



Seattle Pacific University
LIBRARY
Discover, Create, Share

Seattle Pacific University
Digital Commons @ SPU

SPU Works

2015

What Is Moral Reasoning?

Leland F. Saunders

Follow this and additional works at: <http://digitalcommons.spu.edu/works>

 Part of the [Ethics and Political Philosophy Commons](#)

Recommended Citation

Saunders, Leland F. (2015) "What is Moral Reasoning?" *Philosophical Psychology*, 28 (1); 1-20.

This Article is brought to you for free and open access by Digital Commons @ SPU. It has been accepted for inclusion in SPU Works by an authorized administrator of Digital Commons @ SPU.

What is Moral Reasoning?

1. What Role for Moral Reasoning in Moral Judgment?

What role does moral reasoning play in moral judgment? More specifically, what *causal role* does moral reasoning have in the production of moral judgments? Answering this question is important for understanding the causal mechanisms of moral judgment, but it is also important because it promises to shed light on a range of other disputes with respect to morality, such as whether morality and moral judgments are based in reasons and are rationally assessable, and whether, and under what conditions, moral judgments can be true (see, for example, Cholbi, 2006; Haidt & Bjorklund, 2008; Jones, 2006; Kennett & Fine, 2009; Nichols, 2002, 2004; Prinz, 2007; Roskies, 2003). One way of answering this question, and one that has been very influential among both philosophers and psychologists is to claim that moral judgment just is the application of domain-general reason¹ to moral questions; and consequently that all moral judgments are causally produced by moral reasoning (for example, see Bucciarelli et al., 2008; Hare, 1981; Herman, 1993; Kant, 1785/1996; Kohlberg et al., 1983; Piaget, 1932/1965).² For many theorists who subscribe to this particular psychological picture of moral reasoning and moral judgment, one clear upshot of the view that it provides a straightforward account of the rationality of morality and the rational assessability of moral judgments.

¹ Domain-general reasoning, on this sort of view, envisions reasoning as an all-purpose and general capacity for reasoning in all domains, such as scientific, mathematical, and moral.

² This view of moral judging is often attributed to Aristotle as well, based on his notion of the practical syllogism where practical action issues from the conclusion of deductive, syllogistic reasoning (see, *Nicomachean Ethics* 1147a26-32). However, there are reasons to interpret Aristotle here not as claiming that domain-general reasoning is accurate psychological picture of moral judgment, but that it rather provides a rational reconstruction that serves to justify particular moral judgments (Hughes, 2003).

A growing amount of empirical data, however, appears to challenge the claims that moral judgment just is the application of domain-general reasoning to moral questions and that all moral judgments are causally produced by moral reasoning. Research on moral judgment shows that people are often dumbfounded with respect to their moral judgments—that is, they are unable to provide further plausible reasons for them (Cushman et al., 2006; Haidt, 2001; Hauser, 2006; Hauser et al., 2007), that people’s moral judgments are reliably influenced by their emotions, especially disgust and happiness (Haidt et al., 1993; Inbar et al., 2009; Schnall et al., 2008; Valdesolo & DeSteno, 2006; Wheatley & Haidt, 2005), that emotion centers of the brain light up on fMRI scans when subjects are presented moral sentences or images (Moll et al., 2002; Moll et al., 2001) and some moral dilemmas (Greene et al., 2001), and that people’s responses to moral dilemmas are subject to framing and ordering effects (for an overview see, Sinnott-Armstrong, 2008).³

While this research has done much to illuminate the causal role of emotions in moral judgment (though it leaves many open questions as well), it is less clear what implications it has with respect to the causal role of moral reasoning in moral judgment. Some philosophers and psychologists argue that these data show that domain-general reasoning has little, if any, role in moral judgment (Haidt, 2001; Haidt & Bjorklund, 2008a; Prinz, 2006, 2007), while others argue that emotions and moral reasoning are independent pathways for arriving at a moral judgment (Greene, 2008; Greene, et al., 2001), and still others argue that the data are wholly consistent with the claim that moral

³ Framing effects derive their name from the fact that how certain options are framed (for example, as a gain or a loss) affects people’s responses to those cases. An ordering effect occurs when people’s responses to cases are affected by the order of presentation of the cases under consideration (Kahneman, 2003).

judgments are all products of moral reasoning (Fine, 2006; Horgan & Timmons, 2007; Jones, 2003; Kennett & Fine, 2009).

Which of these views is correct? What do the empirical data tell us about the causal role of moral reasoning in moral judgment, and what, if any, light do these data shed on disputes with respect to the rationality of morality? Answering these questions is more difficult than it initially appears, in part, because of the difficulty in determining what moral reasoning *is*. The literature is full of different and inconsistent conceptions of what moral reasoning is, and so far there has been very little work done to help choose among these competing conceptions. The aim of this paper is to make some progress towards addressing this problem by surveying the various competing conceptions of moral reasoning in the current literature, drawing out the relevant areas of dispute among them, and providing a framework that will aid in choosing among them. The following analysis is not meant to be definitive, but rather, it is meant to provide a starting point for ongoing research and a means for adjudicating the disparate and competing claims that currently abound in the moral psychological literature.

1. Defining The Issue

For any investigation it is important first to characterize the object of study. It is not possible to understand the causal role of moral reasoning in moral judgment without first having some understanding of what moral reasoning is. And, as Harman and colleagues point out, determining what causal role moral reasoning has in moral judgment and what counts as an episode of moral reasoning “depends partly on which [psychological] processes get called moral reasoning” (Harman et al., 2010, p. 206). In

this regard there is wide variability in the literature with respect to which process or processes get called moral reasoning, and how moral reasoning is defined or characterized.

As stated before, many philosophers take it that moral reasoning just is the application of domain-general reasoning to moral questions, and that it is conscious, intentional, and effortful (see, for example, Cohen, 2004; Rachels, 1999; Smith, 1994; Toulmin, 1968; Vaughn, 2008).⁴ This view of moral reasoning was explicitly adopted by psychologists in the early stages of empirical investigation studying the role of moral reasoning in moral judgment (Kohlberg, 1984; Piaget, 1932/1965). In the wake of more recent empirical findings, especially those pointing to some role for nonconscious psychological processes and the emotions in moral judgment, psychologists and philosophers have offered a number of different definitions of moral reasoning. For example, Bucciarelli et al. define reasoning as, “any systematic mental process that constructs or evaluates implications from premises of some sort” (Bucciarelli, et al., 2008, p. 123), and moral reasoning as reasoning that involves “deontic propositions” as premises, which are propositions that “concern what you may, should, and should not do or else leave undone” (p. 124). Importantly, for Bucciarelli et al., moral reasoning can be either an intentional and conscious process or a non-intentional⁵ and nonconscious process so long as the process (conscious or nonconscious) is systematic in the right kind of way; a clear departure from the view that moral reasoning is always a conscious process. Haidt, on the other hand, views moral reasoning as always involving at least two

⁴ This is just a small sample of philosophers who hold this position. The view is so widely assumed that it is rarely explicitly stated.

⁵ I use “non-intentional” as opposed to “unintentional” here because “unintentional” often connotes something done accidentally.

conscious steps, and defines it as a “conscious mental activity that consists of transforming given information about people in order to reach a moral judgment” (Haidt, 2001, p. 818). Moreover, according to Haidt, moral reasoning is “intentional, effortful, and controllable, and that the reasoner is aware that it is going on” (p. 818).⁶ This definition restricts moral reasoning to a partly conscious process involving information about people, but moral reasoning, on his view, also heavily relies on nonconscious psychological processes, in particular, the emotions.

