

Can Morally Ignorant Agents Care Enough?

Daniel J. Miller

Forthcoming in *Philosophical Explorations*

Abstract: Theorists attending to the epistemic condition on responsibility are divided over whether moral ignorance is ever exculpatory. While those who argue that reasonable expectation is required for blameworthiness often maintain that moral ignorance can excuse, theorists who embrace a quality of will approach to blameworthiness are not sanguine about the prospect of excuses among morally ignorant wrongdoers. Indeed, it is sometimes argued that moral ignorance always reflects insufficient care for what matters morally, and therefore that moral ignorance never excuses. Furthermore, quality of will theorists treat their skepticism about excuses for the morally ignorant as a natural implication of their approach. It is therefore unsurprising that, while many have argued for the blamelessness of certain morally ignorant agents on grounds concerning reasonable expectation, the possibility that morally ignorant agents might be blameless even according to quality of will views has not been adequately addressed. I illustrate and explain how it is possible for morally ignorant agents to display sufficient care for the morally relevant features of their wrong behavior. Thus, even if quality of will views are correct, moral ignorance sometimes excuses.

1. Introduction

Sometimes agents are ignorant of the wrongness of their behavior because they are unaware of some feature of their behavior that makes it wrong. I might share a personal anecdote about a friend on social media because I'm unaware that my friend wished it to remain private. A doctor might prescribe medicine to a patient who is allergic to it because the doctor is unaware of the patient's allergy. This type of ignorance of the wrongness of one's behavior is sometimes called *circumstantial ignorance*.¹ However, sometimes agents are ignorant of the wrongness of their behavior, not because they are unaware of the feature or features of the behavior that make it wrong, but rather because they are unaware that those features are wrong-making.² Such agents suffer from *moral ignorance*.³ Gideon Rosen offers a historical example:

[C]onsider an ordinary Hittite lord. He buys and sells human beings, forces labour without compensation, and separates families to suit his purposes. Needless to say, what he does is wrong. The landlord is not entitled to do these things. But of course he thinks he is (2003, 64-65).

Contemporary instances of moral ignorance are widespread and varied. A child who is indoctrinated in white supremacist ideology may come to believe that non-white people aren't deserving of the same treatment as white people, and behave accordingly. A mafioso who was raised in a sub-culture of crime and violence may consequently come to value fidelity to the 'family' as paramount. As a result, he might regard his own violence toward ordinary civilians as undesirable but justified. Notably, however, not all instances of moral ignorance are paradigmatically villainous. For example, many ordinary people are aware that factory farming inflicts intense suffering to animals on a mass scale, but nevertheless judge such practices to be permissible (perhaps due to a failure to regard animal suffering as a sufficiently strong moral reason to refrain from such practices).

It's possible that the relevant beliefs and practices in at least some such cases have been so ingrained in an individual's life and upbringing, and so widely accepted in their cultural context, that it's unreasonable to expect them to know better. Many argue that, if this is so, then an agent is blameless for their moral ignorance and thus blameless for their resulting behavior (Rosen 2003, Levy 2009, FitzPatrick 2017). However, this judgment conflicts with the increasingly popular view that blameworthiness depends, not on an agent's moral awareness, or whether it would be reasonable to expect an agent to know better, but rather on the quality of will reflected in an agent's behavior. According to the *Quality of Will* view,

NQW: An agent is blameworthy for some behavior *iff* it reflects a negative quality of will.^{4,5}

The Hittite slaveowner's behavior plausibly reflects a negative quality of will, blameless ignorance notwithstanding. In judging that his behavior is morally permissible despite the awareness of its effects upon his slaves, his behavior displays insufficient concern (or due regard) for them. Likewise, the white supremacist's racist behavior reflects a lack of regard for the rights and interests of others that he deems inferior. And involvement in factory farming practices, despite awareness of the immense suffering these practices cause, seems to reflect insufficient concern for the welfare of the animals harmed by these practices.

But does *all* wrong behavior performed from moral ignorance reflect a negative quality of will? While quality of will theorists don't often offer explicitly affirmative answers to this question, neither are they sanguine about the prospect for excuses among morally ignorant wrongdoers. Nomy Arpaly, for example, maintains that morally ignorant agents are blameworthy because they fail to respond to morally relevant considerations (2002, 231ff). For similar reasons, Matthew Talbert writes that 'while ignorance of the circumstances and consequences of one's actions often undermines blame, moral ignorance typically does not do so' (2013, 226). While Talbert's statement leaves it open whether any morally ignorant wrongdoers might be excused, he neither offers examples of this nor a reason to think this is ever the case.

