

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Intuition versus Reason: Strategies that People Use to Think about Moral Problems

Permalink

<https://escholarship.org/uc/item/07p605xm>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 35(35)

ISSN

1069-7977

Authors

Fedyk, Mark
Koslowski, Barabara

Publication Date

2013

Peer reviewed

Intuition versus Reason: Strategies that People Use to Think about Moral Problems

Mark Fedyk (mfedyk@mta.ca)

Department of Philosophy, Mount Allison University
Sackville, NB, E4L 1G9, Canada

Barbara Koslowski (bmk2@cornell.edu)

Department of Human Development, Cornell University
Ithaca, NY, 14853, USA

Abstract

We asked college students to make judgments about realistic moral situations presented as dilemmas (which asked for an either/or decision) vs. problems (which did not ask for such a decision) as well as when the situation explicitly included affectively salient language vs. non-affectively salient language. We report two main findings. The first is that there are four different types of cognitive strategy that subjects use in their responses: simple reasoning, intuitive judging, cautious reasoning, and empathic reasoning. We give operational definitions of these types in terms of our observed data. In addition, the four types characterized strategies not only in the whole sample, but also in all of the subsamples in our study. The second finding is that the intuitive judging type comprised approximately 26% of our respondents, while about 74% of our respondents employed one of the three styles of reasoning named above. We think that these findings present an interesting challenge to models of moral cognition which predict that there is either a single, or a single most common, strategy – especially a strategy of relying upon one’s intuitions – that people use to think about moral situations.

Keywords: Moral Judgment; Moral Reasoning; Verbal Behavior; Cognitive Strategies.

Introduction

A significant amount of the recent and most influential research in adult moral psychology has focused on the intuitive processes that are frequently represented as the main sort of cognition that occurs when people are asked to make moral judgments about moral situations. This intuition-focused research frequently relies upon so-called “trolley dilemmas” to elicit this intuitive cognition in research participants. Trolley dilemmas are easy-to-understand fictional moral situations that present an either/or choice between one of two ethically appealing courses of action. Participants are asked to decide one of the two courses of action, and this decision is customarily treated as a paradigmatic representation of a moral judgment. (Mikhail, Sorrentino, & Spelke 1998, Mikhail 2012, Nichols & Mallon 2006, Cushman & Greene 2012)

In one of the better-known lines of research in this area, Mikhail (2007) proposes that an innate and subconscious ‘universal moral grammar’ determines the semantic content of people’s moral judgments. The universal moral grammar is defined using the concepts of deontic logic augmented with a variety of psychological concepts like ‘intentional

action’. Moreover, the computations that the universal moral grammar is hypothesized to perform and that generate moral judgments are entirely intuitive. Thus, evidence for the existence of the universal moral grammar comes from studies in which subject’s intuitive decisions about trolley dilemmas are shown to be consistent with theoretical predictions about the content and logical structure of the moral grammar. (Mikhail 2008, Mikhail 2012)

Similarly, Waldmann and Dieterich (2007) argue that the content of moral judgments about trolley dilemmas and similar moral situations is determined by a different kind of intuitive knowledge, namely knowledge of a causal map. This map symbolizes causal pathways in the moral situation that a person is considering, dividing the causal pathways into those that involve *agents* and *patients* and those that do not. Agents play an active role in realizing a harmful effect, and patients experience the harmful effects caused by agents. Accordingly, the either/or choice at the heart of the situation is intuitively represented as choice between two actions (interventions), each of which can alter in different ways the causal relations that hold between agents and patients. Waldmann and Dieterich test whether a “focus on action may sometimes lead to what we call *intervention myopia*”, such that people will focus primarily on interventions on causal pathways involving patients and agents, and focus less on interventions on causal pathways that do not involve agents and patients. (Waldmann and Dieterich 2007) Their data leads them to conclude, “the locus of intervention is one key factor contributing to moral intuitions.” That may be, but the important point for our purposes is that, like Mikhail’s theory, Waldmann and Dieterich’s proposal is that judgments about moral either/or choices are computed by intuitive cognitive processes.

