

Experimental ethics, intuitions, and morally irrelevant factors

Peter Königs

forthcoming in *Philosophical Studies*

1. Introduction

Experimental ethics continues to challenge conventional notions of moral cognition. This essay is about experimental studies that suggest that our moral intuitions respond to factors that lack moral significance.

In one study, Joshua Greene and colleagues have offered a solution to the notorious trolley problem. Rising from the armchair from which trolleyological research has traditionally been carried out, they have conducted experiments to find out which principles govern our intuitions in trolley dilemmas.¹ They found that deontological intuitions are triggered by the conjunction of two factors. One factor is personal force. A dilemma involves personal force if “the force that *directly* impacts the other is generated by the agent’s muscles” (e.g. when one pushes the heavy person off the footbridge).² The other factor is the intention to kill the victim as a means to saving the other people (as in the footbridge case) rather than as a side-effect (as in the standard switch case, in which hitting a switch redirects the trolley onto another track).

Another experimental study by Greene and Jay Musen, which was inspired by Peter Singer’s drowning child scenario, suggests that people’s sense of moral obligation towards people in need varies depending upon spatial distance. Mere spatial distance determines the extent to which we feel morally obliged to, for instance, rescue a dying child.³

In yet another study, Eric Uhlmann and colleagues found that some people’s readiness to sacrifice a small number of people for the greater good depends on the ethnicity of the to-be-sacrificed individuals and on their nationality. Liberals, but not conservatives, were found to be more likely to give consequentialist responses to a sacrificial dilemma when the to-be-sacrificed individual was white than when he was black. A different effect was found for conservatives: They, but not liberals, deemed civilian collateral deaths of a military strike more morally defensible when the victims were Iraqis than when they were Americans.⁴

¹ Greene et al., 2009. The study offers a solution to the descriptive part of the trolley problem (which principles govern our responses to trolley scenarios?), not to its moral part (which responses are correct?).

² Greene et al., 2009, p. 365.

³ Musen & Greene, MS; see Singer, 1972.

⁴ Uhlmann, Pizarro, Tannenbaum, & Ditto, 2009.

What these studies have in common is that they reveal that people's case-specific moral intuitions are sensitive to factors that lack intrinsic moral significance.⁵ We respond differently to moral scenarios due to the presence of (what seem to be) morally irrelevant factors, such as personal force, distance, ethnicity or nationality.⁶

An observation made by Greene nicely captures the idea of identifying morally irrelevant factors: "Were a friend to call you from a set of trolley tracks seeking moral advice, you would probably not say, 'Well, that depends. Would you have to push the guy, or could you do it with a switch?'"⁷ The same test can be used for any other factor the relevance of which we would like to determine.

While experimental investigation into which factors of a scenario elicit our intuitions is still in its infancy, it has been suggested that findings of this sort possess considerable normative significance. More precisely, it has been argued that such findings debunk the case-specific intuitions that have been found to be sensitive to morally irrelevant factors. The fact that we have different intuitions about similar cases due to the presence or absence of a factor that does not justify having different intuitions means that we should distrust these intuitions. Most prominently, Greene has appealed to empirical findings of this sort to attack deontology.

One objective of this paper is to offer a general analysis and assessment of the merits of such empirically informed arguments from moral irrelevance, thereby contributing to the debate surrounding the normative significance of experimental ethics.⁸ But I will also keep an eye specifically on Greene's attempt to use arguments from moral irrelevance to refute deontology. We will see that the success of arguments from moral irrelevance depends crucially on how we conceive of the relationship between case-specific intuitions and intuitions at a higher level of generality. I will argue that such arguments are dialectically useless if we assume that case-specific intuitions are, as a rule, subordinate to intuitions at a higher level of generality (section 3). If we do not make this assumption, the dialectical force of arguments from irrelevance increases. But the prospects of particularly sweeping debunking challenges, such as the one put forth by Greene, decrease (section 4). Finally, I will offer a discussion specifically of the experimental aspect of arguments from moral

⁵ For another study of this sort, see Gino, Shu, & Bazerman, 2010.

⁶ The moral (ir)relevance of intention, the second of the two factors identified by Greene and colleagues, is less clear (see Greene, 2014, pp. 720-721). Note also that Greene and Musen's study was challenged by Nagel & Waldmann, 2013. They found that the apparent distance effect is really due to various confounding variables, such as informational directness, group membership and the relative efficaciousness of one's helping efforts. But as Greene rightly points out, at least some of these factors are morally irrelevant, too (Greene, 2013, p. 378 n261).

⁷ Greene, 2016a, p. 176.

⁸ See Berker, 2009; Greene, 2010, 2014; Kahane, 2013; Kumar & Campbell, 2012; Rini, 2013; Sauer, 2012. I am not in this article concerned with the normative significance of experimental *neuroscientific* findings for moral theory (on this see Berker, 2009; Kamm, 2009; Königs, 2018a).

irrelevance. Arguments from moral irrelevance based on experimental findings such as the above have been presented as examples of how experimental methods can advance moral theory. But, as I will argue, this way of vindicating the normative significance of experimental moral psychology is not unproblematic (section 5). I will begin this essay with a more general characterization of arguments from moral irrelevance (section 2).

One final preliminary remark is in order: Whether a given factor really is morally irrelevant is, of course, to some extent debatable. While I am inclined to agree that the above-mentioned factors are morally irrelevant, I will not in this article defend their moral irrelevance. I am primarily concerned with general methodological issues, which are independent from whether any of the above-cited factors really are morally irrelevant or not.⁹

2. Arguments from moral irrelevance

Normally, when we have different intuitions about similar moral cases, we take this to indicate that there is a moral difference between these cases. This is because we take our intuitions to have responded to a morally relevant difference. But if it turns out that our case-specific intuitions are responding to a factor that lacks moral significance, we no longer have reason to trust our case-specific intuitions suggesting that there really is a moral difference. This is the basic logic behind arguments from moral irrelevance.

Two different types of moral intuitions play a role in arguments from moral irrelevance. The target of such arguments are case-specific intuitions, that is, intuitions about what is the right thing to do in a concrete case. An argument from moral irrelevance might, for instance, target our case-specific intuitions in the switch case and the footbridge case. But they also rely on an intuition about a moral principle at a higher level of generality, namely on an intuition about whether a given feature of a scenario matters from a moral point of view. For instance, an argument from moral irrelevance might rest on the intuition that the involvement of personal force in a trolley scenario does not matter from a moral point of view. An intuition of this sort does not, on its own, tell us what to do in a concrete case, but it tells us whether two cases, e.g. the switch case and the footbridge case, should be treated differently due to the presence or absence of this feature.¹⁰

Arguments from moral irrelevance come in two versions, a precise and an imprecise version. They differ in what conclusion is drawn from the finding that we respond differently to scenarios due to the presence or absence of a morally irrelevant factor. A

⁹ Note that Uhlmann and colleagues confirmed in a pre-test that ethnicity and nationality are typically not considered morally relevant.