Elizabeth Harman takes the most explicit stance. Harman maintains that moral ignorance itself always reflects a negative quality of will, and thus that behaving from moral ignorance typically reflects this negative quality of will.⁶ More specifically, Harman maintains that moral ignorance always involves insufficient care for the morally relevant features of some morally wrong behavior. In conjunction with the *Quality of Will* view expressed in *NQW*, this leads Harman to maintain the following:

No Excuse: Moral ignorance never excuses an agent for behaving wrongly (2011, 2014).⁷

If *No Excuse* is true, then the morally ignorant agents discussed above are all blameworthy for their behavior, despite their ignorance, and regardless of whether their culture or lack of moral sensitivity makes it unreasonable to expect them to know better. Some will find this a hard pill to swallow. Others will readily accept these implications as coinciding with intuitions that slaveowners, racists, violent criminals, and meat-eaters are blameworthy for their false beliefs and morally ignorant behavior.

But the implications of *No Excuse* extend beyond this. *No Excuse* also seems to imply that anyone who lives in accordance with a commitment to a mistaken ethical theory is also often blameworthy for their beliefs and behavior. Thus, either the compliant Utilitarian, who shirks a prior commitment to prevent some harm, or the dutiful Kantian, who refuses to do so, is blameworthy for his behavior. According to *No Excuse*, whichever of these individuals is incorrect about the importance of the relevant moral features will end up behaving in a way that reflects a negative quality of will (i.e., insufficient concern or regard) with respect to one of those features.⁸ If Utilitarianism is false, then the compliant Utilitarian's behavior reflects a lack of concern for promise-keeping or fidelity, and if Kantianism is false then the dutiful Kantian's behavior reflects lack of concern for harm to others.⁹

But, as anyone who has spent time wrestling with ethical problems can recognize, ethics is hard.¹⁰ Consequently, people of good faith and sincerity may nevertheless come to hold mistaken ethical views. Ignorance in ethical reasoning needn't be the result of laziness, arrogance, or a stubborn refusal to consider opposing views. So, while some may find it intuitive that slaveowners, racists, and meat-eaters are blameworthy for their beliefs and behavior, far fewer will find it intuitive

that *everyone* who behaves impermissibly as a result of an error in moral reasoning is blameworthy for their ignorance and resulting ignorant behavior.

Accepting *No Excuse* involves some serious bullet-biting. And while quality of will theorists seem willing to bite this bullet, it remains to be seen whether they need to. Indeed, focusing on cases of deeply objectionable wrongdoers (e.g., slaveholders, racists) might seem to bolster the intuitive plausibility of *No Excuse*, but only by overlooking the undoubtedly larger subset of cases of moral ignorance that are more common in everyday experience.

While many have argued for the blamelessness of certain morally ignorant agents on grounds concerning reasonable expectation (e.g., Rosen 2003, Levy 2009, FitzPatrick 2017), whether morally ignorant agents might be blameless even according to quality of will views has yet to be adequately addressed. But, how might this be possible, given that agents who behave from moral ignorance seem always to do so with a lack sufficient regard, concern, or care for what matters morally? Isn't *No Excuse* a straightforward consequence of the *Quality of Will* view?

In this paper I argue for a negative answer to this question, illustrating how an agent who behaves from moral ignorance can nevertheless display sufficient care for the morally relevant features of his behavior. If this is correct, then even if *NQW* is true there are cases in which an agent's blameless moral ignorance excuses.¹¹ I proceed to argue that the scope of the relevant type of excuse extends both to agents in (what I refer to as) 'close-call' cases as well as some good faith proponents of false ethical theories.

As I will explain, showing that *NQW* does not entail *No Excuse* requires more than establishing that morally ignorant agents can care about the morally relevant features of their behavior. Rather, it must be shown that such agents can display *sufficient* care for those features. My argument begins from the observation that some morally ignorant agents incorrectly judge that some morally wrong behavior is permissible because they mistakenly weigh the moral considerations for

and against that behavior. This doesn't entail, however, that such agents have insufficient concern for the wrong-making features of their behavior. Instead, it may be that they overestimate the relative importance of the moral considerations that count in its favor. Thus, the behavior of such agents needn't reflect a lack of concern for what matters morally, but instead too much concern for something else that also matters morally.