We mention these studies in some detail because we want to use them as evidence for the following claim: if you want to “nudge” subjects towards using intuitive cognitive processes to produce moral judgments about moral situations, one approach which seems likely to have this effect is to present subjects with a moral situation that embeds an explicit either/or choice – that is, a moral *dilemma*. After all, both studies described above are examples of experimental subjects producing moral intuitions in response to situations that contain explicit either/or choices, and there are many other studies like this. (Lombrozo 2009)

But there is another body of literature in recent moral psychology that suggests that different way of inducing intuitive responses is to ask subjects to engage with *affectively salient* moral situations, whether or not they are technically moral dilemmas. (Borg, Hynes, Van Horn, Grafton, & Sinnott-Armstrong, 2006; Ciaramelli, Muccioli, Ladavas, & di Pellegrino, 2007; Damasio, 1994; Greene, Nystrom, Engell, Darley, & Cohen, 2004) Thus, for example, Jonathan Haidt's well-known social intuitionist model (SIM) holds that intuitively generated affective states almost always fix the content of people's moral judgment. (Haidt 2001) The basic idea is that things that we intuitively feel that we like are judged to be morally permissible, and things we intuitively feel dislike for are judged to be morally impermissible. So creating feelings along the like/dislike continuum is another potential way of encouraging uses of moral intuition; and it certainly should be if the social intuitionist model is correct.

We believe that the various lines of research that focus on the intuitive aspects of moral cognition have produced a number of novel and important scientific insights into the relationship between intuition, affect, and moral judgment. However, we are skeptical of the idea that – as per Haidt, Mikhail, and others – intuitive processes are the very nearly the *only* processes by which people form moral judgments. Accordingly, we present here the results of an experiment designed to identify some of the different cognitive strategies that people use when thinking about moral situations. And one of the questions we were most interested in answering was just how *frequently* people rely on their moral intuitions when responding to a moral situation. Because of this, our experiment was designed to maximize the likelihood that some of our participants use their moral intuitions to respond to the moral situations we asked them to consider. Our test conditions were *Dilemma Non-Affect*, *Dilemma Affect*, *Problem Non-Affect*, and *Problem Affect*, and all but the last condition was constructed so as to try to “push” subjects in these conditions towards the use of intuitive cognitive processes. Specifically, our *Dilemma* conditions were designed to replicate closely the trolley dilemmas discussed above by presenting our subjects with a clear either/or choice. Similarly, our *Affect* conditions – which included language designed to elicit feelings of either disgust or sympathy – were designed to target the intuitively-mediated emotional processes that are posited by theories like the social intuitionist model. We provide a fuller description of our test conditions below – for now, we simply want to make the point that the rationale for our test conditions was our goal of trying to encourage subjects in some of our test conditions to use moral intuition. Thus, *Dilemma Non-Affect* targets intuitive systems like those posited by Mikhail, Waldmann, Dieterich and others; *Problem Affect* targets intuition systems like those posited by the social intuitionist model; *Dilemma Affect* targets *both* kinds of intuitive systems; and *Problem Non-Affect* acts as a control condition, insofar as it is not designed to target any specific intuitive process that has been described in the recent moral psychology literature.

Furthermore, we want stress that our aim was not to show that people are more likely to use intuitive cognitive processes in any of our test conditions. Although we designed our experiment to maximize the chances that participants in our dilemma and/or our affect conditions would be more likely to use moral intuition than those in the problem non-affect condition, our fundamental aim was to identify cognitive strategies that did *not* vary across our test conditions. In other words, we wanted if possible to identify any *condition invariant* cognitive strategies, while at the same time employing test conditions that worked against this end by making it more likely that participants in some of these test conditions would use *different* cognitive strategies.

Method

We collected our data using person-to-person interviews rather than online interviews. We did this because in a separate experiment (Koslowski and Fedyk, *in prep*) we observed that in person-to-person interviews subjects produce a richer expression of the cognitive processes they use when responding to moral problems than they do in online-only sampling contexts.

In all of our conditions, participants were simply asked for a judgment about a moral decision faced by a character in a fictional vignette.