¹⁰ The distinction between these two types of intuitions is common in the literature, see e.g. Greene, 2014, p. 724; Kagan, 1998, pp. 13-14; Kamm, 1993, pp. 5-7; McMahan, 2013; Sandberg & Juth, 2011, p. 213.

finding of this sort implies that we have two different case-specific intuitions that cannot both be sound, because we know that they are responding to a factor which does not justify having two different responses. But it does not yet tell us which of the two case-specific intuitions is the problematic one.

A precise argument from moral irrelevance is precise in that it states which of the two case-specific intuitions is to be dismissed. This requires an additional assumption as to how the intuitional conflict is to be resolved. A precise argument from moral irrelevance proceeds by asserting that one of the two case-specific intuitions can be assumed to be sound. A more formal statement of such an argument looks as follows:

P1. Empirical evidence suggests that we have different case-specific intuitions (I_1 and I_2) about similar moral cases due to the presence or absence of some factor F .

P2. Factor F is morally irrelevant.

C1. We have different case-specific intuitions (I_1 and I_2) about similar moral cases due to the presence or absence of a morally irrelevant factor.

P3. If we have different case-specific intuitions (I_1 and I_2) about similar moral cases due to the presence or absence of a morally irrelevant factor, these intuitions cannot both be sound.

C2. I_1 and I_2 cannot both be sound.

P4. I_1 is sound.

C3. I_2 is not sound.

The experimental findings concerning the factors that our case-specific intuitions are responsive to, such as personal force, are contained in premise 1. These case-specific intuitions (at least one of them) are the target of the argument. Premise 2 is a conventional normative premise stating that the factor that triggers the case-specific intuitions is morally insignificant. It rests on a general, rather than case-specific, intuition about which properties matter from a moral point of view. Premise 3 states the core idea behind arguments from moral irrelevance explained above. We usually regard our having different case-specific intuitions about similar moral cases as showing that these cases differ morally, because we assume that our case-specific intuitions are responding to such a morally relevant difference. But if we learn that our case-specific intuitions are responsive to a morally irrelevant factor, which does not justify having different case-specific intuitions, there must be something wrong with one of these case-specific intuitions. We can draw the imprecise conclusion C2 that our case-specific intuitions cannot both be sound. Premise 4 is the additional premise that allows drawing the more

precise conclusion C3, which states that one rather than the other case-specific intuition is the culprit.¹¹

Greene's argument against deontology is an example of a precise argument from moral irrelevance. It specifies exactly how the intuitional conflict is to be resolved, namely in favor of consequentialism. The mere finding that our case-specific responses are sensitive to the morally irrelevant factor personal force is, as it were, symmetrical. It does not yet entail that the deontological response ("It is impermissible to shove the person off the footbridge.") rather than the consequentialist one ("It is permissible to hit the switch.") is flawed. It merely means that our case-specific responses to scenarios involving or lacking personal force should be the same (C2).¹² But Greene seems to take it that our consequentialist case-specific intuitions in scenarios without personal force – e.g., that it is permissible to hit the switch – can be assumed to be correct (P4). The consequentialist judgment about scenarios without personal force is regarded as a moral fixed-point, so to speak.¹³ Given then that personal force is not a morally relevant consideration, the deontological case-specific intuition in personal dilemmas must be erroneous. The deontological case-specific intuition does not give us reason to move away from the consequentialist default. Greene claims "that once all of the inner workings of our judgments are revealed by science, there will be nothing left for deontologists. All of the factors that *push us away from consequentialism* will, once brought into the light, turn out to be things that we will all regard as morally irrelevant."¹⁴ Greene, that is, treats consequentialist case-specific responses as the moral default any deviation from which needs to be justified. His argument from moral irrelevance is meant to establish that deontological case-specific intuitions fail to provide reasons to deviate from consequentialism as they are triggered by morally irrelevant factors. For the purposes of his argument, he defines deontological intuitions as intuitions "in favor of characteristically deontological conclusions (e.g., 'It's wrong despite the benefits')" and consequentialist ones as those "in favor of characteristically consequentialist conclusions

¹¹ I am here expanding and, I hope, improving on previous characterizations of this type argument (Berker, 2009, p. 321; Kumar & Campbell, 2012; Sauer, 2018, p. 43).

¹² As previously observed by Kumar & Campbell, 2012, pp. 317-318.

¹³ At one point, Greene seems reluctant to explicitly endorse this assumption (Greene, 2014, p. 713), but he *has* to if the argument is to be a precise argument from irrelevance against deontology. Due to the mentioned symmetry, he cannot attack the deontological intuition directly by claiming that it, but not the consequentialist response, is triggered by the irrelevant factor. For whenever our case-specific responses vary in response to an irrelevant factor, *both* responses are sensitive to this irrelevant factor. The deontological intuition is triggered by the irrelevant fact that the victim is pushed rather than killed by hitting a switch (or black rather than white). But the consequentialist intuition is likewise triggered by the irrelevant fact that the victim is killed by hitting a switch rather than pushed (or white rather than black).

¹⁴ Greene, 2010, p. 21, my emphasis.

(e.g. ‘better to save more lives’).¹⁵ He thus seeks to challenge deontology by undermining the case-specific intuitions that favor deontology over consequentialism.

Imprecise arguments from moral irrelevance entail a more modest conclusion. Proponents of imprecise arguments from moral irrelevance content themselves with observing that the targeted case-specific intuitions cannot simultaneously be sound while leaving open which of the case-specific intuitions must be given up. An imprecise argument from moral irrelevance looks just like a precise one, except that it stops at C2.

Such a more modest variation of Greene’s approach has been defended by Victor Kumar and Richmond Campbell. They object to its precise version on the grounds that the necessary additional assumption (P4) is too controversial. As they point out, some people have denied that the consequentialist case-specific response to the switch scenario is correct, instead considering it evident that one must not push the heavy person off the footbridge.¹⁶ If one shares this view, one would have to conclude that one has no reason to move away from the *deontological* case-specific judgment. In light of such controversies, it is not clear which way to resolve the conflict between the pair of conflicting case-specific intuitions. Kumar and Campbell therefore offer a qualified defense of Greene’s approach, emphasizing that empirical research can indeed be relevant to moral theory: “Empirical studies can indicate that what accounts for our divergent responses to apparently similar cases does not justify those responses, and therefore that we should withhold from drawing a moral distinction between the cases.”¹⁷ But we cannot draw a more precise conclusion along the lines of C3.