My discussion of moral ignorance is restricted to cases in which an agent is aware of the wrong-making features of his behavior but nevertheless falsely believes that his behavior is permissible.¹² As I use the term, agents who behave *from* moral ignorance (rather than merely *while* morally ignorant) would have behaved otherwise had they been aware of the wrongness of their behavior (Rosen 2003, Wieland 2017).

2. Understanding the 'No Excuse' View

2.1 De Dicto Care and De Re Care

By employing the distinction between *de dicto* and *de re* care for morality, Elizabeth Harman illustrates how a quality of will approach to blameworthiness might lead one to accept *No Excuse*:

Behaving in a certain way is blameworthy just in case (and to the degree that) the behavior results from the agent's caring inadequately about what is morally significant – where this is not a matter of *de dicto* caring about morality but *de re* caring about what is in fact morally significant (2014, 13, emphasis original).

Following Nomy Arpaly (2003), Harman explains that '[a]n agent cares *de dicto* about morality if the agent wants *to act morally*. An agent cares *de re* about what is in fact morally significant if the agent cares about the features of her actions that actually do matter morally' (2014, 13, emphasis

original). An agent may care *de dicto* about morality while ignorant that his behavior is morally wrong because he fails to care *de re* about what is morally significant. To illustrate, the Hittite slaveowner might desire to avoid wrongdoing but fail to see the suffering of his slaves as a wrong-making feature of his keeping slaves, and thus falsely judge that keeping slaves is morally permissible. Again, many people might desire to live in accordance with morality, but fail to see animal welfare as a sufficiently important moral consideration, and as a result believe that factory farming practices are morally permissible.

Indeed, it would seem that any instance of moral ignorance must reflect insufficient *de re* care about what matters morally. As Harman writes, ‘Believing a certain kind of behavior is wrong on the basis of a certain consideration is a way of caring about that consideration’ (2011, 460). This suggests that failing to believe that some behavior is wrong despite being aware of some wrong-making feature is a way of failing to care about that feature to the degree that morality requires. But if so, then moral ignorance can never excuse an agent from acting from it, since, all else equal, acting from this ignorance would reflect this lack of care.¹³ Thus, Harman maintains that moral ignorance never excuses an agent for acting from it. As she writes, ‘[B]eing caught in the grip of a false moral view is not exculpatory’ (2014, 1).

2.2 *De Re Care and Negative Quality of Will*

Since failures of *de re* care for morality are instances of negative quality of will, the view that theorists like Arpaly and Harman hold concerning the relationship between blameworthiness and an agent’s *de re* care for what matters morally is an instance of *NQW*. For this reason, more general remarks are warranted concerning what it means for an agent’s behavior to reflect a negative quality of will. An agent’s morally wrong behavior, *X*, can reflect a negative quality of will toward a person

with respect to some wrong-making feature F either by reflecting *ill will* or by reflecting a lack of *good will*.¹⁴

An agent's X -ing reflects *ill will* toward a person with respect to some wrong-making feature, F , only if the agent either desires or intends to X (at least partly) because it has F , or if the agent sees F as a reason to X . To illustrate, suppose that Adrian publicly exposes Brooke's secrets, which results in Brooke's humiliation. Adrian's behavior plausibly reflects ill will toward Brooke with respect to her humiliation only if he desires to publicly expose her secrets at least partly because doing so will humiliate her, or if he sees her humiliation as a reason to expose her secrets.

In contrast, an agent's X -ing reflects a *lack of good will* toward a person with respect to some wrong-making feature, F , only if either he lacks a (sufficiently strong) desire that F *not* obtain or fails to see F as a reason *not* to X (where F is an instance of harming another agent or where F is a failure to benefit or fulfill some other positive duty).¹⁵ To illustrate, suppose that Candace tells an insensitive joke about Dante that hurts his feelings. Candace's behavior reflects a lack of good will toward Dante with respect to hurting his feelings only if she fails to see Dante's hurt feelings as a reason not to tell the insensitive joke, or has little (or no) desire to refrain on that basis.¹⁶

My subsequent discussion of morally ignorant wrongdoers focuses primarily on lack of good will, which I discuss in terms of having insufficient (*de re*) care or concern for what matters morally.¹⁷ In the following section I consider an extant attempt to divorce *NQW* from *No Excuse* by appealing to the fact that morally ignorant agents can have (*de re*) care for what matters morally. As I will argue, this attempt is inadequate because it fails to show what is needed, viz., that morally ignorant agents can have *sufficient* care for what matters morally. In Section 4 I argue for this stronger claim, showing that even if *NQW* is true there are cases in which an agent's moral ignorance excuses. In Section 5 I discuss the scope of excuses for morally ignorant agents.

