figure 30: The red triangles on the right indicate the 4 control cue words used. Only one of these cues (happy) elicited any moral memories. The green triangles on the right indicate the 5 cues that elicited the most moral memories.



Figure 31: Cue Frequency Distribution in Presentation Order




Figure 31: The dashed, red lines show the division between each section of the task. Section 1 asked subjects to recall 5 memorable events, section 2 asked subjects to recall memories based on emotional words, section 3 asked subjects to recall memories about actions in their past and section 4 asked subjects to remember the best and worst choices in their lives.

Figure 32: Gender Effects on Cue Word Frequencies




Figure 32: Female (in pink marked with diamonds) and male (in blue marked with squares) subjects generated similar numbers of memories to each of the cue words.

Figure 33: Effect of Age on Cue Word Frequencies




Figure 33: The youngest subjects (aged 40 – 45 years old) (shown in blue, marked with diamonds) and the youngest subjects (aged 55-60 years old) (shown in pink, marked with squares) show similar distribution of cue word frequencies.

Figure 34: Frequency in Lexicon Does Not Determine Number of Elicited Memories




Figure 27: The frequency of a word’s appearance in written English (shown in pink, marked with squares), as measured by the Kucera-Francis standards, does not influence the number of memories that cue words elicit (shown in blue, marked with diamonds). Cue words marked with asterisks use some alternate form of the word for the Kucera-Francis frequency. Also, as seen in Table 1, the cue words *took something*, *embarrassed*, and *qualms* are not included in the graph because frequencies are not available for them. The control words *tired*, *exercised* and *funny* are not included because they did not elicit any moral memories. Finally, the superlative cues were omitted because there are no K-F frequencies available for concepts.

Figure 35: Distribution of Positive and Negative Words in Narratives

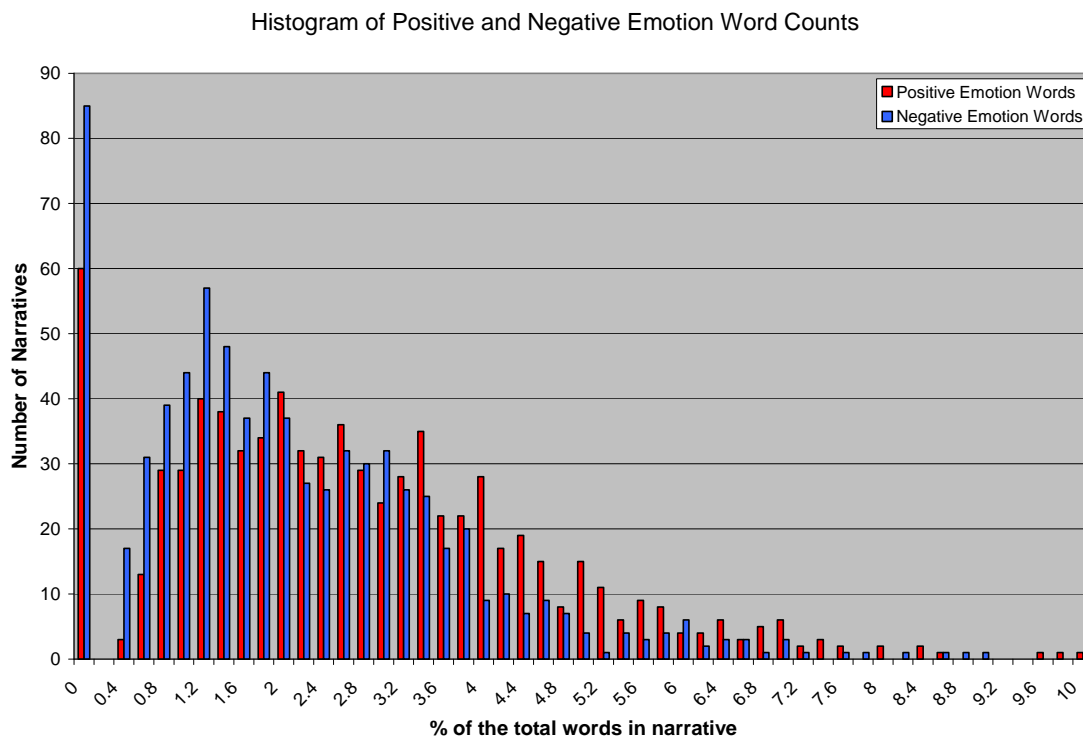


Figure 35: LIWC2007 text analysis was used to calculate the percentage of positive and negative emotion words in each narrative. Two outliers have been removed from this graph. The category 'Positive Emotion Words' has a single outlier at 15.22% and 'Negative Emotion Words' has a single outlier at 11.59%.

Figure 36: Narrative Categorization: What Did They Do?




Figure 36: Raters selected all of the possible categories about what actions occurred in each narrative. The category labels in this graph have been shorted for display; see Appendix G for the complete category labels.

Figure 37: Narrative Categorization: Why Did They Say They Did It?




Figure 37: Raters selected all of the possible categories about why the narrators said that they made the choices described in each narrative.

Figure 38: Narrative Categorization: How Did They Feel?




Figure 38: Raters selected all of the possible categories about how the narrators said that they felt about their choices in each narrative.



Table 10: Location of Action in the Narratives

Scene of the Action	Number of Ratings
Family/Close Friend	813
Home/Neighborhood	633
Job/Business Situation	363
No Scene*	251
School	187
Abroad/Vacation	87
Car	56
Animal	40
Store**	37
Party/Club/Bar	36

Table 10: This table shows the 9 most common locations or situations in which the narratives occurred. Other scenes with fewer than 10 ratings included: Church, City, Clinic, Court, Crowd, Focus Group, Funeral/Death, Gang, Homeless Shelter, Hospital, Internet/Online, Jail, Just Married, Leisure, Little League, Military Training, Movie Set, Movie Theatre, Rehab, and Restaurant.

\*Subjects were not required to choose a scene. 'No Scene' is used when the rater did not choose a scene.

\*\*The scene 'Store' was the only scene that was written in by enough raters to have more than 10 ratings.

Figure 39: Rater Judgments of Narrative Morality




Figure 39: The blue bars (left) show the number of ratings that received the rating score shown on the x-axis. The red bars (right) show the number of narratives with that score as a mean rating. The score -5 was designated as ‘wrong’ on the rater sheet; the score 0 was given as ‘gray area’; and the score +5 was given as ‘right’.

Figure 40: Distribution of Change in Moral Ratings




Figure 40: 45% of participants did not change their feelings about how right or wrong their actions were between the time of the narrative and the time of the rating. Those who did change their feelings did so in a normally distributed way across the entire range of possible changes. The mean change was to lower the rating by 0.5 points, but the variance in this was quite large ( $SD = \pm 2.93$  points).

Figure 41: Change in Moral Behavior Ratings Over Time




Figure 41: A calculation of the difference between the ratings of how participants felt at the time of the event and how participants felt at the time of the rating, averaged over 5 year periods, shows that most participants do not revise their feelings about a situation over time.

*Chapter 4: Figures and Tables*  
Section 3

“I know only that what is moral is what you feel good after and what is immoral is what you feel bad after.”  
-- Ernest Hemingway

Figure 42: Dendrogram of Hierarchical Cluster Analysis of All Narratives

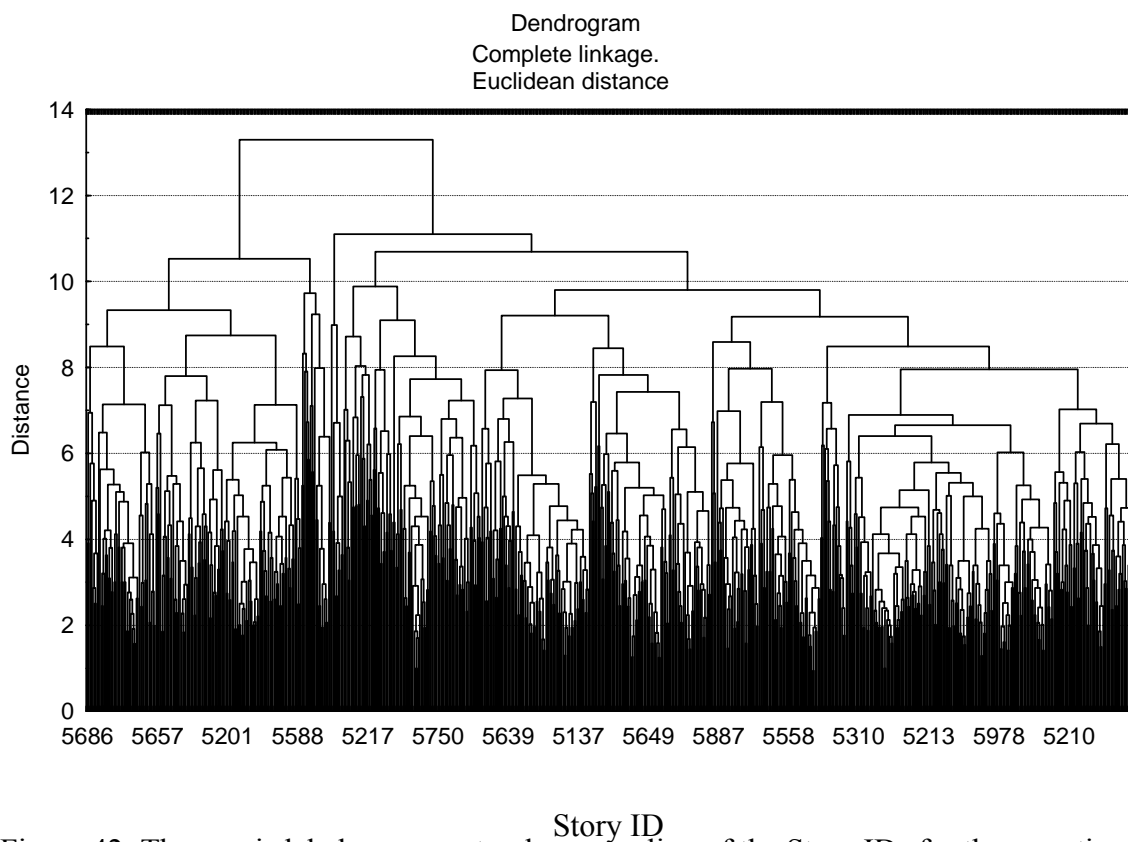


Figure 42: The x-axis labels represent only a sampling of the Story IDs for the narratives. In total, 749 of the narratives are represented across the x-axis. (9 narratives were left out of this analysis.)

Figure 43: Plot of Scree Test

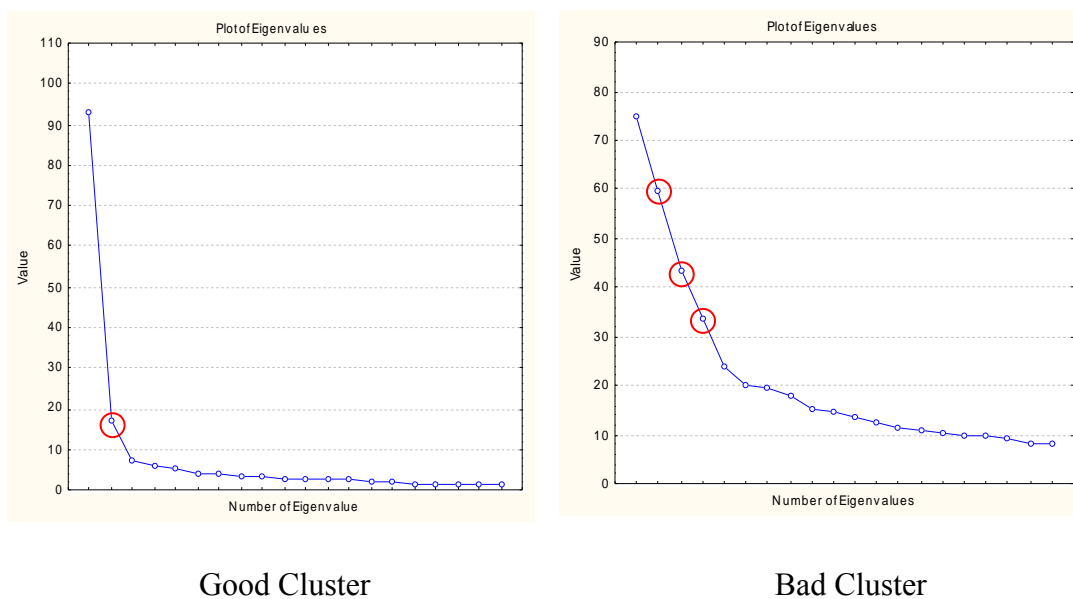


Figure 43: One factor is circled in red for the Good cluster; 3 factors are circled for the Bad cluster.

Figure 44: Factor Loadings for the 3 Bad Clusters




Figure 44a.

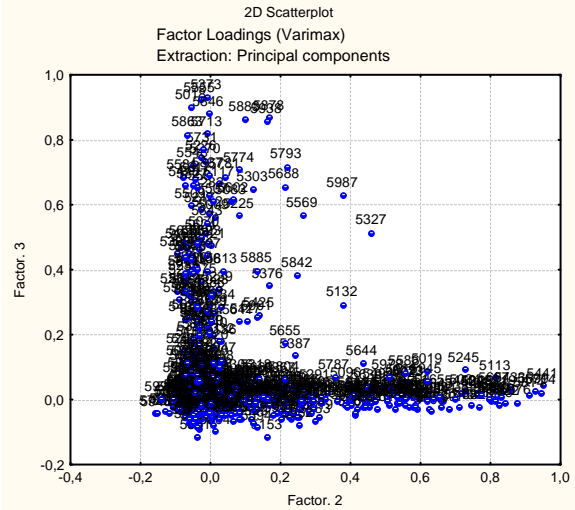


Figure 44b.




Figure 44c.

