THOMSON'S TURN, DUAL PROCESS THEORIES OF MORAL JUDGMENT, AND THE EPISTEMIC STATUS OF ETHICAL INTUITIONS

By all accounts, The Trolley Problem has proven intractable. After over 25 years, no widely held, satisfactory solution to the problem has been found. But why should philosophers find this at all surprising? Most philosophical problems behave in this way. Such obstinacy is just what we have come to expect of philosophical problems. In fact, some might even say that The Trolley Problem's intractability is what makes it a genuine philosophical problem. Judy Thomson, however, insists in an unpublished paper that we shouldn't take The Trolley Problem's intractability lightly; it should cause suspicion of our initial intuitions about the problem. Thomson goes on to argue that our initial intuitions were mistaken, offering an explanation of her own for why so many were led astray.

In many ways, Thomson's paper marks a significant turn in her philosophical views. But more than simply documenting Thomson's own progress of thought, her paper provides an illuminative case study of the epistemic status of ethical intuitions and, in particular, their vulnerability to error due to the interference of emotions. In this paper, I explore what import Thomson's paper has for the epistemic status of intuitions, drawing heavily on recent work in moral psychology on dual process theories of moral judgment. Along the way, I argue that Thomson's argument is inconclusive and that her explanation is only a partial explanation of our intuitions. I begin with an exposition and appraisal of Thomson's paper. This is followed by an overview of dual process theories of moral judgment, concluding with a consideration of how such theories might bear on the epistemic status of ethical intuitions.

I. THOMSON'S TURN

Thomson begins her paper with a summary of The Trolley Problem. At the root of the problem is the Letting Five Die Vs. Killing One Principle, according to which a person must let five die if saving them requires killing one.[1] This principle, she remarks, is based on the "intuitively very plausible" idea that our negative duties (what we owe people in the form of non-interference) are weightier than our positive duties (what we owe people in the form of aid).[2] By the locution "intuitively very plausible," I take Thomson to mean that she has the intuition that this idea, expressed in the Letting Five Die Vs. Killing One Principle, is plausible.[3] Now consider the following two scenarios:

Bystander's Two Options:

> The brakes of a trolley have failed, and it's careening down the track toward five workmen. A bystander happens to be standing by the track, next to a switch that can be used to turn the trolley off the straight track, on which the five men are working, onto a spur of track to the right on which only one man is working. The bystander has only two options: (i) he can do nothing, letting five die, or (ii) he can throw the switch to the right, killing one.[4]

Fat Man:

> The brakes of a trolley have failed, and it's careening down the track toward five workmen. This time, however, there is no alternate spur of track onto which to turn the trolley. You happen to be standing on a footbridge over the track beside a fat man. You have only two options: (i) you can do nothing, letting five die, or (ii) you can push the fat man off the footbridge down onto the track, thereby killing him, but also, since he's very heavy, stopping the trolley and saving the five.[5]

---

[1] Thomson leaves it to the reader to insert an implicit *ceteris paribus* clause to the Letting Five Die Vs. Killing One Principle.

[2] Judith Jarvis Thomson, "Turning the Trolley," (Unpublished, August 2008), 2.

[3] Not much turns this reading of Thomson; it simply helps formulate The Trolley Problem in an interesting way.

[4] Ibid., 3.

[5] Ibid., 4.

When presented with these scenarios, most people have the intuition that choosing option (ii) in Bystander's Two Options is permissible and the intuition that choosing option (ii) in Fat Man is impermissible. But what explains these divergent intuitions? According to the Letting Five Die Vs. Killing One Principle, option (ii) is impermissible in both scenarios. The problem, then, is how to explain away the conflict of two intuitions: the intuition that negative duties are weightier than positive duties, expressed in the Letting Five Die Vs. Killing One Principle, and the intuition that option (ii) in Bystander's Two Options is permissible.

In a startling revision of her former views, Thomson presently maintains that our intuition about Bystander's Two Options is seriously mistaken.[6] We should have never said that it was permissible for the bystander to throw the switch, killing one. Once we recognize this, the conflict of intuitions vanishes; The Trolley Problem becomes a non-problem. To motivate this critical move, Thomson puts forward a new scenario she calls "Bystander's Three Options."