3. An Inadequate Argument Against No Excuse

Jan Willem Wieland (2017) argues against *No Excuse* by maintaining that agents who behave from moral ignorance needn't do so from a negative quality of will.¹⁸ Wieland's argument involves an appeal to a series of cases that culminate in 'case 3**':

Cleo keeps two slaves and forces them to work for her without pay. Cleo is ignorant that this is wrong because she is ignorant that slavery is wrong. She has made a serious attempt to determine whether slavery is wrong, and collected all the non-moral facts about the issue (she knows that they suffer, that she could have been a slave herself if she were unlucky enough, etc.). It is not the case that Cleo keeps slaves because she wants them to suffer. Nor is she indifferent to their suffering: she is aware of it and feels sympathy with the slaves. Still, she did not succeed in drawing the inference that slavery is wrong because of the limited social context, and concluded that it was permissible (156-157).

Wieland argues that, since she neither wants them to suffer nor is indifferent to their suffering, Cleo *de re* cares about what matters morally in this case (i.e., the slaves' suffering), and thus does not have a negative quality of will according to the account maintained by Harman and Arpaly.

Wieland then anticipates an objection from a defender of *No Excuse*: Cleo may care, but she doesn't care *enough*: "Cleo doesn't care enough about her slaves' welfare – if she cared for her slaves in the way that we would regard as morally sufficient, then typically she would treat them in a way that we would regard as morally acceptable" (157). To claim that Cleo doesn't care *enough* in this context seems to imply caring as much as morality demands, which requires that one's degree (and perhaps kind) of care for some consideration is at least proportional to the value or importance of that consideration.

This objection has initial plausibility. Consider a mother who ordinarily provides her son with his basic needs—if his needs don't conflict with her own desires. If she had to choose between

spending money on a new pair of shoes for herself and paying for an annual doctor's appointment for her son, she would spend the money on the shoes, judging it more important to stock her wardrobe than to 'waste' her money on the medical checkup. Reflection upon this case yields a verdict similar to Cleo's case: while it's clear that this mother cares (to some extent) for her son, it's equally clear that her care is morally deficient. It seems that if the mother cared enough for her son's well-being then she would realize that she ought to prioritize the doctor's appointment over her wardrobe, and would behave accordingly. Indeed, it isn't obvious how one might care enough about what matters morally and still act wrongly from moral ignorance.

As Wieland understands it, the objection at hand suggests that acting from a good quality of will requires acting permissibly. Wieland responds:

In my view, this is too strong. Generally, it is not the case that S does X with good will only if X is the morally right or acceptable course of action. There should be cases where good will and wrongful conduct go together, and in my view 3** is a case in point (157).¹⁹

Wieland's response implies that Cleo does care enough in Case 3**: she acts with good will even though her action is morally wrong. However, although Wieland assures us that 3** is a counterexample to *No Excuse*, he doesn't explain how this might be. For example, it's certainly not sufficient for caring enough that Cleo feels the measure of sympathy for her slaves one might feel for a caged animal, or that Cleo desires her slaves to be healthy *slaves*. Accordingly, it's difficult to believe that the objector would be satisfied by what seems little more than an announcement on Wieland's part that the objection fails. An adequate response to the objection at hand must illustrate *how* it's possible for a morally ignorant agent to perform an impermissible action (i) from no ill will

and (ii) without lacking good will. And since it seems that Cleo doesn't care *enough*, 3** does not convincingly illustrate (ii).

Wieland proceeds to argue that moral ignorance itself needn't reflect a negative quality of will. As Wieland explains, Harman's defense of *No Excuse* requires the truth of what he calls 'Necessity':

Necessity: if you care about [some wrong-making feature] Y, then you believe that X is wrong because of Y. (thus, if you care about the slaves' suffering, then you believe that keeping slaves is wrong because of this suffering) (158).