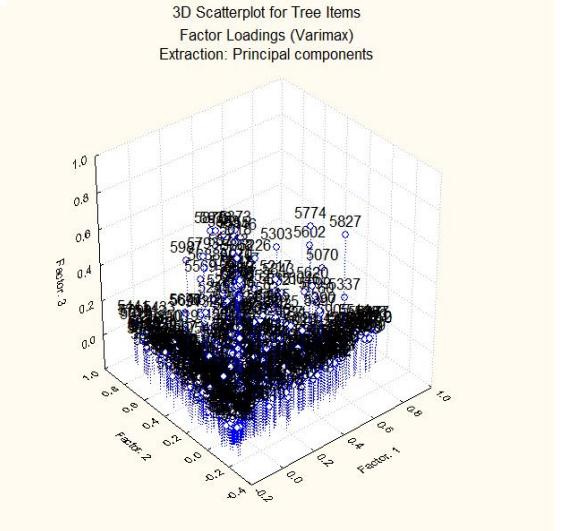


Figure 44d.

Figure 44: Figures a-c are varimax rotations of each of the three factors found from the Bad cluster. Figure d shows the three-dimensional plot of these same data. All plots demonstrate the orthogonality of these three factors.



Figure 45: Varimax Rotation Good Cluster




Figure 45: The varimax rotation of the first two factors from the Good cluster shows a uniform cloud; these data cannot be further separated.

Figure 46: Word Frequency Analysis Bad Factor 1



	Word Frequency
Said	101
Get	101
Know	92
Go	91
One	86
Lie	81
Going	80

Figure 46: 128 narratives were used to construct this word cloud. The most frequent words (after removal of common words) are shown largest in the cloud and quantified in the table.



Figure 48: Word Frequency Analysis Bad Factor 3




Figure 48: 74 narratives were used to construct this word cloud. The most frequent words (after removal of common words) are shown largest in the cloud and quantified in the table.

	Word Frequency
took	84
something	63
money	61
one	52
back	51
really	51
take	51

Figure 49: Word Frequency Analysis Good Narratives




Figure 49: 174 narratives were used to construct this word cloud. The most frequent words (after removal of common words) are shown largest in the cloud and quantified in the table.

	Word Frequency
money	147
know	145
felt	126
people	126
get	123
back	119
one	115
something	112
good	100
feel	98

Figure 50: Distribution of Selected Narratives from each Bad Cluster

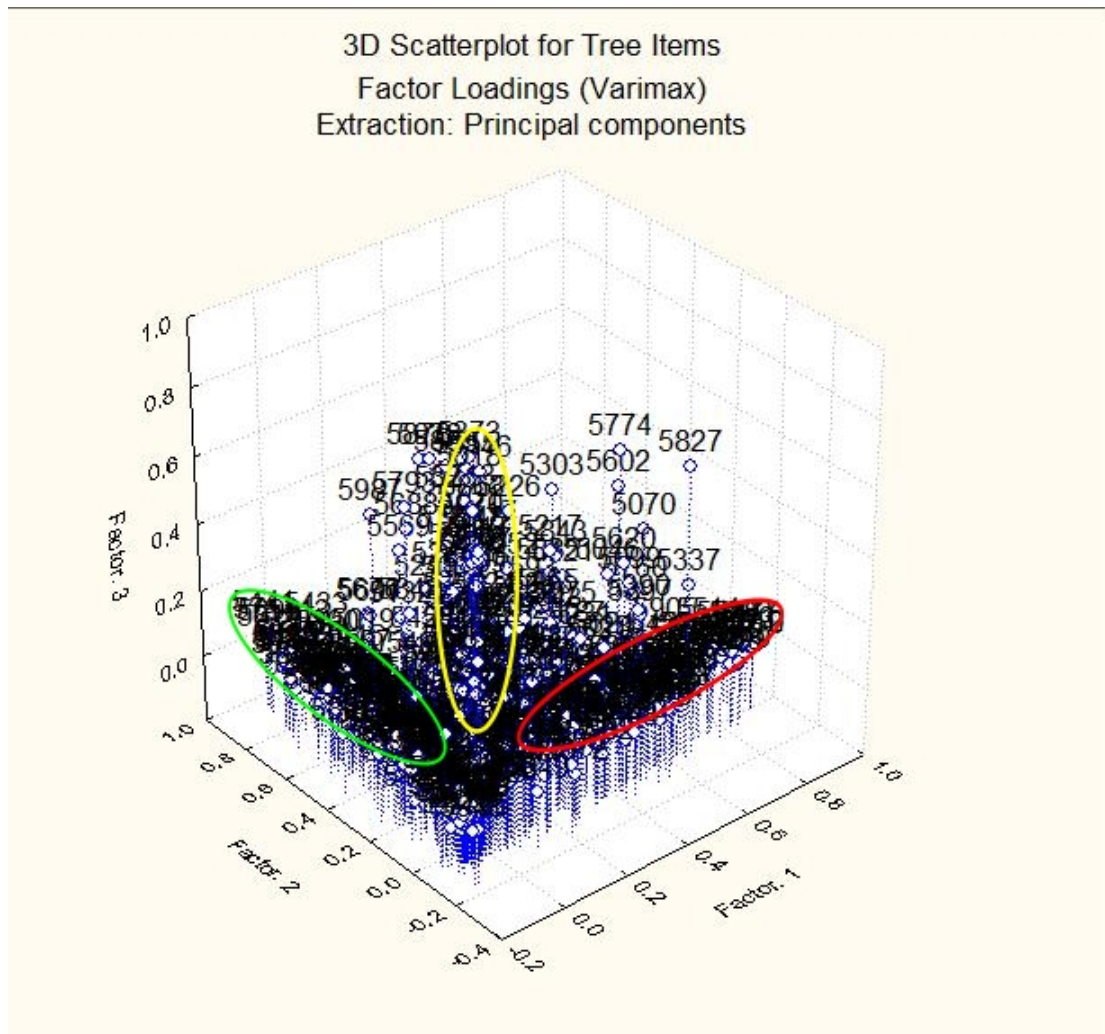


Figure 50: This figure shows the same 3-dimensional factor loading for the bad clusters as above, with the approximate areas selected for further word frequency analysis shown in the colored circles. (The actual areas selected would be represented by cubes in this plot.)

Table 11: Rating Categories that Contribute to Each Bad Factor

Factor 1	*	Factor 2	*	Factor 3	*
Selfishness	124	Regret about behavior	157	Regret about doing the wrong thing	171
Not taking responsibility for actions	99	Regret about doing the wrong thing	146	Regret about behavior	170
Lying	92	Shame/ashamed	131	Breaking a rule/law	141
Tricking someone else	88	Being hurtful to someone	121	Shame/ashamed	135
Doing the wrong thing	84	Being unkind to someone	103	Doing the wrong thing	134
Breaking a rule/law	79	Knowledge of personal moral weakness	101	Knowledge of personal moral weakness	123
Knowledge of personal moral weakness	74	Doing the wrong thing	98	Selfishness	111
Uncertainty about moral correctness of action	68	Taking responsibility for actions	93	Stealing/Spending someone's money	92
Taking advantage of someone	68	Learning a lesson	76	Learning a lesson	84
Because of money	65	Selfishness	75	Taking responsibility for actions	83

\*This the sum of all of the rating values on this item in this factor. Larger numbers reflect more influence of this rating on this cluster.

Table 11: This table depicts the contribution of each of these rating categories to each factor. Shown in green are the key concepts for each factor.

Table 12: Top 5 Eliciting Cue Words

Factor 1 (n = 128)	# Narratives	Factor 2 (n = 121)	# Narratives	Factor 3 (n = 74)	# Narratives
Lied	49	Hurtful to Someone	27	Took Something	46
Sneaky	26	Regretful	22	Guilty	9
Cheated	22	Worst Thing	14	Sneaky	7
Unfaithful	11	Guilty	11	Cheated	6
Guilty	8	Unfaithful	11	Most Afraid	2

Good Narratives (n = 174)	# Narratives
Compassionate	39
Honest	34
Virtuous	26
Tempted	14
Responsible	11

Table 12: These tables show the cue words that elicited the most narratives in each of the four factors (good, bad factor 1: lying, bad factor 2: hurting someone else and bad factor 3: stealing.)



*Chapter 4: Figures and Tables*  
**Section 4**

“Wisdom, compassion, and courage are the three universally recognized moral qualities of men.”  
-- Confucius

Figure 51: Select Query Structure for Gender Ratings

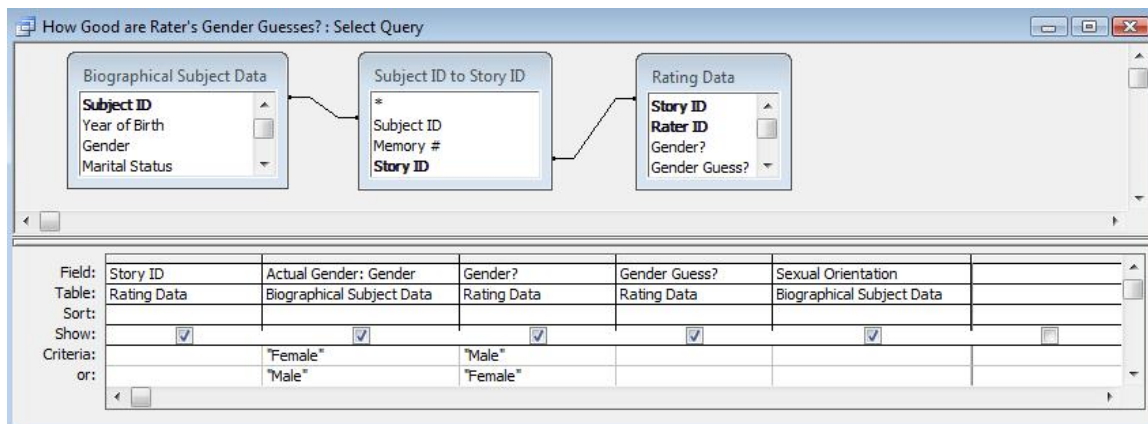


Figure 51: The design view for the select query “How Good are Rater’s Gender Guesses?” shows the fields needed to implement this query as well as the linking done through the Subject ID to Story ID table.

Table 13: Outcome from Data Mining Example

Table 13A

Ratings	All Raters Wrong	Raters Mixed	
No Guess	19 ratings	104 ratings	
Mixed Guess	108 ratings	438 ratings	
<b>Totals</b>	127 ratings	542 ratings	669 ratings

Table 13B

Narratives	All Raters Wrong	Raters Mixed	
No Guess	2 narratives	7 narratives	
Mixed Guess	56 narratives	165 narratives	
<b>Totals</b>	58 narratives	172 narratives	230 narratives

Table 13: These tables show the results of the raters' gender guesses on the rating tool. (A) 19 ratings map onto 17 narratives; 15 of those narratives have other raters who were unsure (guessers), so they move into the mixed guess category. (B) 104 ratings map onto 69 narratives. Only 7 of those have the unique distinction of having raters who were all sure. This only includes ratings/narratives told by self-identified heterosexual narrators.

*Chapter 4: Figures and Tables*  
Section 5

“Hence naturally flows the great variety of opinions concerning moral rules which are to be found among men, according to the different sorts of happiness they have a prospect of, or propose to themselves; which could not be if practical principles were innate, and imprinted in our minds immediately by the hand of God.”

-- John Locke

Table 14: Background and Neuropsychological Data

Participant Group	Number in Group	Mean Age at Testing	IQ (WAIS-III FSIQ)	Wechsler Memory Scale – III	Education (years)	Gender	Handedness	Chronicity (years since onset at time of testing)
Bilateral Amygdala Lesion	1	45	88	89	12	1 F	RH	Developmental Damage
Amnesic (Anoxia and Encephalitis)	5 (3,2)	52.8 ± 3.96	95.6 ± 7.73	61.6 ± 12.22	15.6 ± 3.58	4 M, 1 F	3 RH, 2 LH	8.5 ± 4.51
Frontal Lobe Damage	5	61.4 ± 6.43	107 ± 26.08	98.75 ± 30.5	12.3 ± 2.87	4 M, 1 F	5 RH	11 ± 11.90
Developmental Frontal Lobe Damage	2	26 ± 7.07	Unk	Unk	14.5 ± 2.12	1 M, 1 F	Unk	Developmental Damage
Brain Damaged Controls (Temporal Lobectomies)	8	43.8 ± 14.14	96.13 ± 10.23	92.23 ± 14.5	13.3 ± 1.83	2 M, 6 F	6 RH, 2 LH	4.6 ± 3.82
Iowa Normals	6	Unk	Unk	Unk	Unk	Unk	Unk	No Damage

Unk=Unknown

Table 14: The demographic and neuropsychological details about each participant group in the Iowa testing procedure are shown above.

Figure 52: Number of Words per Narrative




Figure 52: Number of Words per Narrative: All participants in the Iowa testing procedure produce significantly shorter narratives than the participants tested in the main study ( $t[779] = 4.40, p < 0.0001$ ).

Figure 53: Number of Moral Narratives Told




Figure 53: Participants in the Iowa testing procedure recalled fewer moral memories than participants in the main study. There may be large individual differences that contributed to this finding; the main study had a range from 3 to 17 moral narratives produced.

Figure 54: LIWC Analysis of Positive and Negative Emotion Words




Figure 54: The Iowa testing procedure did not affect the emotional content of participants' recollections. The moral narratives that were told by participants in the Iowa study contained similar numbers of positive and negative emotion words to the main study.



*Chapter 4: Figures and Tables*  
**Section 6**

“The first step in the evolution of ethics is a sense of solidarity with other human beings.”  
-- Albert Schweitzer

Figure 55: Cue Word versus Mean Age of Memory




Figure 55: Moral memories elicited by positively valenced cues are more recent. Distribution of the temporal remoteness of memories as evoked by specific cue words (means). Positive and negative memories evoked by the itself valence-neutral cue “most memorable” are highlighted in red.