Participants. Participants were eighty-three undergraduate students (m=43, f=40) at Cornell University. Participants were enlisted using the university's internet-based recruitment tool, and all participants received course credit for their participation.

Interview. Our interviews took place in a quiet lab room at Cornell University. No one but the interviewer and the participant was in the room at the time of the interview. Each interview was recorded using either a digital audio recorder or a tape recorder, and tapes from the latter were subsequently digitized. Data for 1 female participant was excluded from our analysis because of a tape-recorder failure that occurred during this participant's interview and thus prevented her interview from being transcribed.

Stimuli. Each participant was presented with 6 different vignettes that described a situation in which the main character in the vignette faces a moral choice. The moral situation described by our vignettes intentionally resembled the situations described by vignettes that have been used in previous research. For example, our “Smith” vignette is version of the classic runaway trolley case, albeit involving people trapped in a subway tunnel. We also used an updated version of Kohlberg's famous “Heinz” vignette, and a very simple vignette derived from Peter Singer's famous article about moral obligations towards people experiencing a devastating famine in a distant country. (Singer 1972)

All of our vignettes were written in plain English, and one of the ways in which they differed from other vignettes used by some other moral psychologists is that they described

situations that either have occurred or at least could likely occur. We did this in order to increase the ecological validity of the study, as some vignettes used by other researchers require subjects to engage in deeply counterfactual thinking. The “fat man” trolley case, for example, requires subjects to believe that, despite the laws of physics and human physiology, there exists a man fat enough to stop a runaway trolley car. (cf. Pinker 2008)

Our vignettes varied in their length, where the shortest vignette was 85 words long, and the longest 303 words long, with an average length of approximately 196 words. Each vignette introduced a main character with a gender-neutral name (like “Smith”, “Davis”, or “Parker”) and described a situation faced by the main character that called for a moral decision.

Here are more explicit definitions of the four types of vignettes we used:

Dilemma Vignettes – for these vignettes the same language as for the *Problem* condition is used, except that a short phrase (like “Smith can either...”) or sentence is added to the vignette that stipulates that the main character faces an either/or choice.

Affect Vignettes– in this condition, 1 or 2 short sentences were added to either the *Dilemma* or *Problem* vignettes. The sentences were designed to elicit mild feelings of either sympathy or disgust in our participants. Examples of these sentences are:

- a. “The cancer is very painful, and the woman cries most days.”
- b. “Relief workers are trying desperately to treat children who are suffering a range of painful and eventually fatal illnesses caused by malnutrition.”
- c. “The boss is dirty and smells bad. He tells Adams that the sandwich he is eating is a horse-meat and pickles sandwich...”
- d. “Lisa is one of the members of Smith’s team. She works to support her two high-school aged children after her husband died of cancer several years ago.”

Non-Affect Vignettes – for these vignettes, the sentence designed to elicit affect is omitted.

Problem Vignettes – in these vignettes, the language describing the “either/or” choice is omitted.

We constructed four test conditions by crossing these two variables: decision type (problem vs. dilemma) and affect type (affect vs. non-affect). Thus, the four test conditions were *Dilemma Non-Affect*, *Dilemma Affect*, *Problem Non-Affect*, and *Problem Affect*. We did this so that – as explained in more detail above – *Dilemma Affect* was the condition that for theoretical reasons was most likely to push our participants towards using their moral intuitions.

Participants were randomly assigned to one of the four test conditions. Subjects therefore only ever responded to one type of vignette. Within each condition, the order in which the six vignettes were presented was random.

Procedure

After participants had settled into the interview room and consented to participate, participants were given an unbound stack of 6 pieces of paper, where each piece of paper had written upon it one of the 6 vignettes. Participants were then asked to read along silently while the interviewer read the vignette out loud. After the interviewer was done reading the vignette, he or she then asked the participant, “What is your judgment about what X should do?”, where “X” stands for the name of the main character in the relevant vignette. Participants were provided with no further instructions or feedback. They were not asked to produce any other judgments than a judgment about what the main character should do. Neither were they asked *explicitly* to reason about the vignettes they were presented with. Participants were therefore free to respond to our question however they wanted, which means the reasoning we observed (see below) was produced spontaneously. Once each participant concluded his or her response to our question, the interviewer moved on to the next vignette. This process was repeated until each subject had responded to all six of the vignettes.