Greene takes a different approach, and defines moral reasoning (or “cognition” in his terminology) behaviorally, by contrasting the behavioral effects of reason with the behavioral effects of emotions (Greene, 2008). On his view, “‘cognitive’ representations are inherently neutral representations, ones that do not automatically trigger particular behavioral responses or dispositions, while ‘emotional’ representations do have such automatic effects, and are therefore behaviorally valenced” (Greene, 2008, p. 40). Greene further elaborates that cognition is for “reasoning, planning, manipulating information in working memory, controlling impulses, and ‘higher executive functions’ more generally,” whereas emotions are “subserved by processes that in addition to be valenced, are quick and automatic, though not necessarily conscious” (pp. 40-41). In more recent work, however, Greene and colleagues have dropped this behavioral strategy, and instead define moral reasoning very specifically as a “Conscious mental activity through which one evaluates a moral judgment for its (in)consistency with other moral commitments,

⁶ This definition is repeated, almost verbatim, in Haidt & Bjorklund, (2008a).

where these commitments are to one or moral principles and (in some cases) particular moral judgments” (Paxton & Greene, 2010, p. 6).⁷

There are other theorists who do not provide a definition of moral reasoning, *per se*, but rather employ more general characterizations of it throughout their arguments, though these characterizations often reveal certain assumptions with respect to what they take moral reasoning to be. Prinz, for example, characterizes moral reasoning as necessarily involving the manipulation of “affect free” propositional attitudes (Prinz, 2007, p. 99), or as simply “dispassionate” (Prinz, 2006, pp. 37-40) without attempting to define what dispassionate process(es) constitutes moral reasoning. However, it seems as though Prinz is assuming that moral reasoning just is domain-general reasoning, and that domain-general reasoning necessarily involves manipulating “affect free” propositional attitudes, namely, beliefs, in a dispassionate manner. Others, however, characterize moral reasoning as a metacognitive process of some sort. Craigie (2011), for example, characterizes moral reasoning as an effortful, reflective, metacognitive process that can endorse or overturn moral intuitions, and Kennett & Fine (2009) similarly characterize moral reasoning as “a capacity for reflective shaping and endorsement of moral judgments” (p. 77).

The broad variability in definitions and characterizations of moral reasoning in the literature leads to two problems. First, it is not clear that people are all talking about the same thing when they argue that moral reasoning does or does not play a causal role in the production of moral judgments. Because there are many different theorists working with many different definitions or characterizations of the object of study, it is quite

⁷ This also appears to be the definition that Greene and colleagues are using in Paxton et al. (2011), though they do not explicitly define it.

possible for theorists to mean entirely different things when they conclude that moral reasoning does or does not play a causal role in the production of moral judgments. For example, both Bucciarelli et al. and Haidt argue that moral reasoning has a pervasive and important causal role in the production of moral judgments, but it is clear from the differences in their definitions of moral reasoning, and how their arguments proceed from their respective definitions, that they mean something entirely different by this claim. Given Bucciarelli et al.'s definition of moral reasoning as "any systematic mental process that constructs or evaluates implications from premises of some sort," they conclude that moral reasoning is causally implicated in every particular moral judgment,⁸ even those that are arrived at without going through any conscious steps, because moral reasoning, on this definition, can be entirely nonconscious so long as it is systematic and inferential in the right kind of way. They argue that this condition is satisfied, and thus that moral reasoning, in this sense, has a pervasive and important causal role in moral judgment.

Haidt, who defends the highly influential Social Intuitionist Model of moral judgment, also concludes that moral reasoning has a pervasive and important causal role in moral judgment, but he means something entirely different by this claim. According to his Social Intuitionist Model, moral judgments are largely caused by moral intuitions, which he defines as:

the sudden appearance in consciousness, or at the fringe of consciousness, of an evaluative feeling (like-dislike, good-bad) about the character or action of a person, without any conscious awareness of having gone

⁸ A particular moral judgment is judgment whose object is a particular action, agent, or state-of-affairs, e.g., it was wrong of Smith to take Jones's wallet, as opposed to a more general judgment that, for example, stealing is wrong.

through steps of searching, weighing evidence, or inferring a conclusion (Haidt & Bjorklund, 2008a, p. 188).

In the ordinary case, Haidt claims that these moral intuitions “lead directly” (p. 188) to a moral judgment, which on this view is “the conscious experience of blame or praise, including a *belief* in the rightness or wrongness of the act” (p. 188, emphasis in original).⁹ The psychological processes that produce moral intuitions are in sharp contrast to Haidt’s definition of moral reasoning as an intentional, effortful, and controllable “conscious mental activity that consists of transforming given information about people in order to reach a moral judgment.” For Social Intuitionists, moral reasoning plays no role in production of moral intuitions, so how can they claim that moral reasoning is pervasive and causally important to moral judgment? The explanation is that, according to the Social Intuitionist Model, moral judgment is also a social process, and one person’s expressed moral judgments can affect the intuitions and judgments of others. On this view, people are often pressed by others to provide reasons for their moral judgments, and they then confabulate reasons after the fact. These post hoc reasons have the potential to influence the moral judgment of others, either through direct social pressure, or because the reasons cause others to view the situation in a different light, causing them to have a new and different moral intuition leading to a new and different moral judgment.

This back-and-forth process between two or more people satisfies Haidt’s definition of

⁹ The precise informational and causal processes of the capacity for moral judging are not spelled out by Social Intuitionists. They claim that there is a “tight connection between flashes of intuition and conscious moral judgments” (p. 188), but a tight connection is not necessarily a causal one. This is an important problem, because it is difficult to understand what, precisely, Social Intuitionists take the role of emotions to be in moral judging. And as Dwyer notes, failing to spell out the nature of the link between intuition and judgment renders the supposed role of emotions in moral judging mysterious, unless Social Intuitionists simply want to stipulate that “moral judgments are moral intuitions made conscious” (Dwyer, 2009, p. 277).

moral reasoning (however loosely), and because he argues that this social aspect of moral judgment is pervasive, he concludes that moral reasoning has a pervasive and important causal role in moral judgment. This claim is very different than Bucciarelli et al.'s, even though they both claim that moral reasoning has a pervasive and important causal role in moral judgment.

Moreover, both Bucciarelli et al. and Haidt mean something quite different when they claim that moral reasoning has a pervasive and important causal role in moral judgment than when other theorists such as Paxton & Greene, Craigie, and Kennett & Fine argue that moral reasoning has a more episodic causal role in moral judgment, because the theorists in this latter group all view moral reasoning as having a metacognitive component involving reasoning about one's own moral judgments and reflectively endorsing or rejecting them (though these theorists' views differ in other important ways). If this is what moral reasoning consists in, then such reflective and metacognitive episodes will be rarer than the sorts of processes that either Bucciarelli et al. or Haidt consider to be moral reasoning, though such processes might still have an important causal role in moral judgment, namely, in shaping, evaluating, or overturning some moral judgments.

