

ARTICLE



Why do ethicists eat their greens?

Andrew Sneddon

Department of Philosophy, University of Ottawa, Ottawa, Canada

ABSTRACT

Eric Schwitzgebel, Fiery Cushman, and Joshua Rust have conducted a series of studies of the thought and behavior of professional ethicists. They have found no evidence that ethical reflection yields distinctive improvements in behavior. This work has been done on English-speaking ethicists. Philipp Schönegger and Johannes Wagner (2019) replicated one study with German-speaking professors. Their results are almost the same, except for the finding that German-speaking ethicists were more likely to be vegetarian than non-ethicists. The present paper devises and evaluates eleven psychological hypotheses (along with one from Schönegger and Wagner) aimed at explaining why ethical reflection might have motivational influence for this topic but not for others. Three hypotheses are judged to be plausible at this initial stage: generic emotional support, perception of cost as a source of emotional obstacles, and social categorization.

ARTICLE HISTORY

Received 15 May 2019
Accepted 2 January 2020

KEYWORDS

Ethical reflection; emotion; reason; categorization; behavior of ethicists; vegetarianism; eric Schwitzgebel

1. Introduction

Since 2009, Eric Schwitzgebel, along with Joshua Rust and Fiery Cushman, has conducted a series of studies aimed at assessing the psychology of professional ethicists (Schwitzgebel & Cushman, 2012, 2015; Schwitzgebel & Rust, 2009, 2010, 2011, , 2014). One thread of this research examines whether the sort of reflection that is encouraged and cultivated through the study of ethics has any sort of measurable effect of behavior. One way or another, Schwitzgebel and colleagues have collected data about the beliefs and behavior of both professional ethicists and non-ethicist philosophers. The result has either been surprising or unsurprising, depending on who you ask: no evidence has been found to suggest that ethicists act in morally better ways than non-ethicists. Hence, no evidence has been found suggesting that the reflection that is ostensibly performed and developed during and after an education in ethics has any effect on behavior.

Schwitzgebel, Rust, and Cushman have performed their studies on English-speaking professors primarily in North America. In the wake of this work, Schönegger and Wagner (2019) attempted a replication of one of these studies (i.e., Schwitzgebel & Rust, 2014), but with German-speaking professors. Their results were almost exactly the same, but not quite: putting it roughly, Schönegger and Wagner found that all of the philosophers were more likely than non-philosophers to judge that we should be vegetarian, and that the ethicists were more likely to be vegetarian than both non-ethicists and non-philosophers (Schönegger & Wagner, 2019, pp. 544–5, 551). Not only is the behavior of ethicists arguably better here, this is the only topic for which there was marked consistency in ethicists' judgments and behavior. Of particular comparative importance is charity: ethicists judged giving to charity to be more important than the non-philosophers did, but they did not give more than the non-philosophers (Schönegger & Wagner, 2019, p. 545).

These findings about ethicists and vegetarianism are an interesting development in this series of studies, but also in the probing of the links, or lack thereof, between ethical reflection and behavior. Supposing that the findings of Schönegger and Wagner hold up – it's one study so far, so they might not – they suggest that the question framing this work has been posed too broadly. Instead of asking, in general, whether ethical reflection shapes behavior, we should be looking in a more fine-grained fashion for topics for which there do seem to be links and for ones for which there seem to be no links, and then, we should be asking why: why are ethicists more likely to be vegetarians than non-ethicists, and why aren't they more likely to give relatively higher amounts to charity? Why do ethicists eat their greens? Why might reflection have effects on behavior here, but not for other topics?¹

Clearly, more studies are needed to sort out topics for which there are and aren't apparent reflection–behavior links. However, it's not too early to take stock. Given the appearances in the wake of the Schönegger and Wagner study, why might this be the case? My purpose is to frame some hypotheses and to evaluate their initial plausibility in the light of other things that have been revealed in the last generation's work in moral psychology.²

2. The landscape so far

First, let's get a clearer view of what the studies of the behavior of ethics professors have probed and found. Schwitzgebel and colleagues have found no interesting differences in behavior between ethicists and non-ethicists. Schwitzgebel and Rust (2014, pp. 295–296) formulated four possible models of the relation between ethical reflection and behavior in general:

1. **Booster view:** ethical reflection tends to yield discovery of moral truths, and such discoveries influence behavior positively.

2. **Rationalization view:** ethical reflection changes beliefs to fit preexisting patterns of behavior or behavioral inclinations.

3. **Inert discovery view:** ethical reflection tends to yield discovery of moral truths, but with no significant effect on behavior.

4. **Epiphenomenalist view:** ethical reflection produces no significant changes in either beliefs or behavior.

Schwitzgebel and Rust (2014, p. 320) think that the extant data pretty much rule out the booster view, since it would require that ethicists behave significantly better than non-ethicists which, to that date, had not been found. This leaves the other models in play. Schwitzgebel and Rust (2014, p. 320) are careful to entertain the possibility that ethical reflection can have both good and bad effects on behavior, which cancel each other out. In other words, they gesture to the possibility that more fine-grained studies would vindicate the booster view for specific topics.

This might be what Schönegger and Wagner (2019) have found. They surveyed more than 400 German-speaking professors on a variety of topics. The professors were divided virtually evenly amongst ethicists (151), non-ethicist philosophers (133), and non-philosophers (133). Following the study methodology of Schwitzgebel and Rust (2014), Schönegger and Wagner assessed relative stringency of ethical beliefs across several topics, similarities and differences in behavior relevant to these topics, and consistency between beliefs and behavior for these topics. For most topics, the ethicists did not voice more stringent beliefs than the others. Indeed, it was the non-philosophers who tended to have the most stringent moral beliefs. Ethicists had more stringent beliefs than the other two groups about giving to charity (pro), and both philosopher groups had more stringent beliefs about vegetarianism (pro) than non-philosophers (Schönegger & Wagner, 2019, p. 540). There was no consistent difference in behavior across the groups except with regard to vegetarianism (Schönegger & Wagner, 2019, p. 543). Specifically, when asked about their last meal, 24% of ethicists reported eating meat, compared with 40% of non-ethicists and 39% of non-philosophers. With regard to meals per week containing meat, ethicists reported an average of 2.1 meals, non-ethicist philosophers reported 2.8, and non-philosophers reported 3 (Schönegger & Wagner, 2019, p. 544). Accordingly, it is only in the case of vegetarianism that there's significant belief-behavior consistency for ethicists as compared with the two non-ethicist groups (Schönegger & Wagner, 2019, p. 545). While both philosopher groups voiced relatively more stringent attitudes in favor of vegetarianism, only ethicists had significantly lower meat-eating behavior (Schönegger & Wagner, 2019, p. 551). Despite having stronger in-favor attitudes about giving to charity than the other groups, ethicists did not report giving a higher proportion of their income to charity, thereby

exhibiting relative inconsistency on this topic (Schönegger & Wagner, 2019, pp. 545, 551).