As Wieland puts it, 'it seems plausible to think that one might care about the slaves' suffering without the belief that keeping slaves is wrong' (158). As mentioned above, though, Cleo may care about her slaves (and their suffering) in the way one might care about a pet animal. So, Wieland's claim is obviously true. But its truth is also beside the point, since it doesn't address the objection he raises for his own argument. The relevant objection is *not* that an agent cannot care about some wrong-making feature unless he judges an action to be wrong because of it; a response to *that* objection is easy to come by. Again, caring comes cheap. Cleo may care, but only to a degree that obviously falls short of what morality demands. Rather, the relevant objection is that an agent cannot care *enough* about some wrong-making feature unless he judges an action to be wrong because of it (and refrains from so acting). In the following section I offer a reply to this objection that shows both that moral ignorance sometimes excuses and how this is possible. This result has significant implications concerning the scope of excuses for various types of morally ignorant agents, which I discuss in Section 5.

4. Caring Enough

Although Wieland's case 3** illustrates how behaving from moral ignorance is consistent with caring about the morally relevant features of one's behavior, Wieland fails to show that behaving from moral ignorance is consistent with caring *enough* about those features. For that, we need a clearer case, as well as a more thorough explanation of how such a case is possible. Consider the following:

TRUTH-TELLING: Clark works on a production line at a bottling plant. Through no fault of his own, one of the machines he monitors begins to malfunction. Clark promptly fixes the machine, but the malfunction results in a minor setback in production for the day. Given the normal variability in the daily bottling rate, no one will find out about this unless Clark reports it. Clark knows that his irascible manager (who has a policy requiring employees to report machine-malfunctions) will likely punish Clark by cutting his overtime hours if he finds out. Clark's family is on a tight budget and are already doing everything they can to increase their income. Without Clark's overtime pay, he and his wife won't be able to afford the cost of dental work for their two children, who have toothaches due to cavities. Clark cares very deeply about his children. However, he also believes that it is an exceptionless moral principle that one must be forthcoming with the truth even when the consequences of doing so are severe, and for this reason is strongly committed to truth-telling. Clark can barely stand the thought of doing something that will harm his children. But he sincerely believes that failing to tell his manager the truth would be wrong. He correctly believes that doing the right thing sometimes comes at a cost, and believes that his situation is an instance of this. Given the strength of his moral commitment to truth-telling, Clark tells his manager. His manager cuts his overtime hours, and Clark's children have to go without the much-needed dental work.

Clark must choose between telling the truth and prioritizing his children's welfare. Each of these options is supported by genuine moral considerations. Clark's failure to tell his manager the

truth would be an instance of deception and of breaking an implicit commitment to abide by his manager's policies. On the other hand, Clark has a special obligation to take care of and provide for his children, and failing to prioritize their interests will result in suffering and hardship. Given his special obligation to his children and the fact that withholding the truth from his manager would be a fairly minor instance of deception and breach of commitment (with negligible negative consequences), it's plausible that the considerations in favor of prioritizing his children's interests outweigh those in favor of truth-telling.²⁰ Withholding the truth from his manager is not merely permissible, but obligatory. If so, then Clark acted wrongly from moral ignorance.²¹

Many will be inclined to criticize Clark for his failure of moral judgment. Some who are so inclined will argue that Clark's ignorance is culpable since he could have reasonably been expected to know better. If his ignorance *is* culpable, then it does not excuse his impermissible behavior. But, what exactly does reasonable expectation amount to?

According to a prominent account defended by William FitzPatrick (2008, 2017), it's reasonable to expect Clark to know better only if Clark had opportunities to fulfill certain procedural epistemic obligations. Such epistemic obligations may include Clark rethinking or reflecting upon his moral beliefs, discussing the issue with others who disagree, perhaps reading the work of ethicists who argue against the existence of exceptionless moral principles (like those Clark believes in), and so on. Suppose, however, that Clark is meticulously careful in the management of his moral beliefs in all of these ways. Clark nevertheless ends up finding himself with the view that one must always tell the truth (when failing to do so constitutes deception). There may be additional steps that Clark could take that might reduce the likelihood that he will mistakenly weigh moral considerations in the future. Regardless, it's certainly possible for Clark to fulfill all of his epistemic obligations without thereby coming to avoid or eliminate his moral ignorance.²²

FitzPatrick points out that even those who fulfill their procedural epistemic obligations may nevertheless fail to improve their epistemic position if the way that they fulfill these obligations is distorted by epistemic vices, such as ‘failing, for example, to see the merit in opposing views even while ‘considering them’ at length, due to greed, a self-serving sense of entitlement, overconfidence, arrogance, or cowardly resistance to the prospect of uncomfortable conclusions’ (2017, 43).