Thus, what one means in claiming that moral reasoning is or is not pervasive in moral judgment, or what one means in claiming any sort of causal role for moral reasoning in moral judgment (or denying such a role) importantly depends on how moral reasoning is defined or characterized. Indeed, how one even attempts to answer these questions, and what sorts of empirical evidence one takes to be relevant will depend on how moral reasoning is defined or characterized.

The second problem is more methodological, though related to the first, and it is that what gets called moral reasoning is generally a matter of *stipulation*. Most theorists simply stipulate what process or processes they consider to constitute moral reasoning, and then attempt to answer the question as to what causal role moral reasoning has in moral judgment *from that stipulation*, and problematically, the definitions or characterizations of moral reasoning offered are simply stipulated with little or no argument for the choice. This is not always the case, of course, and there are attempts within the literature to refine stipulated definitions by drawing upon relevant empirical findings. Paxton and Greene's definition of moral reasoning, for example, is intended to be a modification of Haidt's definition of moral reasoning based on a survey of relevant empirical findings. I shall return to this point later, but the typical strategy being pursued is to simply designate a particular process as the one that constitutes moral reasoning, and then ask whether that designated process plays any causal role in moral judgment by looking to the empirical literature. Call this the stipulation without argument strategy. This problem is all the more acute when considering how these claims are sometimes used to draw incredibly bold and often skeptical conclusions with respect to morality in general. For example, Haidt argues, based on his view of moral reasoning, that moral judgments cannot be rational in the sense that some are better supported by reasons than others (Haidt & Bjorklund, 2008a, 2008b).¹⁰

The stipulation without argument strategy has historical precedent. Hume very famously adopts such an approach in Book III of *A Treatise of Human Nature*. He defines reason, and moral reason by extension, as, "the discovery of truth or falshood. Truth or

¹⁰ Some moral judgments may be socially better than others, on his view, meaning that a particular moral judgment may make one more likely to secure social allies and participate in cooperative behavior.

falshood consists in an agreement or disagreement either to the real relations of ideas, or to real existence and matter of fact” (Hume, 1740, Book III, Part I, Section I). From this definition he concludes that moral distinctions, or judgments, are not the products of moral reasoning, and that moral reasoning has a very limited causal role in the production of moral judgments (it is limited to discovering facts that may then play a role in our emotional responses).

However, Hume’s famous skepticism with respect to the role of moral reasoning in moral judgment is derived almost exclusively from his limited and skeptical definition of what constitutes reasoning, and he offers little argument for this definition (Korsgaard, 1986). Similarly, most contemporary moral psychologists offer few, if any, arguments to support their definitions of moral reasoning; they simply designate a certain psychological process as the one that should be called moral reasoning, and then determine whether the empirical data indicate that that process has any causal role in moral judgment. But in order for any of these arguments to be successful, it must be the case that the process they have designated as moral reasoning really is what moral reasoning consists in. And it is at precisely this point that there is good reason to be skeptical that any of them have done so.

Indeed, each of the proposed definitions offered by contemporary psychologists and philosophers face serious difficulties as full accounts of moral reasoning. For example, Greene’s definition simply fails to draw a meaningful behavioral distinction between reasoning and emotions. Emotions, on his view, trigger stereotyped responses, such as the so-called “fight-or-flight” response some consider to be characteristic of fear. And this sort of stereotypical response of emotions contrasts with reasoning, which does

not trigger stereotypical behavioral responses. The problem here is that neither do all emotions. Being afraid of a test, for example, does not necessarily, or even generally trigger a “fight-or-flight” response. Likewise, sympathizing with those suffering from a long and protracted civil war on the other side of the world does not necessarily, or even generally trigger a behavioral response that might be considered stereotypical of sympathy, such as intervening in some way to alleviate their suffering. Similar problems can be raised for any morally significant emotion, such as empathy, indignation, anger, and disgust, and so there are no clear behavioral differences between reason and emotions of the sort that Greene relies on in his definitions. Haidt’s definition of moral reasoning as the “conscious mental activity that consists of transforming given information about people in order to reach a moral judgment” is far too limited. Why think that moral reasoning is restricted to transforming information *about people*? He provides no reason for thinking that moral reason cannot include thinking with or about moral principles, or about states-of-affairs, or about what is intrinsically good, or the moral status of animals, or environmental issues. Similarly, characterizing moral reasoning as reflective metacognition is quite limited. Metacognition is generally limited to thinking *about* one’s judgments, e.g., whether to act on them, and does not include thinking *with* one’s judgments or reasoning to a first-order moral judgment. That is simply what the “meta” in metacognition means. But again, why think moral reasoning is limited to just reasoning about moral judgments?

These objections are not meant to be definitive, and the point here is not to spend time discussing how these views of moral reasoning fall short, but rather, it is to make progress towards understanding how we can get the question right. And the general

problem with the stipulation without argument strategy is that one is simply begging the question from the outset—the conclusion one arrives at will simply depend on which psychological process one calls moral reasoning. This is hardly a promising start. What is needed is a way to specify what moral reasoning is in a non-question-begging way; in a way that does not amount to *merely* stipulating a particular psychological process as the process of moral reasoning. This makes it seem as though designating any particular psychological process as the psychological process that deserves to be called moral reasoning is something of a *radical choice*—a choice for which there are no reasons, and, moreover, no possible means of adjudication between competing choices and designations.

2. How to Make Progress

The first step in making progress towards this goal is to get clearer on how to go about specifying what moral reasoning is. Many theorists have attempted to specify what moral reasoning is by defining or characterizing it in terms of a particular psychological process (or a particular kind of psychological process), but another way of specifying what constitutes moral reasoning is not, in the first instance to specify a psychological process, but to specify a *capacity* of a certain sort, much like the capacity to read, or the capacity to multiply, or the capacity to see depth. A capacity, in this sense, is understood in terms of *what* it enables a person who possesses it to do (e.g., read, multiply, see depth), and a capacity can be understood simply as a complex dispositional property (Cummins, 2000). Capacities, in this sense, are subserved by psychological processes,

and what explains *how* a capacity works are the operations and structures of the subserving psychological processes (Cummins, 2000).

This distinction between capacities and psychological processes is similar to Marr's distinction between the computational level and the algorithmic level when analyzing an information-processing system (1982/2010 pp. 22-27). According to Marr, an analysis aimed at the level of computation specifies what the goal of a computational system is (what it does and why), while an analysis aimed at the level of an algorithm specifies how the computations are implemented in an algorithm, i.e., what the precise computations are that transform information within the system from input to output. The specification of a capacity, in the sense used here, can be understood as providing a computational level analysis of the target system in question, in this case, that of moral reasoning. Specifying the particular psychological processes that subserve moral reasoning, their internal operations and structure, can be understood as providing an algorithmic level explanation. Thus, one way of making progress in the case of moral reasoning is to first specify what the capacity for moral reasoning is (a computational level analysis) and then explaining how the capacity operates in terms of subserving psychological processes (an algorithmic level of analysis).