These numbers provide initial support for the more fine-grained analysis indicated by Schwitzgebel and Rust: there's reason to think that ethical reflection yields a difference in behavior with regard to diet, but not with regard to charity, at the very least. Schönegger and Wagner (2019, p. 552) suggest that the tradition of studies on the moral–conventional distinction might be pertinent to explaining this (on the moral–conventional distinction, see, e.g., Smetana, 1981, 1993; Turiel, 1983, 1998, 2003). The idea is that, perhaps, implicitly or explicitly, the respondents grouped the topics they were asked into the “properly moral” (i.e., vegetarianism and charity) and the “not-so-clearly-really-moral” (i.e., the rest). However, there are two shortcomings of this suggestion. The first is that sensitivity to differences between moral and conventional rules – and hence, moral and conventional topics – has been suggested to appear very early in development (e.g., Smetana, 1981; Turiel, 1998). This means that it is found throughout the population before exposure to, and training in, the sort of reflection that is thought to be characteristic, or even definitive, of the expertise of professional ethicists. Hence, it ought not to explain what makes ethicists different from others. Second, and more important, even if such a psychological tendency explained the grouping of charity and vegetarianism together as morally different from the other topics, it would not answer our question, which is why ethicists are more likely than non-ethicists and non-philosophers to eschew eating meat while not being more generous givers to charity. More fine-tuned hypotheses are needed if we are going to shed light of why it seems that ethical reflection has effects on behavior for this topic but not for the other ones that were asked about, especially charitable giving. Let's turn to some possibilities.

3. Why do ethicists eat their greens: Some hypotheses

3.1. Two false starts

I shall start with two hypotheses that are worth considering at the outset, but should be replaced by others pretty quickly. While I have suggested that it appears that the Schönegger and Wagner study reveals that ethical reflection can lead to behavioral improvement with regard to vegetarianism, the first thing to consider is whether this is all that it is: appearance. One thing that might explain the dietary differences among the three groups is dietary preferences in general. The first hypothesis, then, is that the pre-education distribution of vegetarian tendencies explains the questionnaire results. Perhaps vegetarians are somewhat more likely to choose philosophical and specifically ethical studies than non-ethical and non-philosophical ones.

Schwitzgebel and Rust (2011, p. 6) are open about presuming that there are no such differences among the groups while acknowledging the empirical wager being made. They are correct on both counts, as far as I can tell. There is no specific reason that I know of to think that there are such pre-education differences, but we do not have detailed information about this. Pre-education interest in nonhuman animals might be indirectly revealing. For instance, data on pet ownership would tell us something about the involvement of nonhumans in the lives of people surveyed. One might expect greater personal attachment to nonhumans among vegetarians. However, I know of no data about this. For what it's worth, Peter Singer somewhat famously voices lack of interest in pets despite being a foremost advocate for vegetarianism and animal rights more generally (Sommers, 2016, pp. 156–157). Strictly speaking, it would be desirable to have information about dietary preferences at three stages: before post-secondary education, before specifically philosophical education (which may or may not coincide with the beginning of post-secondary education), and before specialization in ethics. Only if there are no differences across the groups at all three stages can we take the Schönegger and Wagner data at face value.

It is theoretically possible that pre-education differences mask practical effects of ethical reflection: if the people who become ethicists have pre-education below-average behavioral tendencies about the topics addressed in this body of work, then practically, efficacious ethical reflection that brought them up to average would not stand out as remarkable. Of course, there is no evidence for this. Barring evidence to the contrary, the plausible assumption is that there are no significant pre-education dietary differences among the three groups of professors.

A bit of reflection about the content of studies of ethics provides even more *prima facie* reason to discount this hypothesis. Why might we think that those with vegetarian proclivities would be attracted to specialization in ethics? If animal rights issues were particularly prominent as topics of study in ethics programs, then those whose vegetarianism was based on values (rather than taste or health) might find the field attractive. Alternatively, if ethicists generally were vegetarian, and much more so than non-ethicist philosophers and non-philosophers, then those who valued their own vegetarianism might be attracted to a field dominated by like-minded eaters. However, the first is certainly not the case, and there is no specific reason to think that the second is the case. Lots of issues are studied in an ethics education. Besides the theoretical aspects of the discipline, practical ethics courses can be taken on many topics: business ethics, social justice, ethics of war, medical ethics, and more, all in addition to animal welfare issues. Vegetarianism does not figure particularly prominently in ethical studies, let alone dominate them. As for the dietary preferences of ethicists, no data exist that I know of. However, my impression is that more ethicists are

omnivores than vegetarian. If this is correct, then pre-education vegetarians would not be entering a field of studies characterized by like-minded eaters. They might find more vegetarians among the ethicists than in other fields, but they would still be in the dietary minority.

On the assumption that there are no significant pre-education differences in dietary tendencies across the three groups, we can move to the second nonstarter hypothesis. It takes as its starting point one important point of any education (or, at least, of any theoretical education, as opposed to practical education such as in performance of a sporting activity): cultivating an understanding of the subject matter. Sometimes, the development of such understanding will deepen one's grasp of what she already knows; other times, however, it will change one's outlook on a topic. The ethicist develops an understanding of ethics – of values, of certain sorts of reasons to act and not to act, of desirable character traits, and of the empirical details about actions, persons, and the wider world relevant to these things. The present hypothesis is that acquiring such understanding tends to change ethicists' view of their relations to nonhumans more than their view of our relations to each other. Such transformation can happen in various ways – via thought about the moral reasons to act and not to act in certain ways, via acquisition of information about the nature of both humans and nonhumans, and more. If ethical reflection changes ethicists' perceived relation to nonhumans more than their perceived relation to humans (and if, which is likely, charitable donations are overwhelmingly made to organizations aiding other humans), then the cultivation of ethical understanding should change ethicists' behavior more with regard to nonhumans than with regard to humans. Those who don't acquire such transformative understanding will not exhibit the same pattern of behavioral change. This, then, explains the behavioral differences among the three groups surveyed by Schönegger and Wagner.

There are two problems with this, and both have to do with the nature of understanding. Consider this: in what does the understanding of some theoretical subject matter consist? Suppose that such understanding is primarily a matter of the beliefs one holds. This is plausible, but as a psychological hypothesis, it encounters a problem in the survey data. Schönegger and Wagner found that ethicists and non-ethicist philosophers had very similar beliefs about the importance of being vegetarian. It seems to be their behavior that differs, not their understanding (construed in terms of beliefs), so appealing to differences in understanding is a nonstarter for shedding light on what differentiates these groups.