Bystander's Three Options:

> The brakes of a trolley have failed, and it's careening down the track toward five workmen. A bystander happens to be standing by the track, next to a switch that can be thrown in *two* ways. If he throws it to the right, then the trolley will turn onto the spur of track to the right, thereby killing one workman. If he throws it to the left, then the trolley will turn onto the spur of track to the left. The bystander himself stands on that left-hand spur of track, and will himself be killed if the trolley turns onto it. The bystander has only three options: (i) he can do nothing, letting five die,

[6] Both this section and the next are divided into two parts: the first part addresses the normative question concerning whether choosing option (ii) in Bystander's Two Options is permissible and the second part addresses the descriptive question concerning what explains our mistaken intuitions.

or (ii) he can throw the switch to the right, killing one, or (iii) he can throw the switch to the left, killing himself.[7]

Thomson suspects that most people (besides calling for a massive recall of trolley brakes) will not choose option (iii), electing for option (ii) instead. Option (ii), though, hardly seems permissible. If the bystander is unwilling to kill himself, what morally sufficient reason could he have for presuming that the one workman on the right-hand spur is willing to be killed? As Thomson notes, "What the bystander does if he turns the trolley onto the one workman is to make the one workman pay the cost of his good deed because he doesn't feel like paying it himself."[8] So option (ii) is off limits. And since option (iii) is not obligatory, Thomson concludes that it is permissible for the bystander to choose option (i).[9]

With this in mind, Thomson returns to Bystander's Two Options. The bystander now mentally runs through Bystander's Three Options before choosing between his two options. He asks himself, "If I were faced with Bystander's Three Options, what would I choose?" If the bystander would not choose option (iii) were he faced with Bystander's Three Options, then Thomson insists that he's not entitled to choose option (ii) in Bystander's Two Options. If successful, this subjunctive exercise elicits the intuition that option (ii) in Bystander's Two Options is impermissible. The Trolley Problem turns out not to be a problem after all.

But if it's not a problem after all, why were so many of us led astray? Why did we have the intuition that option (ii) is impermissible in Fat Man but not in Bystander's Two

[7] Ibid., 6.

[8] Ibid., 7.

[9] Of course, Thomson acknowledges that the rare bystander might choose option (iii), but she argues that he too has no morally sufficient reason to presume that the one workman would perform such an altruistic action. As a result, the altruistic bystander must assume that the workman is unwilling to be killed.

Options? Thomson needs an explanation, an error theory, for our initial intuition in Bystander's Two Options. Her explanation is straightforward enough: "the more drastic the means, the more strikingly abhorrent the agent's proceeding."[10] In other words, the more extreme the agent's means are, the more likely it is that we'll have the intuition that his actions are impermissible. Thomson suggests that this is due to the fact that more drastic means make the infringement of a negative duty to the one more obvious.[11] The agent's drastic means in Fat Man (namely, pushing the fat man off the footbridge) make the infringement of a negative duty toward him more obvious, thereby eliciting the intuition that option (ii) is impermissible. Relative to Fat Man, the agent's means in Bystander's Two Options are significantly less drastic—the bystander merely throws the switch. As a result, we overlook the infringement of a negative duty to the one, and the intuition that option (ii) is impermissible fails to be elicited. Thomson then accounts for our initial intuition that option (ii) is permissible by simply asserting that we were "overly impressed by the fact that if [the bystander] proceeds, he will bring about that more live by merely turning a trolley."[12]

II. Derailing Thomson's Project

The force of Thomson's argument, I think, chiefly lies in its consideration of whether the persons involved in the various scenarios consent to being killed. We're supposed to put ourselves in the shoes of the one workman and ask ourselves whether we would be willing to be killed in order to save the five. Regardless of whether we would consent to being

[10] Ibid., 16. By "means," I understand Thomson in this context to mean whatever is required for getting the trolley to threaten the one. See Judith Jarvis Thomson, "The Trolley Problem," *Yale Law Journal* 94 (1985), 1410.

[11] Thomson, "Turning the Trolley," 16.

[12] Ibid.

killed, Thomson contends that it's impermissible for us to assume that the one workman *would* consent. Likewise, assumptions about whether the one consents could be influencing our intuitions in Fat Man. The fact that the agent in Fat Man must push the fat man off the footbridge at least suggests that the fat man does not consent to being killed. Of course, the fat man could have invited you, the agent, to push him off the footbridge, but that seems rather unlikely. It might be the case, then, that most people simply assume that the fat man does not consent to being pushed, which in turn pumps the intuition that option (ii) in Fat Man is impermissible. In short, the introduction of consent complicates things.