Coding and Analysis

The audio recordings from each of our interviews were transcribed by a professional transcription service that specializes in legal and academic work. The transcriptions were made using an “absolute verbatim” style, which means that every utterance, pause, “hmmm”, and so on was transcribed and, importantly, done so using a standardized notation.

Two coders who were blind to our study’s hypotheses, aims, and methods then coded these transcriptions independently. Disagreements between our coders were very infrequent, and were resolved through discussion. 82 interviews were coded, where each response to the question “What is your judgment about what X should do” was treated as an individual case. This means that our data set consisted of 492 discrete cases.

Coding Categories. We created 11 coding categories that describe easy to observe speech-acts or other kinds of verbal behavior. For example, one of our categories was “Subject asks at least one question about the vignette”. Only one of our 11 categories explicitly referred to verbalized reasoning (see 4. below). The remaining 10 categories were derived from examining our transcripts for reoccurring speech-acts. We used this approach so that we did not render it a priori that our data analysis would find either an intuition / reason distinction, or find different types of reasoning in our cases. Thus, we had prior to running our analysis as much reason to expect that our analysis would sort responses into, for instance, questioning and non-questioning responses as we did for reasoned and intuitive responses.

Each of the coding categories was defined as a categorical variable, and no coding categories were treated as exclusive of any other. This permits our coding categories to nest within one another, and this property allows us to logically

construct the definitions of the cognitive strategies out of the definitions of our coding categories.

Analysis. We used a two-step cluster analysis algorithm to find natural clusters formed in the cases that comprise our observational data. Specifically, we looked for clusters that occurred in *all* of the different populations of cases that we could create by sorting according to gender or test condition. We used the two-step algorithm because it is able to find natural clusters in categorical variable data.

The algorithm looks for cases that have the same coding category values, and creates a preliminary cluster out of any set of cases that have the same values. It then scores a number of different “models” of the clusters identified in first step of the analysis according to their Bayesian Information Criteria (BIC). Importantly, this second step is able to resolve any borderline cases: if a particular case *c* is similar to those cases in a group of cases *G* which all have exactly the same values, and *c* is also different than many other cases that, in the first step, the algorithm did not put into any clusters, the algorithm may then put *c* into *G* if the model which places *c* in *G* has the best BIC score. However, if the case data is too heterogeneous – as may occur if there are nearly as many clusters as there are cases – then the algorithm in the second step will delete some or all of the preliminary clusters. Because of this, it is also possible that two-step cluster analysis will find no natural clusters in some data sets. (cf. Norušis 2011)

We examined combinations of 11 different coding categories applied to 11 different populations of cases (see below). More explicitly, we looked for a combination of coding categories which the two-step algorithm was able to use to find clusters that (a) occurred in all 11 populations of cases we analyzed and that (b) had a silhouette coefficient greater than 0.6. We also wanted to identify a set of clusters that (c) were the *only* natural clusters in all 11 populations of cases.

The 11 populations of cases we analyzed were: all participants, male / female participants, participants in each of the four test conditions (*Dilemma Affect*, *Dilemma Non-Affect*, *Problem Affect*, *Problem Non-Affect*), and participants in each of the types of situations used to construct our test conditions (*Dilemma*, *Affect*, *Problem*, *Non-Affect*). Thus, we used the two-step algorithm to determine if the coding category “Subject says what the main character should do” picked out any natural clusters in the populations listed above. We also used the cluster analysis algorithm to determine if the two coding categories “subject says what the main character should do” and “subject asks for more information about the vignette” *together* picked out any natural clusters in the 11 populations of cases above. And we also asked the cluster analysis algorithm to determine if the *three* coding categories consist of the previous two plus “subject uses moral language in their response” *together* picked out any natural clusters of cases in the 11 populations of cases listed above. And so on.

