

The Baltic International Yearbook of
Cognition, Logic and Communication

October 2012
pages 1-14

Volume 7: *Morality and the Cognitive Sciences*
DOI: 10.4148/biyclc.v7i0.1777

EDOUARD MACHERY
University of Pittsburgh

DELINEATING THE MORAL DOMAIN

ABSTRACT: Moral psychology and naturalistic moral philosophers should strive to understand how moral and non-moral norms differ. Using Patricia Churchland's recent book *Braintrust* as an example, I describe the pitfalls of failing to take this issue seriously, and I introduce a new research program about the delineation of the moral domain.

Which norms and normative judgments belong to the moral domain, and which do not? In recent years, moral philosophers (ethicists as well as meta-ethicists) and moral psychologists have largely overlooked this question (for two exceptions, see [Gert 2011](#) and [Haidt 2012](#)). Philosophers' and psychologists' lack of attention to this question is puzzling since many important debates in moral philosophy and moral psychology depend on how the moral domain is characterized (Machery and Stich, forthcoming). In this brief article, I will show that failing to clarify the distinction between moral and non-moral norms can lead one's theorizing about moral psychology astray, and I will describe the outlines of a research program about the delineation of the moral domain. To fulfill these goals, I will use Patricia Churchland's recent book on moral psychology, *Braintrust*, as a springboard, arguing that this otherwise fine book fails to cast much light on moral psychology because it never attempts seriously to delineate the moral domain. Here is how I will proceed. In Section 1, I will briefly review *Braintrust* before discussing in Section 2 the arguments advanced by Churchland

to justify her refusal to delineate the moral domain in any detail. In Section 3, I will look at psychologist Elliot Turiel's well-known attempt to delineate the moral domain, and I will discuss the recent research casting doubts on this attempt. Finally, in Section 4, I will describe the outlines of a research program meant to delineate the moral domain.

1. *BRAINTRUST*

Braintrust is an enjoyable tour of the recent findings in psychology, systems and cognitive neuroscience, neuroeconomics, evolutionary theory, etc., about human social behavior. Patricia Churchland corals an impressive amount of information—a feat that, I suspect, no other philosopher would have been able to do—and her treatment of the scientific literature is, unsurprisingly, much more sophisticated than is common for this kind of book: She brings a critical eye to the scientific literature she is engaging with, and she honestly acknowledges the shortcomings and limits of our current scientific knowledge. This is refreshing.

Surprisingly, however, with the exception of a fairly superficial discussion of the linguistic analogy defended by Harman, Mikhail (whose important work goes unmentioned), Dwyer, and Hauser, Churchland does not engage with the exciting debates that have enlivened moral psychology for about ten years (see, e.g., [Doris & the Moral Psychology Research Group 2010](#)). For instance, there is no serious discussion of the exact role emotions play in moral judgments, about the existence and nature of virtues (about which, curiously, she has written elsewhere), about the evolution of morality, about the nature of agency, about the nature of altruism, about the nature and extent of moral variation, etc. Furthermore, she does not examine at any length the significance of the empirical work she discusses so carefully for traditional philosophical concerns in ethics and meta-ethics. As Churchland knows full well, much has been made in recent philosophy of the kind of work she is discussing—the evolution of morality is often held to have significant meta-ethical implications ([Ruse 1986](#); [Joyce 2006](#); [Street 2006](#); [Machery & Mallon 2010](#)), classic findings in social psychology are said to undermine virtue theory ([Flanagan 1993](#); [Doris 1998, 2002](#); [Harman 1999](#)), some take cross-cultural variation in norms to undermine

moral realism (Machery et al. 2005; Doris & Plakias 2008), the alleged sensitivity of moral intuitions to biases is sometimes held to raise questions about the nature of justification in ethics (Appiah 2008; Nadelhoffer & Feltz 2008; Sinnott-Armstrong 2008; Machery 2010)—and her silence is puzzling. Perhaps, this is because philosophers and other academics were not her intended audience. (And if that's not the reason, what is it then?)