Figure 56: Valence of Cue Word Predicts Age of Memory




Figure 56: Emotional valence of the cue words is associated with the age of the memory they evoke.

Figure 57: Valence of Rating Impacts Recency of Memory




Figure 57: Memories with positive moral content are more recent. We plotted the mean (and SEM) recency of all those memories that were classified by independent raters as falling into moral categories with positive or negative valence. Contrasts between each of the oppositely valenced pairs of categories were significant ( $p < 0.001$ , for all three pairs). The contrast when comparing all three negative categories to all three positive categories was also significant ( $p < 0.001$ ).

Table 15: Demographics and Neuropsychological Profiles

Age	Gender	Ethnicity	Education (Highest Level Attained)	Handedness	Sexual Orientation	Current Religion <sup>o</sup>
Mean 48.93 ± 5.91 years	Male n=47	Caucasian n=50	Elementary School n=1	Right Handed n=83	Heterosexual n=85	Christian n=49
Range 40 to 60	Female n=53	Black* n=11	Middle School n=0	Left Handed n=14	Homosexual n=4	Jewish n=3
		Asian/Pacific Islander n=10	High School n=9	Ambidextrous n=3	Bisexual n=3	Buddhist n=3
		Hispanic** n=25	Some College n=41		Polyamorous n=1	Agnostic n=1
		Multiracial n=3	Bachelor's Degree n=31		Other n=1	Multiple Religious Practices n=5
		American Indian n=1	Master's Degree n=13		No Answer n=6	Other n=14
			Graduate Degree n=5			None n=25

Table 15 Continued

IQ	NEO raw score (mean, SD)	PANAS (mean, SD)	Ethical Position Questionnaire (mean, SD)
Mean 110.52 ±13.42	Neuroticism 31.55 ± 4.75	Positive Affect 33.97 ± 7.37	Idealism 3.74 ± 0.63
Range 77 to 138	Extraversion 39.95 ± 3.97	Negative Affect 13.61 ± 6.57	Relativism 3.076 ± 0.66
	Openness 35.95 ± 3.57		
	Agreeableness 36.96 ± 5.12		
	Conscientiousness 41.35 ± 3.27		

Table 15: The overall demographic and neuropsychological profiles of the participant population is shown above. °No participants were practicing Hindus or Muslims. \*Black included participants of African and African-American descent. \*\*Hispanic was given as a separate ethnicity choice because of the large Latino community in the Los Angeles area. Participants were free to mark multiple ethnicities (e.g., Caucasian and Hispanic) if they felt that best described themselves.

Table 16: Cue Words Used in Recall Protocol

Cue	Number of Moral Memories	ANEW Valence Rating [Mean, (SD)]	Independent Valence Ratings [Mean, (SD)]	Kucera-Francis Word Frequency
Best Thing	4	*	8.80 (0.41)	*
Bittersweet	3		5.32 (1.06)	1
Cheated	66		1.23 (0.47)	<i>Cheat:</i> 3
Compassionate	42		8.39 (0.78)	2
Doubtful	7		3.07 (0.41)	22
Embarrassed**	1	3.03 (1.85)	1.64 (0.67)	8
Guilty	51	2.63 (1.98)	1.64 (0.67)	29
Happy	7	8.21 (1.82)	8.59 (0.47)	98
Honest	45	7.70 (1.43)	8.18 (0.67)	47
Hurtful to Someone	34	<i>Hurt:</i> 1.90 (1.26)	0.82 (0)	<i>Hurt:</i> 37
Lied	56	<i>Lie:</i> 2.79 (1.92)	1.64 (0.67)	<i>Lie:</i> 59
Memorable	18	----	----	----
Most Afraid Others Will Find Out	20	* <i>Afraid:</i> 2.00 (1.28)	1.23 (0.47)	* 57
Most Like Others to Know	5	*	7.77 (0.47)	*
Most Like to Change	7	*	1.43 (0.78)	*
Most Proud	5	*	8.59 (0.47)	*

Table 16 Continued

Cue	Number of Moral Memories	ANEW Valence Rating [Mean, (SD)]	Independent Valence Ratings [Mean, (SD)]	Kucera-Francis Word Frequency
Proud	5	8.03 (1.56)	8.80 (0.41)	50
Qualms	18		2.66 (1.23)	--
Reckless	10		2.25 (1.93)	9
Regretful	36	2.28 (1.42)	1.64 (0.67)	1
Relieved	9		7.57 (0.78)	<i>Relieve:</i> 13
Responsible	13		7.77 (1.42)	71
Sneaky	56		2.25 (0.78)	2
Tempted	39		3.68 (1.06)	<i>Tempt:</i> 2
Took Something	78		1.43 (0.78)	¥
Unfaithful	41	2.05 (1.55)	0.82 (0)	1
Virtuous	29	<i>Virtue:</i> 6.22 (2.06)	7.98 (0.41)	<i>Virtuous:</i> 6 <i>Virtue:</i> 30
Worst Thing	29	*	0.82 (0)	*

Table 16: This table shows a list of the cue words and the number of moral memories that each cue word elicited. (Three control cues (tired, exercised and funny) are not included in the table.) The second column gives the valence ratings for each available cue taken from the Affective Norms of English Words (ANEW)(Bradley and Lang 1999). These can be compared to the independent rating data that were collected for this study in column three. The final column gives the frequencies of each word in standard written English using the Kucera-Francis frequencies (Kucera & Francis, 1967).

\*Superlative cues were phrases and therefore do not have an ANEW rating.

\*\* Embarrassed was excluded from the analyses because it was an outlier.

° The word “Memory” does have an ANEW (6.62 (1.50)) and K-F rating (76) but the testing cue ‘memorable’ was somewhat different from the other cues and not intended to evoke recollections based on the cue itself.

¥ Kucera Francis Frequency for steal: 5.



*Appendices*

“It is in everybody's interest to seek those [actions] that lead to happiness and avoid those which lead to suffering. And because our interests are inextricably linked, we are compelled to accept ethics as the indispensable interface between my desire to be happy and yours.”

-- Dalai Lama

## Appendix A: The Trolley Problem

In 1985, Judith Jarvis Thomson<sup>287</sup> revived a philosophical thought experiment originally posed by Philippa Foot<sup>288</sup>:



Suppose you are the driver of a trolley. The trolley rounds a bend, and there come into view ahead five track workmen, who have been repairing the track. The track goes through a bit of a valley at that point, and the sides are steep, so you must stop the trolley if you are to avoid running the five men down. You step on the brakes, but alas they don't work. Now you suddenly see a spur of track leading off to the right. You can turn the trolley onto it, and thus save the five men on the straight track ahead. Unfortunately, Mrs. Foot has arranged that there is one track workman on that spur of track. He can no more get off the track in time than the five can, so you will kill him if you turn the trolley onto him. Is it morally permissible for you to turn the trolley?

Thomson and Foot felt that this action, although terrible, was morally permissible. The trolley driver is saving five lives by sacrificing one. Foot originally posed the trolley problem an example of the doctrine of double effect, a Catholic doctrine which claims that it is never acceptable to do intentional harm, but actions that cause unintentional bad side-effects are acceptable. She was unsure whether this doctrine could be supported, ethically, and suggested a number of other thought experiments for comparison<sup>289</sup>. Utilitarianism would also suggest that it is always better to kill one person to save five. However, Thomson posed a number of alternative cases, where our intuition is just the opposite – killing one to save five is morally impermissible. The most frequently compared example is one that Thomson called the ‘Fat Man’. (Typically this is now called the ‘Footbridge’.)

<sup>287</sup>Thomson, J. J. (1985). "The Trolley Problem." *Yale Law Journal* **94**: 1395-1415.

<sup>288</sup>Foot, P. (1978). *The Problem of Abortion and the Doctrine of the Double Effect in Virtues and Vices and Other Essays in Moral Philosophy*. Berkeley and Los Angeles, University of California Press.

<sup>289</sup>Foot and Thomson each designed comparative thought experiments but only Thomson's thought experiments will be discussed here, as they are the ones have been used more generally.

You are standing on a footbridge over the trolley track. You can see a trolley hurtling down the track, out of control. You turn around to see where the trolley is headed, and there are five workmen on the track where it exits from under the footbridge. What to do? Being an expert on trolleys, you know of one certain way to stop an out-of-control trolley: Drop a really heavy weight in its path. But where to find one? It just so happens that standing next to you on the footbridge is a fat man, a really fat man. He is leaning over the railing, watching the trolley; all you have to do is to give him a little shove, and over the railing he will go, onto the track in the path of the trolley. Would it be permissible for you to do this?

Thomson's intuition (which is supported by experimental work by Hauser<sup>290</sup>) is that it is not permissible to push the fat man off the footbridge, even though by killing one person, five can be saved. This intuitive distinction between the morally appropriate action in these two cases has prompted moral philosophers and moral neuroscientists alike to investigate.

A number of other variants on the trolley problem exist, including:

'The Bystander at the Switch' – in which the reader is a bystander next to a switch for the trolley, rather than the driver choosing a track;

'The Loop' – in which the reader is standing next to a looped track with a fat workman, where switching the train will kill the fat workman, saving the five further down the line;

"The Man in the Yard" – in which you can prevent the trolley from killing five people by colliding it with another trolley, derailing both and subsequently causing them to fall down a hill on top of a man sleeping in his hammock in his own yard<sup>291</sup>.

---

<sup>290</sup> Hauser, M. D., F. Cushman, et al. (2007). "A dissociation between moral judgments and justifications." *Mind & Language* 22(1): 1-21.

<sup>291</sup> Unger, P. (1996). *Living High and Letting Die*. Oxford, Oxford University Press.

There are also a number of other moral dilemmas that are used for contrast with these trolley problems. Thomson's original contrast for the trolley problem involved a doctor choosing whether to take the organs from one healthy patient to cure five sick patients, another situation where killing one person is not justified by saving five others.

Additional examples of these dilemmas can be seen in Appendix D.

## Appendix B: Testing Procedure for Session Two

### Part I:

#### *Read to the subject:*

First, we would like to know about 5 things that you remember doing in your past. These can be memories from any time in your life – early childhood right up until a year or so ago. They should be memories that particularly involve you, and something that you did, a decision you made, or an action you took. Don't spend too much time thinking about all the different memories. We are interested in the first 5 that come to mind as vivid, memorable events. Try to pick events that you can describe in detail. We'd like you to try to think of both positive and negative events, but whatever stands out most in your mind is the best. Take about 3 minutes to tell each story. You don't have to fill the whole 3 minutes and you can take longer if you need more time. *Instruct the subject about how they will know when three minutes are done. (i.e., a small clock or timer).* You will have as long as you like to break between memories. Now you can have a few minutes to think of 5 things that you have done that you remember particularly well.

Do you have any questions?

#### *Memory One*

Please tell us about a very memorable thing that you have done.

#### *Memory Two*

Please tell us about another very memorable thing that you have done.

#### *Memory Three*

Please tell us about another very memorable thing that you have done.

#### *Memory Four*

Please tell us about another very memorable thing that you have done.

#### *Memory Five*

Please tell us about another very memorable thing that you have done.

Good work! You've finished Part I.

**Part II:**

Now, we would like to know about things that you have done that made you feel an emotion very strongly. The computer will give you an emotion word and you will tell us about a time when you did something that made you feel that way. For example, if the word is “joy”, you might tell a story about how you felt joyful when you got married. These should be memories of things that YOU have done (decisions that you made or actions you took) so you wouldn’t want to tell a story about how you felt joyful when Sonny married Cher, unless you are Sonny or Cher.

Just like in the last section, you should try to make your memory about 3 minutes long.

Do you have any questions?

*Memory One*

Please talk about a time when you did something that made you feel relieved.

*Memory Two*

Please talk about a time when you did something that made you feel doubtful.

*Memory Three*

Please talk about a time when you did something that made you feel embarrassed.

*Memory Four*

Please talk about a time when you did something that made you feel regretful.

*Memory Five*

Please talk about a time when you did something that made you feel happy.

*Memory Six*

Please talk about a time when you did something that made you feel proud.

*Memory Seven*

Please talk about a time when you did something that made you feel guilty.

*Memory Eight*

Please talk about a time when you did something that made you feel tired.

*Memory Nine*

Please talk about a time when you did something that made you feel virtuous.

*Memory Ten*

Please talk about a time when you did something that made you feel compassionate.

*Memory Eleven*

Please talk about a time when you did something that made you feel responsible.

*Memory Twelve*

Please talk about a time when you did something that made you feel bittersweet.

Good work! You've finished Part II.

**Part III:**

In this part of the study, we would like to know about moments in your life where you acted in a particular way. We will give you the actions that you might have done and ask you to tell us the most important memory that comes to your mind of doing this. Again, the memory could be from any time in your life. If you have never done the action that is presented, you may skip the question. However, please try not to skip more than one question. As before, please try to make each of your stories about 3 minutes long.

Do you have any questions?

*Memory One*

Please talk about a time when you did something you have qualms about.

*Memory Two*

Please talk about a time when you cheated.

*Memory Three*

Please talk about a time when you took something that didn't belong to you.

*Memory Four*

Please talk about a time when you were unfaithful.

*Memory Five*

Please talk about a time when you were funny.

*Memory Six*

Please talk about a time when you did something sneaky.

*Memory Seven*

Please talk about a time when you were reckless.

*Memory Eight*

Please talk about a time when you exercised.

*Memory Nine*

Please talk about a time when you were tempted.

*Memory Ten*

Please talk about a time when you were honest.

*Memory Eleven*

Please talk about a time when you lied.

*Memory Twelve*

Please talk about a time when you were hurtful to someone else.

Good work! You've finished Part III.

**Part IV:**

In this final section, we would like to hear about personal situations in your life. These may be very private memories but remember that everything you say here is confidential. If you do not want to answer a question, you may say so and skip it. As before, we would like you to take about 3 minutes for each story.

Do you have any questions?

*Memory One*

Please talk about the thing that you are most proud of having done.