Still, it's not clear that its introduction renders option (ii) in Bystander's Two Options impermissible. At least in Fat Man, my intuition is that option (ii) is impermissible whether or not the fat man consents to being pushed off the footbridge and killed. The impermissibility of option (ii) is in some sense independent of the fat man's consent. Why couldn't the permissibility of option (ii) in Bystander Two Options be independent of the one workman's consent in a similar manner? It's not implausible to think that our intuition about Bystander's Two Options is that option (ii) is permissible, regardless of whether the one workman consents to being killed. Even if he would not consent, presumably the five workmen on the straight track would likewise not consent to being allowed to die. Because consent muddles the scenarios, empirical studies would be helpful in clarifying what impact the explicit introduction of consent has on our intuitions. Pending such studies, Thomson's argument that option (ii) is impermissible in Bystander's Two Options appears inconclusive.

Furthermore, Thomson's explanation of our initial intuition in Bystander's Two

Options has its own shortcoming. Consider the following variations on Bystander's Two Options and Fat Man.

Looped Bystander's Two Options:

> The brakes of a trolley have failed, and it's careening down the track toward five workmen. A bystander happens to be standing by the track, next to a switch that can be used to turn the trolley temporarily onto a side track that rejoins the straight track at a point before the five workmen. On the side track there is large weight and one workman standing in front of the weight. The weight is heavy enough to stop the trolley, but in order to hit the weight the trolley must first hit the workman, killing him. The bystander has only two options: (i) he can do nothing, letting five die, or (ii) he can throw the switch, killing one.

Looped Fat Man:

> The brakes of a trolley have failed, and it's careening down the track toward five workmen. You happen to be standing by the track, next to a switch that can be used to turn the trolley temporarily onto a side track that rejoins the straight track at a point before the five workmen. Only a fat man stands on the side track, but he is heavy enough to stop the trolley. You have only two options: (i) you can do nothing, letting five die, or (ii) you can throw the switch, killing the fat man.

In both of these scenarios, the agent's means of threatening the one is the same—throwing the switch. Thomson claims that the more drastic the means, the more obvious it is that a negative duty to the one has been infringed. If Thomson is right, we should expect our intuitions about the scenarios to coincide, since the agent's means are the same in both scenarios. Yet empirical studies by John Mikhail have shown that our intuitions do not coincide. According to Mikhail, 62 percent of subjects said that option (ii) in Looped Bystander's Two Options is permissible, while only 48 percent said that option (ii) in Looped Fat Man is permissible; moreover, similar web-based studies with several thousand

subjects drawn from over 120 countries have confirmed these findings.[13] This 14 percent gap indicates that our intuitions about the Looped scenarios diverge in a statistically significant way. And given that the agent's means are the same in both scenarios, Thomson's explanation cannot account for this divergence of intuitions.

Nevertheless, Thomson's explanation can account for a great deal of our intuitions about the original Bystander's Two Options and Fat Man. Unlike the Looped scenarios in which intuitions are roughly evenly divided, 85-95 percent of subjects agree that option (ii) is permissible in Bystander's Two Options, whereas the same percentage of subjects agree that option (ii) is impermissible in Fat Man.[14] This gulf of 70-90 percent is greatly diminished in Looped scenarios like those above to a gap of 14 percent. These findings strongly suggest that something like Thomson's explanation is at least partially right. More drastic means are more likely to elicit the intuition that option (ii) is impermissible.

III. Dual Process Theories of Moral Judgment

Thomson's explanation closely relates to recent work in moral psychology on dual process theories of moral judgment. For the greater part of its history, moral psychology has been dominated by single process theories of moral judgment. On such theories, moral judgments are the product of one and only one process, normally involving conscious principled reasoning. Over the last decade or so, however, many moral psychologists have come to think that moral judgment is also the product of another process—a process that involves rapid, emotional responses, typically called "intuitions" in the literature. Our moral

[13] John Mikhail, "Moral Cognition and Computational Theory," (Unpublished), 88-89.
[14] Ibid.

judgments are therefore products of both conscious principled reasoning and rapid, emotional responses. But dual process theorists are not stuck on two processes. As Daniel Gilbert explains:

> Few [moral psychologists] would come undone if their models were recast in terms of three processes, or four, or even five. Indeed, the only number they would not happily accept is one, because claims about dual processes in dry psychology are not so much claims about how many processes there *are*, but claims about how many processes there *aren't*. And the claim is this: There aren't one.[15]

