Which causes of moral beliefs matter?

Which Causes of Moral Beliefs Matter?
Author(s): Elizabeth O'Neill
Source: *Philosophy of Science*, Vol. 82, No. 5 (December 2015), pp. 1070–1080
Published by: University of Chicago Press on behalf of the Philosophy of Science Association
Stable URL: http://www.jstor.org/stable/10.1086/683441
Accessed: 24-12-2015 07:48 UTC

Which Causes of Moral Beliefs Matter?

Elizabeth O'Neill*†

I argue that information about the distal causes of moral beliefs, such as evolution, is only relevant for assessing the epistemic status of moral beliefs in cases where we cannot determine whether the proximal processes producing these beliefs are reliable just by examining the properties of these proximal processes. Any investigation into the epistemic status of moral beliefs given their causes should start with a look at proximal causes—not at evolution. I discuss two proximal psychological influences on moral beliefs—disgust and sympathy—to demonstrate the feasibility of drawing epistemic conclusions from an examination of proximal causes alone.

1. Introduction. The philosophical literature on the causes of moral beliefs can be divided into two broad categories: works dealing with more distal causes of human morality, especially evolution (whether biological or cultural), and those dealing with more proximal influences on moral belief, such as emotions. There has been much discussion recently about what the evolutionary origins of the psychological processes underlying moral beliefs signify for the epistemic status of moral beliefs, especially in response to Richard Joyce's (2001, 2006) and Sharon Street's (2006, 2008) work (e.g., Copp 2008; Enoch 2010; Griffiths and Wilkins 2010; Machery and Mallon 2010; Schafer 2010; Brosnan 2011; Kahane 2011; Fraser 2014). There has also been some discussion about what more proximal causes of moral beliefs, such as emotions, signify for the epistemic status of moral beliefs (e.g., Singer

*To contact the author, please write to: Philosophy and Ethics, School of Innovation Science, Eindhoven University of Technology, Eindhoven, The Netherlands; e-mail: erh.oneill@gmail.com.

2005; Greene 2008; Berker 2009; Kahane and Shackel 2010; Kelly 2011; Kahane 2012).

In this article I argue that distal causes of moral beliefs are relevant for assessing the epistemic status of moral beliefs only if we cannot determine whether a given proximal process is reliable just by looking at the properties of that proximal process.[1] The influence of evolution will be relevant in some cases, but my thesis implies that any investigation into the epistemic status of moral beliefs given their causes should start with proximal causes. I describe two cases where information about properties of a proximal cause is sufficient to conclude that moral beliefs influenced by that cause have a poor epistemic status. The first case is disgust. I offer an argument for the unreliability of disgust that owes much to Daniel Kelly's work on disgust. Like Kelly, I conclude that disgust is an unreliable influence on moral beliefs, but I offer a different argument to support the conclusion. When we feel disgusted with something, we can say that it has a property of being disgusting. Call this disgustingness. I argue that disgustingness's property of "stickiness"—that disgustingness is easily transmitted and hard to get rid of—is inconsistent with plausible general features of moral facts. Consequently, we should decrease our confidence in moral beliefs influenced by disgust.

My second case is sympathy. On the assumption that if sympathy contributes to truth tracking in the moral domain, it does so by picking out entities in pain, and given empirical evidence that sympathy produces false positives and false negatives about the presence of pain at a significant rate, I argue that we should reduce our confidence in moral beliefs that rely on the influence of sympathy. In each of these cases, we reach conclusions about the epistemic status of moral beliefs influenced by proximal causes without examining the distal causes of moral beliefs, such as the evolutionary origins of disgust or sympathy. Further, my arguments illustrate two general types of etiological debunking argument that do not rely on substantial normative commitments. One strategy is to show that properties of the cause are inconsistent with plausible general features of moral facts. A second strategy is to identify a component of the process leading to a moral belief for which we have independent methods for assessing reliability and show that the component is unreliable.

Much of the literature on the evolutionary causes of moral beliefs, following Street (2006), has focused on the epistemic property of truth tracking. I too will be concerned with externalist varieties of epistemic status in this article, and in particular I have in mind Sherrilyn Roush's (2005) account of

1. By a distal cause of a belief, I refer here specifically to a higher-order process that produces the proximal process that produces the belief. Evolution is a distal cause of our beliefs in this sense because it produces the psychological processes that produce our beliefs.

truth tracking, which is formulated in terms of conditional probabilities, but I will also talk in terms of the reliability of processes that produce moral beliefs. I should also note that in this article I assume moral cognitivism, but the arguments about epistemic status that I offer should apply to both social constructivists and realists about moral facts.