Suppose instead that 'understanding' is not exhausted merely by the beliefs one holds. It is instead a name for a complex psychological phenomenon involving belief, inferential tendencies, behavioral dispositions, affective propensities, and perhaps more. Now, the problem is that the

hypothesis is too vague. It amounts to a mere relabeling of what we're trying to explain. There is some psychological difference among these groups that makes ethicists relatively more likely to be vegetarian – what is it? To answer this question with “the difference lies in the understanding: a complex psychological phenomenon involving belief etc.” hardly says more than that there exists some psychological difference between these groups. The real question asks what the psychological difference is: what, more specifically, accounts for the behavioral pattern? What difference or differences exist in understanding? The upshot is that a hypothesis focused on understanding as a complex psychological phenomenon is trivially true, and hence uninformatively so.

3.2. *Emotion-based hypotheses*

Assuming that there are no significant pre-education differences in dietary tendencies across the three groups, we need semi-fine-grained hypotheses that address why vegetarianism might be psychologically different from the other topics addressed in the survey, especially charitable donation. One general possibility is that emotion is at the core of the difference. The work over the last generation on moral psychology in general and on moral judgment in particular has revealed emotion to be particularly important in a variety of ways (for a tiny sampling see, e.g., Haidt, 2001; Nichols, 2004; Prinz, 2007; Greene, 2008; for a probing overview, see Tiberius, 2015, Chapters 5–6). Emotion is of particular interest in the present context because our question concerns the practical efficacy of such judgments (e.g., eating meat is morally wrong), and emotion is renowned for its motivational power. It's one thing to believe that, for example, tax evasion is wrong; it's quite another thing to be angry about it (or disgusted, or distraught, etc.). Because of this, emotion is a good candidate for explaining why ethicists tend to act in accordance with their beliefs about the importance of being vegetarian.

A handful of emotional hypotheses can be generated. One family of emotional hypotheses focuses on emotional support for vegetarianism-relevant action. According to one version, the vegetarian tendencies of ethicists are due to specific emotional support for such action. For example, perhaps the vegetarian ethicists have an emotional concern for nonhumans; the combination of this concern with such beliefs (that eating them is wrong), tends to yield a refusal to eat them. The crucial thing here is the relation of such concern to the ethical reflection that the ethicists perform due to their academic specialization. If the emotion in question is independent of this reflection, this hypothesis risks collapsing into the nonstarter hypothesis about pre-education differences: it's the pre-education, or at least, the independent-of-education psychology of the people surveyed

that end up as ethicists that accounts for their behavior. This leaves the other option, that ethical reflection produces the concern. This, however, is not promising. The reason is that it leaves unexplained why reflection can produce a motivationally efficacious emotion in the case of eating animals but not in the case of giving to charity (and perhaps not in other cases either). More explanation is needed, lest again we have a mere relabeling of the phenomenon to be explained.

A second sort of emotional support hypothesis cites generic emotional support. For instance, rather than a concern for nonhuman animals specifically, perhaps the vegetarian ethicists have a general care for morally relevant others. This, however, suffers from the same problem as the previous hypothesis: it does not shed light on the difference between vegetarianism and charity (not to mention other issues), since both involve doing good and preventing harm for morally relevant others. Another version is more plausible. We are deeply sensitive to the difference between actions and effects perceived as personal, and others perceived as impersonal (see, e.g., Greene, 2008). Arguably, this is part of the psychology of typical responses to trolley cases (on the pattern of responses, see Hauser, 2006, pp. 112–121, 127–128; for this explanation, see Greene, 2008, pp. 41–44, 2013, Chapter 4). Most of us, by far, say that it's morally permissible to pull a lever to switch a runaway train from a track on which it will hit and kill five people to a different track where it will kill just one. However, most of us, again by far, say that it's impermissible to push a single person in front of the train, thereby killing them, even when doing so will save five people. The personal involvement of putting your hands on somebody, by hypothesis, is psychologically different from the impersonal involvement of pulling a lever, despite the fact that the life–death math is the same.

Here's how this might pertain to the present issues. Giving to charity is famously impersonal. Our money goes somewhere, then something else happens somewhere else, all without us seeing any effects of our giving in the vast majority of cases. Maybe we're paying for vaccinations, or clean water, or education – we typically see none of this. However, we are personally involved in what we eat. Very rarely is it a complete mystery as to what we're putting in our mouths. We often do not know all of the ingredients of our meals, but we usually know the main ones. Animal matter is often quite prominent in the foods made of it. We ingest this into our physical persons, typically with our own hands; it's hard to get more personally involved than that. While this hypothesis is, I think, clearly more plausible than the previous ones, study is needed to probe it, especially in connection with other topics that involve comparably personal involvement.

Other possible emotional hypotheses turn away from emotional support to emotional obstacles. It's one thing to think that, for example, you should

treat people equally, and another to think this while hating brown people. In the latter case, there's an emotional obstacle in the way of acting in accordance with your moral judgment. In the case of biological omnivores with wide-ranging tastes, the adoption of a vegetarian diet need not come with costs. One way or another, we have to eat; if you have a wide array of tastes, it makes little gustatory difference whether you eat animals or vegetables. However, charity is all cost, *prima facie*: without choosing to do so, we can avoid giving to charity, and a safe starting assumption is that people prefer to spend their limited resources in ways that benefit themselves directly or indirectly. There is no claim to necessity here. We can admit that for omnivores with narrow tastes, vegetarianism will involve a cost, potentially a very high one, and that for those with other-minded sentiments, the costs of charity can be very low, to the point of being negligible. Nevertheless, the general point stands. Empirical studies aimed at contexts in which the costs of charitable donations are varied would be particularly illuminating for this hypothesis. I am selfish enough to be gratified to see, every spring, the benefit on my income tax return of my annual charitable giving. If these tax benefits diminished, I would be less happy about giving to charity – that is, there would be a greater emotional obstacle to such donation. However, if the tax benefits increased (are you listening, government of Canada?), I would readily consider giving more. Generalizing, contexts with great and well-known benefits of charitable giving lessen the differences between vegetarianism and charitable donation. Contexts with small or hardly known benefits increase this perceived gap. Data about giving habits sensitive to such factors might shed light on the role of emotional obstacles in the practical efficacy of ethical reflection.