In fact, *Braintrust's* subtitle notwithstanding (“What neuroscience tells us about morality”), much of the book has only a loose connection with moral cognition (see, e.g., the chapter about mindreading). Under the guise of theorizing about moral cognition, it is really about social cognition in general. It is unclear what led Churchland to write about social cognition in general rather than moral psychology more specifically. She may just find the former more interesting; or she may find the scientific study of social cognition better than moral psychology; or she may have expected her readers to be more excited by the science of social cognition in general. All this may well be true, but I think that there is another explanation for Churchland's somewhat surprising focus in a book explicitly about moral cognition: because she does not seriously attempt to delineate the moral domain by identifying what distinguishes moral and non-moral norms, her discussion is insufficiently focused on the moral domain.

2. CHURCHLAND ON THE MORAL/NON-MORAL DISTINCTION

Philosophers routinely distinguish various kinds of norms, such as moral norms, etiquette norms, coordination norms, and prudential norms. While the resulting classifications of norms are often orthogonal to one another (compare, e.g., Nichols 2004 and Bicchieri 2006), these always end up distinguishing *moral* norms from other kinds of norms. Psychologists, neuroscientists, and some evolutionary biologists often follow suit, sometimes only implicitly. For those moral psychologists who endorse the linguistic analogy, the hypothesis of a moral grammar is meant to explain the acquisition of *moral* norms and their deployment in judgment (in contrast to other kinds of norms); Turiel, Smetana, and others hold that, universally and from an early age on, the violation of *moral* norms elicits reactions that differ from the violation

of other kinds of norms (Nucci & Turiel 1978; Smetana 1981; Sousa et al. 2009, e.g.); other psychologists argue that *moral* norms have a distinctive content—e.g., they prohibit harmful actions (Gray et al. 2012). Similarly, evolutionary biologists (e.g., Waddington, Alexander) and psychologists (e.g., Haidt) who speculate about the evolution of morality are often concerned with the evolution of the capacity to learn, apply, and comply with a particular kind of norm—viz. *moral* norms—and not with norms in general.

The distinction between moral and other norms raises at least four distinct clusters of empirical questions. First, is it really the case that lay people distinguish moral norms from other kinds of norms and that they treat the former differently? Second, if lay people in our modern, Western societies do draw this distinction, is this distinction universal and ancient? Is it a feature of the human mind that it distinguishes a subset of norms—moral norms—and treats them differently? Third, how is this distinction drawn? What distinguishes moral norms from other kinds of norms? Finally, where does this distinction come from? Is it an adaptation? If so, what is its function and what is its phylogeny? Or, rather, is it a cultural invention?

One may wonder whether focusing on the way people draw the distinction between moral and non-moral norms is the best way to proceed. Wouldn't it be better to focus on how moral and non-moral norms *really* differ, which may, or may not, be identical to the distinction drawn by lay people? I will not answer this question fully here, but I will say this: Which approach is optimal depends on the context in which this distinction is relevant. When one is interested in moral psychology, understanding how lay people themselves distinguish moral from non-moral norms is crucial because how motivated they are to comply with norms, how conservative they are in changing norms, how willing they are to punish norm violators, etc., depend on *their* understanding of the distinction.

Churchland briefly engages with the empirical issues singled out above. In the introduction of *Braintrust* (p. 9), she characterizes “ethics or morality” as “a four-dimensional scheme for social behavior that is shaped by interlocking brain processes: (1) *caring* (rooted in attachment to kin and kith and care for their well-being), (2) *recognition of others' psychological states* (...), (3) *problem solving in a social context*

(e.g., how we should distribute scarce goods [...]), and (4) *learning social practices*.” This is a vague characterization of morality, which does not identify what is distinctive about morality. Much of human social behavior, including playing football and going to a movie theater, involves mindreading, problem solving in a social context, and learning social practices. Furthermore, many actions that belong to the moral domain have nothing to do with caring, and many actions that involve caring do not belong to the moral domain (e.g., caring for one’s plants). From the get go, this vague and non-diagnostic characterization of morality misdirects the book toward an examination of human social behavior in general.