Thus, we supplied approximately 121 different

combinations of coding categories to the two-step algorithm. As we said, we were searching for clusters of cases that occurred in all 11 populations of cases and which scored a high silhouette coefficient (> 0.6). Any such clusters would therefore represent types of responses that were invariant across conditions and populations.

Results

We found 4 such clusters. Specifically, we found that a combination of 4 coding categories defined four different natural clusters that occurred in all 11 populations listed above. What’s more, these 4 clusters were the *only* natural clusters in 9 of the 11 populations. The coding categories that define these clusters are:

1. Subject says what the main character should do.
2. Subject uses the word “might” or “probably” or a similar word to express hesitation when verbalizing their judgment.
3. Subject says something indicating that they are imagining themselves in the situation of the main character of the vignette (such as “Well, what I would do in that situation is...” or “If it was me there, I think that I...”).
4. Subject expresses at least one inference when making their response (such as “if ... then ...” or “... because ...”).

Because these were treated as categorical variables, each of these coding categories can take the only the values “true” or “false”. Each of our 492 cases will therefore have a value of “true” or “false” for each of these categories. This means that there are 16 logically possible clusters that the cluster analysis algorithm could have found using these coding categories, although it is also possible that the algorithm find could have found no clusters at all.

Here are the clusters that the algorithm found. Note that each cluster is operationally defined out of logical values for the four coding categories listed directly above:

Simple Reasoning = subject expresses a judgment and expresses an inference, but does not use language indicative of hesitation and does not imagine themselves in the position of the main character.

Intuitive Judging = subject expresses a judgment, and does not express an inference, does not use language indicative of hesitation and does not imagine themselves in the position of the main character.

Cautious Reasoning = subject expresses a judgment, expresses an inference, and does use language indicative of hesitation, but does not imagine themselves in the position of the main character.

Empathic Reasoning = subject expresses a judgment and expresses an inference and imagines themselves in the position of the main character, but does not use language indicative of hesitation.

We think that these four natural clusters – or, if you prefer, types of response – represent four *cognitive strategies* that people use to respond to moral situations. And just to be clear, these are not types of *people*; they are *cognitive*

strategies that occurred in all of our test conditions.

Figure 1 presents the proportions of these clusters in our total respondents – that is, for all 492 cases taken together. For this population, the cluster analysis algorithm placed every case into one of the four clusters, meaning that no case was excluded. Importantly, there was no significant variation in the relative proportion of these four clusters across all of the populations of cases that we analyzed. This means that the ratio of reasoners to intuitors was approximately 3:1 in all of our populations. It also means that the proportion of the four clusters in, for instance, the male population of cases looks very similar to the proportion of clusters in our total population of cases; for illustration, please see Figure 2.

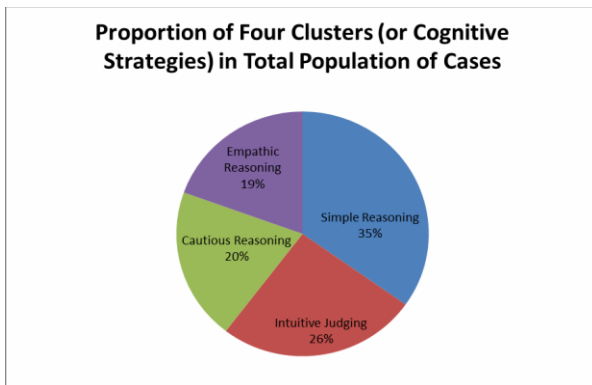


Figure 1: The proportion of four natural clusters found in the population consisting of all 492 of our cases. Each of the 492 cases was placed into one of these four clusters. Proportion is expressed as a percentage, rounded to the nearest whole number. (Silhouette measure of cohesion and separation = 0.7)




Figure 2: The proportion of four natural clusters found in the population consisting of only our male cases. As in Figure 1, each case was placed into one of these four clusters. Proportion is expressed as a percentage, rounded to the nearest whole number. (Silhouette measure of cohesion and separation = 0.7)