There is another way in which there can be an emotional obstacle to charitable donation but not to vegetarianism. Early formulations of what is now referred to as Moral Foundations Theory (Graham et al., 2011; Haidt, 2012) posited three domains of values and three psychological roots of these domains. This was known as the CAD triad: each evaluative domain and each psychological root was tied to a particular letter in this acronym (Shweder et al., 1997). The 'A' stands for the autonomy domain of values, associated with anger. The domain of autonomy "relies on regulative concepts such as harm, rights, and justice ... and aims to protect the discretionary choice of 'individuals' and to promote the exercise of individual will in the pursuit of personal preferences" (Shweder et al., 1997, p. 138). The correlative psychological hypothesis is that anger – as opposed to contempt, the 'C' root associated with community values, and disgust, the 'D' root associated with divinity value – underlies our ideas about this domain and is typically experienced with violations of its rules. Let's put aside more recent formulations of this idea (involving more domains and more psychological roots) and consider autonomy and anger. Giving to charity, with the exception of

environmental and especially nonhuman animal charities, involves contributing to the lives of other humans. Many, but not all, of these people count as autonomous, broadly speaking. They have the power, if not the opportunity, to exercise discretionary choice and pursue individual preferences. If there is a psychological sensitivity to autonomy, it involves not just anger at its violation, but respect for its presence and authority. Indeed, such respect is the flip side of anger: if it's important that we become angry at violations of autonomy, we had better have inclinations to notice and respect autonomy lest we become apt targets of anger (and its repercussions) ourselves. Therefore, most of the topics involved in charitable donation can be predicted to activate the psychology of sensitivity to autonomy. By contrast, vegetarianism is psychologically different. Here's my empirical wager: even while many of us will distinguish between more and less autonomous nonhumans, we typically, if implicitly and unconsciously, lump them all together into the nonautonomous category. This is just to say that thoughts about nonhumans do not typically trigger the autonomy emotions that seem to be involved in our thoughts about each other. These autonomy emotions can constitute an emotional obstacle: by hypothesis, we face a psychological barrier with regard to giving to human-involving charities that is not involved in choosing a vegetarian diet.

This hypothesis is as plausible, at most, as the notion of autonomy emotions. Work on the CAD triad has largely been superseded by Moral Foundations Theory, which recognizes at least five domains and associated psychological roots (the exact number varies among exponents). The autonomy domain of values has not been preserved intact in Moral Foundations Theory. While, for example, harm certainly figures as part of the moral landscape according to this view, there is less reason to see in it a distinct psychological tendency to recognize and respect humans as capable of having authority and responsibility over their own lives. Thus, while not outright implausible, the plausibility of the present autonomy emotions hypothesis is weakened to the extent that its empirical and theoretical foundations have been replaced.

3.3. Categorization

So much for emotion; let's turn to other aspects of our psychology. I shall start with categorization. Presumably, the ways in which we frame the things we have to deal with affect subsequent thoughts and actions regarding them. Think about categorizing something as food versus garbage, private versus shared, a threat versus a neighbor, as just a few examples. Two overlapping categorizations are pertinent to the present topic and provide materials for our attempts to explain it. First, suppose that we focus on the life and death aspect of both diet and charities. There are two basic and famous ways in

which we can categorize our involvement in a death. On one hand, we can see ourselves as participating in a killing; alternatively, we can see ourselves as letting someone or something die. The distinction between killing and letting die has been a focal point in applied ethics since at least the 1970s (see, e.g., Rachels, 1975), so presumably, it is psychologically salient for people in general. It might be, at least, part of what accounts for the typical pattern of responses to trolley problems: when we don't see ourselves as killing, but instead, as "merely" letting die, we tend to see intervention as permissible; but we tend to balk at allowing ourselves to kill. Consequentialist moral reasoning tends to cast the killing–letting die distinction as morally irrelevant, whereas deontological patterns of reasoning tend to preserve its moral relevance.

Here is how all of this might apply to the present issue. If we come to conceive of ourselves as participating in the killing of our food, then our psycho-normative sensitivity to the killing–letting die distinction will be brought to bear on it. To the extent that we are more inclined to hesitate at killing than at participating in a non-killing taking advantage of a death, this sensitivity ought to produce reticence to eat that which is killed. The intimate relation that we have to our food makes it relatively easy to see ourselves as participating in killing it. By contrast, the causal chain between us and most humans in need of lifesaving aid makes it a stretch to see ourselves as participating in killing when we refrain from giving to a lifesaving charitable endeavor. At most, it seems that we are letting die. Therefore, the psycho-normative sensitivity to the killing–letting die distinction will tend to frame not giving (at all or more) to such charities as more permissible than killing.

There are at least two potential sources of evidence readily at hand for assessing this hypothesis. First, we should look to the dietary preferences of people who, either historically or currently, slaughter nonhuman animals. These people cannot help but see themselves as killing, I imagine. To the extent that they are not more inclined than others, under the same cultural and socioeconomic conditions, to be vegetarian, the present hypothesis lacks support. I do not have good information about this class of people, but my impression is that they are not more vegetarian than those who do not slaughter nonhumans. Second, we should canvass avowed consequentialists who endorse the moral irrelevance of the killing–letting die distinction. To the extent that these people are nonetheless vegetarian, then we should think that some other psychological phenomenon is at work; and indeed, broadly utilitarian cases for vegetarianism have been fairly prominent in the last couple of generations of discussion of the moral status of nonhumans. Peter Singer is once again relevant (see Singer, 1975, 2011, Chapter 3 on animal rights, pp. 181–184 on killing, letting die). Again, the available evidence suggests that this version of categorization hypothesis is not especially plausible.

The broader version of this sort of hypothesis widens the scope of the purportedly germane distinction from killing versus letting die to doing versus allowing harm. As with the killing–letting die distinction, this broader distinction has received explicit philosophical attention (e.g., Bennett, 1995; Woollard, 2015). It, again, seems relevant to the pattern of results found in trolley cases. Greene (2013, pp. 240–245) has argued for why it should be psychologically salient: it’s computationally tractable to keep track of what we do, and hence, of the harms that we perform, but it’s essentially impossible to track all of the things that we do not do, and hence, of all the harms that we allow through inaction. This does not mean that the distinction is morally relevant, of course, but the present issue is psychological, not moral. If we come to see our dietary practices as involving the bringing about of harm, then we should trigger the subsequent thought and action tendencies associated with harming. For many, this will be a psychological obstacle, in part or in whole, to continued participation in such practices. It’s much more difficult to see our inaction with regard to charity as active harming; it’s more likely to be categorized as allowing harm, thereby activating other thought and action tendencies, ones that typically function as less of an obstacle to going on as we were. However, so far as I can tell, the slaughterers and utilitarians are just as relevant here as they were to the previous categorization hypothesis, making it comparably implausible (*prima facie*, given my rough assessment of this evidence).