Despite the difficulty pointed out by Kumar and Campbell, not all precise arguments from moral irrelevance are doomed to fail. Cases in which it is clear in which way such an inconsistency is to be resolved are certainly conceivable. Whether we can construct a precise or only an imprecise argument from moral irrelevance depends on the plausibility of the necessary supporting assumption (P4).

The general idea behind arguments from moral irrelevance is not new. Singer’s suggestion that there is no morally relevant difference between a starving child in a remote country and a nearby drowning child as well as Peter Unger’s book-length elaboration of Singer’s insight are two notable applications of this approach.¹⁸ Greene’s reasoning is clearly

¹⁵ Greene 2008, p. 39. Elsewhere, he defines deontological judgments “as ones that are naturally justified in deontological terms (in terms of rights, duties, etc.) and that are more difficult to justify in consequentialist terms”, and consequentialist judgments “as ones that are naturally justified in consequentialist terms (by impartial cost-benefit reasoning) and that are more difficult to justify in deontological terms” (Greene, 2014, p. 699). In this paper, I go along with Greene’s definition.

¹⁶ Kumar & Campbell, 2012, p. 318.

¹⁷ Kumar & Campbell, 2012, p. 322.

¹⁸ Singer, 1972; Unger, 1997. Singer’s argument may be an example of a successful *precise* argument from moral irrelevance. Surely, no one would want to suggest that we should resolve

inspired by their work. But the current revival of this approach is different in that it is driven by experimental investigations into which factors trigger our responses. Earlier such arguments were armchair-based.¹⁹ Regina Rini, another advocate of this approach, asserts that the primary use of experimental moral psychology for normative moral theory is precisely that of informing such arguments from moral irrelevance: “[E]mpirical investigation allows us to identify psychological factors that influence our moral judgments, yet which we do not reflectively regard as surviving normative abstraction.”²⁰

3. Asserting the priority of general intuitions over case-specific intuitions

The target of arguments from moral irrelevance are case-specific moral intuitions, that is, intuitions about what is the right thing to do in a particular case. Greene, for instance, attacks the intuitions that it is morally acceptable to let faraway children starve by not donating to charity and that one ought not push the heavy person off the footbridge to stop the trolley. Arguments from moral irrelevance do not target moral principles at a higher level of generality. To be sure, Greene ultimately seeks to refute deontology as a general theory. But he proceeds by debunking case-specific deontological intuitions such as the above rather than deontological principles at a higher level of generality.

Arguments from moral irrelevance are thus ‘liberationist’ in spirit. They suggest that we should discount many of our case-specific intuitions. The liberationist approach contrasts with the ‘preservationist’ approach, which seeks to preserve our case-specific intuitions. Unger, who coined these terms and who is himself a chief proponent of liberationism, explains:

On [the] *Liberationist* view, folks’ intuitive moral responses to many specific cases derive from sources far removed from our Values and, so, they fail to reflect the Values, often even pointing in the opposite direction. So even as the *Perservationist* seeks (almost) always to *preserve* the appearances promoted by these responses, the *Liberationist* seeks often to *liberate* us from such appearances.²¹

Precise and imprecise arguments from moral irrelevance are both liberationist in this sense, the only difference being that the former specify exactly which case-specific intuitions we should distrust. Arguments from moral irrelevance ‘liberate’ us from case-

the conflict by concluding that we must not save the drowning child. See Campbell & Kumar, 2012 for a related discussion.

¹⁹ Unger has, however, “[i]nformally and intermittently [...] asked many students, colleagues and friends” for their intuitions (Unger, 1997, p. 31).

²⁰ Rini, 2013, p. 267. Notice, though, that the morally irrelevant factors that Rini discusses also include factors that are external to the moral scenarios themselves, such as framing effects and psychological manipulation. I return to this point below.

²¹ Unger, 1997, pp. 11-12.

specific intuitions by showing that they clash with an intuition at a higher level of generality. A general intuition about which properties are morally relevant (P2) is invoked in an attempt to undermine case-specific intuitions. This means that for an argument from moral irrelevance to go through, the intuition at the higher level of generality must have priority over the case-specific ones.²² By this I mean that the more general intuition must be taken to be more reliable than the case-specific ones. If the case-specific intuitions were taken to be more reliable than the general intuition, we would have to question the latter's correctness (that is, P2) rather than the former's, when these intuitions clash. For instance, rather than to reconsider our case-specific intuitions about the strength of our obligation to help nearby and faraway people in need, we would have to conclude that spatial distance *is* a morally relevant factor after all, rejecting premise 2. The question concerning the relative status of case-specific and general intuitions is hence of great importance. Proponents of arguments from moral irrelevance can construe the relationship between case-specific and general intuitions in two different ways.

One option, which I discuss in this section, would be to provide some principled rationale for why we should always give priority to general intuitions over case-specific ones, whatever this rationale might be. According to this view, case-specific intuitions can, as a rule, be assumed to be significantly less reliable than general intuitions, or indeed completely unreliable. Whenever there is a clash between case-specific intuitions and a general intuition, it is always the latter that prevails, even when the former are intuitively very compelling.

The other option, considered in more detail in the next section, would be to hold that how much confidence we should have in case-specific and general intuitions is simply a function of the intrinsic strengths of these intuitions. As ethical intuitionist Michael Huemer explains, “[s]ome appearances are stronger than others – as we say, some things are ‘more obvious’ than others – and this determines what we hold on to and what we reject in case of a conflict.”²³ Which intuitions we should trust would depend primarily on their strengths, not their level of generality, and the way to adjudicate a conflict would be to attend to the relative strengths of these intuitions.²⁴

²² See Bruni, Mameli, & Rini, 2014, p. 170; Ernst, 2007, p. 136; Kamm, 2007, p. 417.

²³ Huemer, 2005, p. 100; similarly DePaul, 2006, pp. 599-600. Huemer points out elsewhere, however, that case-specific intuitions may eventually turn out to be more susceptible to debunking explanations (Huemer, 2008, p. 383).