FitzPatrick maintains that such agents are culpable for their ignorance, fulfillment of their epistemic obligations notwithstanding. But Clark needn’t have or exercise any epistemic vices in order to maintain his mistaken moral beliefs. We can stipulate that Clark is sincere in his search for the truth, and humbly open to the possibility that he might be wrong (though he is ultimately unconvinced in this matter). Clark’s situation is therefore what FitzPatrick calls a ‘hard case,’ in which ‘people can deliberate thoughtfully, follow procedural norms, and even do all of this in a way that manifests all of the relevant virtues while *still* failing to get to the truth’ (2017, 43, emphasis original). If the conditions on culpable ignorance should be understood in terms of reasonable expectation along these lines, then Clark is blameless for his moral ignorance, and thus blameless for his morally ignorant behavior.

This reasoning will be unpersuasive for theorists like Talbert, Arpaly, and Harman, who maintain that a negative quality of will is necessary and sufficient for culpable ignorance or behavior.²³ Such theorists may contend that Clark’s false moral judgment itself reflects a lack of care for his children’s interests, regardless of whether he has fulfilled his epistemic obligations. Because of this, they might argue, Clark is blameworthy for his moral ignorance and his action performed from this ignorance.

But Clark simply doesn’t lack sufficient care for his children. On the contrary, he agonizes over the prospect of doing what he believes is morally right precisely because doing so will harm his

children. If blameworthiness for acting wrongly requires acting in a way that reflects a negative quality of will, then Clark is blameless for acting from moral ignorance, and *No Excuse* is false.²⁴

Some will object that, since the interests of Clark's children in this case were in fact more important than telling his manager the truth, Clark failed to care *enough* about his children (or assign sufficient weight to them as a source of moral reasons). Accordingly, one might maintain that his belief and behavior therefore reflect a lack of care: his care for his children's interests should have outweighed his care for truth-telling in this case. Of course, since his action is (we have supposed) morally wrong, it's *trivially* true that Clark incorrectly weighed the moral reasons at play. But this doesn't settle the issue at hand. It's possible that Clark cared *enough* for his children's interests, but cared *too much* for truth-telling. Thus, Clark's false belief and morally ignorant behavior don't reflect lack of care for his children, but rather excessive concern with telling the truth.

It may be helpful here to distinguish between two senses of 'caring enough.' First, one might care enough about something when one's degree of care is causally sufficient for recognizing some relevant fact (e.g., that an action is morally wrong). Second, one may care enough in the sense that one cares as much as morality demands.²⁵ Harman seems to assume that if one cares as much as morality demands, then one would care enough in the causal sense. Reflection upon Clark's case shows this assumption to be false. Clark's care for the interests of his children is not causally sufficient to make him aware of the fact that his action is morally wrong, but this does not entail that Clark fails to care as much for their interests as morality demands. Caring as much as morality demands doesn't require that one care enough to judge that one's behavior is morally wrong.

One might instead agree that Clark cared for his children's interests *as much as morality demands*, but deny that he cared for their interests *as much as they can reasonably expect*. But these two standards can come apart only if the standard of reasonable expectation at hand isn't a *moral* standard. If it isn't, it can't be relevant to any sense of reasonable expectation at issue with respect to

moral blameworthiness. While one could maintain that Clark is an appropriate target of some non-moral blame, this is beside the point of this debate, which concerns whether *morally* ignorant wrongdoers can be *morally* blameless according to the quality of will view expressed in *NQW*.

One who is still unconvinced that Clark cares enough is invited to consider Mark, who finds himself in a similar situation. Mark also cares very deeply for his children's interests—at least as much as anyone else does. Mark, however, doesn't care as much for truth-telling as Clark does. Because of this, Mark judges correctly that he should behave in the interests of his children instead of telling his manager the truth, and refrains from telling his manager the truth.

We are free to stipulate that Clark has an equal degree and kind of care for his children as Mark. And, if *anyone* cares enough about one's children, then Mark cares enough (since, by stipulation, Mark cares at least as much as anyone else). But if Mark cares enough for his children, and Clark cares as much as Mark, then it seems to follow that Clark also cares enough for his children. This implies that Clark's action doesn't reflect a negative quality of will; in particular, it doesn't reflect a lack of good will (or insufficient care). Assuming *NQW*, it follows that Clark is blameless for acting from moral ignorance and thus that *No Excuse* is false.

The defender of *No Excuse* has a possible recourse at this point: she can maintain that caring enough about something requires caring enough *relative* to one's other cares, maintaining the following principle:

CER: Caring enough about something *F* requires caring about *F* more than competing factors, where *F* is more important (or a weightier consideration) than these factors.