Killing, and so on, are not the only materials from which to make a hypothesis about categorization, of course. Let’s turn, instead, to the ways in which we think about minds. Daniel Dennett (1987) famously explained this in terms of the “intentional stance.” A stance is an explanatory approach to something; it can be taken deliberately or deployed intuitively. The intentional stance involves the interpretation of things in terms of psychological states with mental content, such as beliefs and desires. To see something via the intentional stance is to see it as a planner, a reasoner, a chooser. Dennett contrasts interpretation using these ideas with the kinds of explanations and predictions that happen when we deploy the cognitive resources of either the “physical stance,” which involves seeing something as a physical object, or the “design stance,” which involves seeing something in terms of functions which it has been produced to perform (either by a minded maker or by mindless nature). Evidence from neuroscience suggests that these stances rely on specific and separate neural networks (Martin & Weisberg, 2003; Robbins & Jack, 2006). Other work indicates that the intentional stance can be triggered automatically by simple suggestions of goal-directed behavior, the understanding of which seems to require the positing of such mental states as beliefs, desires, goals, plans, intentions, and the like (Arico et al., 2011; Hamlin et al., 2007; Heider & Simmel, 1944).

More recent empirical work has added another stance. Robbins and Jack (2006) argue that we also explain each other in terms of the “phenomenal stance.” When we adopt the phenomenal stance toward something, we understand it by attributing a conscious life of feelings, emotions, and other phenomenologically salient states (see Theriault & Young, 2014, pp. 108–110 for a reflective review of this, related evidence). Much the same distinction has been made in terms of “dimensions” of mental state attributions. Kurt Gray and colleagues argue that mental state attributions are made along “agency” and “experience” dimensions (e.g., K. Gray et al., 2011; K. Gray & Wegner, 2009; K. Gray et al., 2012; H. M. Gray et al., 2007). Of particular importance to present concerns is a study in which researchers asked subjects to rate the mental capacities of a variety of things, including both humans and nonhumans. Subjects were asked such questions as whether a five-year-old human was more or less likely to feel pain than something else, such as a chimpanzee (H. M. Gray et al., 2007, p. 619). Patterns in responses were analyzed, resulting in the identification of the agency and experience dimensions. “Experience” included such items as hunger, pain, pride, and joy; overall, the items linked under “experience” accounted for 88% of variance in responses. Another 8% was linked to “agency” factors, which included such things as self-control, memory, and planning. The overall result is that there is both theoretical and empirical reason to think that we understand minds in two psychologically distinct ways. One of these involves classic propositional attitudes such as belief and desire, whereas the other involves affective states.

The researchers working on the psychology of mental state attributions link the ways in which we cognize minds to moral cognition. Gray and colleagues argue that the experiential dimension of mental state attributions is central to our seeing-minded creatures as bearers of rights, whereas the agency dimension is central to attributions of responsibility (e.g., K. Gray et al., 2012). Sytsma and Machery (2012) dispute some of these details and claim that both the agential and experiential dimensions, the intentional and phenomenal stances, are needed for the attribution of moral standing, in the sense of the recognition of things as bearing rights.³ Either way, there is empirical reason to think that, at least, the resources of the phenomenal stance are needed, psychologically, to see something as mattering morally. This yields material for explaining why ethicists eat their greens, but, relative to others, don’t give more to charity. If ethical education leads to more people adopting the phenomenal stance toward nonhumans, then it should also increase the number of people who see nonhumans as morally relevant. The general mechanism for this is aptly put in terms of categorization: people undergoing this change in outlook now categorize nonhumans as experiencing affective states (at least). This categorization change leads to another: the nonhumans are moved from the “morally irrelevant” category

to the “morally relevant” one. People without education in ethics will tend to adopt the phenomenal stance toward nonhumans in relatively smaller numbers, and hence, they will be more likely to see them as not mattering morally. Such a categorization change would not change our charitable inclinations.

While it may well be the case that seeing things as being bearers of rights is intimately tied to our propensity to attribute affective states to them, the present hypothesis is nonetheless implausible. It is unlikely that ethical reflection of the sort cultivated in post-secondary ethical education would bring about a change in tendency to adopt the phenomenal stance toward nonhumans. We have already seen why: this topic – the psychology of nonhumans – is not likely to be a particularly central part of an ethics education. Such an education will be primarily concerned with such things as values, the nature of moral reasons to act and not to act, and the classification of character traits as virtues or vices, to name just a few things. The nature of nonhuman minds would fall into the category of empirical information relevant to the direct topics of ethical expertise, rather than being constitutive of such expertise itself. Moreover, it is unlikely that those without an education in ethics lack the tendency to adopt the phenomenal stance toward nonhumans. This is an empirical issue about which I lack data, but such an assumption strikes me as the appropriate assumed baseline. Data about changes in pet ownership might be relevant here: my impression is both that it has increased (in North America) in the last generation, and that people are more interpersonally (if that’s the right word) engaged with their pets than prior generations. If this is correct, then it suggests no general lack of ability among people in general to see nonhumans as having such mental states as affective ones. The more likely effect is that an ethical education will lead people to see nonhumans and/or their experiences as morally relevant. This is less aptly put in terms of recategorization of nonhumans from outside of the domain of the phenomenal stance to inside of it. Instead, it’s better understood as an adjustment internal to this stance. Still, it’s a promising idea, so let’s turn to some hypotheses that give specific content to it.

3.4. Reason

Some might have found the earlier sections oddly off topic. After all, the overarching topic of inquiry in the whole series of studies is whether ethical reflection makes any difference to our behavior. Surely, then, it makes sense to focus on ethical reflection. Indeed, it does, so let’s do so by turning our attention to reason. While emotion has figured prominently in recent moral psychology, reasoning has not been neglected. For instance, Jonathan Haidt (2001) has famously argued for a “social intuitionist” model of moral

judgment. The intuition component – the role of fast processes such as emotional ones – gets most of the attention, but the social component is our present concern. Simply put, Haidt argues that moral judgments tend to be quickly produced by emotion, but that they are often shaped by subsequent reasoning processes. The reasoning in question tends to take place interpersonally rather than privately. While such reasoning often aims at defending the products of intuition rather than at assessing whether these products are true or, indeed, worth defending, it sometimes modifies or even rejects these products. The crucial point, for present purposes, is that Haidt's model offers some reason to think that social reasoning about a topic changes the relations people have to their ideas about this topic. Whether it deepens or erodes their endorsement, interpersonal reasoning has psychological effects. Let's call this the "social performance" view of the psychological power of reasoning.