²⁴ There is the third possibility of claiming that case-specific intuitions have priority over general ones. Preservationists like Frances Kamm are associated with this view. Ethical particularists, who do not believe in moral principles in the first place, are naturally inclined towards this view, too. While in principle conceivable, this would entail that arguments from moral irrelevance do not work at all, as they require that an intuition at a higher level of generality can override the case-specific intuitions. The third possibility is therefore not an option for proponents of arguments

I will first consider the first option, that of holding that there is some principled reason why case-specific intuitions should be considered less reliable than more general intuitions. Howard Nye, for instance, contends that case-specific intuitions have no justificatory force at all, only the ancillary function of clarifying and illustrating more general principles.²⁵ Distrust towards case-specific intuitions is also common among utilitarian thinkers. Singer seems to endorse the priority of general judgments in his famous *Famine, Affluence, and Morality*. People's case-specific verdicts militate against his view that indifference towards the plight of people in remote countries is morally unacceptable. But "the way people do in fact judge has nothing to do with the validity of my conclusion. My conclusion follows from the principle which I advanced earlier, and unless that principle is rejected, or the arguments shown to be unsound, I think the conclusion must stand, however strange it appears."²⁶ Unger's development of Singer's argument is also informed by the liberationist view that general judgments should correct our case-specific judgments, rather than vice versa.²⁷ Building upon this tradition, proponents of arguments from moral irrelevance could maintain that intuitions at a higher level of generality should be taken to be more reliable than case-specific ones *as a matter of principle*, whatever the rationale behind this principle might be.

The problem with this first way of understanding the relationship between the two types of intuitions is that it would defeat the purpose of arguments from moral irrelevance. Arguments from moral irrelevance would then target precisely the sort of intuitions that are claimed to be rather irrelevant.

Before I explain this in more detail, it is worth noting that particularly ambitious debunkers, who seek to undermine a very large swathe of case-specific intuitions, might in fact be forced to embrace this conception of the relative status of general and case-specific intuitions. Greene, for instance, hopes that *all* case-specific deontological intuitions can ultimately be dismissed as sensitive to morally irrelevant factors. As mentioned earlier, he suggests 'that once all of the inner workings of our judgments are

from moral irrelevance. These three ways of understanding the relation between general and case-specific intuitions are also distinguished by Kagan, 1998, pp. 13-14; Kamm, 1993, pp. 5-7.

²⁵ Nye, 2015.

²⁶ Singer, 1972, p. 236. The principle Singer refers to is: "[I]f it is in our power to prevent something very bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it." (Singer, 1972, p. 231). This raises an interpretative question: Is the duty to give to charity entailed by our obligation to save the drowning child and the moral irrelevance of spatial distance? Or is it directly entailed by the above principle (which would render the other argument obsolete)? I won't address this interpretative question here (refer e.g. to Nye, 2015, p. 630). Singer is inspired by Henry Sidgwick (1981), another utilitarian who favors intuitions at a high level of generality.

²⁷ Unger, 1997. Another skeptic about case-specific intuitions is Shelly Kagan (1989, pp. 13-15; 2016)

revealed by science, there will be nothing left for deontologists. All of the factors that push us away from consequentialism will, once brought into the light, turn out to be things that we will all regard as morally irrelevant.’ But this very ambitious plan is only promising if we have some principled reason to maintain that case-specific intuitions should not be taken very seriously. For if how much trust we should place in intuitions depended instead on how strong, or how ‘obvious’, these intuitions are, it would be a surprise if the case-specific deontological intuitions would always be overridden by the conflicting general intuitions. Surely, we should expect that at least some case-specific deontological intuitions are so compelling as to override conflicting general intuitions.

Things would be different if Greene’s aim were to undermine only the small set of case-specific intuitions that have already been shown by him and his colleagues to be responsive to irrelevant factors. It would then suffice for him to claim that *some* general intuitions (about the irrelevance of personal force or spatial distance) are stronger, or ‘more obvious’, than *some* case-specific ones (e.g. that one must not shove the heavy person off the footbridge). In my view, this is not an implausible claim to make, which is why Greene’s argument deserves to be taken seriously. But his larger aim is ultimately to undermine *all* case-specific deontological intuitions that support deontology and count against consequentialism. And such case-specific deontological intuitions are quite numerous, as are the factors that they are responsive to. Whichever ethical debate we consider – abortion, cloning, eugenics, euthanasia, lying, organ sales, prostitution, (capital) punishment, etc. – we find that people have strong case-specific intuitions that support characteristically deontological conclusions rather than consequentialist ones. People intuit that it is just plain wrong to gene-edit babies to immunize them against HIV, to give away one’s kidney for money, and so forth. It may turn out that such deontological intuitions are ultimately responsive to factors that we would not regard as morally relevant. This remains to be seen. But even if they are, it can hardly be expected that, whenever there is a conflict between intuitions, the general intuitions about the factors’ moral irrelevance are always more compelling or ‘obvious’ than the targeted case-specific ones. It is just very likely that sometimes the opposite is the case. To be sure, Greene’s argument *could* go through, but this would come as a surprise, given how many case-specific deontological intuitions there are.²⁸ This, then, is why ambitious debunkers like Greene might in fact be forced to make the strong assumption that general intuitions

²⁸ This is not to deny that there are many uncontroversial cases, in which Greene’s deontological/consequentialist distinction does not apply. He explains: “When it comes to uncontroversial moral questions, these terms have little meaning. Disapproving of child abuse is both ‘characteristically consequentialist’ and ‘characteristically deontological’, or neither – take your pick. [...] These two ‘characteristically’ labels are not very meaningful outside the context of moral dilemmas in which considerations about rights and duties, at least superficially, appear to conflict with an impartial cost-benefit analysis.” (Greene 2016b, p. 179) But ethical inquiry is all about truly controversial questions that involve precisely such a conflict, and there are therefore plenty genuine ‘characteristically deontological’ intuitions that challenge consequentialism.

should be assumed to be more reliable as a matter of principle, independently from their intrinsic intuitive force.

Now, the problem with this assumption is that it would take the sting out of arguments from moral irrelevance. It makes mounting arguments from moral irrelevance easier, but at the same time it renders them dialectically ineffective. To see why, note that for arguments from moral irrelevance to be dialectically forceful, it must be assumed that the targeted intuitions play a crucial role in the construction of moral theories. For instance, the fact that Greene attacks deontology by attacking case-specific deontological intuitions commits him to the view that these case-specific intuitions constitute the decisive evidence for deontology. It would be pointless to attack case-specific deontological intuitions if deontology rested primarily on more general intuitions. This has also been noted by Kumar and Campbell, who point out that Greene's argument rests on the hidden "assumption [...] that the principal evidence for moral theories is our first order intuitions about concrete cases. One moral theory is more justified than another principally insofar as it better explains and systematizes our first order intuitions."²⁹ Presumably, what Greene has in mind is deontological theory that relies heavily on evidence from thought experiments involving the sort of dilemmas examined by Greene. The prototypical deontologist of this sort is Frances Kamm, who has practiced and encouraged the study of case-specific intuitions like no other. Her deontological views are based on a large set of carefully examined case-specific intuitions, and they are precisely the sort of case-specific intuitions that Greene seeks to undermine in an attempt to refute deontology.