Yet dual process theorists in moral psychology generally claim more than just "There aren't one." They also claim that our moral judgments are *primarily* the product of rapid, emotional responses. The more careful theorists, of course, restrict this claim to a particular class of moral judgments, namely, those concerning physically harmful behavior. And most are quick to add qualifications like, "the moral domain encompasses much more than reactions to and prohibitions against causing bodily harm," and, "Our focus on two systems that are important for judgments concerning harm is by no means presented as a complete account of moral psychology."[16] Additionally, moral psychologists claim that these emotional responses are shaped largely by morally irrelevant factors; these include more immediate social factors such as family, friends, and culture but also more distant evolutionary factors such as natural selection.[17]

---

[15] Daniel T. Gilbert, "What the Mind's Not," 4. Emphasis in original. Because the name "dual process" can be somewhat misleading, many theorists prefer to use "multi-process" or "multi-system" instead.

[16] Fiery Cushman, Liane Young, and Joshua D. Greene, "Our Multi-System Moral Psychology: Towards a Consensus View," (Forthcoming), 2-3.

[17] For a fuller discussion of this, see Jonathan Haidt, "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment," *Psychological Review* 108:4 (2001), 814-834.

One main area of research in moral psychology for dual process theorists is The Trolley Problem. When presented with Bystander's Two Options and Fat Man, the overwhelming majority of people say that option (ii) is permissible in the former scenario but not in the latter. In an effort to explain these intuitions, Joshua Greene has introduced a distinction between "personal" and "impersonal" harm scenarios: scenarios in which the agent is involved in the one's death in an "up close and personal" manner (like Fat Man) and scenarios in which the agent is involved in the one's death in a more impersonal manner (like Bystander's Two Options).[18] Greene hypothesized that personal scenarios trigger a rapid, emotional response that impersonal scenarios do not. Thus, the thought of killing someone by means of pushing them off a footbridge is "more emotionally salient" than the thought of killing someone my means of throwing a switch. The difference in our intuitions about Bystander's Two Options and Fat Man, Greene proposes, is explained by this difference in emotional response.[19] Greene's hypothesis has since been confirmed by several studies. In one study subjects presented with personal scenarios displayed increased neural activity in brain regions associated with emotional response, while the same subjects displayed relatively less activity in the relevant brain regions when presented with impersonal scenarios. At the same time, subjects presented with impersonal scenarios displayed relatively greater activity in brain regions associated with conscious principled reasoning. And as expected, the

---

[18] Greene holds that the rationale for this distinction is largely evolutionary, a fact that will become relevant below. For a more detailed explication, see Joshua D. Greene, "The Secret Joke of Kant's Soul," *Moral Psychology* (Vol. 3), ed. Walter Sinnott-Armstrong (Cambridge: MIT Press, 2008), 43.

[19] Ibid. Interestingly, Greene now admits that this difference in emotional response is only a partial explanation of our intuitions about The Trolley Problem and that it cannot account for the 14 percent divergence on Looped scenarios. My thoughts on Thomson's explanation where reached quite independently of Greene's statements. See Ibid. and also "Reply to Mikhail and Timmons," *Moral Psychology* (Vol. 3), ed. Walter Sinnott-Armstrong (Cambridge: MIT Press, 2008), 113.

reaction times of subjects when presented with personal scenarios were markedly faster than when presented with impersonal scenarios.[20]

It should be clear how Greene's hypothesis relates to Thomson's explanation of our initial intuitions about The Trolley Problem. To put his hypothesis in Thomson's words, pushing the fat man of the footbridge in Fat Man is more "strikingly abhorrent" than throwing the switch in Bystander's Two Options. Similarly, Thomson's maxim about more drastic means can be easily mapped onto Greene's distinction between personal and impersonal harm. The agent's more drastic means in Fat Man triggers an emotional response that is not triggered by the agent's relatively less drastic means in Bystander's Two Options.

Given that the emotional response driving our intuition about Fat Man is shaped by morally irrelevant factors like family and natural selection, it's rather curious that Thomson should embrace this intuition and dismiss our initial intuition about Bystander's Two Options. If one was tempted to dismiss an intuition about The Trolley Problem at all, the more natural one to dismiss seems to be our intuition about Fat Man. This is because such emotional responses are not truth-aimed; they have no special connection to ethical truths. To be sure, it's possible over time to cultivate emotional responses in accordance with what one believes to be true, but these emotional responses are only reliable insofar as one's beliefs are true. Whatever connection they have to the truth would be parasitic on the truth

[20] Ibid., 43-44. For some interesting figures related to these studies, see Joshua Greene and Jonathan Haidt, "How (and Where) Does Moral Judgment Work?" *TRENDS in Cognitive Science* 6:12 (December 2002), 517-523.

of one's beliefs or else purely coincidental.[21] This is not to say that Thomson's turn cannot be motivated—just to say that the direction of her turn is somewhat unintuitive.