We can make the social performance view a bit more precise. When people are called upon to reason interpersonally about their already existing beliefs and behavior about a topic, this will have psychological repercussions. While the reasoning will often entrench the ideas and habits already in place, sometimes it will lead to changes in them. This yields a prediction: imagine two groups of people. Suppose that ideas about topic X, either *yea* or *nay*, are distributed fifty-fifty in these two groups. Suppose that the people in one group perform public interpersonal reasoning about X. This can be expected to produce, eventually, at least a small change in the distribution of endorsement of X in this group. However, by hypothesis, there is no reason to expect a change in such distribution in the other group – or, alternatively, no reason that does not apply to both groups.

In the light of this, we can construct a social performance reason-based hypothesis about ethical reflection, ethicists, and vegetarianism. A post-secondary education in ethics is going to include, pretty much by definition, performance of interpersonal reasoning about some range of topics. It will be very common but not strictly necessary for this range to include diet and whether it is morally permissible to eat nonhumans. Therefore, any group of ethicists can be expected to have gone through the belief- and habit-changing powers of the social performance of reasoning with regard to this issue. By contrast, a much smaller proportion of any group of non-ethicist philosophers can be expected to have gone through such reasoning processes. Until very recently, the proportion of non-philosophers that we could expect to have gone through such social reasoning about vegetarianism would have been even smaller; dietary concerns and fashions are presently changing, however. Therefore, even if the people now in these three groups had the same distribution of dietary tendencies before they went through the post-secondary education system, the effects of their educations on their dietary beliefs and habits should be different. In

particular, it would not be surprising to find a pattern much like the one found by Schönegger and Wagner.

There is much that I like about the social performance hypothesis.⁴ However, it is deficient in the present context. While it might explain differences in belief and action with regard to vegetarianism among the three groups, it does not explain why ethicists eat their greens but don't give more to charity. After all, the responsibilities of the relatively rich to the relatively poor can be expected to be a topic of interpersonal reasoning in an ethics education at least as much as diet. It is both of long-standing pedigree (e.g., Kant, 1997 considered such aid as an example of an imperfect duty to others in *Groundwork*) and of present concern (e.g., it is at the heart of discussion of effective altruism [MacAskill, 2017]). Since we have no reason to expect exposure of the beliefs and habits about these topics to interpersonal reasoning pressures to differ from those about nonhumans, we have no grounds, here, to expect the effects of these pressures on these beliefs and habits to differ across these topics; yet they do. The social performance hypothesis seems inadequate to explain the difference in behavior of ethicists regarding vegetarianism and charity.

There is another way in which the social performance of reasoning can be used to devise a hypothesis about ethical reflection. Let's return to the topic of the previous section: categorization. Suppose that the general idea that the ways in which we categorize topics have effects for subsequent thoughts and actions about those topics is plausible (as it surely is). Suppose, also, that Haidt is correct that moral reasoning's natural home is interpersonal. The social nature of such reasoning can be surmised to have effects on the categorization of topics. Consider your beliefs about two topics – one you expect never to discuss with anyone else; the other you expect to discuss with others. Indeed, suppose that you expect not only to discuss that topic with other people at some time, but that you expect that such conversation is likely to recur. These conversations will involve probing of your tastes regarding the topic and of your ideas about the worth of your involvement with this topic. I surmise that people with these sorts of beliefs will tend to have two different sorts of stances toward them. We will tend not to reason about, nor even to be prepared to reason about, the first sort of topic, unless circumstances prompt us to. By contrast, we will be much likely to be prepared to reason interpersonally about the second sort of topic, and a way of being prepared to reason interpersonally about some topic is to reason about it personally, in advance of such public activity. Let's call this the "social categorization" hypothesis.

Now, once again, consider ethicists, vegetarianism, and charity. While both topics will probably be the subject of social reasoning as part of an ethics education, they differ in other ways. Students can be expected to sit through lectures about these topics and, maybe, to write something about

them, and maybe, but less certainly, to say something about them in their classrooms. This much is in common between the two topics. Outside of the classroom, however, there are important differences. Giving – or not giving – to charity is easily and usually done privately. While it might be done in a face-to-face manner, it is more often, and in more significant amounts of money, done using an envelope or a website in the absence of other people. Of course, people sometimes give money to people on the street or at the door, and of course, some people participate with their time and effort in charitable endeavors. Nevertheless, much more is done privately rather than socially (I expect – I’m making an empirical wager here). Eating, however, while often done privately, is also often done interpersonally. Pretty much all of us eat very often with people we know well, and semi-often with people we don’t know well. Joint breaking of bread is probably the most famous way for people to get together to learn about each other a little bit. Shared lunchtimes among coworkers are renowned for developing collegial camaraderie. While we eat, one thing we’re likely to talk about is eating. It’s not the only thing, of course, but it is a pretty likely topic. According to the social categorization hypothesis, charitable donation and diet are likely to be categorized differently according to our expectations about interpersonal conversation. We can expect it to be more likely that we will talk with others about diet, on more than one occasion, than about charity. Hence, we should be more apt to categorize eating as something we’ll need to reason about than charity. This will invite preparation for such social reasoning, which is likely to take the form of pre-conversation solitary reasoning.

These points hold, in general, but now consider the education of an ethicist. Two things can be added. First, the interpersonal conversation about eating in the educational setting while eating can be expected to be about its ethical dimensions more than it would be among, for example, data processors. Second, the conversation can be expected to be more regular, more sustained, and more probing about these aspects, partly because of the interests in ethics of the interlocutors, partly because that’s what academic conversation is like, perhaps especially among philosophers (see, e.g., Brandreth, 2012; Leonard, 2013). All of this gives us reason to think that the social aspects of reasoning will have effects on the beliefs and habits of ethicists’ dietary choices that they do not have on their charitable donation tendencies.

I do not see an easy way to test the social categorization hypothesis. Historical sources might be the best place to look for data. If we have records of the experiences of people who have gone through a social change characterized by a topic that was once treated as a private concern becoming a matter of public interest, then we would have information about what happens in both thought and behavior when people recategorize something from “private,” or at least, “not social” to “social.” Esthetic tastes and

religious beliefs strike me as the right sort of topic, but there could easily be others. Without this or other data, I'm inclined to rule this hypothesis in as somewhat plausible. Much turns on the adequacy of the view of reasoning as predominantly social. If this is not the case, then we lose the present reason to expect these downstream differences in categorization.

4. Conclusion

Why do ethicists eat their greens? We don't yet know. Does ethical reflection make a difference in behavior? In some ways, no, in other ways, maybe yes. As studies of behavior are performed, my bet is that we'll need more fine-grained analysis of the relations between moral reasoning and behavior than has been performed as part of this work so far.