2. The Priority of Proximal Causes. My thesis says that we should give priority to proximal causes in our investigation of the epistemic status of moral beliefs. Here I present a general argument for this thesis. If we can ascertain whether the proximal process tracks the truth on the basis of features of that process, looking at the distal process that produced that proximal process provides us with no additional information about the epistemic status of our beliefs. By contrast, if we ascertain the probability that the distal process would produce a truth-tracking proximal process, we can still gain additional information about the epistemic status of our beliefs by looking at the features of the actual proximal process that the distal process ended up producing. For instance, even if we determine that the distal process has a low probability of producing truth-tracking proximal processes, our proximal process may nonetheless, against the odds, be one that tracks the truth. In such a case, we can potentially gain additional information by investigating the proximal process directly. The evolutionary origins of the faculties that produce moral beliefs can provide us with some information about whether our moral beliefs track the truth, but there is an asymmetry in how much they can tell us compared with how much proximal processes can tell us. I propose that, in principle, information about proximal causes is sufficient for ascertaining whether our moral beliefs track the truth, but that information about distal causes is not sufficient for the same purpose.

If we could determine just by looking at features of our moral psychological processes that we track moral facts, we would not need to investigate how probable it was that evolution would produce truth-tracking moral psychological processes. A parallel point holds if we can determine on the basis of features of our moral psychological process that we do not track the truth of moral facts. In such a case, we need not investigate the probability that a truth-tracking psychological process would evolve. I am skeptical that we can conclude with confidence on the basis of information about the features of moral psychological processes that we track the truth of moral facts, but, as I show with the cases of disgust and sympathy, it is possible to conclude with confidence that we do not track the truth of moral facts without looking at distal processes.

In cases where we cannot determine the reliability or unreliability of our proximal processes on the basis of information about their features alone, we can use information about the evolutionary origins of these processes to inform our assessment of the epistemic status of moral beliefs. For instance,

if we cannot gather enough information about whether our current proximal process tracks the truth because there is conflicting evidence or we have an inadequate understanding of the proximal process, the probability that evolution would have produced a process that tracked the facts in question can provide additional information. Alternately, if we were interested in the epistemic status of moral beliefs across multiple generations and the relevant psychological processes would be subject to continued influence from evolution, information about evolution could improve our estimation of whether those psychological processes will track the truth in future generations. These are the situations to which the evolutionary debunking arguments advanced by Street, Joyce, and others are useful—situations where we cannot already be confident about whether proximal processes track truth just by looking at the features of those processes. But if we can determine with confidence the reliability of proximal processes on the basis of the features of those processes, we need not look at distal causes in order to draw conclusions about truth tracking. Thus, when we set out to ascertain whether our moral beliefs track the truth, our first stop should be the empirical work on proximal psychological processes rather than work on the evolution of those proximal processes.

This point generalizes to beliefs outside morality. That is, information about proximal processes is more valuable than information about distal processes even when we are assessing the epistemic status of nonmoral beliefs. Suppose that there are two people, Juana and Marta, and we want to evaluate the epistemic status of their beliefs about certain mathematical propositions. Suppose that we are confident that each person acquired reliable mathematical reasoning processes early in life that regularly produce true beliefs about some set of mathematical propositions. We are confident that each person will continue to hold these mathematical reasoning processes for the rest of her life. If we were only interested in these two people and their beliefs about the relevant set of mathematical propositions, and we could be confident that their mathematical reasoning processes were reliable, then we could be confident that each is tracking the truth of these mathematical propositions, regardless of where their reasoning processes came from.

However, if we were unsure about whether each person possessed a reliable mathematical reasoning process, we could obtain some information about whether they track the truth by looking at the distal process by which they obtained their mathematical reasoning processes. Suppose that the distal process by which Juana obtained her mathematical reasoning process was reading a reputable mathematics textbook, whereas the process by which Marta obtained her mathematical reasoning process was listening to a local mystic who usually disseminates mathematical falsehoods and ineffective tricks. Looking at those distal processes provides us with information about the probability that each woman acquired a reliable mathematical reasoning

process—high in Juana's case, low in Marta's case. Of course, the information we obtain by this route misleads us here, since in fact Marta, against the odds, acquired a reliable proximal process and does track the truth of the relevant propositions. We would be better off with good information about the proximal reasoning processes that Juana and Marta actually possess.