According to *CER*, Clark doesn't care enough about his children because he doesn't care about them more than telling his manager the truth, and his children's interests are more important than

that. But *CER* certainly isn't obvious, and reflecting upon an analogous case of caring casts doubt on it:

FRIENDSHIP VS. CHARITY: Kairah cares greatly for the poor, often spending her spare time on volunteer work and donating much of her expendable income to charity. However, Kairah is also deeply devoted to her friends, and values bringing them happiness with surprise gifts. This leads Kairah to spend the remaining income from her recent paycheck on an amusing knickknack for a friend instead of spending it on a meal for the impoverished and starving person she noticed upon entering the store.

A natural judgment to draw here is that Kairah cares *too much* about bringing happiness to her friends, since (in this case) her care for this shouldn't exceed her care for the needs of the impoverished person. But why should we accept the judgment that Kairah cares too much about bringing happiness to her friends instead of the alternative judgment that Kairah just cares *too little* for the poor? One reason, of course, is that Kairah's usual behavior consistently reflects a care for the poor that exceeds that of ordinary people. But here's another reason: suppose that if Kairah cared about the poor even more than she does, this would result in other bad decisions (e.g., spending so much time and money helping the poor that she is unable to adequately provide for her own family). These reasons make it plausible that her care for the needs of the poor is at least proportional to their actual importance. If something is askew with Kairah's caring, it's plausibly not her level of care for the poor. It's her caring too much about making her friends happy. But if this is true then *CER* is false, since Kairah's caring enough for the needs of the poor doesn't require that she cares about them more than something less important (e.g., bringing happiness to her friends). And, if *CER* is false, as reflection on this case suggests, then it's possible for an agent like Clark to have sufficient care for his children despite his mistaken judgment that truth-telling is more important than preventing their harm. So again, we have good reason to reject *No Excuse*.

CER is a general principle concerning what is required to care sufficiently for something. Although I've offered reasons to doubt that the requirement it expresses holds across the board, there may yet be cases where caring sufficiently about something does require caring more than competing considerations. I consider this in Section 5.

5. The Sufficient Care Excuse

If the foregoing is correct, then *No Excuse* isn't a consequence of the *Quality of Will* view expressed in *NQW*. Morally ignorant wrongdoers can be blameless even if *NQW* is true, since they can display sufficient care for the morally relevant features of their behavior. Call the sort of excuse at issue (the sort that morally ignorant agents like Clark have) the *Sufficient Care Excuse*. While the Sufficient Care Excuse applies to at least some morally ignorant agents, it's worth asking how wide its scope is.

5.1 Close Call Cases

The Sufficient Care Excuse applies most clearly to what I refer to as 'close call' cases. Close call cases involve competing (sets of) moral considerations that are, objectively speaking, close in terms of their relative moral importance. While disagreement about *which* particular cases are close calls is unavoidable, it's uncontroversial that there are such cases. And it's plausible that, in many such cases, agents come to hold incorrect beliefs about the relative importance of the competing considerations while nevertheless caring sufficiently about each. These agents act from moral ignorance and yet are excused, given *NQW*. For example, one needn't care insufficiently about preventing harm to others simply because one mistakenly judges that, on a particular occasion, telling the truth is more important. And one needn't display a lack of concern with honesty just because one incorrectly believes that a loved one's benefit is slightly more important in a given case.

Again, one might care greatly for the needs of the poor while mistakenly thinking that keeping a promise to a friend matters a bit more in a particular situation. That each of these mistaken judgments would in fact be correct if applied to other very similar cases makes this all the more plausible: in some cases telling the truth *is* more important than preventing the relevant harm, and sometimes keeping a promise to a friend (depending on the nature or content of the promise) *does* matter more than attending to the needs of the poor.

As we've seen, all that it takes to make such a mistake is to place *too much* importance on one consideration (rather than placing too little importance on a competing consideration). The behavior of such agents doesn't reflect a lack of care for what matters morally, but instead too much care for something else that *also* matters morally. If this is correct, then the Sufficient Care Excuse extends to a considerable range of cases where genuinely caring people are mistaken about the moral permissibility of their behavior.²⁶

5.2 Good Faith Proponents of False Ethical Theories

In Section 1 I pointed to a troubling implication of *No Excuse*, namely that all those who live in accordance with their commitment to a false ethical theory are often blameworthy for their actions. Since *No Excuse* is false (on *NQW*), it's worth revisiting this issue. Recall the concern: agents who fail to judge that some moral consideration is important due to the implications their ethical theory seem to display a lack of *de re* care for what matters morally. The Kantian seems to display a lack of concern for harm to others when he judges that it's morally right to keep a promise even though doing so will bring about such harm. The Utilitarian may display a lack of concern for fidelity when he acts according to his judgment that he ought to break a promise to bring about some good outcome. In response to this worry, I will argue that the *Sufficient Care Excuse* plausibly

extends to at least some cases where good faith proponents of false ethical theories behave from moral ignorance, and that not all of these are close calls.