An example may make this strategy clearer, and illustrate the difference between a capacity and psychological processes. Vision is a good example. It makes little sense to say that there is a psychological process for vision. Rather, vision is a capacity—a capacity that enables us to see certain objects in certain lights, etc.—and what explains how that capacity works is a complicated and carefully orchestrated set of psychological

and physiological processes (Marr, 1982/2010).¹¹ That the capacity for vision requires a set of psychological processes is easily seen when we consider that it is possible to develop any number of distinct problems with the capacity for vision (Pinker, 1997). For example, some people cannot see edges well, or letters, or they lack the ability to see objects as those objects (Sacks, 1998).¹² Thus, how the capacity for vision works, and how it breaks down, is explained (in some sense) by the operations and structures of the psychological processes that subserve it. Determining whether a psychological process helps explain how a given capacity works requires providing an adequate specification of the capacity in the first place, and doing so is rarely straightforward, in part because specifying what a capacity is cannot always be completely separated from how the capacity operates.¹³

One way to proceed in the case of moral reasoning then, is not first to stipulate a psychological process, or set of psychological processes, or a sort of psychological process that constitutes moral reasoning, rather, it is to specify what the capacity for moral reasoning amounts to, which is done by specifying what it is moral reasoning enables us to do. This is not a simple task, and there can and will be substantive disagreements that arise here, but this does not undercut the point that it is important to specify what moral reasoning is as a *capacity*.

Adopting this strategy raises two related points. First, the psychological processes that constitute moral reasoning cannot be specified until there is an adequate specification

¹¹ Physiological processes are important in explaining the capacity for vision because the operations of the eyeball, for example, are not physiological, not psychological in nature.

¹² The eponymous patient in Sacks's book, for example, can see objects, but conceptualizes those objects incorrectly. Thus, he sees his wife, but mistakes her for a hat.

¹³ Cummins (2000), for example, points out that it is not always clear whether an observed psychological effect is incidental to the operation of a capacity, or part of the capacity itself. He uses the example of rotation and scanning effects in vision.

of the capacity for moral reasoning. That is, what constitutes an adequate *explanans* of the capacity for moral reasoning cannot be determined without an adequate specification of the *explanandum*, namely, the capacity for moral reasoning. Very likely, the reason that there are so many disparate definitions and specifications of the psychological processes that constitute moral reasoning is that there has not been sufficient attention paid to what the capacity for moral reasoning is and what an adequate specification of it would amount to. In the absence of any particular specification of the capacity for moral reasoning, it will be all but impossible to determine whether a particular psychological process explains (or helps explain) that capacity. The upshot of the capacity approach to moral reasoning is that it provides a method for adjudicating among competing specifications of the psychological processes of moral reasoning, namely, by determining which of them provides the best explanation of the capacity for moral reasoning. But doing that, of course, requires having some specification of the capacity for moral reasoning in the first place.

Secondly, on the capacity approach to moral reasoning it is very unlikely that the capacity for moral reasoning is going to be explained by a single psychological process. Just as vision can only be explained by a set of psychological and physiological processes working in concert—some of which may be unique to vision, and others of which may be shared with other capacities¹⁴—it is very likely, indeed, it would be highly unlikely if it were not the case, that moral reasoning as a capacity is explained by several

¹⁴ This sort of analysis is further supported by fMRI studies. Very few tasks rely on a single area of the brain, but rather draw on many different areas. And, assuming that there is some correspondence between brain activity and psychological processes, this supports the claim that capacities are not a single psychological process, but many psychological processes working together. fMRI studies with respect to moral reasoning reveal similarly diffuse patterns, with some brain areas being more active than others in responding to certain moral dilemmas (Greene et al., 2001; Moll et al., 2002).

psychological processes, some, none, or all of which may be unique to it. But again, in order to determine which psychological processes best explain the capacity for moral reasoning, on this approach, requires some substantive specification of what the capacity for moral reasoning is.

The way to provide such a substantive specification of the capacity for moral reasoning is by understanding what it is moral reasoning enables us to do. This is no different than how any other capacity is specified, such as reading, multiplying or seeing. While specifying what the capacity for moral reasoning is will rely on empirical data, it will also require careful consideration of what it is creatures like us are doing when we engage in it. Ideally, this sort of specification will be theory neutral, but that may not be entirely possible, and substantive theoretical disagreements will likely emerge. The aim of what follows, then, is not to provide a full specification of what the capacity for moral reasoning is, nor is it to be taken as a final analysis of what the capacity is; rather, the aim is to begin the task specifying what the capacity for moral reasoning is and offer an initial specification of it by focusing on a few of the things moral reasoning enables us to do. I shall conclude by drawing out some implications this limited specification of the capacity for moral reasoning has for the question of what role moral reasoning has in moral judgment.

3. Moral Reasoning *In Situ*

In the most general terms, moral reasoning is a capacity that enables us to think consciously and deliberately about morality; it involves thinking with moral and non-moral considerations that often, but not always, leads to a moral conclusion. This is a

substantive claim, of course, and one that requires some defense, in particular, the claim that moral reasoning, as a capacity, is conscious and deliberate. However, there is good reason for thinking of moral reasoning in this way, which is that, to all appearances, the question of what role moral reasoning plays in moral judgment is framed as it is in order to distinguish among other sorts of capacities that could play a role in moral judgment, such as the capacity for moral intuition, and the capacity for moral emotions. There is little doubt, for example, that Hume intends to distinguish moral reasoning from the moral emotions, and Prinz (2006, 2007) and Greene (2008) explicitly draw the same distinction. Haidt (2001; Haidt & Bjorklund, 2008a) introduce a further distinction among moral intuitions and moral reasoning, drawing on a body of research that indicates that many judgments are arrived at through processes that are quick, automatic, and nonconscious, while other judgments are arrived at through processes that are slower, controlled, and conscious. Processes of the first kind are often called System 1 processes, and processes of the second kind are often called System 2 processes (Evans & Over, 1996; Kahneman, 2002, 2003; Sloman, 1996; Stanovich, 1999; Stanovich & West, 2000; Tversky & Kahneman, 1983). In psychology these theories of cognitive architecture are generally referred to as Two Systems or Dual-Process views, and many moral psychologists employ the distinction between System 1 processes and System 2 processes and identify moral reasoning as a System 2-type process (Craigie, 2011; Haidt & Bjorklund, 2008; Kennett & Fine, 2009; Saunders, 2009).

Thus, when moral psychologists ask the question of what role moral reasoning has in moral judgment, the question is intended to ask very specifically what role does conscious and deliberate moral thinking have in producing moral judgments. For this

reason the capacity for moral reasoning can be appropriately specified, initially, as the capacity that enables us to think consciously and deliberately about morality; it involves thinking with moral and non-moral considerations that often, but not always, leads to a moral conclusion, because this initial specification differentiates the capacity for moral reasoning from other capacities that could play a role in moral judgment, such as the capacity for moral intuition and the capacity for moral emotions.¹⁵ However, this initial specification of the capacity for moral reasoning should not be taken as the final word, or even near final. It is subject to revision from further empirical investigation and *in situ* observation, and there will no doubt be areas of continuing substantive disagreement. Even with these limitations, however, this initial specification can help clarify what the capacity for moral reasoning is.

One further point of clarification is necessary before moving on to specify the capacity for moral reasoning in more detail. The claim that moral reasoning, as a capacity is conscious and deliberate does not imply that every and all of the psychological processes that subserve moral reasoning are conscious and deliberate, nor that every individual psychological process that subserves moral reasoning is consciously accessible. The capacity, as a whole, can have these qualities without attributing them to each and every subserving psychological process. Moral reasoning can be conscious and deliberate and still be subserved by some psychological processes that themselves are not, such as, for example, a Moral Faculty.