The social categorization hypothesis is explicitly presented here as a hybrid – it brings together concerns that I have elsewhere divided under the labels 'categorization' and 'reason.' This indicates a certain amount of artifice in my divisions throughout. For example, the discussion of the so-called "autonomy" emotions clearly involves categorization; the psychological resources of the phenomenal stance involve both emotion and reason; and more, I'm sure. To a certain extent, we should blur these divisions. Still, we should not blur them altogether. The specific hypotheses canvassed above are not all equally plausible. The ones that look best at this point involve generic emotional support, perception of cost as a source of emotional obstacles, and social categorization (albeit extra provisionally compared to the others). The rest look less promising, but more data are needed for firmer conclusions.

Those with particular theoretical commitments might grumble about my array of relatively more plausible hypotheses. None of the three casts reasoning as directly motivationally powerful. One explicitly relies on emotion for motivational oomph, and the other two get reason in only indirectly. This range of candidates hardly looks like fodder for a vindication of the conative capacities of ethical reflection, understood as centrally involving reason. True enough, I suppose. At this early point, it strikes me as misguided to insist on reason being productive of action in a very specific manner. For one thing, I'm inclined toward psychological pluralism (see Sneddon, 2011), so, for instance, I'm happy with a broad Humeanism about motivation for at least some topics, such as vegetarianism. Such pluralism is apt to keep in mind for the fine-grained analysis of action tendencies in the light of ethical reflection, and perhaps, for the varied results that could follow in the wake of the Schönegger and Wagner study. I'm also inclined to think that the project of shedding light on the powers of reason (and emotion, and more) is one best pursued through a combination of conceptual analysis, empirical data collection, and explicit psychological

theorizing. In other words, it's a hard job to be done after paying attention to empirical details, not before.

The re-blurring of divisions among emotion, reason, categorization, and more can take a different form. More complex hypotheses putting together two or more of these psychological states and/or processes while preserving their conceptual distinctness could be formed. Fruitful complex hypothesis formation needs more data than there are for the present topic, so I won't construct any examples. The result would have implications for our understanding of, well, understanding. My initial worry about relying on "understanding" as a psychological phenomenon was that the idea is too vague to shed any light on the psychological differences among ethicists, other philosophers, and non-philosophers. However, if we articulate conceptual differences between emotion, reason, and so on, and try to find the best theoretical and empirical reasons to explain the ways in which we think and act in terms of specific relations among these phenomena, we may earn ourselves the right to claim to have charted the nature of understanding, and thereby, to have vindicated the very idea rather than turned our backs on it.

Notes

1. Strictly speaking, an even more fine-grained question is appropriate: why is this pattern of judgment and behavior found for German-speaking ethicists but, apparently, not for English-speaking ones? In what follows, I will put aside cultural-linguistic specificity and ask only about the psychology of ethicists, but the ruminations should be understood to apply first to this specific subset and only second to a wider group of ethicists. In this, I follow Schönegger and Wagner, whose own speculation about what might account for their results is not framed in cultural or linguistic terms. I discuss this in the next section. The more that German-speaking ethicists turn out to have an education and a profession integrated with English-speakers, the less significant the linguistic and cultural differences.
2. What follows is a tour of some psychological hypotheses, but note, "psychological" does not here mean "individualistic." In [Section 3.4](#), explicitly social, and hence widely distributed, processes are considered. See [Sneddon \(2007, 2008, 2011\)](#) for discussion of some forms that externalist hypotheses can take. That said, there might be externalist possibilities not well represented by the array of topics addressed below. In particular, what might be thought of as "cultural" possibilities – albeit ones realized in professional "cultures" such as the study of ethics – are worthy of exploration. My thanks to an anonymous reviewer for pressing the significance of this point.
3. The work of Buckwalter and Phelan ([2013](#)), [Buckwalter & Phelan \(2014\)](#)), which suggests that the phenomenal and intentional stances are not as dissociable as other work suggests, accords with this position.
4. No surprise here: see [A. Sneddon \(2007, 2008, 2011\)](#)).

Disclosure statement

No potential conflict of interest was reported by the author.

Notes on contributor

Andrew Sneddon is professor of philosophy at the University of Ottawa. He studies ethics and philosophical psychology. His books include *Like-Minded: Externalism and Moral Psychology* (MIT 2011), *Autonomy* (Bloomsbury 2013), and *Offense and Offensiveness: A Philosophical Account* (Routledge, forthcoming).

References

- Arico, A., Fiala, B., Goldberg, R. F., & Nichols, S. (2011). The folk psychology of consciousness. *Mind & Language*, 26(3), 327–352. <https://doi.org/10.1111/j.1468-0017.2011.01420.x>
- Bennett, J. (1995). *The act itself*. Oxford University Press.
- Brandreth, B. (2012, February). How to talk to a philosopher. *Wired*. Retrieved May 14, 2019, from <https://www.wired.co.uk/article/talk-to-a-philosopher>
- Buckwalter, W., & Phelan, M. (2013). Function and feeling machines: A defense of the philosophical conception of subjective experience. *Philosophical Studies*, 166(2), 349–361. <https://doi.org/10.1007/s11098-012-0039-9>
- Buckwalter, W., & Phelan, M. (2014). Phenomenal consciousness disembodied. In J. Sytma (Ed.), *Advances in experimental philosophy of mind* (pp. 45–74). Bloomsbury Academic.
- Dennett, D. (1987). *The intentional stance*. MIT Press.
- Graham, J., Nosek, B. A., Haidt, J., Iyer, R., Koleva, S., & Ditto, P. H. (2011). Mapping the moral domain. *Journal of Personality and Social Psychology*, 101(2), 366–385. <https://doi.org/10.1037/a0021847>
- Gray, H. M., Gray, K., & Wegner, D. M. (2007, February). Dimensions of mind perception. *Science*, 315(5812), 619. <https://doi.org/10.1126/science.1134475>
- Gray, K., Jenkins, A. C., Heberlein, A. S., & Wegner, D. M. (2011). Distortions of mind perception in psychopathology. *Proceedings of the National Academy of Sciences of the United States of America*, 108(2), 477–479. <https://doi.org/10.1073/pnas.1015493108>
- Gray, K., & Wegner, D. M. (2009). Moral typecasting: Divergent perceptions of moral agents and moral patients. *Journal of Personality and Social Psychology*, 96(3), 505–520. <https://doi.org/10.1037/a0013748>
- Gray, K., Young, L., & Waytz, A. (2012). Mind perception is the essence of morality. *Psychological Inquiry*, 23(2), 101–124. <https://doi.org/10.1080/1047840X.2012.651387>
- Greene, J. (2008). The secret joke of Kant's soul. In W. Sinnott-Armstrong (Ed.), *Moral psychology vol 3: Neuroscience of orality: Emotion, brain disorders, and development* (pp. 35–80). MIT Press.
- Greene, J. (2013). *Moral tribes: Emotion, reason, and the gap between us and them*. Penguin Books.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108(4), 814–818. <https://doi.org/10.1037/0033-295X.108.4.814>
- Haidt, J. (2012). *The righteous mind: Why good people are divided by politics and religion*. Pantheon Books.
- Hamlin, J. K., Wynn, K., & Bloom, P. (2007). Social evaluation by preverbal infants. *Nature*, 450(7169), 557–559. <https://doi.org/10.1038/nature06288>