*Memory Two*

Please talk about the best thing you ever did.

*Memory Three*

Please talk about the thing that you would most like to change if you could go back and do it differently.

*Memory Four*

Please talk about the thing that you are most afraid that other people will find out that you did.



*Memory Five*

Please talk about the thing that you would most like other people to know that you did.

*Memory Six*

Please talk about the worst thing you ever did.

Good work! You've finished!

*For the experimenter:*

*Please feel free to use the space below to make any relevant notes about the testing session.*

## Appendix C: Session Three Testing Questionnaire

### Part I:

1. When did this event occur? Please give the month, date and year to the best of your ability. If you are unsure, please give your best guess of your age at the time this event occurred. \_\_\_\_\_

2. On a scale from 1 to 10, how pleasant was this experience?

←-----→  
 not pleasant at all      neither pleasant nor unpleasant      extremely pleasant

3. On a scale from 1 to 10, how vivid was this experience?

←-----→  
 not vivid at all      somewhat vivid      extremely vivid

4. On a scale from 1 to 10, how intense was this experience?

←-----→  
 not intense at all      somewhat intense      extremely intense

5. On a scale from 1 to 10, how emotional was this experience for you?

←-----→  
 not emotional at all      somewhat emotional      extremely emotional

6. Please choose the emotion from the following list that best describes how this experience made you feel. You may choose to rate up to three emotions -- the emotion that best describes how you feel should be rated '1', the second '2' and the third '3'.

Happy	Angry	Embarrassed	Proud
Surprised	Disgusted	Guilty	Jealous
Afraid	Sad	Ashamed	Regretful

7. On a scale from 1 to 10, how significant was this experience in your life?

←-----→  
 not significant at all    neither significant nor insignificant    extremely significant

8. On a scale from 1 to 10, how confidential do you consider this memory to be?

Think in terms of how much you would want to share this memory with anyone else. A rating of 10 would mean that you would never tell this to anyone else, even your closest friend, and a rating of 1 would mean that you are willing to tell everyone this story and/or it doesn't matter if other people know about it.

←-----→  
 not confidential at all      somewhat confidential      extremely confidential

9. On a scale from 1 to 10, how much do you regret this experience now?

←-----→  
 do not regret at all      regret somewhat      regret a great deal

10. On a scale from 1 to 10, how guilty do you feel now about your behavior during this experience?

←-----→  
 not guilty at all                      somewhat guilty                      extremely guilty

11. On a scale from 1 to 10, how embarrassed are you now about your behavior during this experience?

←-----→  
 not embarrassed at all              somewhat embarrassed              extremely embarrassed

12. On a scale from 1 to 10, how proud are you now about your behavior during this experience?

←-----→  
 not proud at all                      somewhat proud                      extremely proud

#### Part II:

Next we will ask you to make some moral judgments about your memory. These are judgments about how right or wrong, good or bad, your actions were. Please read each question carefully.

1. **At the time**, how wrong or right did your actions in this episode seem to you? Use a scale from 1 to 10. If, at the time you did it, you felt this was a very wrong thing to do, you would give it a 1. If, at the time you did it, you felt this was a very right thing to do, you would give it a 10.
2. **At the time**, did you feel like you had options about what you could do in the situation? Use a scale from 1 to 10. If it seemed at the time that you could not really do anything else, then you would give it a 1. If it seemed, at the time, that you had other options and could just as easily have done something else, then you would give it a 10. If it seemed like you could have done something else, but it would have been difficult, then you might give it a 5 or a 6.
3. To the best of your recollection, approximately how many distinctly different things could you have done at the time?
4. **Immediately after** you made the decision and did these actions, how content were you with your actions? If, for example, you thought that you made a hard decision but were completely satisfied with your actions, you would give it a 1 on a scale from 1 to 10. If you were completely discontent with your actions, you would give it a 10.
5. **At the time**, how hard was it, morally, for you to do what you did? We are interested in how much of a psychological struggle it was for you to take that particular action. If, at the time, you found it very easy and did it without having to think or worry about it much, you would give it a 1. If, at the time, you found it to be one of the most difficult decisions in your life and you deliberated and agonized over it before doing it, then you would give it a 10.
6. **At the time**, did you second-guess your decision?

7. To the best of your recollection, how many times in the year following this episode did you think about this decision?
8. To the best of your recollection, how many times in the year following this episode did you wish that you'd done something different?

Now we are going to ask you some similar questions but instead of thinking about how you felt back then, please think about how you feel today, sitting in this room, when you read the memory.

9. **Reading this today**, how wrong or right do your past actions in this episode seem to you? Use a scale from 1 to 10. If now you feel this was a very wrong thing to do, you would give it a 1. If now you feel this was a very right thing to do, you would give it a 10.
10. **Reading this today**, do you feel like you had options about what you could have done in the situation? Use a scale from 1 to 10. If it seems now that you could not really have done anything else, then you would give it a 1. If it seems now that you had other options and could just as easily have done something else, then you would give it a 10. If it seems like you could have done something else, but it would have been difficult, then you might give it a 5 or a 6.
11. **Thinking about the situation today**, approximately how many distinctly different things could you have done at the time?
12. **Today**, do you think there was something morally better that you could have done?
13. **Today**, do you think there was something morally better that you should have done?
14. **Today**, how content are you with your past actions? If, for example, you thought that you made a hard decision but you are still completely satisfied with your actions, you would give it a 1 on a scale from 1 to 10. If you are now completely discontent with your past actions, you would give it a 10.
15. **If you were in the same situation again**, how hard would it be, morally, for you to decide what to do? If, today, you would find it very easy and would not have to think or worry about it much, you would give it a 1. If, today, you would find it to be one of the most difficult decisions in your life and you would have to deliberate and agonize over it before doing it, then you would give it a 10.
16. Do you still think about whether you made the right decision in this situation?
17. How many times in this past year have you thought about the decision you made in this episode? (Do not include thinking about it since you began this study.)
18. When you think about this episode **today** do you wish that you'd done something different?

Now we're going to ask you to imagine what most other people might think if they were reading this exact memory. These people are strangers to you. They are just other normal people who might be in a research study like you and the ONLY information they have is the description of the situation that you are going to read. Please try to evaluate the next questions from most other people's perspective.

19. **If other people read this situation**, how wrong or right would the actions in this episode seem to them? On a scale from 1 to 10, if you think people would generally feel

this was a very wrong thing to do, you would give it a 1. If you think people would generally feel this was a very right thing to do, you would give it a 10.

20. **If other people read this situation**, would they feel like you had options about what you could do in the situation? On a scale of 1 to 10, if you think people reading this would think that you could not really have done anything else, then you would give it a 1. If you think people would think that you had other options and could just as easily have done something else, then you would give it a 10. If other people would think that you could have done something else, but it would have been difficult, then you might give it a 5 or a 6.

21. Would **other people** think there was something morally better that you could have done?

22. Would **other people** think there was something morally better that you should have done?

23. **If other people reading this story were rating on a scale of 1 to 10**, how hard would they say that it was, morally, for you to decide what to do? If other people would think that you would find it very easy and would not have to think or worry about it much, you would give it a 1. If other people would think that you would find it to be one of the most difficult decisions in your life and you would have to deliberate and agonize over it before doing it, then you would give it a 10.

24. Do you think **other people** reading this would say that you made the right decision in this situation?

25. Do you think that **other people** reading this situation would say that they would make the same choice?

### Part III:

Now we are going to ask you some more general questions about yourself and your reflections on the episode you've just read. These questions may require you to type or speak answers into the computer. If you have any problems, please come let the experimenter know.

1. In the previous section, you had to imagine what other people would think of this story. If you were able to give someone reading this more information about the situation, what would you say?

2. How would you explain your choices to someone reading this situation?

3. If other people only read the situation you described, do you think they would say that they would do the same thing you did<sup>292</sup>?

4. If you want to explain your answer to the last question, you may record it here.

5. Do you think a very strictly ethical person (for example Gandhi) would do the same thing?

---

<sup>292</sup> There was an additional question intended here: "If other people were allowed to read your explanation AND the extra information you just gave, do you think that they would do the same thing that you did?" This question was not included in the Session Three testing done with our participants. It would be informative to include it in future experiments, however.



## Appendix D: Moral Dilemma Questionnaire

Instructions: Please answer Yes or No to the following items. There are no right or wrong answers.

### Item 1

While picking turnips you arrive at two paths, one going to the left, one to the right. If you go to the left you will harvest ten bushels of turnips. If you go to the right you will harvest twenty bushels of turnips.

Is it okay for you to go to the right in order to harvest twenty bushels of turnips instead of ten?

### Item 2

An old friend has invited you to visit at his summer home. There are two ways to get there by car: the highway and the coast road. The highway will get you there in three hours, but the scenery is boring. The coast road will get you there in three and a half hours, but the scenery is beautiful.

Is it okay for you to take the coast road instead of the highway in order to see the beautiful scenery?

### Item 3

You get a letter from a highly respected charity. The letter asks you to make a donation of two hundred dollars. The letter explains that a two hundred-dollar donation will allow them to give badly needed food and medicine to poor people on the other side of the world.

Is it okay for you to not make a donation to this organization in order to save money?

### Item 4

You are the owner of a small business trying to save money. You see that you could lower your taxes by pretending that some of the things you bought for yourself were bought for your business. For example, you could pretend that the stereo in your bedroom is being used at the office.

Is it okay for you to pretend that your personal things are being used by your business in order to lower your taxes?

### Item 5

You have a friend who has not been able to find a job. He figured that he would be more likely to get hired if he had better things to say on his job applications. He decided to put some false information on his job applications in order to make him seem like a better person for the job. By doing this he got a job, beating out other people who were better for the job than him.

Was it okay for your friend to put false information on his job application in order to help him get a job?

Item 6

You are driving along a country road when you hear someone yelling for help by the side of the road. You pull over and find a man who has had an accident. His legs are covered with blood. The man asks you to take him to the hospital. You would like to help him. He will probably lose his leg if he does not get to the hospital soon. However, if you take him to the hospital, his blood will ruin the leather seats in your car.

Is it okay for you to leave this man by the side of the road in order to preserve the leather seats in your car?

Item 7

You have to go from New York to Boston for a meeting. You can take the train or the bus. The train will get you there right before the meeting no matter what. The bus is supposed to arrive an hour before the meeting, but it is sometimes several hours late due to traffic. It would be nice to get there an hour early, but it would be very bad if you were late.

Is it okay for you to take the train instead of the bus in order to make sure that you are not late for the meeting?

Item 8

You have a party in a few hours, and before you go you will need to take a shower. You have some yard work to do before the party, and doing the yard work will make you sweaty. A shower would be nice right now, but if you shower now you will have to shower again after you do the yard work. You are worried about your water bill and want to use as little water as possible.

Is it okay for you to shower before doing your yard work in order to enjoy a shower now?

Item 9

A person running a telephone survey calls while you are having a quiet dinner at home. If you do the survey you will get \$200 but you will have to interrupt your dinner.

Is it okay for you to interrupt your dinner in order to get \$200?

Item 10

You are visiting the sculpture garden of a rich art collector. The garden overlooks a valley containing a set of train tracks. A railway workman is working on the tracks, and an empty runaway trolley is heading down the tracks toward the workman. The only way to save the workman's life is to push one of the art collector's prized sculptures down into the valley so that it will roll onto the tracks and stop the trolley. Doing this will destroy the sculpture.



Is it okay for you to destroy the sculpture in order to save this workman's life?

Item 11

You are on vacation by the sea. You see a group of people get on a small boat and start sailing for a nearby island. After they leave, you hear that there is a big storm coming. The storm will hit them if they are not told to come back. The only way to protect them from the storm is to quickly go out and warn them. To do this you would have to borrow a nearby speedboat. The speedboat belongs to a man who would not like it if you borrowed his boat.

Is it okay for you to borrow the speedboat in order to warn the people about the storm?

Item 12

This afternoon you will do two things: go for a jog and do some work at your desk. It's sunny now, but it is supposed to rain later in the afternoon. You don't like to jog in the rain, but you don't care what the weather is like when you are working at your desk.

Is it okay for you to work at your desk now and jog later in order to get your work done before you jog?

Item 13

You are the leader of a group of soldiers. You are on your way back from a secret mission in enemy territory. One of your men has stepped in a trap and is badly hurt. The trap sets off an alarm so the enemy will be coming for you right away. If the enemy finds the injured soldier they will torture him and kill him. He begs you not to leave him behind, but if you try to take him with you the entire group will move too slowly and be captured. The only way to prevent this injured soldier from being tortured is to shoot him yourself.

Is it okay for you to shoot this soldier in order to prevent him from being tortured by the enemy?

Item 14

You are walking down the street when you find a wallet lying on the ground. You open the wallet and see that it contains several hundred dollars in cash. You can tell that the wallet belongs to a rich person. You have been hit by hard times and need extra money. It would help you if you sent the wallet back to the owner without the cash, keeping the cash for yourself.

Is it okay for you to keep the money you found in the wallet in order to have more money for yourself?

Item 15

You are the leader of a small army that consists of warriors from two tribes, the hill tribe and the river tribe. You belong to neither tribe. During the night a hill tribe warrior murdered a river tribe warrior. The river tribe will attack the hill tribe unless the

murderer is put to death, but the hill tribe refuses to kill one of its own warriors. The only way for you to avoid a war between the two tribes that will cost hundreds of lives is to kill the murderer in front of everyone by cutting off his head with your sword.

Is it okay for you to cut off this man's head in order to prevent the two tribes from fighting a war that will cost hundreds of lives?

Item 16

You are at the wheel of a runaway trolley heading for a place where the tracks split in two. On the tracks going to the left is a group of five railway workmen. On the tracks going to the right is one railway workman. If you do nothing the trolley will go to the left, and it will run over the five workmen and kill them. The only way to save the five workmen is to hit a switch that will make the trolley go to the right, but if you do this it will run over and kill one workman instead.