¹⁵ Some theorists, notably Bucciarelli et al. (2008) following Johnson-Laird (2006) hold that the capacities for moral intuition, moral emotions, and conscious moral reasoning are all properly analyzed as forms of reasoning and all properly referred to as moral reasoning. What is needed, on this view, are simply different accounts of moral reasoning. This may be quite right, and nothing claimed above undercuts this understanding of the various capacities that can contribute to the production of a moral judgment. However, it is often useful to draw verbal distinctions to correspond to substantive distinctions so as to avoid misunderstandings and the very real possibility that some disagreements may persist simply as a result of such misunderstandings.

With the initial specification of the capacity for moral reasoning in place, it is possible to attempt to specify it in greater detail, which will require careful observation of its operations *in situ*. What this means is that it is necessary to pay careful attention to what episodes of moral reasoning function to do within our overall moral psychology in order to specify what the capacity is. In brief, moral reasoning is the capacity that enables us, for example, to consciously weigh various moral considerations to come to an all-things-considered conclusion regarding what should or should not be done. It enables us to consciously apply moral principles to new or novel situations to come to a conclusion about what should be done. It enables us to consider the moral implications of complex policies, and make a decision about whether they should be endorsed. It enables us to consider our own moral judgments to determine whether they are appropriate, correct, or the moral judgments that we should have, and it enables us (in some situations) to change our minds with respect to the moral permissibility of some act or action, e.g., the permissibility of a given war, or the permissibility of eating meat. Moral reasoning enables us to consider and evaluate moral arguments, and in light of those, evaluate our own moral commitments, attitudes, and judgments.

Moral reasoning allows us to come to moral conclusions, though it is important to point out that not all of those conclusions will be moral judgments, in the sense of being first-order judgments about the permissibility of some act or action, or the goodness or badness of something or someone (e.g., war is wrong, lying is bad). Some conclusions of moral reasoning will be second-order judgments about one's first-order judgments (e.g., I was mistaken to think that war is always wrong), or hypothetical judgments (e.g., If certain conditions are met in the right sort of way, then some wars could be permissible),

or coming to a view with respect to some general principle of morality (e.g., that the principle of utility is incorrect), or what is intrinsically good, or something similarly abstract. Moral reasoning will not always result in a moral conclusion, however. This can occur when the issue or question one is reasoning about is sufficiently complicated and where one recognizes good countervailing reasons favoring different and exclusive conclusions, such as may occur when considering abstract questions or complicated and controversial questions of social policy. Moreover, any of the foregoing sorts of conclusions of moral reasoning can be held with varying degrees of confidence; one can hold a moral conclusion quite tentatively, subject to further reflection and analysis, or very confidently, letting it serve as the basis of further moral thought and reflection.

Some may object that this characterization of moral reasoning begs the question, namely, that it asserts rather than argues for the view that moral reasoning produces moral judgments, and thus begs the question against those who argue that moral reasoning cannot, or rarely, produce moral judgments. But this is to misunderstand the difference between a capacity and a psychological process, and confuses what the real set of issues and disagreement are about. The real disagreement here is not about what the capacity for moral reasoning is, but *how* it operates. What those who claim that moral reasoning cannot produce a moral judgment by itself typically mean is that the psychological processes that subserve domain-general reasoning, by themselves, cannot produce a moral judgment—that there must be some other (typically emotive)¹⁷ process that also plays a role—but none of them deny that humans typically possess the capacity for moral reasoning as described above.

¹⁷ Haidt (2008a) claims that emotions are typically involved in moral judgment (greater than 95% of the time), while Prinz (2007) claims that emotions are necessarily involved in moral judgment.

One of the benefits of characterizing moral reasoning as a capacity is that it brings to the fore the actual substantive issues and disagreements with respect to moral reasoning. This is an important point that bears emphasizing: the capacities that humans typically possess are obvious, and there is no dispute as to whether humans are genuinely capable of moral reasoning as described above; the real dispute is about how best to explain that capacity, that is, *how* it operates. Seen in this light, one of the substantive points of disagreement with respect to moral reasoning is whether there is anything special about it. If moral reasoning just is domain-general reasoning applied to moral questions, then there is not really anything special or significant about our capacity for it, and the capacity does not require any sort of special explanation; it is just like scientific reasoning, or mathematical reasoning, or any other domain of reasoning.¹⁸ If, however, moral reasoning cannot be fully explained by domain-general reasoning, but requires the addition of other psychological processes, such the emotions or a Moral Faculty (Dwyer, 2006, 2009; Hauser, 2006; Hauser et al., 2008; Mikhail, 2000, 2007, 2009),¹⁹ then there is something genuinely different and special about the causal mechanisms of moral reasoning than those of domain-general reasoning. This also highlights a second related dispute, which is whether the capacity for moral reasoning fully explains the capacity for moral judgment, or whether the capacity for moral judgment requires an explanation beyond the capacity for moral reasoning. These two points are some of the substantive issues with respect to understanding moral reasoning, and are the genuine areas of disagreement with respect to what moral reasoning is, namely, what best explains the

¹⁸ On the assumption that all such domains of reasoning are subserved by the identical set of psychological processes, which is itself a substantive claim (see, Carruthers, 2006).

¹⁹ According to these theorists, moral judgment is best explained by a Moral Faculty analogous to the Language Faculty in linguistics.

capacity for moral reasoning, and whether moral reasoning fully explains the capacity for moral judgment.²⁰

A significant upshot of organizing the dispute as one about how best to explain the capacity for moral reasoning is that it provides a method and criteria for choosing among competing definitions of moral reasoning. Most theorists typically designate a particular psychological process as *the* process of moral reasoning, and Harman and colleagues suggest that choosing among such competing definitions and the psychological processes they designate is just a matter of what one wants to call moral reasoning. But this method just allows people to assert a favored view of moral reasoning and does nothing to help settle the original question of what causal role moral reasoning has in moral judgment. But, by viewing moral reasoning as a capacity, the method for choosing among competing definitions of moral reasoning is quite clear: a proposed definition of moral reasoning is (or entails) a substantive claim with respect to what psychological processes explain how the capacity works, and thus any proposed definition can be and should be assessed against its actual explanatory burdens, that is, how well it explains the capacity for moral reasoning.

For example, Haidt's claim that moral reasoning is primarily the process of the back-and-forth trading of moral intuitions and *post hoc* reasons between two people is one possible explanation of the capacity for moral reasoning, and whether it counts as a good one depends on how well it actually explains the capacity for moral reasoning. And the test here is straightforward: how well does the proposed explanation capture the

²⁰ Specifying the capacity for moral judgment (that is, the capacity to make moral judgments) is a much broader topic, and goes well beyond the aim of the current paper, which is to begin to specify the capacity for moral reasoning. However, it may not be possible to provide a fully specification of the capacity for moral reasoning without also specifying what a moral judgment is.

phenomena of moral reasoning? Does it explain the various uses of moral reasoning? Does it explain the various sorts of moral conclusions that moral reasoning can produce? Does it explain how moral reasoning functions in our broader moral psychology? If it does not and cannot, then this definition of moral reasoning should be modified to do so or else be rejected. The same is true for any proposed definition of moral reasoning. Thus, the test for any proposed definition or account of the psychological processes of moral reasoning is an explanatory one, namely, how well it explains the capacity for moral reasoning. Moreover, taking the capacity for moral judgment as the *explanandum* not only provides a method for determining when definitions with respect to moral reasoning require empirical revisions, it also helps determine which empirical findings are relevant in revising such definitions. This is because the empirical revision of definitions with respect to moral reasoning requires determining which empirical data are relevant to that revision, but determining whether and how such data are relevant depends upon taking a certain view with respect to what the capacity for moral reasoning is. That is, determining whether and how a particular set of empirical findings sheds light on the operations of moral reasoning requires first having some understanding of what moral reasoning is, and taking moral reasoning to be a certain sort of capacity provides such an understanding.