As we indicated above, we also found that these clusters were the *only* natural clusters that were present in 9 of the 11 case populations we analyzed. The two exceptions were the subpopulation of cases in the *Problem Non-Affect* condition (where these four clusters accounted for 75% of the natural clusters in the population) and the female respondents (where these four clusters accounted for 85% of

the natural clusters in the populations). Remember, the algorithm can resolve borderline cases by placing statistically similar though not identical cases together in the same cluster, but borderline cases will not always be resolved in the same way across different populations of cases. This is because the treatment of a borderline depends partly on what the statistical properties of other borderline cases in the population under analysis.

Finally, we would like to report that that the coding category “subject uses moral language” failed to figure in any of the condition invariant natural clusters. We find this result particularly intriguing.

Impact of the Conditions and Other Objections

A natural worry with our claim that the four cognitive strategies that we observed are condition invariant is that our test conditions simply failed to have any experimentally meaningful impact on our subjects – even though three of our four test conditions were designed on theoretical grounds to try to push subjects in those conditions towards intuitive responses.

So as a control for this possibility, we analyzed the cases in the different test conditions for any significant differences, and we found several. For instance, subjects in the two affect conditions were more likely to ask our interviewer for information about the vignette than subjects in the two non-affect conditions ($x^2 = 6.54, p=0.0105$). We also found that, when we scored the coherency of the reasoning on a 7 point scale derived from a grading rubric used in a critical reasoning course, the coherency of reasoning of the subjects in the *Problem Non-Affect* condition was significantly higher than all other conditions (e.g., for *Problem Non-Affect* versus *Dilemma Affect*, $x^2 = 19.05, p = 0.0019$). These data indicate that our test conditions did have different psychological effects, and this speaks to the strength of the clusters we found in our data.

We would also like to speak to the assumption that differences in people’s verbal responses can be read without further experimental constructs as evidence of differences in the underlying cognitive processes. This assumption is often implicit in the analysis of experimental data in moral psychology, and it is most prominent in the work of researchers who have taken the view that moral cognition is largely driven by intuition processes (cf. Haidt 2001). Our position is that this assumption is warranted as a premise in an abductive inference for our conclusion – namely, that the best explanation of the differences we observed in our subject’s verbal responses is that these differences reflect different underlying cognitive strategies.

In sum, we claim that the four natural clusters we defined above represent four different cognitive strategies that people use to respond to moral situations. Sometimes people are simple reasoners, intuitive judgers, cautious reasoners, or empathic reasoners – and, importantly, these four strategies are used whether or not people are asked explicitly to think about an either/or dilemma, and whether or not they read a vignette designed to induce mild feelings of either disgust or sympathy.

Discussion

Some of the categories we used in our analysis did not yield any condition invariant clusters, and this result provides an interesting independent confirmation of some of the claims made by the social intuitionist model. The operational definition of “intuitive judger” above is nearly identical to the theoretical definition of an “intuitive judgment” given by Haidt (2001). And because in our experiment there was no a priori reason to think that the intuitive judging cluster would be one of the four clusters found in all of our test conditions, our observation that a large number of our respondents behaved in a way that very nearly exactly satisfies the social intuitionists’ definition of “intuition” is therefore evidence of the accuracy of their theoretical definition for the concept. This result, moreover, comports very well with dual-process approaches to cognition.

However, we failed to observe significantly more intuitive judgers in the three conditions designed to induce intuitive moral cognition. The ratio of intuitive judgers to reasoners held steady across all of our test conditions. Remember: participants were given no explicit instruction to reason about the moral situations we read to them; we asked each participant for *only* a judgment about what the main character in the vignette should do. So the fact that we were unable to “push” subjects in some conditions to rely more frequently on moral intuition is a challenging result to intuition-based models of moral cognition like the social intuitions model and Mikhail’s universal moral grammar. Despite our attempts to maximize the likelihood that participants would use intuitive cognitive processes in some of our test conditions, participants were in all of our conditions at least three times more likely to use some kind of reasoning than to use intuition.