Thus, distal processes can supply useful information for our assessment of the status of beliefs in cases where we are unsure about the reliability of proximal processes. In other cases, distal processes do not provide information for assessing epistemic status beyond what can be gained from information about proximal processes.

3. Two Proximal Causes and Their Implications. Although their work has not attracted as much debate as Street's, some philosophers, such as Peter Singer, Joshua Greene, and Daniel Kelly, have looked at proximal causes to debunk moral beliefs. Many of the attempts to debunk using proximal causes involve an appeal to significant normative assumptions, such as that poisons and pathogens are not morally relevant (e.g., Kelly and Morar 2014, 22) or that there is no morally significant difference between killing someone in a personal or impersonal way (Singer 2005, 348). By contrast, I suggest two distinct strategies for assessing the reliability of proximal processes while minimizing assumptions about which properties or actions are right or wrong. I will illustrate these strategies with two cases. Note that in each of these cases, we conclude that the proximal cause influencing our moral belief is unreliable without looking at the evolutionary origins of the proximal cause. Thus, these cases also illustrate the priority of proximal causes that I argued for in the previous section.

3.1. Disgust. The first strategy I propose operates by establishing that a cause exhibits properties that are inconsistent with plausible general features of moral facts. The premise required here is not that there are moral facts with these features, but rather that if there are any moral facts, then they have these features. Such general features might include that the time of day during which one considers a moral proposition does not affect the truth-value of that proposition or that the order in which one considers moral propositions does not affect the truth-values of those propositions. This strategy does not require assuming the truth of any particular moral proposition, such as the proposition that pain is generally bad.

Empirical evidence shows that a person's level of disgust can affect his or her moral judgments on a variety of questions (Haidt and Hersh 2001; Wheatley and Haidt 2005; Schnall et al. 2008; Horberg et al. 2009; Inbar et al. 2009; Zhong, Strejcek, and Sivanathan 2010; Helzer and Pizarro 2011; Chapman and Anderson 2013; however, for objections see also May 2014; Landy and Goodwin 2015). Two key features of disgust are the following:

First, entities that we consider disgusting can contaminate other entities that become associated with them—disgustingness can be transmitted quite easily to things we did not previously consider disgusting. Second, once one associates disgust with an entity, it is hard to stop thinking of that entity as disgusting even when the situation changes or when one acquires new evidence. Together, we can call these two features of disgustingness *stickiness*—disgustingness is easily transmitted between entities, and once disgustingness is attached to an entity, it is hard to revise that association.

Classic cases that support the stickiness of disgust judgments include subjects' unwillingness to drink juice stirred with a new comb or a sterilized cockroach and unwillingness to wear the clothes of an amputee. One might want to object that these are nonmoral contexts, and that perhaps disgustingness is not sticky in the same way in moral contexts. However, there is reason to think that the stickiness of disgustingness extends to putatively moral contexts as well. For instance, Haidt's work on moral dumbfounding suggests the difficulty people have in revising both their disgust response and their resulting moral judgments about a variety of cases, such as (ostensibly) harm-free incest (Haidt, Bjorklund, and Murphy 2000).

Kelly (2011, 2014) elucidated these two features of disgustingness, which I am calling stickiness, and used them for moral epistemological purposes, reaching the same conclusion that I am advancing, that disgust is an unreliable influence on moral beliefs. But his argumentative strategy is different from the one I am proposing. I argue that the stickiness of disgustingness is inconsistent with plausible general features of moral facts. Given the ease with which we come to associate disgust with a broad array of entities and actions and the difficulty we have dissociating feelings of disgust from those entities and actions, disgust is unlikely to be tracking some sort of moral property. The only assumption that this argument requires about the nature of moral facts is that the moral facts (if there are any) must be static in a certain sense—they do not vary wildly, with a moral property such as *badness* having the potential to apply to actions or objects as easily as disgustingness can apply to actions or objects. This argument does not involve claims about what things should be treated as disgusting, what morally relevant features of the world disgust is picking out, or what sorts of things are likely to be right or wrong, which is what makes the argument unusual.[2]

Here is a hypothetical case of sticky moral disgust that highlights the unreliability of the influence: if we feel disgust in response to an outgroup, we might come to feel disgust in response to various traditions that the outgroup engages in (whatever those traditions are) and symbols associated with the

---

2. Admittedly, it does involve denying that there is a constitutive relationship between feeling disgust toward an action and that action being wrong.

group (whatever those symbols are), leading us to indelibly classify those traditions and symbols as wrong or to be avoided. Thus, we have reason to reduce confidence in moral beliefs that have been substantially influenced by disgust because one of the properties of disgust makes it unlikely to be tracking morally relevant features of the world.