IV. The Epistemic Status of Ethical Intuitions

But what does any of this mean for the epistemic status of ethical intuitions? One response to dual process theories of moral judgment, exemplified by Peter Singer, has been to treat most ethical intuitions as epistemically worthless. That is to say, dual process theories are seen to give us good reason to think that ethical intuitions are, by and large, of no evidential value. For Singer, ethical intuitions just are rapid, emotional responses. As such, they have been primarily shaped by evolution and natural selection. Their effectiveness in survival, though, in no way guarantees their accuracy. "The direction of evolution," Singer stresses, "neither follows, nor has any necessary connection with, the path of moral progress. 'More evolved' does not mean 'better.'"[22] Consequently, ethical intuitions do not count as *prima facie* evidence, and we should not seek to develop our moral theories around them.[23] Philosophers should rather seek to develop what Singer calls "more reasoned" moral theories.

---

[21] Indeed, Greene thinks such a coincidental connection tells against Thomson's deontological ethic. See "The Secret Joke of Kant's Soul," 66-72.

[22] Peter Singer, "Ethics and Intuitions," *The Journal of Ethics* 9 (2005), 342.

[23] Stephen Stich advances a similar line of argument, contending that intuitions are the product of cognitive processes deeply dependent on cultural heritage. To think these count as prima facie evidence, thinks Stich, you'd have to be an "epistemic xenophobe." See his "Reflective Equilibrium, Analytic Epistemology and the Problem of Cognitive Diversity," *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, ed. Michael R. DePaul and William Ramsey (Lanham, Maryland: Rowman & Littlefield Publishers, 1998), 95-112.

Singer also stops to consider what this means for The Trolley Problem. Following Greene, he writes at length about the evolutionary rationale for the distinction between personal and impersonal harm. "For most of our evolutionary history," he explains,

> human beings have lived in small groups, and the same is almost certainly true of our pre-human primate and social mammal ancestors. In these groups, violence could only be inflicted in an up-close and personal way—by hitting, pushing, strangling, or using a stick or stone as a club. To deal with such situations, we have developed immediate, emotionally based responses to questions involving close, personal interactions with others. The thought of pushing the stranger off the footbridge elicits these emotionally based responses. Throwing a switch that diverts a train that will hit someone bears no resemblance to anything likely to have happened in the circumstances in which we and our ancestors lived. Hence the thought of doing it does not elicit the same emotional response as pushing someone off a bridge.[24]

Because our ethical intuitions are simply evolved responses formed by morally irrelevant factors, the philosophical project of seeking to justify our diverging intuitions about Bystander's Two Options and Fat Man is fruitless. "Very probably," Singer concludes, "there is no morally relevant distinction between the cases."[25]

Of course, Singer is aware of one problem that might beset such a dismissal of ethical intuitions: it is susceptible to self-defeat arguments similar to those put against Quinean empiricism by George Bealer. Our "more reasoned" moral theories are "still based on an intuition, for example the intuition that five deaths are worse than one, or more fundamentally, the intuition that it is a bad thing if a person is killed."[26] Yet he emphasizes that these are not the kind of intuition he intends to dismiss; they are different from the intuitions to which moral psychologists like Greene refer. These might be better called

---

[24] Ibid., 347-348.
[25] Ibid., 348.
[26] Ibid., 350.

"rational intuitions," he says.[27] So Singer confesses that there might be ethical intuitions that count as *prima facie* evidence, but that the intuitions involved in The Trolley Problem are not of this kind.

Another possible response dual process theories of moral judgment, likely to be adopted by rationalists like Bealer, is to insist in the teeth of the evidence, as it were, that all ethical intuitions count as *prima facie* evidence. All dual process theories show is that ethical intuitions are particularly vulnerable to error due to the interference of emotions. It's easy to imagine a rationalist like this quipping, "We didn't have to bother with dual process theories of moral judgment to figure *that* out." Our moral theories should continue to be developed around our ethical intuitions, although we should exercise caution to avoid the undue interference of our emotional responses to various scenarios.