Is it okay to hit the switch in order to save the five workmen?

Item 17

You, your husband, and your four children are crossing a mountain range on your return journey to your homeland. You have inadvertently set up camp on a local clan's sacred burial ground. The leader of the clan says that according to the local laws, you and your family must be put to death. However, he will let you, your husband, and your three other children live if you will kill your oldest son yourself.

Is it okay for you to kill your oldest son in order to save your husband and your other three children?

Item 18

You are on a big boat at sea. There is a fire on the boat and everyone has to get off. People get into the lifeboats. All the lifeboats, including yours, have too many people in them. The sea is getting rough, and water is coming in over the sides. If nothing is done the lifeboat will sink and everyone on board will die. However, there is an injured person who will not survive in any case. If you throw that person overboard the boat will stay afloat and the remaining passengers will be saved.

Is it okay for you to throw this person overboard in order to save the lives of the remaining passengers?

Item 19

You work for a large company, and you know secret things about that company. You have a friend who plays the stock market, and you owe her a lot of money. By telling your friend secret things about your company you can help your friend make money in the stock market. If you do this she will forget about the money you owe her, but helping people make money by telling them company secrets is against the law.

Is it okay to tell your friend secrets about your company so that she will forget about the money you owe her?

Item 20

A runaway trolley is heading down the tracks toward five workmen who will be killed if the trolley isn't stopped. You are on a footbridge over the tracks, in between the trolley and the five workmen. Next to you on this footbridge is a stranger who happens to be very large. The only way to save the lives of the five workmen is to push this stranger off the bridge and onto the tracks below where his large body will stop the trolley. The stranger will die if you do this, but the five workmen will be saved.

Is it okay to push the stranger on to the tracks in order to save the five workmen?

Item 21

You are the leader of a mountain climbing group that is stranded in the wilderness. Your group includes a family of six that needs a certain vitamin very badly. Some people's kidneys contain large amounts of this vitamin. There is one person in your party who has the vitamins in his kidneys. The only way to save the lives of the six members of this family is to remove one of this man's kidneys, get the vitamins out of it, and give them to the six people. The man will not die if you do this, but his health will suffer. The man is opposed to this plan, but you have the power to do as you see fit.

Is it okay for you to take this man's kidney in order to save the lives of the six people who need vitamins?

Item 22

You are a doctor. You have five patients who need different kinds of organ transplants. They will die without new organs. You also have another patient who is a healthy young man. The only way that you can save the five patients is to steal the organs out of the healthy young man's body and transplant them into the five patients. If you do this, the young man will die, but the other five patients will live.

Is it okay for you to perform this transplant in order to save the five patients?

Item 23

A deadly virus is spreading around the world, killing millions of people. You are a scientist. You have invented two chemicals. One of them kills the virus. The other one is a deadly poison. You don't know which chemical is which because the labels on your containers got mixed up. Once you figure out which is which you can use the good chemical to save millions of lives, but the only way to find out which is which is to test each chemical on someone. If you do this, one person will live, and the other will die, but you will then be able to start saving many other people's lives.

Is it okay for you to kill one of these people with a deadly chemical in order to save millions of lives?

## Item 24

The enemy has taken over your country. You have two children, ages five and eight. There is an enemy doctor who performs painful experiments on humans that always lead to death. He intends to perform experiments on one of your children, but he will allow you to choose which of your children will be experimented upon. You have twenty-four hours to bring him one of your children. If you refuse to bring one of your children to him he will find them both and experiment on both of them.

Is it okay for you to bring one of your children to this doctor in order to avoid having them both die?

## Item 25

You are at the wheel of a runaway trolley heading for a place where the tracks divide in two. On the tracks going to the left is a group of five railway workmen. On the tracks going to the right is a group of seven railway workmen. If you do nothing the trolley will go to the left, killing the five workmen. The only way to save the five workmen is to hit a switch that will make the trolley go to the right, killing the seven workmen instead.

Is it okay to hit the switch in order to save the five workmen?

## Item 26

You've been offered two jobs and you are deciding which one to take. Company A will pay you \$40,000 and will give you ten days of vacation per year. Company B will pay you \$20,000 and will give you twelve days vacation per year. Everything else about the two jobs is the same.

Is it okay for you to take Company B's offer in order to have two more days of vacation per year?

## Item 27

Your plane has crashed in the mountains. The only survivors are you, another man, and a young boy. It is very cold and you have no food. To survive, you must reach a small village on the other side of the mountain. It will take you several days to get there. The boy has a broken leg and cannot walk. He will almost certainly not survive the journey. Without food, you and the other man will probably die as well. The other man suggests that you kill the boy and eat his remains over the next few days.

Is it okay to kill this boy so that you and the other man will survive?

## Item 28

You are on vacation by the sea. You see a group of people get on a small boat and start sailing for a nearby island. After they leave, you hear that there is a big storm coming. The storm will hit them if they are not told to come back. The only way to protect them from the storm is to quickly go out and warn them. To do this you would have to borrow a nearby speedboat. The speedboat belongs to a man who would not like it if you

borrowed his boat, and he has hired a guard to make sure no one uses it. To get to the speedboat you will have to lie to the man guarding the boat.

Is it okay for you to lie to the guard in order to borrow the speedboat and warn the people about the storm?

Item 29

You are part of a group of scientists studying animals in the jungle. The group includes eight children. The group is captured by a gang of terrorists who keep you and the others locked up for several days. One of the terrorists likes you. He tells you that the leader plans to kill you and the others the next day. He says he will help you escape, but that he needs to be sure you won't go to the police. To make sure you won't go to the police, he wants to videotape you killing one of the other scientists. If you accept his offer, the other adults will die but you and the eight children will escape.

Is it okay for you to kill a member of the group in order to escape from the terrorists and save the eight children?

Item 30

You take your broken VCR to a repair shop. It will cost \$100 to have it fixed. You see in the newspaper that a newer and slightly better version of your VCR is on sale for \$100.

Is it okay for you to have your old VCR fixed in order to avoid spending money on a new one?

Item 31

Enemy soldiers have taken over your village. They have orders to kill everyone they find. You and some others are hiding in a basement. You hear the soldiers enter the house above you. Your baby begins to cry loudly. You cover his mouth to block the sound. If you remove your hand from his mouth he will cry loudly and the soldiers will hear. If they hear the baby they will find you and kill everyone, including you and your baby. To save yourself and the others you must smother your baby to death.

Is it okay for you to smother your child in order to save yourself and the other people?

Item 32

You are making brownies from a book of recipes. According to the book, the brownies are supposed to have walnuts. You don't like walnuts, but you like peanuts.

Is it okay for you to use peanuts instead of walnuts in order to avoid eating walnuts?

Item 33

You go to the store in order to buy headache medicine. The store is out of the brand you wanted to buy. The woman working at the store suggests that you buy a different brand that is exactly the same as the one you were looking for.

Is it okay for you to buy the other brand of medicine instead of waiting to buy the brand you were looking for?

### Appendix E: Table of Cue Words in Presentation Order

These cues were presented in the following randomized, fixed order to every subject.

<p>Warm-up Cue: Presented 5 times</p> <p>Please tell us about a very memorable thing that you have done.</p>
<p>Emotion Cues (3 Positive, 4 Intermediate, 3 Negative, 2 Controls)</p> <p>Please talk about a time when you did something that made you feel:</p> <p>Relieved (I)  Doubtful (I)  Embarrassed (N)  Regretful (N)  Happy*  Proud (P)  Guilty (N)  Tired*  Virtuous (P)  Compassionate (P)  Responsible (I)  Bittersweet (I)</p>
<p>Action Cues (10 Cues, 2 Controls)</p> <p>Please talk about a time when you:</p> <p>Did something you have qualms about  Cheated  Took something that didn't belong to you  Were unfaithful  Were funny*  Did something sneaky  Were reckless  Exercised*  Were tempted  Were honest  Lied  Were hurtful to someone else</p>

Superlative Cues (6 Cues)

Please talk about:

The thing you are most proud of having done

The best thing you ever did

The thing that you would most like to change if you could go back and do it differently

The thing that you are most afraid other people will find out that you did

The thing that you would most like other people to know that you did

The worst thing you ever did

\*indicates control item



## **Appendix F: Description of NEO Personality Types**

Neuroticism (N) is the most dominant personality factor and broadly it measures an individual's emotional stability. The NEO professional manual describes individuals with low N scores as "calm, even-tempered, and relaxed, and they are able to face stressful situations without becoming upset or rattled." By comparison, subjects with high N scores tend to have difficulty coping with stress and experience more negative feelings. They tend to be poorly emotionally adjusted. Facets of N: Anxiety, Angry Hostility, Depression, Self-Consciousness, Impulsiveness, Vulnerability.

Extraversion (E) contrasts the commonplace 'types' of extraverts and introverts. However, it is important to note that these should be thought of as two ends of a range, not necessarily polar opposites. That is, introversion is the "absence of extraversion". Thus, while extraverts may be described as "liking people and preferring large groups and gatherings", introverts do not dislike people and gatherings. Rather, they prefer to be alone in the same way extraverts prefer to be with people. Extraverts "are also assertive, active and talkative. They like excitement and stimulation and tend to be cheerful in disposition." Introverts are characterized as reserved, independent and even-paced. Introverts "are not unhappy or pessimistic". Facets of E: Warmth, Gregariousness, Assertiveness, Activity, Excitement-Seeking, Positive Emotions.

Openness (O) refers to an individual's openness to experience. Individuals with low O scores are referred to as 'Closed' and present themselves as more "conventional in behavior and conservative in outlook" than more Open individuals. Open individuals are more "willing to entertain novel ideas and unconventional values". They also tend to feel emotions (both positive and negative) more strongly than their Closed counterparts. Open individuals tend to be more exploratory of the world, both in concept and in

practice. Openness is also associated with creativity, aesthetic sensitivity, an active imagination and intellectual curiosity. Closed individuals should not be thought of as lacking these traits, but rather that they tend to stay close to home with their expression of them. One particularly interesting point about Open individuals is that they are “willing to question authority and entertain new ethical, social and political ideas. . . . An open person may apply his or her evolving value system as conscientiously as a traditionalist does.” This may make people with high O scores particularly interesting in studying moral decision making and behavior. Facets of O: Fantasy, Aesthetics, Feelings, Actions, Ideas, Values.

Agreeableness (A) primarily looks at an individual’s reaction toward others. People with high A scores are “fundamentally altruistic. [They] are sympathetic towards others and eager to help them, and believe that others will be equally helpful in return.” People with low A scores are often referred to as ‘Disagreeable’ or ‘Antagonistic’ but this may unfairly characterize these individuals. People with low A scores are often ‘egocentric, skeptical of others’ intentions and competitive.” They are also “ready to fight for [their] own interests” and good at engaging in critical thinking. Individuals at the far end of both sides of this dimension may suffer from personality disorders. Facets of A: Trust, Straightforwardness, Altruism, Compliance, Modesty, Tender-Mindedness.

Conscientiousness (C) primarily refers to an individual’s self-control. It is based on “individual differences in planning, organizing, and carrying out tasks”. People with high C scores are “purposeful, strong-willed and determined” as well as high achievers, but they may also tend towards “annoying fastidiousness, compulsive neatness or workaholic behavior”. Low C scorers may often be characterized as lazy or sloppy,

because they are not as goal-directed or purposeful as their counterparts. However, they are “not necessarily lacking in moral principles, but they are less exacting in applying them”. People with low C scores may also be more hedonistic and interested in sex.

Facets of C: Competence, Order, Dutifulness, Achievement Striving, Self-Discipline, Deliberation.

## Appendix G: Rating Worksheet

### Who?

Male (Guess?)

Female (Guess?)

-----

Child

Race?

Pre-teen

Religion?

Adolescent

Other?

Adult

### What Did They Do? (0-3)

- Helping someone
- Protecting someone
- Taking (physical) care of someone
- Being kind to someone
- Self-sacrifice
- Cheating in school/work
- Cheating in business
- Sneaking into a movie theatre
- Falsifying a document (resume, paperwork)
- Tricking someone else
- Physically unfaithful (to a partner, with someone else's partner)
- Mentally unfaithful (to a partner, with someone else's partner)
- Neglecting someone
- Being hurtful to someone
- Being unkind to someone
- Physical harm to someone
- Lying
- Telling a white lie
- Stealing/Spending someone's money
- Stealing an idea/plagiarized
- Shoplifting
- Returning lost money/wallet/things
- Bank error/too much change/undercharged
- Enlisting someone else's help in doing wrong action
- Enlisting someone else's help in doing right action
- Fighting for someone else's rights
- Taking advantage of someone else
- Breaking a rule/law
- Vandalism/Damaged Property

### Where? (Scene of the Action)

Job/Business Situation      School

Family/Close Friend      Animal

Abroad/Vacation      Party/Club

Home/Neighborhood      Car

### Why Did They Say They Did It? (0-3)

- Taking the high road
- Doing the wrong thing
- Doing the right thing
- Making extra effort to do right thing
- Making extra effort to do wrong thing
- Doing lesser harm to prevent greater harm (consequentialist)
- Means do NOT justify the ends (deontological)
- Fear of retribution/karma
- Cheating "the system"/ "the man"
- Righting a wrong
- Peer Pressure/Social Pressure to act
- Altruism
- Teaching a Lesson/ Revenge
- Self-Preservation
- Getting out of a tight spot
- Selfishness
- Taking a chance
- Because of money

### How Did They Feel? (0-3)

- Regret about doing the right thing
- Regret about doing the wrong thing
- Regret about behavior
- Knowledge of personal moral weakness
- Know. of personal moral strength
- Uncertainty about moral correctness of action
- Not taking responsibility for actions
- Taking responsibility for actions
- Learning a lesson
- Putting self in someone else's shoes
- Shame/Ashamed
- Powerful

### Did they: (0-3)

- Rationalize an action/behavior?