To counter this kind of objection, Churchland notes (p. 9-10) that she “generally [shies] away from trying to cobble together a precise definition of ‘moral,’ preferring to acknowledge that there is a spectrum of social behaviors, some of which involve matters of great seriousness, and tend to be called moral, such as enslaving captured prisoners or neglecting children, while others involve matters of more minor moment, such as conceptions for behavior at a wedding.” This is a misguided response. How moral norms differ from other norms is not to be answered by stipulation. Rather, empirical evidence is needed to determine whether and how lay people single out moral norms, and whether those have a specific place in people’s minds (in terms of motivation, etc.), and seriousness is unlikely to be a very enlightening characterization of the distinction between moral and non-moral norms.

Finally, to justify her reluctance to examine what, if anything, distinguishes moral norms, Churchland insists (p. 10) that “the boundaries of the concept ‘moral,’ like the boundaries of ‘house’ and ‘vegetable,’ are fuzzy.” This is, however, a confused line of reasoning. That a concept has a vague extension is no obstacle to characterizing its content (see, e.g., Labov 1973 on the concepts expressed by “mug” and “bowl”).

3. TURIEL’S MORAL/CONVENTIONAL DISTINCTION

So, how do moral and non-moral norms differ from a psychological point of view? Some, including perhaps Churchland, may think that there is already a substantial body of evidence bearing on this ques-

tion. In particular, philosophers and psychologists often treat Turiel and colleagues’ work as having shown that in many different cultures even young (2 and 1/2 year-old) children distinguish between moral and conventional norms (e.g. Nucci & Turiel 1978; Smetana 1981). To investigate this question, Turiel and colleagues used a task, commonly called “the moral/conventional task,” in which participants are presented with a short scenario describing an action—for instance, a child leaving the classroom without asking permission or a child hitting another child in a classroom. Participants are then typically asked four questions (adapted from Smetana 1981, except for the fourth question not asked in this study):

1. Permissibility question: “How bad is [the event]?”
2. Authority-dependence question: “Would [the depicted event] be OK if there was no rule about it here?”
3. Universality question: “Would [the event] be OK at home or in another school?”
4. Justification question: “Why is [the event] wrong?”

Turiel and colleagues have repeatedly found that children distinguish two kinds of actions that are judged wrong. Those actions (e.g., hitting another child) that are judged worse are also judged to be authority-independent (they would still be wrong, even if the relevant authority allowed people to act this way) and to be wrong everywhere; finally, people justify their opinion that these actions are wrong by appealing to considerations of harm, justice, or rights. Turiel and colleagues call the norms prohibiting these actions “moral norms.” By contrast, those actions (e.g., leaving the classroom without asking permission) that are judged to be less wrong are judged to be authority-dependent (they would not be wrong if the relevant authority allowed people to act this way) and to be wrong only locally; finally, people justify their opinion that these actions are wrong by appealing to authority and convention. Turiel and colleagues call the norms prohibiting these actions “conventional norms.” This substantial body of findings seems to suggest that the distinction between moral and non-moral norms is a universal and, plausibly, ancient feature of the human mind, and that

its existence calls for an evolutionary explanation (Nichols 2004; Joyce 2006).

However, while widely agreed upon, this conclusion should be viewed with skepticism. Recent findings suggest that the neat separation of wrong actions into two kinds is an artifact of the restricted class of actions used in the moral/conventional task. When a larger class of actions is used, the different features that are meant to characterize moral and non-moral norms come apart. In a famous study, Haidt et al. (1993; see also Shweder et al. 1987) presented American and Brazilian participants with vignettes describing disgusting actions. For instance, participants read the following vignette (617):

A family's dog was killed by a car in front of their house. They had heard that dog meat was delicious so they cut up the dog's body and cooked it and ate it for dinner.