But the assumption that case-specific intuitions constitute the principal evidence for deontology is in direct tension with the assumption that case-specific intuitions, as a rule, are much less reliable than more general ones. If the latter are more reliable than the former, it is them that provide the principal evidence for moral theories, rather than the unreliable case-specific ones. Thus, assuming that general intuitions trump case-specific ones makes it easy to construct arguments from moral irrelevance and to maybe even challenge entire swathes of case-specific intuitions, such as all deontological intuitions. But it undermines the very force of these arguments, as they target intuitions that, ex hypothesi, should not play an important role in the construction of moral theories to begin with. It does not render these arguments invalid, but dialectically toothless. What Greene would have to do instead is attack deontological intuitions at a higher level of generality, such as, say, that the separateness of persons, the signing of contracts, or the giving of promises are morally relevant factors.³⁰ It would be them that are crucial for the justification of deontology, not the case-specific intuitions that he is actually attacking.

²⁹ Kumar & Campbell, 2012, p. 313.

³⁰ See also Berker, 2009, p. 325. Note that these intuitions possess independent intuitive plausibility, unlike perhaps the doctrine of double effect (Greene, 2014, p. 721).

To see the problem more clearly, consider how a deontologist might respond to Greene's challenge. A deontologist who accepts that case-specific intuitions are less reliable than general intuitions need not be too concerned about the finding that case-specific deontological intuitions are responsive to morally irrelevant factors. Her endorsement of deontology is motivated by intuitions at a higher level of generality anyway, precisely because she considers them more trustworthy than case-specific ones. By contrast, a deontologist who rests her case for deontology on case-specific intuitions, such as Kamm, will be loath to accept the assumption that we should always listen to our general intuitions rather than our case-specific ones when they are shown to clash.

It is, however, important to observe that there is a not yet mentioned type of argument from moral irrelevance that this second response (rejecting the priority of general intuitions) cannot defuse. The focus of the present article are arguments from moral irrelevance that identify morally irrelevant factors that are *internal* to the moral scenarios. For instance, the involvement of personal force in a trolley scenario is an internal feature of the scenario itself. By contrast, some morally irrelevant factors that our intuitions are sensitive to are *external* to the moral scenarios. There are, for instance, studies suggesting that moral intuitions are responsive to the order of presentation, to hypnotic suggestion and to whether one is seated in a clean or dirty environment.³¹ These findings differ from the ones mentioned at the beginning of this paper. They, too, identify morally irrelevant factors that our intuitions are sensitive to, but these factors are not features of the scenarios themselves. Rather, they pertain to the *presentation* of the dilemmas and are in this sense external. Now, a preservationist response to findings of this sort strikes me as hopeless. No one could plausibly insist on the reliability of *these* case-specific intuitions, arguing that the sensitivity to external factors of this sort is indicative of *these* factors' moral significance. Surely, no one would want to claim that, say, order effects indicate that the order of presentation must be morally relevant. The moral irrelevance of such external factors is beyond doubt. This means that arguments from moral irrelevance that show that our intuitions are triggered by morally irrelevant *external* factors are more promising than regular arguments from moral irrelevance.³²

4. Taking seriously both case-specific and general intuitions

³¹ Liao, Wiegmann, Alexander, & Vong, 2012; Petrinovich & O'Neill, 1996; Schnall, Haidt, Clore, & Jordan, 2008; Schwitzgebel & Cushman, 2012, 2015; Tobia, Chapman, & Stich, 2013; Wheatley & Haidt, 2005; Wiegmann, Okan, & Nagel, 2012. In light of the replication crisis, findings like these should be taken with a grain of salt (see e.g. (Landy & Goodwin, 2015).

³² As already mentioned, many but not all of the examples discussed by Rini (2013) concern external factors. On the unreliability of moral intuitions due to sensitivity to morally irrelevant external factors, see also Sinnott-Armstrong, 2008.

Returning to debunking arguments from morally irrelevant *internal* factors, we found that the first way of conceiving of the relation between case-specific and general intuitions leads to an impasse. Arguments from moral irrelevance that debunk case-specific intuitions, even if successful, would do little to advance moral theory, because they would target the wrong set of intuitions. The alternative option, anticipated above, would be to take seriously both types of intuitions and to attend to the intrinsic force of intuitions to settle conflicts between intuitions at different levels. We must then decide on a case-by-case basis which way the tension between case-specific intuitions and general intuitions is to be resolved. In some cases, particularly strong general intuitions may prompt us to dismiss conflicting case-specific intuitions. In other cases, case-specific intuitions may be so compelling as to require a revision of more general principles. What matters, then, is how compelling or ‘obvious’ the individual intuitions are, not their level of generality as such. This is a plausible method of moral theory construction, which a significant portion of ethicists are explicitly or implicitly committed to. It is also in line with the method of reflective equilibrium, which takes seriously judgments at different levels of abstraction and allows that judgments of both types may have to be revised in light of the others.³³ And it means that arguments from moral irrelevance may be dialectically effective and advance moral theory. Their targets, case-specific intuitions, play an important role in the construction of moral theories, as they are regarded as one important source of ethical knowledge. This makes these intuitions worthwhile targets for debunkers.

Of course, arguments from moral irrelevance that target case-specific intuitions that are not very compelling to begin with again risk to be uninteresting, since these intuitions will not play a major role in the justification of moral claims. The power of arguments from moral irrelevance increases with the strength of the targeted intuitions. It is best to target strong and therefore frequently appealed-to case-specific intuitions and to show that they clash with even stronger intuitions at a higher level of generality. One would then advance moral theory by targeting dialectically important (namely strong and frequently invoked) moral intuitions.

A sweeping debunking program of the sort proposed by Greene is now, however, unlikely to go through for reasons that are already familiar. Greene may be able to demonstrate that *some* intuitively compelling case-specific deontological intuitions must be relinquished as they clash with even more compelling intuitions at a higher level of generality and plausible consequentialist judgments. Indeed, Greene’s argument so far

³³ In fact, even Kamm, the leader of the preservationist camp, acknowledges, if only by lip-service, that a moral principle derived from case-specific intuitions may stand in need of further validation: We must “consider the principle on its own, to see if it expresses some plausible value or conception of the person or relations between persons. This is necessary to justify it as a *correct* principle, one that has normative weight, not merely one that makes all of the case judgments cohere.” (Kamm, 2007, p. 5, see also pp. 346, 379). In practice, Kamm shows relatively little interest in whether a principle considered on its own is plausible (Nye, 2015, p. 627).

may be a case in point. The intuition that one ought not shove the heavy person off the footbridge is an important and compelling moral data point. But the more general principle that personal force is morally irrelevant and the judgment that it is permissible to hit the switch may be sufficiently compelling as to override this intuition.³⁴ But, as noted above, it would be a surprise if *all* case-specific deontological intuitions can eventually be undermined in this way. It is to be expected that there are many very strong deontological case-specific intuitions, which will sometimes force us to reconsider the general principles with which they clash, rather than vice versa.