**Appendix G: Rating Worksheet continued****Cue Words (Circle All That Apply)**

Relieved  
 Doubtful  
 Embarrassed  
 Regretful  
 Happy  
 Proud  
 Guilty  
 Tired  
 Virtuous  
 Compassionate  
 Responsible  
 Bittersweet  
 Did something you have qualms about  
 Cheated  
 Took something that didn't belong to you  
 Were unfaithful  
 Were funny  
 Did something sneaky  
 Were reckless  
 Exercised  
 Were tempted  
 Were honest  
 Lied  
 Were hurtful to someone else  
 The thing you are most proud of having done  
 The best thing you ever did  
 The thing that you would most like to change if you could go back and do it differently  
 The thing that you are most afraid other people will find out you did  
 The thing that you would most like other people to know you did  
 The worst thing that you ever did

**Lazarus' Core Relational Themes (Circle All That Apply)**

Anger: A demeaning offense against me and mine  
Anxiety: Facing uncertain, existential threat  
Fright: Facing an immediate, concrete and overwhelming physical danger  
Guilt: Having transgressed a moral imperative  
Shame: Having failed to live up to an ego-ideal  
Sadness: Having experienced an irrevocable loss  
Envy: Wanting what someone else has  
Jealousy: Resenting a third party for loss or threat to another's affections  
Disgust: Taking in or being too close to an indigestible object or idea (metaphorically speaking)  
Happiness: Making reasonable progress toward the realization of a goal  
Pride: Enhancement of one's ego-identity by taking credit for a valued object or achievement, either our own or that of someone or group with whom we identify  
Relief: A distressing goal-incongruent condition that has changed for the better or gone away  
Hope: Fearing the worst but yearning for better  
Love: Desiring or participating in affection, usually but not necessarily reciprocated  
Compassion: Being moved by another's suffering and wanting to help

Please Judge:



## Appendix H: Directions for Raters

These directions may seem overly elaborate (or elementary) at times, but I hope that that will help prevent any miscommunication. With a rating project of this size, it's pretty important that everyone do things the same way. Please try very hard to make everything legible. Unreadable or missing story numbers will make your ratings unusable.

When you rated the stories before, we asked you to rate them on paper. This time we'd like to ask you to rate them on the computer, if possible. If that's not possible, paper is still fine.

The point of this rating is to put each story in EVERY category that it belongs in. If you come across a story which you do not think should be in the database, please let me know. Stories should meet two criteria: 1. they should be episodic and personal; 2. either the reader or the speaker should feel that the story involves some kind of right or wrong action.

Stories may span a significant range of behaviors and still be in the same category. For example, a child who takes a piece of candy and a professional jewel thief who robs a museum are both stories about "taking something that doesn't belong to you". Equally, a story about someone who decides NOT to steal something should still be categorized under stealing. Please note these stories with the greek letter alpha  $\alpha$  (to signify that it's an anti-stealing story). If you are rating online, you can use a negative sign (-) to signify this instead.

1. The rating sheet
  - a. Label EVERY rating sheet with
    - i. your unique user ID – this was emailed to you
    - ii. the story number
2. Category 1: Who?
  - a. Gender:
    - i. This is the **only** category where you should guess.
    - ii. ALWAYS guess which gender the speaker is and mark if it is a guess or if you are certain. (We are using this as a pilot for a gender study at the same time.)
  - b. Age:
    - i. Only circle this if it is relevant to the story.
    - ii. Typically, the subject being an adult is not relevant to the story.
  - c. Race/Religion
    - i. If this category is relevant to the story, please note it.
    - ii. If there is some other category that comes up that is similar and relevant to the story, please write it in. (For example, if being in the Girl Scouts is a crucial element of the story, please write that in.)
3. Category 2: Where?

- a. This category is to highlight important story elements that may “change the rules” of morality. For example, you might lie in business, but never to your family. There are a few items in this category that are not places but may change the moral rules.
  - b. Only circle these if they are part of the story. Do not guess.
  - c. Family and Animals are important categories here even though they’re not places.
4. Categories 3-5: What, Why and How?
- a. All of these categories should be rated on a scale from 0 to 3.
    - i. 0 = does not apply to this category at all
    - ii. 3 = great exemplar of this category
  - b. If you do not want to write lots of zeros, that’s fine. I will assume that blanks are zeros. Be careful not to miss any lines if you’re doing this.
  - c. Remember: Acts and Omissions count in these stories. Situations where the subject chose NOT to do something are just as important. Mark these with an ‘alpha’ or negative sign and your rating number.
5. Category 6: Rationalization
- a. How strongly does rationalization play a role in this story?
  - b. Use the same 0 to 3 rating scale as before.
6. Cue Words
- a. One cue word was used to elicit this story, however, many cue words may apply.
  - b. If a subject tells a story about cheating on their spouse, they may feel guilty, like it’s the worst thing they ever did, and regretful. This story should be categorized under all of those plus unfaithful.
  - c. Do not spend too much time over-thinking this category; just do it quickly and move on.
7. YOUR judgment
- a. How right or wrong do YOU think the action in the story is?
    - i. Scale from -5 to +5
    - ii. -5 = completely wrong
    - iii. +5 = completely right
    - iv. 0 DOES NOT equal “not moral”; 0 means equally right and wrong (completely morally ambiguous)
    - v. If you come across a story that you think is not moral (even though the subject might), write “not moral” instead of a rating.
  - b. Unlike the rest of these ratings, this is purely subjective – tell us what YOU think.
  - c. Just to reiterate – please make things legible. Your negative signs will be very important here.
8. Core Relational Themes
- a. Choose the core relational themes that best fit the story, based on the description of the theme NOT the emotion word.
  - b. If none of them fit, don’t circle anything.

### **Similarity Ratings**

We are using similarity ratings to help identify "categories" that might not be intuitively obvious. The way to fill out this for is as follows:

1. Read the first story and do the rating sheet as described above.
2. Read the next story. Before you do the rating sheet, decide how similar the two stories you just read are. This could be for any number of reasons -- and you don't have to have a reason. Just do it by "feel" if you will.
3. Put your rating on the similarity sheet and then do the second rating sheet, etc...
4. In this fashion you'll compare each story to the next one.

Finally, if you have any questions or problems, please call or email. Most people have at least a few questions about rating or categories and the more questions that you ask, the better your ratings will be. Please don't hesitate to ask.

Thank you so much for helping with our research project. You are an invaluable resource.



## Appendix I: Sample Narratives from Lesion Subjects

Participant with developmental frontal lobe damage:

My sophomore year of high school was my sister's sophomore year of college. At the time, her boyfriend and I were really good friends. My sister got very jealous of that and she wanted me to quit emailing him and quit talking to him. I didn't care what she wanted, so I didn't stop. That was hurtful to her. Eventually I did stop talking to him but that's not the point.

Participant with frontal lobe damage:

Last winter, my church started a clothing giveaway for people that don't have enough. About the same time, I went to the thrift store looking at things for my grandson. I saw a coat on the 50 cent rack and I thought, "Well, I'll take that for the church". I told the lady who worked there that I was buying this for the church and she said, "Oh, would you take all these others?" So, I left with half a dozen bags of coats to take to the giveaway at my church, because they were coats that weren't selling at the thrift store. For a free giveaway, they were really nice looking coats. I got a variety of children's through men's coats.

Participant with bilateral amygdala damage:

The only time that I've ever lied was a long, long time ago. There was a lady I knew who was badly burned. She felt she was ugly and had a low self-esteem. She hated herself actually. I had to lie to this woman and I told her she was absolutely beautiful. She was badly scarred, but I told her she was beautiful. I wanted her to feel good about herself. I wanted to put a smile on her face, to boost her self-confidence.

Participant with amnesia:

I was dating this girl one time and one night I couldn't get a hold of her. I wanted to go to this concert, so I brought another girl to the event. I knew right away, that probably was not a good deal. The girls started talking to each other and it went downhill after that. I don't really want bad deals.

Participant with temporal damage:

I remember this one time when I was at a store and someone dropped some bills on the floor. I picked the money up to give it to them instead of keeping it myself. I don't even know what size bills they were. I know there's other people that would just walk off and take money like that with them but I just wasn't raised that way.

Normal participant from Iowa:

For the last 5 years, I've been going with my church on a service project in Appalachia. The very first year was very life changing for me. I was going through marriage problems, which I didn't realize until I got home and I found that he'd moved out. It's just a small group of about 30 people from our church that go. We go down and do hard labor -- making homes safer, warmer and drier for people. They live in much more unfortunate circumstances than we live in, in our day to day life. It was very eye opening and gave me a new understanding of other people's teenagers. Each experience is different, and

each family is different, but I try to keep in touch with those families at least once or twice a year. It just makes you realize that their culture is totally different in Appalachia. They're so appreciative, though, and they have a strong faith just like you or I and that's been very inspiring to me.

## Appendix J: Selection of Narratives from the Database

### Memory 5144

This is a recollection that I have of about 47 years ago -- it is a time when I was more honest than I should have been, at least in terms of the way my father looked at it. My father was an executive in the hearing aid industry. He traveled back and forth between England and countries like Germany, Holland and Belgium, attending various trade fairs and marketing the goods of his particular company to German importers. One time he took my sister and me and my mother on a combination of business and pleasure trip to Germany. While he was there he bought himself some very expensive, state-of-the-art cameras. I remember when we came back to England, my father gave me a camera and gave my sister one. He told me very distinctly and clearly what to say to the customs officer if I was asked. If the customs officer said, "Do you have anything to declare, have you brought anything into the country?" I was supposed to take the camera that was around my neck and give it to the customs officer and say, "Here we brought this camera." Meantime the second camera was very carefully hidden away where it was unlikely that they would find it. Well, the customs agent came to our car and asked for the declaration. My father looked at me and said, "Give him the camera." Without batting an eyelid I piped up (in my rather squeaky high pitched voice of the time) to the customs officer and said, "Which one, Daddy? You've got two." I got my father in quite a bit of trouble. He got fined some \$400 for trying to bring something in the country that he shouldn't have done. While the customs officer praised me for my honesty, I was persona non grata with my dad for quite some time after that. I got punished for being honest and forthright. I guess at the time I felt that honesty was the best policy and I'm not certain why I disobeyed my dad; maybe I was intimidated by the authority of the uniforms at the customs office.

### Memory 5201

I hate sneakers; I wear shoes at all times. I cannot stand sneakers. A friend of mine gave me twenty pairs of sneakers because he was buying some more, but I wear 12's and these were 10's. I passed downtown in my car from time to time and if I saw somebody who looked like he needed a pair of tennis shoes I would open my trunk and give him a pair. One particular man, one time, was standing in front of Church's Chicken. He had some old shoes on and was walking on the back of them. I knew that was uncomfortable, because I had been in that position. I said, "Would you like to have a pair of tennis shoes?" But I was in the wrong car; the car that had the tennis shoes was at the shop. I said, "If you wait here 5 minutes I'm going to go and get you a nice pair of tennis shoes. Just wait." Normally if you catch a person in that position they are on drugs or something, but not this guy. It seemed like he had to be mentally ill. I think that is wrong of the government just to put all the mentally ill people on the street. They're not all able to fend for themselves. Anyway, I came back with the tennis shoes. I thought that when I came back I would buy him a dinner or something. When I got back though, there was a lady who had already bought him a dinner. Some people will still help each other. So I

think he came up that day. He wound up getting a dinner and a new pair of tennis shoes. He threw his old ones in the trash and went walking down the road with his new shoes and his chicken.

#### Memory 5257

I lied once a job resume. I put on there that I had this school background and I wanted to put as much information on skills as I could – more than I had. I think I put that I knew Photoshop and you know, other computer skills. I didn't feel comfortable about it because I knew that if I were to get that job, they were gonna ask me to do whatever skills I listed. After I thought about that, I decided just not lie about it and I fixed my resume. I only put down the skills that I have. I wanted to lie because I was applying for this particular job. I felt that maybe putting more skills, more abilities than I had, more experience, would give me a better chance at getting the job. I thought that it would look better on my resume. After I decided not to lie and fixed it, I got the job anyway. I turned out that I don't have to lie. I have pretty good experiences and skills and abilities and that should be enough.

#### Memory 5489

About two months ago I was in the grocery store and I was at the cash register. There was this elderly lady; she must have been maybe 75, 80, I don't know how old. She had just gotten her groceries and she left her wallet on the counter. I picked her wallet up and called her back, because I didn't want her to forget it. I think that was an honest thing to do.

#### Memory 5787

I was seeing a man who was married. I wasn't involved with anyone, so I think most people wouldn't say that I was being unfaithful. I think the thing that would be considered unfaithful on my part, would be that I was doing a dreadful, unjust thing to his wife. Now that it has been done to me, I can really see what a terrible, terrible thing it was to do to his wife. I had fallen madly in love with this man when we were working together. We were working and traveling together and it just happened. That's no excuse; it doesn't make it right, because it wasn't right. I was unfaithful to a relationship between another man and woman. I try to be honest and faithful. I don't want to sound like a Pollyanna, but that's how I am. Maybe I'm square but that's how I live my life. It comes back to my beliefs in a universal law -- the way we treat people, it's going to come back on the other end. If we do something wrong it's going to come back and we may pay really dearly. If we do right, god willing, it will come back to us, so I try to be faithful and go forward with my life in that way.

## Memory 5894

When my mother passed away, oh it's been seven, eight years ago now, I had to wrap up her affairs. In doing so, I learned that she was supporting an African child through the Christian Children's fund. I hadn't known about that and I didn't know anything about that particular charity. As I read the materials that she had there, I learned that she had been making a contribution of twenty five dollars a month that was going a long way to supporting this child, and feeding them and providing them with livestock and those kinds of things. The money really made a difference in the lifestyle and life of that family that that child was part of. As I thought about the fact that my mother was no longer with us, I realized that my mother would no longer be making those contributions to that charity and that child. It seemed like it was not right that just because my mother was no longer there to carry on that vision that she had for helping those people, that that work wouldn't be done and that that sense of giving and caring would somehow disappear. I found my heart going out to the child that my mother had been supporting and to that whole mission. As a result, I did something I hadn't even considered doing up until that point -- I signed up to support a child of my own, one who was assigned to me. I make regular, monthly contributions to this child through this charity so that she has clothes and food and livestock and her family and she are not at risk. That's something that really makes me feel compassionate.