On this view, it's likely that The Trolley Problem remains a problem. Perhaps many of our intuitions about Fat Man can be explained away as the product of unreliable emotional responses, but there still exists a statistically significant divergence of intuitions about Looped scenarios. And it appears that this 14 percent gap cannot be accounted for by the interference of emotions. It might also be suggested that empirical studies by Greene and others have simply failed to locate genuine ethical intuitions about The Trolley Problem, so the problem could remain as intractable as ever.[28]

I suspect the right response to the dual process theories of moral judgment lies somewhere between the two responses mentioned here. Surely, many mental states that

---

[27] Ibid., 351.

[28] Bealer contends that the majority of empirical studies do this very thing. George Bealer, "The Incoherence of Empiricism," *Proceedings of the Aristotelian Society* 66 (1992), 130.

pass as ethical intuitions in philosophy are merely rapid, emotional responses. Philosophers have generally failed to properly discriminate between these sorts of mental states, which has resulted in the loose application of the term "intuition." Such carelessness could be rectified in part with a rich phenomenological description of intuitions and the emotional responses with which they are often confused in ethics.[29] For instance, intuitions are propositional attitudes; when you have the intuition that *P*, it seems to you that *P*. Emotional responses, although they may incline us to have a particular attitude toward a proposition, are not propositional attitudes themselves. In general, emotional responses are intentional states, but their intentional objects are not neatly thought of as propositions. Similarly, ethical intuitions are intellectual seemings. If emotional responses are to be construed as seemings at all, then they should definitely not be characterized as intellectual. Phenomenological descriptions like these, I think, could help allay the temptation to play fast-and-loose with ethical intuitions.

In addition, dual process theories of moral judgment should give us pause about the influence of emotions over our intuitions, especially in ethical thought experiments involving personal harm. Such thought experiments trigger strong emotional responses, which may or may not coincide with our genuine ethical intuitions. And while personal harm is certainly not the only thing that can trigger these emotional responses, it is one of the most recognizable ways in which they are triggered. It pays to be aware of these emotional responses and to be sensitive to the possibility of their leading to mistaken ethical intuitions.

---

[29] In several of his articles, Bealer provides a detailed phenomenological description of intuitions. See, for example, Ibid., 101-104. See also John Bengson, "Intuition and Perception," (Unpublished Draft, November 2008), 9-18.

V. CONCLUSION

Whether or not Thomson can successfully motive her turn, her explanation of our initial intuitions and dual process theories of moral judgment give us considerable insight into how we make judgments about scenarios like Bystander's Two Options and Fat Man. The emotional response triggered by more drastic means explains many of our intuitions about the Trolley Problem, but it cannot explain them all. Nevertheless, these findings in moral psychology at the very least underscore ethical intuition's particular vulnerability to error from the interference of emotions.

REFERENCES

Bealer, George. "The Incoherence of Empiricism." *Proceedings of the Aristotelian Society* 66 (1992): 99-138.

Bengson, John. "Intuition and Perception." Unpublished Draft. November 2008.

Cushman, Fiery, Liane Young, and Joshua D. Greene. "Our Multi-System Moral Psychology: Towards a Consensus View." Forthcoming.

Gilbert, Daniel T. "What the Mind's Not." 3-11.

Greene, Joshua D. "The Secret Joke of Kant's Soul." *Moral Psychology* (Vol. 3), ed. Walter Sinnott-Armstrong. Cambridge: MIT Press, 2008.

———. "Reply to Mikhail and Timmons." *Moral Psychology* (Vol. 3), ed. Walter Sinnott-Armstrong. Cambridge: MIT Press, 2008.

Greene, Joshua D. and Jonathan Haidt. "How (and Where) Does Moral Judgment Work?" *TRENDS in Cognitive Science* 6:12 (December 2002): 517-523.

Haidt, Jonathan. "The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment." *Psychological Review* 108:4 (2001): 814-834.

Mikhail, John. "Moral Cognition and Computational Theory." 81-91.

Singer, Peter. "Ethics and Intuitions." *The Journal of Ethics* 9 (2005): 331-352.

Stich, Stephen. "Reflective Equilibrium, Analytic Epistemology and the Problem of Cognitive Diversity," *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, ed. Michael R. DePaul and William Ramsey. Lanham, Maryland: Rowman & Littlefield Publishers, 1998.

Thomson, Judith Jarvis. "The Trolley Problem." *Yale Law Journal* 94 (1985): 1395-1415.

———. "Turning the Trolley." Unpublished, August 2008.