3.2. Sympathy.   The second strategy I propose identifies a component in the moral belief production process of which we have independent methods for assessing reliability and shows that the component is unreliable.[3] In the case of sympathy, the targeted component is the move from feeling sympathy to the belief that the person is in pain, understood broadly. For the sake of argument, I make the substantial moral assumption that if sympathy contributes to truth-tracking moral beliefs, it does so by picking out entities in pain. I show that sympathy produces both false positives and false negatives at a significant rate and conclude that we should reduce our confidence in moral beliefs that hinge on pain attributions produced by sympathy.[4]

The psychological cause of interest here is the emotion-laden process that picks out entities that are undergoing pain. I use the term "sympathy" to refer to this process. Only for the sake of the argument, I assume that pain is sometimes morally relevant and that making accurate moral beliefs sometimes depends on the accuracy of our attributions of pain to entities in our environment. I take it that some emotional processes related to sympathy or compassion often play a role in the production of our belief that we ought to take steps to reduce the pain of another entity, in part by helping us identify the entity as experiencing pain. If this process does frequently play a role in producing moral beliefs, we should be concerned about how reliably it picks out those entities that are in fact in pain. And indeed, there is reason for concern in two directions: sympathy appears to produce both false positives and false negatives. It sometimes leads us to classify entities as in pain when we have independent reason to believe that they are not in pain, and it sometimes leads us to classify entities as not in pain when we have independent reasons to believe that they are, in fact, in pain. For instance, on the false-negative side this mechanism of pain attribution is less likely to engage when it comes to outgroup members. People often empathize less with outgroup members experiencing pain than ingroup members; for instance, Trawalter, Hoffman, and Waytz (2012) show that white subjects tend to attribute lower levels of pain to black people (see also Cikara and Fiske 2011).

3. This strategy could be considered one way to construct what Nichols calls a "process debunking argument" (2014, 730).

4. Bloom (2013) and Prinz (2011) have recently offered distinct arguments for the conclusion that empathy is an unreliable source for moral beliefs.

On the false-positive side, consider recent work on robots. Rosenthal-von der Putten et al. (2012) found that subjects who watched videos of robots being treated violently exhibited higher physiological arousal, expressed empathy for the robot, and reported more negative emotional reactions than subjects who watched videos of robots being treated in a friendly way. Rosenthal-von der Putten and colleagues also found similar patterns of brain activation when subjects watched videos of humans and robots being treated affectionately (though admittedly somewhat higher brain activation when subjects viewed humans being treated violently than when they viewed robots being treated violently). Anecdotal evidence about the use of companion robots like "Paro," a seal-like robot, and "Keepon," a robot that looks like an Easter peep, also suggests a tendency to respond to the behavior of entities that exhibit certain features, such as responsiveness, as if it indicated something about internal experiences.

A critical link in this argument is the move from the fact that people have certain neurological and emotional responses to the wrong sorts of things to the idea that these responses influence their beliefs about pain and obligations. Some evidence for this link comes from reports that soldiers tend to become fond of the robots they use in battle and that some soldiers have actually risked their own lives to rescue robots—presumably not just because the robot was a valuable piece of machinery (Singer 2009). Also, work by Bartneck and Hu found that although subjects, when instructed to do so, destroyed robots with which they had interacted, the subjects expressed hesitance and regret, saying things like "I didn't like to kill the poor boy," "The robot is innocent," and "This is inhumane!" (2008, 426). (For further discussion of this case and the features that lead subjects to attribute pain to robots, see Fiala, Arico, and Nichols 2014.)

If it is sympathy or some other sort of emotional pain-detection process that leads us to think we should, for instance, avoid stepping on insects, or keep patients in a persistent vegetative state alive, the fact that this emotional pain-detection process so easily produces false positives may be reason to reduce confidence in our beliefs and to look elsewhere for additional evidence about whether the entities in question experience pain. This is not to say that it would be better to ignore sympathy when detecting entities in pain in the world. Rather, the suggestion is that we should reduce confidence in beliefs involving pain attributions that have been influenced by sympathy, at least until we have identified the features that characterize the contexts where our sympathy-influenced pain attribution judgments go astray.