## References

- "Craig's List: Los Angeles." from <http://www.losangeles.craigslist.com>.
- "Gilligan's In a Different Voice." from <http://www.stolaf.edu/people/huff/classes/handbook/Gilligan.html>.
- "Report of Registration: January 2, 2004." from [http://www.ss.ca.gov/elections/ror/regstats\\_01-02-04.pdf](http://www.ss.ca.gov/elections/ror/regstats_01-02-04.pdf).
- STATISTICA. Tulsa, OK, Stat Soft, Inc at <http://www.statsoft.com>.
- "US Census Bureau American FactFinder", from <http://factfinder.census.gov/>.
- "USC Center for Religious and Civic Culture." from <http://www.usc.edu/schools/college/crcc/demographics/lacounty.html>.
- (2000). Religious Congregations and Membership in the United States. Nashville, TN, Glenmary Research Center.
- (2000-2008). "Belief-O-Matic - A personality quiz about your religious and spiritual beliefs." from [www.beliefnet.com/story/76/story\\_7665\\_1.html](http://www.beliefnet.com/story/76/story_7665_1.html).
- (2006). Schadenfreude. Dictionary.com, Random House, Inc. Unabridged (v 1.1).
- (2007). Electronic Textbook: Cluster Analysis. Tulsa, OK, StatSoft. Online at: <http://www.statsoft.com/textbook/stathome.html>.
- (2007). Electronic Textbook: Principle Components and Factor Analysis. Tulsa, OK, StatSoft. Online at: <http://www.statsoft.com/textbook/stathome.html>.
- (2007). "European Social Survey." from <http://www.europeansocialsurvey.org/>.
- (2008). "Moral." from <http://www.merriam-webster.com/dictionary/moral>.
- Abdi, H. (2007). Bonferroni and Sidak corrections for multiple comparisons. Encyclopedia of Measurement and Statistics. N. J. Salkind. Thousand Oaks, CA, Sage.
- Agerstrom, J. and K. Moller (2006). "Moral Reasoning: The Influence of Affective Personality, Dilemma Content and Gender." Social Behavior and Personality **34**(10): 1259-1276.
- Aldrich, D. and R. Kage (2003). "Mars and Venus at Twilight: A Critical Investigation of Moralism, Age Effects, and Sex Differences." Political Psychology **24**(1): 23-40.
- Alexander, L. and M. Moore. (Fall 2008 Edition). "Deontological Ethics." The Stanford Encyclopedia of Philosophy 2008, from <http://plato.stanford.edu/archives/fall2008/entries/ethics-deontological/>.
- Allman, J. and J. Woodward (2007). "Moral intuition: Its neural substrates and normative significance." Journal of Physiology – Paris **101**: 179-202.
- Allman, J. M., Hakeem, A., Erwin, J. M., Nimchinsky, E., Hof, P. (2001). "Anterior cingulate cortex: the evolution of an interface between emotion and cognition." Annals of the New York Academy of Sciences **935**: 107–117.
- Allman, J. M., Watson, K.K., Tetreault, N.A., Hakeem, A.Y. (2005). "Intuition and autism: a possible role for Von Economo neurons." Trends in Cognitive Sciences **9**: 367–373.
- Anderson, S. W., Bechara, A., Damasio, H., Tranel, D., Damasio, A. R. (1999). "Impairment of social and moral behavior related to early damage in human prefrontal cortex." Nature Neuroscience **2**: 1032-1037.

- Anderson, S. W., Bechara, A., Damasio, H., Tranel, D., & Damasio, A.R. (1999). "Impairment of social and moral behavior related to early damage in human prefrontal cortex." Nature Neuroscience **2**: 1032 – 1037.
- Anscombe, G. E. M. (1958). "Modern Moral Philosophy." Philosophy **33**(124): 1.
- Aristotle *Nicomachean Ethics*. H. Rackham.
- Baron, J. (1995). "A Psychological View of Moral Intuition." The Harvard Review of Philosophy **Spring**: 36-40.
- Batson, C. D., Lishner, D. A., Carpenter, A., Dulin, L., Harjusola-Webb, S., Stocks, E. L., Gale, S., Hassan, O., Sampat, B. (2003). "... As you Would have Them Do Unto You: Does Imagining Yourself in the Other's Place Stimulate Moral Action?" Personality & Social Psychology Bulletin **29**(9): 1190.
- Batson, C. D. and E. R. Thompson (2001). "Why don't moral people act morally? Motivational considerations." Current Directions in Psychological Science **10**(2): 54-57.
- Bauman, K. J. and N. L. Graf (2003). *Educational Attainment: 2000*, US Census Bureau.
- Beauchamp, T. L. and J. F. Childress (2001). Principles of Biomedical Ethics, 5th edition. New York, Oxford University Press.
- Bechara, A., Damasio, H., Damasio, A. R. (2000). "Emotion, Decision Making and the Orbitofrontal Cortex." Cerebral Cortex **10**: 295-307.
- Beck, A. T. (1972). Depression: Causes and Treatment. Philadelphia, University of Pennsylvania Press.
- Bem, S. L. (1981). Bem Sex-Role Inventory. Redwood City, CA, Mind Garden; Consulting Psychologists Press, Inc.
- Bentham, J. (1961). An Introduction to the Principles of Morals and Legislation. Garden City, Doubleday.
- Berntsen, D. and D. C. Rubin (2002). "Emotionally charged autobiographical memories across the life span: the recall of happy, sad, traumatic, and involuntary memories." Psychology and Aging **17**: 636-652.
- Berthoz, S., Grezes, J., Armony, J. L., Passingham, R. E., Dolan, R. J. (2006). "Affective response to one's own moral violations." NeuroImage **31**: 945 – 950.
- Bjorklund, F. (2003). "Differences in the justification of choices in moral dilemmas: Effects of gender, time pressure and dilemma seriousness." Scandinavian Journal of Psychology **44**: 459-466.
- Bolender, J. (2001). "A Two-Tiered Cognitive Architecture For Moral Reasoning." Biology and Philosophy **16**: 339-356.
- Borg, J. S., Hynes, C., Van Horn, J., Grafton, S., Sinnott-Armstrong, W. (2006). "Consequences, Action, and Intention as Factors in Moral Judgments: An fMRI Investigation." Journal of Cognitive Neuroscience **18**(5): 803-817.
- Bradley, M. M. and P. J. Lang (1999). *Affective norms for English words (ANEW): Stimuli, instruction manual and affective ratings*. Technical report C-1. Gainesville, FL, The Center for Research in Psychophysiology, University of Florida.
- Broad, C. D. (1930). Five Types of Ethical Theory. New York, Harcourt, Brace and Co.
- Brody, L. R. H., J.A. (2000). *Gender, Emotion and Expression*. Gender Handbook of Emotions. M. H.-J. Lewis, J.M. New York, Guilford Press: 338-347.

- Buchanan, T. W, Tranel, D., Adolphs, R. (2005). "Emotional autobiographical memories in amnesic patients with medial temporal lobe damage." The Journal of Neuroscience **25**: 3151-3160.
- Burke, A., Heuer, F., Reisberg, D. (1992). "Remembering emotional events." Memory and Cognition **20**: 277-290.
- Cahill, L., Haier, R. J., Fallon, J., Alkire, M. T., Tang, C., Keator, D., Wu, J., McGaugh, J. L. (1996). "Amygdala activity at encoding correlated with long-term, free recall of emotional information." Proceedings of the National Academy of Sciences **93**: 8016-8021.
- Casebeer, W. D. (2003). "Moral cognition and its neural constituents." Nature Reviews Neuroscience **4**: 841-846.
- Casebeer, W. D. (2003). Natural Ethical Facts: Evolution, Connectionism, and Moral Cognition. Cambridge, MA, The MIT Press.
- Casebeer, W. D. (2003). "The Neural Mechanisms of Moral Cognition: A Multiple-Aspect Approach to Moral Judgment and Decision-Making." Biology and Philosophy **18**: 169-194.
- Chomsky, N. (1955). The Logical Structure of Linguistic Theory. Cambridge, MA, MIT
- Christianson, S.-A. and E. F. Loftus (1991). "Remembering emotional events: the fate of detailed information." Cognition and Emotion **5**: 81-108.
- Ciamarelli, E., Muccioli, M., Ladavas, E., di Pellegrina, G. (2007). "Selective deficit in personal moral judgment following damage to ventromedial prefrontal cortex." Social Cognitive and Affective Neuroscience **2**(2): 84-92.
- Coltheart, M. (1981). "The MRC Psycholinguistic Database." Quarterly Journal of Experimental Psychology **33A**: 497-505.
- Coltheart, M. (1981a). "MRC Psycholinguistic Database User Manual: Version 1." from [http://www.psych.rl.ac.uk/User\\_Manual\\_v1\\_0.html](http://www.psych.rl.ac.uk/User_Manual_v1_0.html)
- Costa, P. T. and R. R. McCrae (1992). NEO PI-R Professional Manual. Lutz, FL.
- Crovitz, H. and H. Schiffman (1974). "Frequency of episodic memories as a function of their age." Bulletin of the Psychonomic Society **4**: 517-518.
- Cummiskey, D. (1996). Kantian Consequentialism, Oxford University Press.
- Cutler, S. E., Larsen, R.J., Bunce, S.C. (1996). "Repressive Coping Style and the Experience and Recall of Emotion: A Naturalistic Study of Daily Affect." Journal of Personality and Social Psychology **64**(2): 379-405.
- Damasio, A. R. (1994). Descartes' Error: Emotion, Reason, And The Human Brain. New York, Grosset/Putnam.
- de Waal, F. B. M., Ed. (2002). Tree of Origin: What Primate Behavior Can Tell Us About Human Social Evolution, Harvard University Press.
- DePaul, M. "Comments on Walter Sinnott-Armstrong." Retrieved September 16, 2008, from <http://www.slu.edu/colleges/AS/philos/PDF/depaul.pdf>.
- Dupoux, E. and P. Jacob (2008). "Response to Dwyer and Hauser: Sounding the retreat?" Trends in Cognitive Sciences **12**(1): 2-3.
- Dupoux E, J. P. (2007). "Universal moral grammar: a critical appraisal." Trends in Cognitive Sciences **11**(9): 373-378.
- Eichenbaum, H., Otto, T., Cohen, N. (1994). "Two functional components of the hippocampal memory system." Behavioral & Brain Sciences **17**: 449-518.
- Feinberg, J. (2008). "Wordle." from <http://wordle.net/>.



- Fieser, J. (2006). "Ethics." from <http://www.iep.utm.edu/e/ethics.htm>
- Foot, P. (1978). The Problem of Abortion and the Doctrine of the Double Effect in Virtues and Vices and Other Essays in Moral Philosophy. Berkeley and Los Angeles, University of California Press.
- Forsyth, D. (1980). "A Taxonomy of Ethical Ideologies." Journal of Personality and Social Psychology **39**(1): 175-184.
- Forsyth, D. (November 2006). Personal Communication.
- Friedman, R. C. and J. I. Downey (1994). "Homosexuality." New England Journal of Medicine **331**(14): 923-930.
- Gebhard, P. H. and A. B. Johnson (1998 ). The Kinsey Data: Marginal Tabulations of 1938-1963 Interviews Conducted by the Institute for Sex Research Bloomington, IN, Indiana University Press.
- Gigerenzer, G. (2006). Bounded and Rational. Oxford, Blackwell.
- Gilligan, C. (1977). "In a Different Voice: Women's Conceptions of Self and Morality." Harvard Educational Review **47**(4): 1-37.
- Gilligan, C. (1982). In a different voice: Psychological theory and women's development. Cambridge, MA, Harvard University Press.
- Goldstein, D. G. and G. Gigerenzer (2002). "Models of Ecological Rationality: The Recognition Heuristic." Psychological Review **109**(1): 75-90.
- Grafman, J. (1995). Similarities and distinctions among current models of prefrontal cortical functions. Structure and Functions of the Human Prefrontal Cortex. K. J. H. J. Grafman, & F. Boller. New York, New York Academy of Sciences: 337-368.
- Greene, J. D. (2003). "From neural 'is' to moral 'ought': what are the moral implications of neuroscientific moral psychology?" Nature Reviews Neuroscience **4**: 847-850.
- Greene, J. D. (2005). Cognitive neuroscience and the structure of the moral mind. The Innate Mind: Structure and Contents. S. Laurence, P. Carruthers and S. Stich. New York, Oxford University Press.
- Greene, J. D. (2007). "Why are VMPFC patients more utilitarian? A dual-process theory of moral judgment explains." Trends in Cognitive Sciences **11**(8): 322-323.
- Greene, J. D. and J. Haidt (2002). "How (and where) does moral judgment work?" Trends in Cognitive Sciences **6**(12): 517-523.
- Greene, J. D., Morelli, S. A., Lowenberg, K., Nystron, L. E., Cohen, J. D. (2008). "Cognitive load selectively interferes with utilitarian moral judgment." Cognition **107**: 1144-1154.
- Greene, J. D., Nystron, 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. D., Sommerville, R. B., Nystron, L. E., Darley, J. M., Cohen, J. D. (2001). "An fMRI Investigation of Emotional Engagement in Moral Judgment." Science **293**: 2105-2108.
- Haahr, M. (1998-2008). "Random.org." from <http://www.random.org/>.
- Haidt, J. (2001). "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment." Psychological Reviews **108**(4): 814-834.
- Haidt, J. (2007). "The new synthesis in moral psychology." Science **316**: 998-1002.