Of course, doing all this requires an even greater specification of the capacity for moral reasoning than the one provided so far here. The aim of this paper is not to provide a full specification of what the capacity for moral reasoning is, but to point to the task that needs to be undertaken. A fuller specification of moral reasoning will need to involve, *inter alia*, a specification of its role in bringing about changes to a person's moral

attitudes, commitments, and judgments (and under what circumstances it cannot),²¹ a specification of if and when moral reasoning is rationally assessable, and a specification of its causal relationships to other capacities. This is a large undertaking, but a necessary one for understanding what moral reasoning is, as a capacity, and how best to explain it. Just as the capacity for vision cannot be adequately specified by claiming that it is the capacity to see without specifying the manner and substance of what it enables us to see—colors (within a certain spectrum), shapes, edges, movement—the capacity for moral reasoning must be specified in much greater detail in order to ground a viable research program.

It will likely not be possible to provide a full specification of the capacity for moral reasoning while completely bracketing off substantive disagreements with respect to how it operates, or by bracketing it off from other moral capacities, such as the capacity for moral intuition or the capacity for moral emotions. That is, specifying what moral reasoning is will sometimes involve making substantive claims with respect to how it operates, and how it operates with respect to other moral capacities. This is to be expected, but it does not render the task of specifying what the capacity for moral reasoning is impossible, so long as it is possible to flag such disagreements and limitations and to develop methods and criteria for choosing among them. Nor does the possibility of substantive disagreements undercut the central point that the appropriate starting point for understanding the causal role of moral reasoning is not in designating a

²¹ This caveat is important, because there seem to be some situations in which a moral commitment, attitude, or judgment is simply incorrigible to moral reasoning. A full specification of the capacity for moral reasoning needs to address the persistence of some recalcitrant moral commitments, attitudes, and judgments and whether such judgments are rational, or what meta-attitudes may be necessary to render a *person* rational as opposed to their judgments.

psychological process as the psychological process of moral reasoning, but in specifying what the capacity for moral reasoning is.

One further point of clarification is needed here. One way that substantive disagreements arise in specifying the capacity for moral reasoning involves particular metaethical claims, such as whether moral judgments necessarily motivate, or what moral terms mean, or what constitutes a genuine (as opposed to ersatz) moral judgment. There are those who argue that, as a conceptual truth, genuine moral judgments must be based in reasons in the right sort of way (Deigh, 1995; Rachels, 1993; Smith, 1994; Svavarsdottir, 1999),²² and others who argue that genuine moral judgments, as a conceptual truth, must be constituted by the emotions (Prinz, 2007).²³ Some argue that genuine moral judgments necessarily motivate (Dreier, 1990; Nagel, 1979; Prinz, 2007; Smith, 1994), while others deny this claim (Brink, 1989; Svavarsdottir, 1999). There are real and substantive issues at play here, driven primarily by the metaethical question of what moral terms mean, but they also have real implications for understanding what moral reasoning is and how it operates that cannot be settled on purely empirical grounds. For example, if genuine moral judgments necessarily motivate, it implies that moral

²² Rachels puts the point quite clearly. He writes:

If someone tells you that a certain action would be wrong...you might ask why it would be wrong and if there is no satisfactory answer you may reject that advice as unfounded. In this way, moral judgments are different from mere expressions of personal preferences...[because] moral judgments require backing by reasons, and in the absence of such reasons, they are merely arbitrary. This is a point about the logic of moral judgment...One *must* have reasons or else one is not making a moral judgment at all (Rachels, 1993, p. 483, emphasis in original).

²³ Prinz is equally clear that, on his view, moral judgments that are not partly constituted by an emotion are what he calls “verbalized moral judgments,” of which he writes:

We often talk as if verbalizations of moral judgments were moral judgments in their own right [i.e., genuine moral judgments]. I will refer to sentences such as “Pickpocketing is wrong” as a verbalized moral judgment. But this label is really shorthand. “Pickpocketing is wrong” is not a judgment; it’s a string of words. In calling it a verbalized moral judgment, I mean it is a verbal form that might be used to express a moral judgment. Verbalized moral judgments are very useful because they allow us to reason about moral values that we don’t actually hold (Prinz, 2007, p. 100).

reasoning is necessarily a practical as opposed to theoretical capacity, and that any account of how moral reasoning operates must explain this practical upshot. Moreover, the question of whether genuine moral judgments necessarily motivate cannot be settled purely on empirical grounds because the question is precisely what gets to count as a genuine moral judgment, which is not something empirical investigation into the psychological processes of moral judgment can settle in advance.²⁴

A full specification of the capacity for moral reasoning will require wading into these metaethical disputes, but it need not derail the current effort to begin specifying what the capacity amounts to. In part this is because the criteria for choosing among competing metaethical claims is the same as that for choosing among competing claims with respect to how moral reasoning operates, namely, how well the respective claims account for the observed phenomena of moral judgment, which includes how well they can explain the observed phenomena with respect to moral reasoning. There is thus significant overlap between the empirical and conceptual projects with respect to moral reasoning, and similar criteria for choosing among competing psychological claims and competing conceptual claims. Conceptual analysis has a significant role to play in understanding what the capacity for moral reasoning is, but it should not trump empirical investigation, or *vice versa*. Specifying what the capacity for moral reasoning is, thus, requires working at the intersection of moral psychology and metaethics.

²⁴ Take psychopaths as an example of the difficulties here. Some people argue that psychopaths show that moral judgments do not necessarily motivate, because psychopaths seem to make genuine moral judgments but do not act in ways consistent with them (Nichols, 2002; Roskies, 2003; Svavarsdottir, 1999). Others argue that because psychopaths do not seem to be motivated by their moral judgments, it shows that they simply fail to make genuine moral judgments at all (Cholbi, 2006; Hare, 1952; Smith, 1994). Another third way of understanding psychopaths is that they make genuine moral judgments, and are motivated by them, they are just not motivated sufficiently to override other competing motivations, such as those of self-interest (Kennett, 2006). Empirical investigation alone cannot settle these competing claims with respect to psychopaths and what they imply with respect to moral reasoning and moral judgment, because the dispute partly hangs on conceptual, metaethical analysis.

One further point of emphasis is called for at this point: There is no reason to think that empirical moral psychological research will neatly follow the same categories and distinctions that are current in metaethical research. It may turn out that the way certain questions are typically framed in the metaethical literature need to be re-thought, or rendered more consistent with empirical findings. For example, some metaethical disputes presuppose a neat dichotomy between reason and the emotions (such as the dispute between cognitivists and non-cognitivists), but the empirical research suggests that there is no such neat dichotomy *in situ* (Cushman et al., 2006; Greene et al., 2008; Greene et al., 2004; Greene et al., 2001; Haidt, 2001; Haidt et al., 1993; Hauser et al., 2007; Koenigs et al., 2007; Moll et al., 2002; Moll et al., 2001; Schnall et al., 2008; Valdesolo & DeSteno, 2006; Wheatley & Haidt, 2005; Young et al., 2010). Ultimately, answering the question of *how* the capacity for moral reasoning (and the capacity for moral judgment more broadly) operates may require different categories and distinctions, and a different understanding of the mind than the one now prevalent in metaethics.²⁵ That said, the task of specifying what moral reasoning is requires both empirical and conceptual resources, and it is this task that must be undertaken first before attempting to explain how the capacity operates.