The argument I have just presented says that our pain attribution judgments, which influence our moral beliefs about what ought to be done, are unreliable. Similar arguments could be constructed using other types of judgments we make that influence our moral beliefs, such as judgments about actors' intentions, the causal structure of a series of events, and judgments

about costs and benefits of a given action. If we can show that these contributors to moral beliefs are unreliable or subject to biases that we can diagnose independently, we can conclude that the moral beliefs themselves are not tracking the truth.

3.3. The Implications of Proximal Causes.   In each of these two cases, information about a proximal process leads to the conclusion that moral beliefs obtained via that process are not truth tracking. We need not look at the evolutionary origins of disgust or sympathy to inform our assessment of whether moral beliefs produced by processes involving disgust and sympathy track the truth.

One reason we might continue to focus on information about the evolutionary origins of moral beliefs is the thought that if we could assess the reliability of evolution, we could draw general conclusions about whether we track moral facts, because the influence of evolution is so wide reaching. By contrast, it is reasonable to think that disgust and sympathy influence only subsets of our moral beliefs. For this reason, by targeting only moral beliefs subject to influence by disgust or sympathy, the debunking arguments I have offered are of limited scope. In principle, though, one could supply proximal debunking arguments that apply to all moral beliefs: if there are certain components that all moral belief production processes share, then determining that those components were unreliable would permit a global rather than selective proximal debunking argument. At the same time, if we have distinct evolutionary stories for different types of moral beliefs (e.g., one evolutionary explanation for justice and fairness norms and another for purity norms), we can have distal debunking arguments that are themselves of only limited scope.[5] Thus, proximal debunking arguments need not always be of narrower scope than evolutionary debunking arguments.

4. Conclusion. I have argued that any investigation into the epistemic status of moral beliefs given their causes should start with proximal causes. In cases where we are confident in our estimation of the degree to which a proximal process tracks the truth, we need not look at the distal origins of those proximal causes in order to draw conclusions about the epistemic status of the moral beliefs that they influence. I have proposed two distinct strategies for reaching epistemological conclusions from information about the psychological and biological causes of moral beliefs and illustrated them with examinations of two types of proximal causal influences on moral beliefs: disgust and sympathy. The properties of those proximal causes

5. See Kelly (2014) for discussion of selective debunking arguments and an example of an evolutionary debunking argument of limited scope.

alone license conclusions about the epistemic status of the beliefs they influence: we should reduce confidence in moral beliefs produced by processes that involve disgust or that rely on the activation of sympathy.

REFERENCES

Bartneck, Christoph, and Jun Hu. 2008. "Exploring the Abuse of Robots." *Interaction Studies* 9:415–33.

Berker, Selim. 2009. "The Normative Insignificance of Neuroscience." *Philosophy and Public Affairs* 37:293–329.

Bloom, Paul. 2013. "The Baby in the Well: The Case against Empathy." *New Yorker*, May 20.

Brosnan, Kevin. 2011. "Do the Evolutionary Origins of Our Moral Beliefs Undermine Moral Knowledge?" *Biology and Philosophy* 26:51–64.

Chapman, Hanah A., and Adam K. Anderson. 2013. "Things Rank and Gross in Nature: A Review and Synthesis of Moral Disgust." *Psychological Bulletin* 139:300–327.

Cikara, Mina, and Susan T. Fiske. 2011. "Bounded Empathy: Neural Responses to Outgroup Targets' (Mis)fortunes." *Journal of Cognitive Neuroscience* 23:3791–3803.

Copp, David. 2008. "Darwinian Skepticism about Moral Realism." *Philosophical Issues* 18:186–206.

Enoch, David. 2010. "The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with It." *Philosophical Studies* 148:413–38.

Fiala, Brian, Adam Arico, and Shaun Nichols. 2014. "You, Robot." In *Current Controversies in Experimental Philosophy*, ed. Edouard Machery and Elizabeth O'Neill, 31–47. New York: Routledge.

Fraser, Benjamin James. 2014. "Evolutionary Debunking Arguments and the Reliability of Moral Cognition." *Philosophical Studies* 168:457–73.

Greene, Joshua. 2008. "The Secret Joke of Kant's Soul." In *Moral Psychology*, vol. 3, *The Neuroscience of Morality: Emotion, Brain Disorders, and Development*, ed. Walter Sinnott-Armstrong, 35–80. Cambridge, MA: MIT Press.

Griffiths, Paul, and John Wilkins. 2010. "When Do Evolutionary Explanations of Belief Debunk Belief?" http://philsci-archive.pitt.edu/5314/.