An intermediate conclusion is that both ways of understanding the relation between case-specific and general intuitions entail problems for very ambitious debunking projects, such as Greene's anti-deontological project. If we assume the strict priority of general intuitions, arguments from moral irrelevance target uninteresting intuitions. If we deny their strict priority, some of the many targeted case-specific intuitions (e.g. deontological ones) must be expected to prevail over conflicting general intuitions. Besides, for Greene, there would still be the problem that deontology is also supported by intuitions at a higher level of generality. Thus, even if Greene manages to undermine some crucial and compelling deontological case-specific intuitions, he would still be only halfway there. All this does not mean that he will not succeed in refuting some deontological views by debunking some case-specific deontological intuitions. But it is more than uncertain whether the method of sifting through our case-specific intuitions and checking whether they are in harmony with more general intuitions can deal a fatal blow to deontology.³⁵

5. On using experimental findings for arguments from moral irrelevance

So far, I have considered the merits of arguments from moral irrelevance without special consideration of their experimental aspect. Much of what was said applies equally to arguments from moral irrelevance that do not rely on experimental findings. It is, after all, possible to identify the factors that our intuitions are sensitive to from the armchair and to use these findings to construct arguments from moral irrelevance, an observation I will return to shortly.³⁶ Recent proponents of arguments from moral irrelevance, however, have presented these arguments as examples of how experimental moral psychology can advance moral theorizing. Greene has been particularly adamant that arguments from moral irrelevance demonstrate "that interesting scientific facts about moral psychology can, when combined with relatively uninteresting normative assumptions, lead us to relatively interesting normative conclusions."³⁷

³⁴ I am here bracketing the problem that intention *does* seem to be a morally relevant factor.

³⁵ Berker, 2009, p. 325; Königs, 2018a.

³⁶ See again Campbell & Kumar, 2012 for a related discussion.

³⁷ Greene, 2014.

In this final section, I will explain why this attempt at vindicating the normative significance of experimental moral psychology is somewhat underwhelming. The use of experimental methods to construct arguments from moral irrelevance is at worst counterproductive and at best helpful but not game-changing.

The experimental studies that have revealed sensitivity to morally irrelevant internal factors have so far mostly investigated laypeople's moral intuitions. Greene's argument from moral irrelevance, being based on these studies, thus targets in the first instance laypeople's intuitions. But targeting laypeople's intuitions reduces the force of arguments from moral irrelevance. Deontologists can retort that it has not been shown that *their* deontological intuitions are sensitive to irrelevant factors, the intuitions of trained philosophers. And it is these expert intuitions that deontological theories are based upon. That is, deontologists can claim that philosophers' and laypeople's intuitions diverge and that the former can be expected to be less susceptible to morally irrelevant factors. In this case, an argument from moral irrelevance that relies on data about laypeople's intuitions will not put much pressure on advocates of the targeted theory

The notion that philosophers are expert intuiters is, of course, controversial.³⁸ The proponent of an argument from moral irrelevance that relies on data about laypeople's intuitions could reject the expertise claim, insisting that laypeople's and philosophers' intuitions are sensitive to exactly the same irrelevant factors. Whether this is the case is difficult to say. There are studies suggesting that philosophers' moral intuitions are equally sensitive to morally irrelevant *external* factors.³⁹ But as of yet there is no evidence concerning their sensitivity to morally irrelevant *internal* factors. To simply take for granted that the folk moral intuitions revealed by surveys are representative of those held by trained philosophers is risky. The expertise claim may itself be fairly speculative, but it possesses some prima facie plausibility. It therefore seems unwise, from a dialectical point of view, to attack a moral view, such as deontology, by attacking the layperson's version of it.

There is a related and somewhat more technical reason why specifically Greene's anti-deontological argument from moral irrelevance should have targeted philosophers' rather than laypeople's intuitions. One element of his argument against deontology is the allegation that more elaborate deontological theories are the output of confabulatory post-hoc rationalization rather than that of unbiased and careful moral reflection.⁴⁰ This allegation is supported by two observations: First, there is ample empirical evidence that

³⁸ See especially Weinberg, Gonnerman, Buckner, & Alexander, 2010, but also e.g. Kauppinen, 2014, p. 295.

³⁹ Schwitzgebel & Cushman, 2012, 2015; Tobia, Buckwalter, & Stich, 2013; Tobia, Chapman, et al., 2013; Wiegmann, Horvath, & Meyer, forthcoming.

⁴⁰ Greene, 2008; 2014, p. 718.

people routinely engage in confabulatory post-hoc rationalization.⁴¹ Second, more elaborate deontological theories tend to confirm our deontological gut reactions: “For example, [...] there is a complicated, highly abstract theory of rights that explains why it’s okay to sacrifice one life for five in the *trolley* case [i.e. the standard switch case] but not in the *footbridge* case, and it *just so happens* that we have a strong negative emotional response to the latter case but not to the former.”⁴² Greene reckons that the best explanation of this striking coincidence is that deontologists are really just rationalizing their gut reactions. It is important to note, however, that post-hoc reasoning is epistemically problematic only when the relevant intuitions have been shown to be unreliable. If deontological intuitions about different trolley cases have been shown to be unreliable, the deontologist who is engaged in post-hoc reasoning is trying to vindicate intuitions that there is no reason to assume to be correct. By contrast, should the deontological intuitions be reliable, the theories produced by deontologists might well provide accurate systematizations and explications of the intuited truths.⁴³ But this means that Greene’s argument from moral irrelevance, which is supposed to debunk deontological intuitions, must target those intuitions that the more elaborate deontological theories are allegedly rationalizations of. This is necessary for the argument from post-hoc rationalization to work.⁴⁴ And since the philosophers who have developed these elaborate deontological theories must surely be assumed to have rationalized their *own* intuitions, Greene’s argument from moral irrelevance must target the intuitions of these philosophers. Therefore, if Greene’s argument from moral irrelevance undermines laypeople’s deontological intuitions rather than those of deontological philosophers, the attack on more elaborate deontological theories – a pivotal component of his attack on deontology – collapses. This collapse can be prevented by making the assumption that laypeople’s and philosophers’ intuitions are identical and thus sensitive to the same morally irrelevant factors. But this assumption is controversial, and Greene’s argument had better not rest upon it.