4. Conclusions

This paper has argued that in order to answer the question of what causal role moral reasoning has in moral judgment requires a specification of what moral reasoning is. Current definitions and characterizations fall short of providing such a specification, because they attempt to specify what moral reasoning is by designating a particular

²⁵ This is the topic of a paper currently in progress.

psychological process as the process of moral reasoning, instead of attempting to provide a specification of what the capacity for moral reasoning is. However, in order to understand what psychological processes subserve moral reasoning requires first specifying what the capacity for moral reasoning is, as a capacity of a certain sort. Another way of putting the point is that it is necessary to provide an adequate characterization of the *explanandum* before attempting to provide an *explanans*. It is quite likely that the reason so much attention has been focused primarily on psychological processes is because much research into moral psychology attempts to import the categories and distinctions prevalent in metaethics directly into psychological study of moral reasoning. This is a mistake, as there is no reason to think that the categories and distinctions prevalent in metaethics provide the right way of thinking about the psychological study of moral reasoning. This paper is meant to provide a new approach for understanding what moral reasoning is, and to provide a new framework for studying and understanding what moral reasoning is.

References:

- Aristotle. (1999). *Nicomachean Ethics* (T. Irwin, Trans. 2nd ed.). Indianapolis: Hackett Publishing Company.
- Brink, D. O. (1989). *Moral Realism and the Foundations of Ethics*: Cambridge University Press.
- Bucciarelli, M., Khemlani, S., & Johnson-Laird, P. N. (2008). The Psychology of Moral Reasoning. *Judgment and Decision Making*, 3(2), 121-139.
- Carruthers, P. (2006). *The Architecture of the Mind*. New York: Oxford University Press.
- Cholbi, M. (2006). Belief Attribution and the Falsification of Motive Internalism. *Philosophical Psychology*, 19(5), 607-616.
- Cohen, S. (2004). *The Nature of Moral Reasoning*. New York: Oxford University Press.
- Craigie, J. (2011). Thinking and Feeling: Moral Deliberation in a Dual-Process Framework. *Philosophical Psychology*, 24(1), 53-71.

- Cummins, R. (2000). "How Does it Work?" versus, "What Are the Laws?": Two Conceptions of Psychological Explanation. In F. C. Keil & R. A. Wilson (Eds.), *Explanation and Cognition* (pp. 117-144). Cambridge, MA: MIT Press.
- Cushman, F., Young, L., & Hauser, M. D. (2006). The Role of Conscious Reasoning and Intuition in Moral Judgments: Testing Three Principles of Harm. *Psychological Science*, *17*(12), 1082-1089.
- Deigh, J. (1995). Empathy and Universalizability. *Ethics*, *105* (4), 743-763.
- Dreier, J. (1990). Internalism and Speaker Relativism. *Ethics*, *101*(1), 6-26.
- Dwyer, S. (2006). How Good is the Linguistic Analogy? In P. Carruthers, S. Laurence & S. Stich (Eds.), *The Innate Mind Volume 2: Culture and Cognition* (pp. 237-256). New York: Oxford University Press.
- Dwyer, S. (2009). Moral Dumbfounding and the Linguistic Analogy: Methodological Implications for the Study of Moral Judgment. *Mind & Language*, *24*(3), 274-296.
- Evans, J., & Over, D. (1996). *Rationality and Reasoning*. New York: Psychology Press.
- Fine, C. (2006). Is the Emotional Dog Wagging Its Tail, or Chasing It? *Philosophical Explorations*, *9*(1), 83-98.
- Greene, J. (2008). The Secret Joke of Kant's Soul. In W. Sinnott-Armstrong (Ed.), *Moral Psychology Volume 3: The Neuroscience of Morality: Emotion, Brain Disorders, and Development*: MIT Press.
- Greene, J., Morelli, S. A., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2008). Cognitive Load Selectively Interferes with Utilitarian Moral Judgment. *Cognition*, *107*, 1144-1154.
- Greene, J., Nystrom, L. E., Engell, A. D., Darley, J. M., & Cohen, J. D. (2004). The Neural Basis of Cognitive Conflict and Control in Moral Judgment. *Neuron*, *44*, 389-400.
- Greene, J., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI Investigation of Emotional Engagement in Moral Judgment. *Science*, *293*, 2105-2108.
- Haidt, J. (2001). The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment. *Psychological Review*, *108*(4), 814-834.
- Haidt, J., & Bjorklund, F. (2008a). Social Intuitionist Answer Six Questions. In W. Sinnott-Armstrong (Ed.), *Moral Psychology Vol. 2: The Cognitive Science of Morality: Intuition and Diversity* (pp. 181-218). Cambridge, MA: MIT Press.
- Haidt, J., & Bjorklund, F. (2008b). Social Intuitionists Reason, In Conversation. In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Volume 2: The Cognitive Science of Morality: Intuition and Diversity* (pp. 241-254). Cambridge, MA: MIT Press.
- Haidt, J., Koller, S. H., & Dias, M. G. (1993). Affect, Culture, and Morality, or Is It Wrong to Eat Your Dog? *Journal of Personality and Social Psychology*, *65*(4), 613-628.
- Hare, R.M. (1952). *The Language of Morals*. New York: Oxford University Press.
- Hare, R. M. (1981). *Moral Thinking: Its Levels, Method, and Point*. New York: Oxford University Press.
- Harman, G., Mason, K., & Sinnott-Armstrong, W. (2010). Moral Reasoning. In J. Doris & M. P. R. Group (Eds.), *The Moral Psychology Handbook*. New York: Oxford University Press.
- Hauser, M. D. (2006). *Moral Minds*. New York: Ecco.