- Hauser, M. (2006). *Moral minds: How nature designed our universal sense of right and wrong*. Harper Collins.
- Heider, F., & Simmel, M. (1944). An experimental study of apparent behavior. *The American Journal of Psychology*, 57(2), 243–259. <https://doi.org/10.2307/1416950>
- Jack, A. I., Robbins, P., Friedman, J. P., & Meyers, C. D. (2014). More than a feeling: Counterintuitive effects of compassion on moral judgment. In J. Sytma (Ed.), *Advances in experimental philosophy of mind* (pp. 125–180). Bloomsbury Academic.
- Kant, I. (1997). *Groundwork of the metaphysics of morals* (M. Gregor, Trans.). Cambridge University Press.
- Leonard, H. (2013). How to talk to philosophers. *Youtube*; posted Sept. 11, 2013; <https://www.youtube.com/watch?v=3KbEXnGfA2U>. <https://www.youtube.com/watch?v=3KbEXnGfA2U> May 14, 2019
- MacAskill, W. (2017). Effective altruism: Introduction. *Essays in Philosophy*, 18(1), Article 1. <https://doi.org/10.7710/1526-0569.1580>
- Martin, A., & Weisberg, J. (2003). Neural foundations for understanding social and mechanical concepts. *Cognitive Neuropsychology*, 20(3–6), 575–587. <https://doi.org/10.1080/02643290342000005>
- Nichols, S. (2004). *Sentimental rules: On the natural foundations of moral judgment*. Oxford University Press.
- Prinz, J. (2007). *The emotional construction of morals*. Oxford University Press.
- Rachels, J. (1975). Active and passive euthanasia. *New England Journal of Medicine*, 292(2), 78–80. <https://doi.org/10.1056/NEJM197501092920206>
- Robbins, P., & Jack, A. I. (2006). The phenomenal stance. *Philosophical Studies*, 127(1), 59–85. <https://doi.org/10.1007/s11098-005-1730-x>
- Schönegger, P., & Wagner, J. (2019). The moral behavior of ethics professors: A replication-extension in German-speaking countries. *Philosophical Psychology*, 32(4), 532–559. <https://doi.org/10.1080/09515089.2019.1587912>
- Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning? Order effects on moral judgment in professional philosophers and non-philosophers. *Mind & Language*, 27(2), 135–153. <https://doi.org/10.1111/j.1468-0017.2012.01438.x>
- Schwitzgebel, E., & Cushman, F. (2015). Philosopher’s biased judgments persist despite training, expertise and reflection. *Cognition*, 141, 127–137. <https://doi.org/10.1016/j.cognition.2015.04.015>
- Schwitzgebel, E., & Rust, J. (2009). The moral behavior of ethicists: Peer opinion. *Mind*, 118 (472), 1043–1059. <https://doi.org/10.1093/mind/fzp108>
- Schwitzgebel, E., & Rust, J. (2010). Do ethicists and political philosophers vote more often than other professors? *Review of Philosophy and Psychology*, 1(2), 189–199. <https://doi.org/10.1007/s13164-009-0011-6>
- Schwitzgebel, E., & Rust, J. (2011). *The self-reported moral behavior of ethics professors*. <http://www.faculty.ucr.edu/~eschwitz/SchwitzAbs/EthSelfRep.htm>
- Schwitzgebel, E., & Rust, J. (2014). “The moral behavior of ethics professors: Relationships among self-reported behavior, expressed normative attitude, and directly observed behavior. *Philosophical Psychology*, 27(3), 293–371. <https://doi.org/10.1080/09515089.2012.727135>
- Shweder, R. A., Much, N. C., Mahapatra, M., & Park, L. (1997). The “big three” of morality (autonomy, community, divinity) and the “big three” explanations of suffering. In A. M. Brandt & P. Rozin (Eds.), *Morality and health* (pp. 119–169). Routledge.
- Singer, P. (1972). Famine, affluence and morality. *Philosophy & Public Affairs*, 1, 229–243.
- Singer, P. (1975). *Animal liberation*. Random House.
- Singer, P. (2011). *Practical ethics, third edition*. Cambridge University Press.

- Smetana, J. (1981). Preschool children's conceptions of moral and social rules. *Child Development*, 52(4), 1333–1336. <https://doi.org/10.2307/1129527>
- Smetana, J. (1993). Understanding of social rules. In M. Bennett (Ed.), *The child as psychologist* (pp. 111–141). Harvester Wheatsheaf.
- Smetana, J. (1999). The role of parents in moral development: A social domain analysis. *Journal of Moral Education*, 28(3), 311–321. <https://doi.org/10.1080/030572499103106>
- Sneddon, A. (2007). A social model of moral dumbfounding: Implications for studying moral reasoning and moral judgment. *Philosophical Psychology*, 20(6), 731–748. <https://doi.org/10.1080/09515080701694110>
- Sneddon, A. (2008). The depths and shallows of psychological externalism. *Philosophical Studies*, 138(3), 193–208. <https://doi.org/10.1007/s11098-006-9058-8>
- Sneddon, A. (2011). *Like-minded: Externalism and moral psychology*. MIT Press.
- Sommers, T. (2016). *A very bad wizard: Morality behind the curtain, second edition*. Routledge.
- Sytsma, J., & Machery, E. (2012). The two sources of moral standing. *Review of Philosophy and Psychology*, 39(3), 303–324. <https://doi.org/10.1007/s13164-012-0102-7>
- Theriault, J., & Young, L. (2014). Taking an “intentional stance” on moral psychology. In J. Sytsma (Ed.), *Advances in experimental philosophy of mind* (pp. 101–124). Bloomsbury Academic.
- Tiberius, V. (2015). *Moral psychology: A contemporary introduction*. Routledge.
- Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge University Press.
- Turiel, E. (1998). The development of morality. In W. Damon & N. Eisenberg (Eds.), *Handbook of child psychology, 5th edition* (pp. 863–932). John Wiley & Sons.
- Turiel, E. (2003). Resistance and subversion in everyday life. *Journal of Moral Education*, 32(2), 115–130. <https://doi.org/10.1080/0305724032000072906>
- Woollard, F. (2015). *Doing and allowing harm*. Oxford University Press.