Haidt, Jonathan, Fredrik Bjorklund, and Scott Murphy. 2000. "Moral Dumbfounding: When Intuition Finds No Reason." Unpublished manuscript, University of Virginia.

Haidt, Jonathan, and Matthew A. Hersh. 2001. "Sexual Morality: The Cultures and Emotions of Conservatives and Liberals." *Journal of Applied Social Psychology* 31:191–221.

Helzer, Erik G., and David A. Pizarro. 2011. "Dirty Liberals! Reminders of Physical Cleanliness Influence Moral and Political Attitudes." *Psychological Science* 22:517–22.

Horberg, E. J., Christopher Oveis, Dacher Keltner, and Adam B. Cohen. 2009. "Disgust and the Moralization of Purity." *Journal of Personality and Social Psychology* 97:963–76.

Inbar, Yoel, David A. Pizarro, Joshua Knobe, and Paul Bloom. 2009. "Disgust Sensitivity Predicts Intuitive Disapproval of Gays." *Emotion* 9:435–39.

Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.

———. 2006. *The Evolution of Morality*. Cambridge, MA: MIT Press.

Kahane, Guy. 2011. "Evolutionary Debunking Arguments." *Nous* 45:103–25.

———. 2012. "On the Wrong Track: Process and Content in Moral Psychology." *Mind and Language* 27:519–45.

Kahane, Guy, and Nicholas Shackel. 2010. "Methodological Issues in the Neuroscience of Moral Judgement." *Mind and Language* 25:561–82.

Kelly, Daniel. 2011. *Yuck! The Nature and Moral Significance of Disgust*. Cambridge, MA: MIT Press.

———. 2014. "Selective Debunking Arguments, Folk Psychology and Empirical Moral Psychology." In *Advances in Experimental Moral Psychology*, ed. J. C. Wright and H. Sarkissian, 130–47. New York: Continuum.

Kelly, Daniel, and Nicolae Morar. 2014. "Against the Yuck Factor: On the Ideal Role of Disgust in Society." *Utilitas* 26:153–77.

Landy, Justin F., and Geoffrey P. Goodwin. 2015. "Does Incidental Disgust Amplify Moral Judgment? A Meta-Analytic Review of Experimental Evidence." *Perspectives on Psychological Science* 10:518–36.

Machery, Edouard, and Ron Mallon. 2010. "Evolution of Morality." In *The Moral Psychology Handbook*, ed. John Doris, 3–46. Oxford: Oxford University Press.

May, Joshua. 2014. "Does Disgust Influence Moral Judgment?" *Australasian Journal of Philosophy* 92:125–41.

Nichols, Shaun. 2014. "Process Debunking and Ethics." *Ethics* 124:727–49.

Prinz, Jesse. 2011. "Against Empathy." *Southern Journal of Philosophy* 49:214–33.

Rosenthal-von der Pütten, A. M., N. C. Krämer, L. Hoffmann, S. Sobieraj, and S. C. Eimler. 2012. "An Experimental Study on Emotional Reactions towards a Robot." *International Journal of Social Robotics* 5:17–34.

Roush, Sherrilyn. 2005. *Tracking Truth*. Oxford: Oxford University Press.

Schafer, Karl. 2010. "Evolution and Normative Scepticism." *Australasian Journal of Philosophy* 88:471–88.

Schnall, Simone, Jonathan Haidt, Gerald L. Clore, and Alexander H. Jordan. 2008. "Disgust as Embodied Moral Judgment." *Personality and Social Psychology Bulletin* 34:1096–1109.

Singer, P. W. 2009. *Wired for War: The Robotics Revolution and Conflict in the 21st Century.* New York: Penguin.

Singer, Peter. 2005. "Ethics and Intuitions." *Journal of Ethics* 9:331–52.

Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127:109–66.

———. 2008. "Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About." *Philosophical Issues* 18:207–28.

Trawalter, Sophie, K. M. Hoffman, and A. Waytz. 2012. "Racial Bias in Perceptions of Others' Pain." *PLoS ONE* 7:e48546.

Wheatley, Thalia, and Jonathan Haidt. 2005. "Hypnotic Disgust Makes Moral Judgments More Severe." *Psychological Science* 16:780–84.

Zhong, Chen-Bo, Brendan Strejcek, and Niro Sivanathan. 2010. "A Clean Self Can Render Harsh Moral Judgment." *Journal of Experimental Social Psychology* 46:859–62.