They were then asked questions similar (though not identical) to those used by Turiel and colleagues. Haidt and colleagues found that low socio-economic status (but not high socio-economic status) participants in Brazil and in the USA judged that disgusting actions were universally wrong, although they did not involve harm or lack of fairness. Thus, the third feature meant to distinguish moral norms from conventional norms (justification in terms of harm or fairness considerations) can dissociate from the two other features (universality and authority-independence).

More recently, Kelly et al. (2007) presented participants with scenarios describing a harmful action committed either now or in the past. For instance, participants were presented with the following two scenarios (123–124):

Scenario 1

Three hundred years ago, whipping was a common practice in most navies and on cargo ships. There were no laws against it, and almost everyone thought that whipping was an appropriate way to discipline sailors who disobeyed orders or were drunk on duty. Mr. Williams was an officer on a cargo ship 300 years ago. One night, while at sea, he found a sailor drunk at a time when the sailor should have been on watch. After the sailor sobered up, Williams

punished the sailor by giving him 5 lashes with a whip.

Scenario 2

Mr. Adams is an officer on a large modern American cargo ship in 2004. One night, while at sea, he finds a sailor drunk at a time when the sailor should have been monitoring the radar screen. After the sailor sobers up, Adams punishes the sailor by giving him 5 lashes with a whip.

Kelly and colleagues found that people judged that some harmful actions (e.g., whipping) were worse now than in the past and that their badness depended on whether the relevant authority authorized them. Again, it would seem that authority-independence, universality, and harm or fairness considerations dissociate when a broader range of examples is used in the moral/conventional task.

Kelly and colleagues' work has not gone unchallenged. Particularly, Paulo Sousa has criticized their work in two noteworthy articles (Sousa 2009; Sousa et al. 2009; for discussion, see also Stich et al. 2009). Sousa and colleagues (2009) followed closely Kelly and colleagues' studies, but also asked participants to explain their response. Their results were largely similar to Kelly and colleagues' (see Table 2 p. 87): A small, but not negligible minority of participants (around 15–20%) treated some harmful wrongdoings such as whipping and the use of torture in the interrogation of terrorists as being local or as being authority-independent, while about half of the participants judged that harming soldiers in order to train them was fine if it had been authorized. However, Sousa and colleagues do not view their results as providing support for Kelly and colleagues' criticism of the moral/conventional task on the grounds that the explanations provided by participants appeal to the victim's consent or to the possible utility of the norm violation in addition to, or instead of, the conventional nature of norms when they judge that a harmful violation is not universal or is authority-dependent.

We should resist Sousa and colleagues' suggestion that any *prima facie* disconfirming finding can be ignored when participants appeal to utility or informed consent to justify their judgment about the non-universality and authority-dependence of a norm against a kind of harmful action, because this suggestion insulates Turiel's hypothesis from disconfirmation: After all, people are likely to justify their judg-

ment on such grounds. More important, even when it is so justified, a judgment that some harmful violation is permissible if authorized by some authority or is not universally applicable is a counterexample to, and so evidence against, Turiel's hypothesis. Consider the case of torture. It is harmful, and it violates the victim's rights. As shown by Sousa and colleagues' data themselves, people also judge that torture is wrong. So, Turiel predicts that people would find it universally wrong and authority-independent. However, a small, but not negligible minority of participants found it ok, if authorized, a finding that is not very surprising in light of many American conservatives' enthusiastic embrace of torture. That people often justify their answer by appealing to the possible utility of torture does not change the fact that their answer is at odds with Turiel's prediction. Surely, Turiel's view is that actions that are judged wrong because they are harmful (or because they are harmful and violate someone's rights, as Sousa and colleagues propose) are wrong everywhere whether or not they are useful!

Naturally, Kelly and colleagues' findings do *not* show that lay people do not distinguish moral norms from other kinds of norms. However, they suggest that the findings obtained by Turiel and colleagues with the moral/conventional task do not provide definite evidence that across cultures and from a very early age on lay people distinguish moral norms from other kinds of norms. Another approach is needed.