All this suggests that building arguments from moral irrelevance on experimental data about laypeople’s intuitions is, in fact, counterproductive. Proponents of arguments from moral irrelevance, especially Greene, would be well-advised to target philosophers’ intuitions rather than those of laypeople. This leaves them with two options. They can either use data gathered experimentally by surveying philosophers’ intuitions or use data gathered in a collaborative armchair-based enterprise by philosophers. Both options capture philosophers’ intuitions, but only the first option is truly experimental, whereas the second is conventional.

⁴¹ See e.g. Uhlmann et al., 2009; Wheatley & Haidt, 2005; Wilson, 2002.

⁴² Greene, 2008, p. 68.

⁴³ See Königs, 2018b, pp. 388-389; Mihailov, 2016, p. 2.

⁴⁴ Greene’s argument from moral irrelevance is, however, only one of several different attempts to debunk deontological intuitions (see Berker, 2009; Königs, 2018a; Paulo, 2019).

The first approach, that of testing philosophers' intuitions experimentally, strikes me as an interesting and worthwhile research project. Just as one has shown in experiments that philosophers are sensitive to irrelevant external factors (e.g. order effects), so one could now set out to demonstrate that this also applies to irrelevant internal factors. But it must be noted that, given the availability of the second option, such experiments are at least not indispensable for the construction of arguments from moral irrelevance. As mentioned above, this conventional method has in fact already been used to construct arguments from moral irrelevance, in particular by Unger. And, it may be added, in the works of such thinkers as Unger and Kamm, the conventional method of investigating our intuitions from the armchair has reached an admirable level of sophistication and precision.

The experimental approach may be credited with improving on the armchair-based one by correcting for biases and individual idiosyncrasies, thereby guaranteeing greater validity and generalizability. This may be true even taking into account that the armchair-based investigation of moral intuitions is often a collaborative effort by the philosophical community rather than an individual one. But experimental arguments from moral irrelevance still represent only a marginal improvement over traditional, non-experimental arguments from moral irrelevance. They do not constitute anything resembling a methodological breakthrough, contrary to what Greene insinuates.

Contrast this with the idea that empirical moral psychology can reveal the evolutionary origins of morality, which would allow the construction of evolutionary debunking arguments. Greene originally proposed such an evolutionary debunking explanation of deontology, but he appears to have abandoned this project in the meantime.⁴⁵ Unlike arguments from moral irrelevance, evolutionary debunking arguments could not possibly be constructed without the help of empirical methods.

As a final side note, things are again slightly different if we consider irrelevant *external* factors. It is difficult to see how philosophers could reliably find out about the influence of irrelevant external factors upon moral judgment from the armchair, precisely because these factors are external to the thought experiments proper. We need systematic and rigorous scientific experiments to discover that, say, hypnotic disgust or framing effects influence our moral intuitions. It seems, then, that experimental methods are much more indispensable to arguments from morally irrelevant *external* factors.

6. Conclusion

My assessment of the merits of empirically informed arguments from moral irrelevance has led to a mixed verdict. Arguments from moral irrelevance will do little to advance

⁴⁵ On this see Königs, 2018a, p. 197.

moral theory if we assume that, as a matter of principle, case-specific intuitions are less reliable than more general ones. They would target precisely the sort of intuitions that moral theorists should not rely on in the first place. By contrast, if we make the plausible assumption that intuitions of both types should be taken seriously, arguments from moral irrelevance may play a valuable role in moral philosophy. But they are less likely to go through, as the case-specific intuitions may trump the more general one. Very ambitious debunking projects, such as the one proposed by Greene, are likely to fail. There is little reason to expect that, say, no deontological case-specific intuition is ever strong enough to prevail over a general one in case of a conflict. Finally, proponents of arguments from moral irrelevance must say more about the added value of building these arguments on experimental findings, if they want to sell experimental arguments from moral irrelevance as a major methodological innovation rather than as a marginal improvement. In fact, Greene's reliance on experimental data about laypeople's intuitions has significant drawbacks.

May experimental findings about which (internal) factors trigger our case-specific intuitions advance moral theory? Yes, but, it seems, only marginally so. Can they be used to refute deontology? Probably not, although they might undermine *some* deontological case-specific intuitions.

Acknowledgments

I wish to thank the anonymous referee, Chiara Brozzo, Christian Seidel, Emilian Mihailov, Hanno Sauer, Irina Schumski, Katharina Brecht, Leo Menges, Michael W. Schmidt, Nora Heinzelmann, Norbert Paulo and audiences in Karlsruhe, Porto, Saarbrücken and Tübingen for their valuable comments.

References

- Berker, S. (2009). The Normative Insignificance of Neuroscience. *Philosophy and Public Affairs*, 37(4), 293–329.
- Bruni, T., Mameli, M., & Rini, R. A. (2014). The Science of Morality and its Normative Implications. *Neuroethics*, 7(2), 159–172.
- Campbell, R., & Kumar, V. (2012). Moral Reasoning on the Ground. *Ethics*, 122(2), 273–312.
- DePaul, M. (2006). Intuitions in Moral Inquiry. In D. Copp (Ed.), *The Oxford Handbook of Ethical Theory* (pp. 595–623). New York: Oxford University Press.
- Ernst, Z. (2007). The Liberationists' Attack on Moral Intuitions. *American Philosophical Quarterly*, 44(2), 129–142.