This finding is relevant to the methodology of moral psychology. Moral intuition is often defined as the absence of reasoning, and reasoning is a normative ability the manifestation of which varies according to the skill and epistemic context of a subject. By setting the range of permitted reactions to an either/or choice or recording agreement with a proposition on a Likert scale, many moral psychological experiments are automatically designed not to record any reasoning. Yet, these constructs do not cause subjects to not reason during the experiment; they only proactively “screen-off” the expression of any underlying reasoning that may or may not occur. Our experiment was designed to see what subjects would do when this screen was removed, and our findings suggest that reasoning is a very common cognitive strategy used to arrive at moral judgments. But the deeper lesson implied by our findings is that there is more than a single concept of moral intuition employed in contemporary experimental moral psychology. Let an *experimental intuition* be any judgment recorded in an experiment where the subject is prevented from expressing any reasoning that may or may not occur in the production of the judgments, and let a *psychological intuition* be any judgment that is not produced on the basis of any immediately prior reasoning. Our findings suggest that many moral psychologists are studying *only*

experimental intuitions – for exactly the same reason that our “intuitive judging” category might capture only experimental, not psychological, intuitions.

Acknowledgments

We would like to thank Richard Boyd, Tamar Kushnir, Luca Bonatti, Albert Costa, Jane Dryden, Robbie Moser, Roopen Majithia, and the members of the RICO group at the University Pompeu Fabra. We are grateful for the financial support of both Mount Allison University and Cornell University.

References

- Ciaramelli E, Muccioli, M., Làdavas, E, and di Pellegrino, G. (2007) Selective deficit in personal moral judgment following damage to ventromedial prefrontal cortex. *Social Cognitive and Affective Neuroscience*, 2, 84-92.
- Cushman, F.A. and Greene, J.D. (2012) Finding faults: How Moral Dilemmas reveal cognitive structure. *Social Neuroscience*, 7(3), 269-279
- Damasio, A (1994) *Descartes’ Error: Emotion, Reason, and the Human Brain*. Putnam Books.
- Greene, J.D., Nystrom, L.E., Engell, A.D., Darley, J.M., and Cohen, J.D. (2004) The neural bases of cognitive conflict and control in moral judgment. *Neuron*, 44, 389-400.
- Haidt, J. (2001) The Emotion Dog and its Rational Tail. *Psychological Review*, 108, 814-834.
- Koslowski, B, and Fedyk, M. (in prep) Online versus person-to-person surveys in moral psychology
- Lombrozo, T. (2009) The role of moral commitments in moral judgment. *Cognitive Science*, 33, 273-286.
- Mikhail, J., Sorrentino, C.M., and Spelke, E. (1998) Toward a Universal Moral Grammar. *Proceedings of the 12th Annual Conference of the Cognitive Science Society*.
- Mikhail, J. (2011) *Elements of Moral Cognition: Rawls’ Linguistic Analogy and the Cognitive Science of Moral and Legal Judgment*. Cambridge University Press.
- Mikhail, J. (2007) Universal Moral Grammar: Theory, Evidence, and the Future. *TRENDS in Cognitive Science*, 11(4), 143-152
- Nichols, S. and Mallon, R. (2006) Moral Dilemmas and Moral Rules. *Cognition*.
- Norušis, M.J. (2011) *IBM SPSS Statistics 19 Guide to Data Analysis*. Prentice Hall.
- Schaich Borg, J., Hynes, C., Van Horn, J., Grafton, S., and Sinnott-Armstrong, W. (2006) Consequences, action, and intention as factors in moral judgments: an fMRI investigation. *Journal of Cognitive Neuroscience*, 18(5), 803-17
- Singer, P. (1972) Famine, Affluence, and Morality. *Philosophy and Public Affairs*, 1, 229-243.
- Pinker, S. (2008) *The Moral Instinct*. New York Times Magazine.
- Waldmann, M.R. and Dieterich, J.H. (2007) Throwing a Bomb on a Person versus Throwing a Person on a Bomb: Intervention Myopia in Moral Intuitions. *Psychological Science* 18(3), 247-253