- Hauser, M. D., Cushman, F., Young, L., Jin, R. K.-X., & Mikhail, J. (2007). A Dissociation Between Moral Judgments and Justification. *Mind & Language*, 22(1), 1-21.
- Hauser, M. D., Young, L., & Cushman, F. (2008). Reviving Rawls's Linguistic Analogy: Operative Principles and the Causal Structure of Moral Actions. In W. Sinnott-Armstrong (Ed.), *Moral Psychology Volume 2: The Cognitive Science of Morality: Intuition and Diversity* (pp. 107-144). Cambridge, MA: MIT Press.
- Herman, B. (1993). *The Practice of Moral Judgment*. Cambridge, MA: Harvard University Press.
- Horgan, T., & Timmons, M. (2007). Morphological Rationalism and the Psychology of Moral Judgment. *Ethical Theory and Moral Practice*, 10, 279-295.
- Hughes, G. (2003). *Routledge Philosophy Guidebook to Aristotle: On Ethics*. New York: Routledge Press.
- Hume, D. (1740). *A Treatise of Human Nature*. New York: Oxford University Press.
- Inbar, Y., Pizarro, D. A., Knobe, J., & Bloom, P. (2009). Disgust Sensitivity Predicts Intuitive Disapproval of Gays. *Emotion*, 2009, 435-439.
- Johnson-Laird, P. N. (2006). *How We Reason*. New York: Oxford University Press.
- Jones, K. (2003). Emotions, Weakness of Will, and the Normative Conception of Agency. In A. Hatzimoysis (Ed.), *Philosophy and the Emotions* (pp. 181-200): Cambridge University Press.
- Jones, K. (2006). Metaethics and Emotions Research: A Response to Prinz. *Philosophical Explorations*, 9(1), 45-53.
- Kahneman, D. (2002). Maps of Bounded Rationality: A Perspective on Intuitive Judgment and Choice. *Prize Lecture*.
- Kahneman, D. (2003). A Perspective on Judgment and Choice: Mapping Bounded Rationality. *American Psychologist*, 58(9), 697-720.
- Kant, I. (1785/1996). Groundwork of the Metaphysics of Morals. In M. J. Gregor & A. Wood (Eds.), *The Cambridge Edition of the Works of Immanuel Kant: Practical Philosophy*: Cambridge University Press.
- Kennett, J. (2006). Do Psychopaths Really Threaten Moral Rationalism? *Philosophical Explorations*, 9, 69-82.
- Kennett, J., & Fine, C. (2009). Will the Real Moral Judgment Please Stand Up? The Implications of the Social Intuitionist Models of Cognition for Meta-ethics and Moral Psychology. *Ethical Theory and Moral Practice*, 12(1), 77-96.
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M. D., et al. (2007). Damage to the prefrontal cortex increases utilitarian moral judgments. *Nature*, 446, 908-911.
- Kohlberg, L. (1984). *The Psychology of Moral Development: The Nature and Validity of Moral Stages*. New York: Harper & Row.
- Kohlberg, L., Levine, C., & Hewer, A. (1983). *Moral Stages: A Current Formulation and A Response to Critics*. New York: Basel.
- Korsgaard, C. M. (1986). Skepticism about Practical Reason. *The Journal of Philosophy*, 83(1), 311-344.
- Marr, D. (1982/2010). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge, MA: MIT Press.

- Mikhail, J. (2000). *Rawls' Linguistic Analogy: A Study of the "Generative Grammar" Model of Moral Theory Described by John Rawls in A Theory of Justice*. Cornell University, Ithaca, NY.
- Mikhail, J. (2007). Universal Moral Grammar: Theory, Evidence and the Future. *Trends in Cognitive Sciences*, 11(4), 143-152.
- Mikhail, J. (2009). Moral Grammar and Intuitive Jurisprudence. *Psychology of Learning and Motivation*, 50, 27-100.
- Moll, J., de Oliveira-Souza, R., Eslinger, P. J., Bramati, I. E., Mourao-Miranda, J., Andreiuolo, P. A., 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., Eslinger, P. J., & de Oliveira-Souza, R. (2001). Frontopolar and Anterior Temporal Cortex Activation in a Moral Judgment Task. *Arq Neuropsiquiatr*, 59(3-B), 657-664.
- Nagel, T. (1979). *The Possibility of Altruism*. Princeton, NJ: Princeton University Press.
- Nichols, S. (2002). How Psychopaths Threaten Moral Rationalism: Is It Irrational to Be Amoral*? *The Monist*, 85(2), 285-304.
- Nichols, S. (2004). After Objectivity: An Empirical Study of Moral Judgment. *Philosophical Psychology*, 17(1), 5-29.
- Paxton, J. M., & Greene, J. D. (2010). Moral Reasoning: Hints and Allegations. *Topics in Cognitive Science*, 1-17.
- Paxton, J. M., Ungar, L., & Greene, J. D. (2011). Reflection and Reasoning in Moral Judgment. *Cognitive Science: A Multidisciplinary Journal*, 36(1), 163-177.
- Piaget, J. (1932/1965). *The Moral Judgment of the Child* (M. Gabain, Trans.). New York: The Free Press.
- Pinker, S. (1997). *How The Mind Works*. New York: Norton.
- Prinz, J. (2006). The Emotional Basis of Moral Judgments. *Philosophical Explorations*, 9(1), 29-43.
- Prinz, J. (2007). *The Emotional Construction of Morals*. New York: Oxford University Press.
- Rachels, J. (1993). Subjectivism. In P. Singer (Ed.), *A Companion to Ethics*. Oxford: Blackwell.
- Rachels, J. (1999). *The Elements of Moral Philosophy; Third Edition* (3rd ed.). New York: McGraw-Hill.
- Roskies, A. (2003). Are Ethical Judgments Intrinsically Motivational? Lessons from "Acquired Sociopathy". *Philosophical Psychology*, 16(1), 51 - 66.
- Sacks, O. (1998). *The Man Who Mistook His Wife for a Hat: And Other Clinical Tales*. New York: Touchstone.
- Saunders, L. F. (2009). Reason in the Moral Life: A Dual Process Account of Moral Justification. In J. Evans & K. Frankish (Eds.), *In Two Minds: Dual-Processes and Beyond* (pp. 335-354): Oxford University Press.
- Schnall, S., Haidt, J., Clore, G. L., & Jordan, A. H. (2008). Disgust as Embodied Moral Judgment. *Personality and Social Psychology Bulletin*, 34(8), 1096.
- Sinnott-Armstrong, W. (2008). Framing Moral Intuitions. In W. Sinnott-Armstrong (Ed.), *Moral Psychology, Vol. 2: The Cognitive Science of Morality: Intuition and Diversity*. Cambridge, MA: MIT Press.

- Sloman, S. A. (1996). The Empirical Case for Two Systems of Reasoning. *Psychological Bulletin*, 119(1), 3- 22.
- Smith, M. (1994). *The Moral Problem*. Malden, MA: Blackwell.
- Stanovich, K. E. (1999). *Who Is Rational? Studies of Individual Differences in Reasoning*. New York: Psychology Press.
- Stanovich, K. E., & West, R. F. (2000). Individual Differences in Reasoning: Implications for the Rationality Debate? *Behavioral and Brain Sciences*, 23, 645-726.
- Svavarsdottir, S. (1999). Moral Cognitivism and Motivation. *Philosophical Review*, 108 (2), 161-219.
- Toulmin, S. (1968). *Reason in Ethics*. New York: Cambridge University Press.
- Tversky, A., & Kahneman, D. (1983). Extensional vs. Intuition Reasoning: The Conjunction Fallacy in Probability Judgment. *Psychology Review*, 90, 293-315.
- Valdesolo, P., & DeSteno, D. (2006). Manipulations of Emotional Context Shape Moral Judgment. *Psychological Science*, 17(6), 476-477.
- Vaughn, L. (2008). *Doing Ethics: Moral Reasoning and Contemporary Issues*. New York: W.W. Norton.
- Wheatley, T., & Haidt, J. (2005). Hypnotic Disgust Makes Moral Judgments More Severe. *Psychological Science*, 16(10), 780-784.
- Young, L., Camprodon, J. A., Hauser, M. D., Pascual-Leone, A., & Saxe, R. (2010). Disruption of the Right Tempoparietal Junction with Transcranial Magnetic Simulation Reduces the Role of Beliefs in Moral Judgments. *Proceedings of National Academy of Sciences*, 107(15), 6753-6758.