- Hardy, S. A. (2006). "Identity, Reasoning, and Emotion: An Empirical Comparison of Three Sources of Moral Motivation." Motivation and Emotion **30**: 207-215.
- Hare, R. M. (1991). The Language of Morals, Clarendon.
- Hauser, M. "The Moral Sense Test." 2008, from <http://moral.wjh.harvard.edu/index.html>.
- Hauser, M. D. (2006). Moral Minds: How Nature Designed a Universal Sense of Right and Wrong. New York, Harper Collins.
- Hauser, M. D., Cushman, F., Young, L., Jun, R. K-X., Mikhail, J. (2007). "A dissociation between moral judgments and justifications." Mind & Language **22**(1): 1-21.
- Heberlein, A. S. and R. Adolphs (2005). The anatomy of human social cognition. The Cognitive Neuroscience of Social Behaviour. A. E. Easton, Nathan. Philadelphia, Psychology Press: 157-194.
- Heekeren, H. R., Wartenburger, I., Schmidt, H., Schwintowski, H-P., Villringer, A. (2003). "An fMRI study of simple ethical decision-making." Neuroreport **14**(9): 1215-1219.
- Heuer, F. and D. Reisberg (1990). "Vivid memories of emotional events: the accuracy of remembered minutiae." Memory and Cognition **18**: 496-506.
- Hogan, R. and E. Dickstein (1972). "A measure of moral values." Journal of Consulting and Clinical Psychology **39**(2): 210-214.
- Hooker, B. (Fall 2008 Edition). "Rule Consequentialism " The Stanford Encyclopedia of Philosophy, from <http://plato.stanford.edu/archives/fall2008/entries/consequentialism-rule/>.
- Huff, C. "Gilligan's In a Different Voice." 2008, from <http://www.stolaf.edu/people/huff/classes/handbook/Gilligan.html>.
- Hume, D. (2003). A Treatise of Human Nature, Project Gutenberg.
- Hursthouse, R. (2007). "Virtue Ethics." from <http://plato.stanford.edu/entries/ethics-virtue/>.
- Johnson, R. (2008). Kant's Moral Philosophy. The Stanford Encyclopedia of Philosophy. E. N. Zalta.
- Kahane, G. and N. Shackel (2008). "Do abnormal responses show utilitarian bias?" Nature **452**(7185): E5-E5.
- Kant, I. (1993 ). Grounding for the metaphysics of morals; with, On a supposed right to lie because of philanthropic concerns (Third Edition). Indianapolis, Hackett, ed.
- Kemerling, G. (2002). "Utilitarianism." from <http://www.philosophypages.com/hy/5q.htm#util>.
- Knobe, J. (2003). "Intentional action and side effects in ordinary language." Analysis **63**: 190-193.
- Knobe, J. (2005). "Theory of mind and moral cognition: exploring the connections." Trends in Cognitive Sciences **9**(8): 357.
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., Damasio, A. (2008). "Koenigs et al. reply." Nature **452**: E5-E6.
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., & Damasio, A. (2007). "Damage to the prefrontal cortex increases utilitarian moral judgements." Nature **446**: 908–911.
- Kohlberg, L. (1958). The Development of Modes of Thinking and Choices in Years 10 to 16, University of Chicago. Ph. D. thesis.

- Kohlberg, L. (1963). "The development of children's orientations toward a moral order. I: Sequence in the development of human thought." Vita Humana **6**: 11-33(a).
- Kreider, R. M. and T. Simmons (2003). Marital Status: 2000, US Census Bureau.
- Kucera, H. and W. N. Francis (1967). Computational Analysis of Present-Day American English. Providence, Brown University Press.
- Lewis, P. A., Critchley, H. D., Rotshtein, P., Dolan, R. J. (2007). "Neural Correlates of Processing Valence and Arousal in Affective Words." Cerebral Cortex **17**(3): 742-748.
- Loftus, E. F. (1993). "The reality of repressed memories." Journal of the American Psychological Association **48**: 518-537.
- Mah, L. W. Y., Arnold, M.C., Grafman, J. (2005). "Deficits in Social Knowledge Following Damage to Ventromedial Prefrontal Cortex." The Journal of Neuropsychiatry and Clinical Neurosciences **17**: 66-74.
- McFarland, C., and R. Buehler (1998). "The Impact of Negative Affect on Autobiographical Memory: The Role of Self-Focused Attention to Moods." Journal of Personality and Social Psychology **75**(6): 1424-1440.
- McGaugh, J. L. (1990). "Significance and remembrance: the role of neuromodulatory systems." Psychological Science **1**: 15-25.
- McGaugh, J. L. (2000). "Memory-- a century of consolidation." Science **287**: 248-251.
- McGaugh, J. L. (2002). "Memory consolidation and the amygdala: a systems perspective." Trends in Neurosciences **25**: 456-461.
- McGaugh, J. L. (2004). "The amygdala modulates the consolidation of memories of emotionally arousing experiences." Annual Review of Neuroscience **27**: 1-28.
- Meltzoff, A. N. and J. Decety (2003). "What imitation tells us about social cognition: a rapprochement between developmental psychology and cognitive neuroscience." Philosophical transactions of the Royal Society of London **358**(1431): 491.
- Mikhail, J. (2007). "Universal moral grammar: theory, evidence and the future." Trends in Cognitive Science **11**(4): 144-152.
- Mill, J. S. (1998). Utilitarianism New York, Oxford University Press.
- Moll, J. (2006). "Human fronto-mesolimbic networks guide decisions about charitable donation." Proceedings of the National Academy of Sciences **103**(42): 15623.
- Moll, J., and de Oliveira-Souza, R. (2007). "Response to Greene: Moral sentiments and reason: friends or foes?" Trends in Cognitive Sciences **11**(8): 323-324.
- Moll, J., de Oliveira-Souza, R., Eslinger, P. J., Bramati, I. E., Mourao-Miranda, J., Andreiuolo, P. A., Pessoa, L. (2002). "The neural correlates of moral sensitivity: A functional magnetic resonance imaging investigation of basic and moral emotions." The Journal of Neuroscience **22**(7): 2730-2736.
- Moll, J., de Oliveira-Souza, R., Eslinger, P. J. (2003). "Morals and the human brain: a working model." Neuroreport **14**(3): 299-305.
- Moll, J., de Oliveira-Souza, R., Bramati, I. E., Grafman, J. (2002). "Functional networks in emotional moral and nonmoral social judgments." NeuroImage **16**(3): 696-703.
- Moll J., Z. R., de Oliveira-Souza R., Krueger F., Grafman J. (2005). "The neural basis of human moral cognition." Nature Reviews Neuroscience **6**(10): 799-809.
- Moore, G. E. (1903). Principia Ethica, Cambridge University Press.

- Mori, E., Ikeda, M., Hirono, N., Kitagaki, H., Imamura, T., Shimomura, T. (1999). "Amygdalar volume and emotional memory in Alzheimer's disease." American Journal of Psychiatry **156**: 216-222.
- Motulsky, H. (1995). Intuitive Biostatistics. Oxford, Oxford University Press.
- Muuss, R. E. (1988). "Carol Gilligan's theory of sex differences in the development of moral reasoning during adolescence." Adolescence **23**: 229-243.
- Neisser, U. and N. Harsch (1992). Phantom flashbulbs: false recollections of hearing the news about Challenger. Affect and Accuracy in Recall: Studies of "Flashbulb" Memories. E. W. U. Neisser. Cambridge, MA, Cambridge University Press: 9-31.
- Nichols, S. (2002). "How Psychopaths Threaten Moral Rationalism, or Is it Irrational to Be Amoral?" The Monist **85**: 285-304.
- Nichols, S. and R. Mallon (2006). "Moral dilemmas and moral rules." Cognition **100**(3): 530-542.
- Oldfield, R. C. (1971). "The assessment and analysis of handedness: The Edinburgh Inventory." Neuropsychologia **9**: 97-113.
- Overman, W., Graham, L., Redmond, A., Eubank, R., Boettcher, L., Samplawski, O., Walsh, K. (2006). "Contemplation of Moral Dilemmas Eliminates Sex Differences on the Iowa Gambling Task." Behavioral Neuroscience **120**(4): 817-825.
- Pennebaker, J. W., Chung, C. K., Ireland, M., Gonzales, A., Booth, R. J. (2007). The development and psychometric properties of LIWC2007. Austin, TX, LIWC.net.
- Phelps, E. A. (2006). "Emotion and cognition: insights from studies of the human amygdala." Annual Review of Psychology **57**: 27-53.
- Piaget, J. (1977). The Essential Piaget. New York, Basic Books.
- Prinz, J. (2006). "The Emotional Basis of Moral Judgments." Philosophical Explorations **9**(1): 29-43.
- Prinz, J. J. (2007). Is Morality Innate? Moral Psychology, vol. 1: Evolution of Morals. W. Sinnott-Armstrong. Cambridge, MA, MIT Press.
- Rawls, J. (1971). A Theory of Justice. Cambridge, Harvard University Press.
- Rokeach, M. (1973). The Nature of Human Values. New York, Free Press.
- Roskies, A. (2006). "Patients with ventromedial frontal damage have moral beliefs." Philosophical Psychology **19**(5): 617-627.
- Rubin, D. C. and M. D. Schulkind (1997). "The distribution of autobiographical memories across the lifespan." Memory and Cognition **25**: 859-866.
- Ruse, M. (2004). "Book Review: Natural Ethical Facts: Evolution, Connectivism, and Moral Cognition by William D. Casebeer, Cambridge, MA: MIT Press, 2003." Evolutionary Psychology **2**: 89-91.
- Schacter, D. L. (1996). Searching for Memory: The Brain, The Mind, And The Past. New York, Basic Books.
- Schacter, D. L. and S. D. Slotnick (2004). "The cognitive neuroscience of memory distortion." Neuron **44**: 149-160.
- Schnall, S. (2008). "Disgust as embodied moral judgment." Personality & Social Psychology Bulletin **34**(8): 1096-1109.
- Schwartz, S. (1994). "Are There Universal Aspects in the Structure and Contents of Human Values?" Journal of Social Issues **50**(4): 19-45.

- Schwartz, S. H. (1992). "Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries." Advances in Experimental Social Psychology **25**: 1-65.
- Scoville, W. and B. Milner (1957). "Loss of recent memory after hippocampal lesions." Journal of neurology, neurosurgery and psychiatry **20**: 11-21.
- Sharp, F. C. (1898). "An objective study of some moral judgments." American Journal of Psychology **9**: 198-234.
- Simpson, E. (1974). "Moral development research: A case study of scientific cultural bias." Human Development **17**: 81-106.
- Sinnott-Armstrong, W. (2008). Framing moral intuitions. In Moral psychology, volume 2: the cognitive science of morality. W. Sinnott-Armstrong. Cambridge: MIT, Bradford Books: 47-76
- Sinnott-Armstrong, W. (2008). Moral Intuitions and Heuristics. William Bennett Munro Memorial Seminar Pasadena, CA, California Institute of Technology.
- Slote, M. (1992). From Morality to Virtue. New York, Oxford University Press.
- Squire, L. R. (1986). "Mechanisms of memory." Science **232**: 1612-1619.
- Squire, L. R. and S. M. Zola (1996). "Structure and function of declarative and nondeclarative memory systems." Proceedings of the National Academy of Sciences **93**: 13515-13522.
- Swanton, C. (2003). "75 Virtue Ethics." International Encyclopedia of the Social and Behavioral Sciences, Elsevier Science.
- Thomson, J. J. (1985). "The Trolley Problem." Yale Law Journal **94**: 1395-1415.
- Thomson, J. J. (1986). Rights, Restitution, and Risk: Essays, in Moral Theory. Cambridge, MA, Harvard University Press.
- Tulving, E. (1972). Episodic and semantic memory. Organization of Memory. New York, Academic Press: 381-403.
- Tversky, A. and D. Kahneman (1974). "Judgment under Uncertainty: Heuristics and Biases." Science **185**(4157): 1124-1131.
- Unger, P. (1996). Living High and Letting Die. Oxford, Oxford University Press.
- van den Berg, H., Manstead, A.S.R., van der Pligt, J., Wigboldus, D.H.J. (2006). "The impact of affective and cognitive focus on attitude formation." Journal of Experimental Social Psychology **42**: 373-379.
- Walker, L., J. (1989). "A longitudinal study of moral reasoning." Child Development **60**(1): 157-166.
- Watson D, Clark, L., Tellegen A. (1988). "Development and validation of brief measures of positive and negative affect: the PANAS scales." Journal of Personality and Social Psychology **54**(6): 1063-1070.
- Wechsler, D. (1999). Wechsler Abbreviated Scale of Intelligence. San Antonio, Psychological Corporation.
- Welniak, E. and K. Posey (2005). Household Income: 1999, US Census Bureau.
- Wilson, M. (1987). "MRC Psycholinguistic Database: Machine Usable Dictionary. Version 2.00." from [http://www.psy.uwa.edu.au/mrcdatabase/uwa\\_mrc.htm](http://www.psy.uwa.edu.au/mrcdatabase/uwa_mrc.htm).
- Winograd, E. and U. Neisser (1992). Affect and Accuracy in Recall: Studies of "Flashbulb" Memories. Cambridge, MA, Harvard University Press.

- Wuensch, K. L. and G. M. Poteat (1998). "Evaluating the morality of animal research: Effects of ethical ideology, gender, and purpose." Journal of Social Behavior and Personality **13**: 139-150.
- Young, L., Cushman, F., Adolphs, R., Tranel, D., Hauser, M. D. (2006). "Does Emotion Mediate the Effect of an Action's Moral Status on its Intentional Status? Neuropsychological Evidence." Journal of Cognition and Culture **6**: 291-304.
- Zola-Morgan, S., Squire, L. R., Amaral, D. G. (1986). "Human amnesia and the medial temporal region: enduring memory impairment following a bilateral lesion limited to field CA1 of the hippocampus." The Journal of Neuroscience **6**(10): 2950-2967.