- Gino, F., Shu, L. L., & Bazerman, M. H. (2010). Nameless + harmless = blameless: When seemingly irrelevant factors influence judgment of (un)ethical behavior. *Organizational Behavior and Human Decision Processes*, 111(2), 93-101.
- Greene, J. (2008). The Secret Joke of Kant's Soul. In W. Sinnott-Armstrong (Ed.), *Moral Psychology: Volume 3: The Neuroscience of Morality: Emotion, Brain Disorders, and Development* (pp. 35–80). Cambridge, MA: MIT Press.
- Greene, J. (2010). Notes on 'The Normative Insignificance of Neuroscience' by Selim Berker (unpublished manuscript).
- Greene, J. (2013). *Moral Tribes: Emotion, Reason, and the Gap Between Us and Them*. New York: Penguin Press.
- Greene, J. (2014). Beyond Point-and-Shoot Morality: Why Cognitive (Neuro)Science Matters for Ethics. *Ethics*, 124(4), 695–726.
- Greene, J. (2016a). Solving the Trolley Problem. In J. Sytsma & W. Buckwalter (Eds.), *A Companion to Experimental Philosophy* (pp. 175-189). Malden, MA: Wiley Blackwell.
- Greene, J. (2016b). Reply to Driver and Darwall. In S. M. Liao (Ed.), *Moral Brains: The Neuroscience of Morality* (pp. 170-181). New York: Oxford University Press.
- Greene, J., Cushman, F. A., Stewart, L. E., Lowenberg, K., Nystrom, L. E., & Cohen, J. D. (2009). Pushing Moral Buttons: The interaction between personal force and intention in moral judgment. *Cognition*, 111(3), 364–371.
- Huemer, M. (2005). *Ethical Intuitionism*. Basingstoke/New York: Palgrave Macmillan.
- Huemer, M. (2008). Revisionary Intuitionism. *Social Philosophy and Policy*, 25(1), 368–392.
- Kagan, S. (1989). *The Limits of Morality*. Oxford: Oxford University Press.
- Kagan, S. (1998). *Normative Ethics*. Boulder: Westviewga.
- Kagan, S. (2016). Solving the Trolley Problem. In E. Rakowski (Ed.), *The Trolley Problem Mysteries* (pp. 151-168). New York: Oxford University Press.
- Kahane, G. (2013). The armchair and the trolley: an argument for experimental ethics. *Philosophical Studies*, 162(2), 421-445.
- Kamm, F. (1993). *Morality, Mortality: Volume 1*. New York: Oxford University Press.
- Kamm, F. (2007). *Intricate Ethics: Rights, Responsibilities, and Permissible Harm*. Oxford: Oxford University Press.
- Kamm, F. (2009). Neuroscience and Moral Reasoning: A Note on Recent Research. *Philosophy and Public Affairs*, 37(4), 330-345.
- Kauppinen, A. (2014). Ethics and Empirical Psychology – Critical Remarks to Empirically Informed Ethics. In M. Christen, C. Van Schaik, J. Fischer, M. Huppenbauer, & C. Tanner (Eds.), *Empirically Informed Ethics: Morality Between Facts and Norms* (pp. 279-305). Cham: Springer.
- Königs, P. (2018a). On the normative insignificance of neuroscience and dual-process theory. *Neuroethics*, 11(2), 195-209.
- Königs, P. (2018b). Two types of debunking arguments. *Philosophical Psychology*, 31(3), 383-402.
- Kumar, V., & Campbell, R. (2012). On the normative significance of experimental moral psychology. *Philosophical Psychology*, 25(3), 311-330.
- Landy, J. F., & Goodwin, G. P. (2015). Does Incidental Disgust Amplify Moral Judgment? A Meta-Analytic Review of Experimental Evidence. *Perspectives on Psychological Science*, 10(4), 518-536.

- Liao, M. S., Wiegmann, A., Alexander, J., & Vong, G. (2012). Putting the Trolley in Order: Experimental Philosophy and the Loop Case. *Philosophical Psychology*, 25(5), 661-671.
- McMahan, J. (2013). Moral Intuition. In H. LaFollette & I. Persson (Eds.), *The Blackwell Guide to Ethical Theory (2nd Edition)* (pp. 103-120). Malden, MA: Blackwell.
- Mihailov, E. (2016). Is Deontology a Moral Confabulation? *Neuroethics*, 9(1), 1-13.
- Musen, J. D., & Greene, J. (MS). Mere Spatial Distance Weakens Perceived Moral Obligation to Help Those in Desperate Need. (Unpublished Manuscript)
- Nagel, J., & Waldmann, M. R. (2013). Deconfounding Distance Effects in Judgments of Moral Obligation. *Journal of experimental Psychology*, 39(1), 237-252.
- Nye, H. (2015). Directly Plausible Principles. In C. Daly (Ed.), *The Palgrave Handbook of Philosophical Methods*. New York: Palgrave Macmillan.
- Paulo, N. (2019). In Search of Greene's Argument. *Utilitas*, 31(1), 38-58.
- Petrinovich, L., & O'Neill, P. (1996). Influence of Wording and Framing Effects on Moral Intuitions. *Ethology and Sociobiology*, 17(3), 145-171.
- Rini, R. A. (2013). Making Psychology Normatively Significant. *The Journal of Ethics*, 17(3), 257-274.
- Sandberg, J., & Juth, N. (2011). Ethics and Intuitions: A Reply to Singer. *The Journal of Ethics*, 15(3), 209-226.
- Sauer, H. (2012). Morally irrelevant factors: What's left of the dual process-model of moral cognition? *Philosophical Psychology*, 25(6), 783-811.
- Sauer, H. (2018). *Debunking Arguments in Ethics*. Cambridge: Cambridge University Press.
- Schnall, S., Haidt, J., Clore, G. L., & Jordan, A. H. (2008). Disgust as Embodied Moral Judgment. *Personality and Social Psychology Bulletin*, 34(8), 1096-1109.
- 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.
- Schwitzgebel, E., & Cushman, F. (2015). Philosophers' biased judgments persist despite training, expertise and reflection. *Cognition*, 141, 127-137.
- Sidgwick, H. (1981). *The Methods of Ethics*. Indianapolis: Hackett Publishing.
- Singer, P. (1972). Famine, Affluence, and Morality. *Philosophy and Public Affairs*, 1(3), 229-243.
- Sinnott-Armstrong, W. (2008). Framing Moral Intuitions. In W. Sinnott-Armstrong (Ed.), *Moral Psychology: Volume 2: The Cognitive Science of Morality: Intuition and Diversity* (pp. 47-76). Cambridge, MA: MIT Press.
- Tobia, K., Buckwalter, W., & Stich, S. (2013). Moral Intuitions: Are Philosophers Experts? *Philosophical Psychology*, 26(5), 629-638.
- Tobia, K., Chapman, G. B., & Stich, S. (2013). Cleanliness is Next to Morality, Even for Philosophers. *Journal of Consciousness Studies*, 20(11-12), 195-205.
- Uhlmann, E. L., Pizarro, D. A., Tannenbaum, D., & Ditto, P. H. (2009). The motivated use of moral principles. *Judgment and Decision Making*, 4(6), 476-491.
- Unger, P. (1997). *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford University Press.
- Weinberg, J. M., Gonnerman, C., Buckner, C., & Alexander, J. (2010). Are philosophers expert intuiters? *Philosophical Psychology*, 23(3), 331-355.
- Wheatley, T., & Haidt, J. (2005). Hypnotic Disgust Makes Moral Judgments More Severe. *Psychological Science*, 16(10), 780-784.

- Wiegmann, A., Horvath, J., & Meyer, K. (forthcoming). Intuitive Expertise and Irrelevant Options. In T. Lombrozo, J. Knobe, & S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy (Vol. 3)*.
- Wiegmann, A., Okan, Y., & Nagel, J. (2012). Order effects in moral judgment. *Philosophical Psychology*, 25(6), 813-836.
- Wilson, T. D. (2002). *Strangers to ourselves: Discovering the adaptive unconscious*. Cambridge, MA: Harvard University Press.