4. A RESEARCH PROGRAM

4.1. *The Moral Questionnaire*

Recently, Steve Stich, Joe Henrich, and I have started examining empirically whether lay people distinguish moral from non-moral norms, whether they draw this distinction in a cross-culturally consistent manner, and, if they do, on what basis this distinction is drawn. While it is too early to report our results, I would like to describe the kind of experimental tool we have been using so far, which we call "the Moral Questionnaire."

Participants are presented with a list of 20 statements about the appropriateness of specific actions in various cultures, including the two following ones:

Many Americans think that fair-skinned people should use

sunscreen when they are on the beach.

Many members of the Kuna tribe in Panama think that thieves should be punished by having their hands cut off and tied around their neck.

After each statement, participants are asked two distinct questions:

1. Please indicate whether or not you agree with their judgment, using the 7-point scale below.
2. Now consider the judgment you just expressed. Do you consider it a moral judgment or some other kind of judgment, using the 7-point scale below?

The first question asks whether participants endorse the normative judgment about the action described in the statement, while the second question asks whether people think that their own normative judgment bears on a moral matter—viz. whether the action described by the statement belongs to the moral domain. For present purposes, only participants' answers to the second question after each of the 20 statements we used are relevant.

We are particularly interested in how these answers correlate with one another. A high positive correlation between two answers means that people tend to treat the two relevant actions similarly—viz. if someone thinks that the first action belongs to the moral domain, she thinks that the other one belongs to the moral domain too, and if someone thinks that the first action does not belong to the moral domain, she thinks the same thing about the other one. A low correlation between two answers means that the way she thinks about one action is unrelated—it does not give any information—about the way she thinks about the other action. A cluster analysis of participants' answers provides a convenient way to analyze these correlations since, roughly, it groups two actions into the same cluster if people make similar judgments about them, and it groups these resulting clusters into broader clusters based on the same principle.

What kinds of findings would show that people distinguish moral from non-moral norms? If people draw such a distinction, two main clusters of norms would clearly emerge. Plausibly too, these two clusters would divide into sub-clusters, corresponding to different kinds

of moral norms (perhaps norms against actions harming innocent individuals and against actions involving infringing upon other people's freedom) and different kinds of non-moral norms (perhaps, etiquette norms, prudential norms, etc.).

Naturally, it is important to study whether lay people converge in their classifications of norms across cultures, religions, and other demographic groups, and we should be cautious in generalizing findings obtained with Western participants to non-Westerners. Recent cross-cultural research in cultural psychology has shown that culture influences how people process information, make decisions, etc. (Henrich et al. 2010), while experimental philosophers have similarly shown that some intuitions vary across cultures (Machery et al. 2004, 2009 for cross-cultural studies of semantic intuitions). More important, a small body of evidence raises the tantalizing possibility that the distinction between moral and non-moral norms is not a psychological universal. Of particular interest is the fact that, while deontic modals—that is, words translating “ought”—and translations of the normative predicates “good” and “bad” (but not “right” and “wrong”) are found in every language (Wierzbicka 2001, p. 167–169; Wierzbicka 2007), some languages do not have a translation of “moral” and thus do not lexicalize the distinction between morally good and morally bad (Wierzbicka 2007, p. 68). In particular, this word seems to have no translation in Hindu (Henrich, personal communication).

So, what kind of findings with the Moral Questionnaire would suggest that people in some culture fail to distinguish moral from non-moral norms. First, it could be that, when participants' answers to the Moral Questionnaire are analyzed, more than two clusters emerge. That would suggest that people do not distinguish moral from non-moral norms, but classify the prompts they read into three kinds. Second, if people in some culture distinguish two main kinds of norms, these kinds of norms could be so different that it makes little sense to identify one of them with moral norms.

4.2. *The Role of the Moral vs. Non-Moral Distinction in Cognition*

If some population (e.g., Westerners) distinguishes moral from non-moral norms, this distinction may be of psychological importance: Perhaps, people may be more motivated to abide by moral norms than by

other norms; they may be more upset when the latter are violated; and so on.

Some unpublished research by Chiara Lisciandra and colleagues illustrates beautifully how this hypothesis can be tested. Lisciandra and colleagues replicated Ash's well-known conformity experiments with the following twist. Instead of focusing on perceptual judgments, they examined whether people would be disposed to conform their normative assessment to that of the group. In substance, Dutch participants were first asked to assess how wrong various actions were. Later, having been told that their original answers had been lost, they were brought in a lab, where they were asked to assess the same actions again. While the first assessment was done privately, the second assessment was done in public, and before answering participants heard first other participants' (really confederates') answers. The question of interest was whether people's second assessment would differ from their first, showing that their answer had been influenced by others' answers.

Lisciandra and colleagues compared different kinds of wrong actions. Lisciandra and colleagues classified some actions as being morally wrong, while others were wrong, but not morally wrong. Rape illustrates the first kind of action, and not reciprocating illustrates the second kind of action. Participants were also presented with scenarios describing disgusting actions.

Lisciandra and colleagues found that, while people are disposed to conform their assessment to that of other people for all types of wrong actions, they are less likely to do so for the actions that fall in the moral domain (as the researchers define it) than for the other two kinds of actions. So, norms that were classified by the researchers as belonging to the moral domain seem to have a distinctive effect on people's cognition.

The proposal, then, is to apply this experimental paradigm, or similar paradigms, to the kinds of norms identified by the Moral Questionnaire described above instead of relying on the distinction between kinds of norms drawn by the experimenters. This would allow us to determine whether people treat differently the kinds of norms they themselves distinguish.

5. CONCLUSION

Clearly, there is much we do not know about the psychological nature of the distinction between moral and non-moral norms as well as about its role in our cognitive lives. Previous empirical research on the topic by Turiel and colleagues is not satisfying, and the superficial arguments marshaled by Churchland fail. But it would not be wise to just overlook this distinction when theorizing about moral psychology since, as illustrated by Churchland's *Braintrust*, theorizing may then lose its focus. Hopefully, the Moral Questionnaire will be useful in casting some light on the way lay people distinguish moral from non-moral norms.

References

- Appiah, K. A. 2008. *Experiments in Ethics*. Cambridge, MA: Harvard University Press.
- Bicchieri, C. 2006. *The grammar of society*. Cambridge: Cambridge University Press.
- Churchland, P. S. 2011. *Braintrust*. Princeton: Princeton University Press.
- Doris, J. 1998. 'Persons, situations, and virtue ethics'. *Nous* 32: 504–530.
- Doris, J. & the Moral Psychology Research Group (eds.). 2010. *The Moral Psychology Handbook*. Oxford: Oxford University Press.
- Doris, J. M. 2002. *Lack of Character*. Cambridge: Cambridge University Press.
- Doris, J. M. & Plakias, A. 2008. 'How to argue about disagreement: Evaluative diversity and moral realism'. In W. Sinnott-Armstrong (ed.) 'Moral Psychology, Volume 2; The Cognitive Science of Morality: Intuition and Diversity', 303–332. Cambridge, MA: MIT Press.
- Flanagan, O. J. 1993. *Varieties of Moral Personality: Ethics and Psychological Realism*. Cambridge, MA: Harvard University Press.
- Gert, B. 2011. 'The Definition of morality'. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*, URL: <http://plato.stanford.edu/entries/morality-definition/>.
- Gray, K., Young, L. & Waytz, A. 2012. 'Mind perception is the essence of morality'. *Psychological Inquiry* 23: 101–124.
- Haidt, J. 2012. *The Righteous Mind: Why Good People are Divided by Politics and Religion*. New York: Pantheon.
- Haidt, J., Koller, S. & Dias, M. 1993. 'Affect, culture, and morality, or is it wrong to eat your dog?' *Journal of Personality and Social Psychology* 65: 613–628.
- Harman, G. 1999. 'Moral philosophy meets social psychology: Virtue ethics and the fundamental attribution error'. *Proceedings of the Aristotelian Society* 99: 315–331.
- Henrich, J., Heine, S. J. & Norenzayan, A. 2010. 'The weirdest people in the world?' *Behavioral and Brain Sciences* 33: 61–83.
- Joyce, R. 2006. *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Kelly, D., Stich, S. P., Haley, K. J., Eng, S. J. & Fessler, D. M. T. 2007. 'Harm, affect, and the moral / conventional distinction'. *Mind & Language* 22: 117–131.
- Labov, W. 1973. 'The boundaries of words and their meanings'. In C.-J. Bailey & R. Shuy

- (eds.) 'New Ways of Analyzing Variation in English', 340–373. Georgetown: Georgetown University Press.
- Lisciandra, C., Postma, M. & Colombo, M. ms. 'Group conditioning of normative judgment'.
- Machery, E. 2010. 'The bleak implications of moral psychology'. *Neuroethics* 3: 223–231.
- Machery, E. & Mallon, R. 2010. 'Evolution of morality'. In J. M. Doris & the Moral Psychology Research Group (eds.) 'The Moral Psychology Handbook', 3–46. Oxford: Oxford University Press.
- Machery, E. & Stich, S. P. forthcoming. 'You can't have it both ways'. *Behavioral and Brain Sciences*.
- Machery, E., Mallon, R., Nichols, S. & Stich, S. P. 2004. 'Semantics, cross-cultural style'. *Cognition* 92: B1–B12.
- Machery, E., Kelly, D. & Stich, S. P. 2005. 'Moral realism and cross-cultural normative diversity'. *Behavioral and Brain Sciences* 28: 830.
- Machery, E., Olivola, C. Y. & de Blanc, M. 2009. 'Linguistic and metalinguistic intuitions in the philosophy of language'. *Analysis* 69: 689–694.
- Nadelhoffer, T. & Feltz, A. 2008. 'The actor-observer bias and moral intuitions: Adding fuel to Sinnott-Armstrong's Fire'. *Neuroethics* 1: 133–144.
- Nichols, S. 2004. *Sentimental Rules. On the Natural Foundations of Moral Judgment*. New York: Oxford University Press.
- Nucci, L. & Turiel, E. 1978. 'Social interactions and the development of social concepts in preschool children'. *Child Development* 49: 400–407.
- Ruse, M. 1986. *Taking Darwin Seriously*. Oxford: Basil Blackwell.
- Shweder, R. A., Mahapatra, M. & Miller, J. 1987. 'Culture and moral development'. In J. Kagan & S. Lamb (eds.) 'The Emergence of Morality in Young Children', 1–82. Chicago: University of Chicago Press.
- Sinnott-Armstrong, W. 2008. 'Framing moral intuitions'. In W. Sinnott-Armstrong (ed.) 'Moral Psychology, Volume 2: The Cognitive Science of Morality', 47–76. Cambridge, MA: MIT Press.
- Smetana, J. 1981. 'Preschool children's conceptions of moral and social rules'. *Child Development* 52: 1333–1336.
- Sousa, P. 2009. 'On testing the "moral law"'. *Mind & Language* 24: 209–234.
- Sousa, P., Holbrook, C. & Piazza, J. 2009. 'The morality of harm'. *Cognition* 113: 80–92.
- Stich, S., Fessler, D. M. T. & Kelly, D. 2009. 'On the morality of harm: A response to Sousa, Holbrook and Piazza'. *Cognition* 113: 93–97.
- Street, S. 2006. 'A Darwinian dilemma for realist theories of value'. *Philosophical Studies* 127: 109–166.
- Wierzbicka, A. 2001. *What Did Jesus Mean? Explaining the Sermon on the Mount and the Parables in Simple and Universal Human Concepts*. New York: Oxford University Press.
- . 2007. "'Moral sense'". *Journal of Social, Evolutionary, and Cultural Psychology* 1: 66–85.