We can begin from the observation that an agent can *care* about something even though he *judges* it to be morally unimportant. According to the account of quality of will I offer in Section 2, an agent's quality of will may involve either a doxastic component (i.e., a belief or judgment that *F* is a reason to *X*, or a reason not to *X*), a conative component (i.e., a desire to *X* or to refrain from *X*-ing because it has feature *F*), or both. This more general point is applicable when discussing quality of will in terms of *de re* care or concern: while believing is *one way* of caring, it's not the only way.

Both Harman and Arpaly seem to embrace this idea. Recall Harman's claim: 'Believing a certain kind of behavior is wrong on the basis of a certain consideration is a *way* of caring about that consideration' (2011, 460, emphasis added). Elsewhere Harman expands upon her view about the relationship between caring and belief in reference to the widely discussed Huck Finn case, where Huck helps his friend Jim escape from slavery despite believing that doing so is morally wrong:

[T]his view takes there to be a psychologically real phenomenon of *caring about* certain features of one's action, which is not identical to *believing* that those features matter morally; sometimes caring about a feature and believing it matters go together; sometimes they come apart. Huck, for example, cares about Jim's humanity but does not believe that Jim's humanity matters morally in the same way that a white person's does (2014: 19).

This helps make sense of Harman's view that 'Huck is *both* praiseworthy for refraining from turning Jim in and blameworthy for his moral view about Jim' (2014, 15, emphasis original). Huck's morally right behavior reflects sufficient care for what matters morally, despite his moral ignorance. However, imagine an inversion of this case where, just as conflicted about his decision, Huck

behaves in accordance with his false moral belief that helping Jim escape would be wrong, and refrains from helping Jim escape. If this is consistent with Huck's having a similar kind and degree of care for Jim's humanity as in the original case, then it might seem that what Harman says implies that he's excused. After all, if Huck's care for Jim is sufficient to make him praiseworthy when he does the right thing, why wouldn't it be sufficient to excuse him if he instead did the wrong thing?

Harman might respond that, since Huck believes that Jim's humanity doesn't matter as much as that of a white person's, this *belief* reflects insufficient doxastic care for Jim's interests. But, when Huck (in the original case) does the right thing, his behavior doesn't *reflect* his false belief about Jim's interests; rather, it reflects his conative care for Jim's interests (his *desire* to help Jim). Accordingly, Harman might maintain that when Huck (in the inverted case) instead does the wrong thing and refrains from helping Jim, then his morally wrong behavior reflects, not his conative care for Jim, but rather insufficient doxastic care: his belief that Jim's interests don't matter as much as those of a white person.

The question now is whether something similar will always be true of an agent who acts wrongly from moral ignorance due to a commitment to a false ethical theory. A crucial question, then, is whether such agents' morally ignorant wrongful behavior must always reflect either insufficient conative *or* doxastic care for what matters morally. With this in mind, consider a case offered by Nomy Arpaly:

Immanuel concludes that he must never lie—not even to save a life. One day a would-be murderer shows up, asking after Gotlieb's whereabouts. Immanuel, though quite tormented, lies to save Gotlieb. Immanuel believes he did the wrong thing, but feels guiltily relieved, as he cares about human lives at least as much as you and I do, weird views about morality notwithstanding (2015, 141).

Arpaly maintains that Immanuel is praiseworthy for his action even though he (like Huck) acts rightly from moral ignorance, since his action reflects significant conative care for human life.²⁷ Consider an inversion of Arpaly's case. Suppose that Immanuel is (due to his care for human life) just as tormented by his predicament, but instead ends up acting in accordance with his false moral beliefs and tells the would-be murderer the truth. Arpaly apparently grants that, in the original case, Immanuel doesn't merely display care, but sufficient care: he cares about human life at least as much as you or I. However, even if Arpaly granted that Immanuel can have a similar kind and degree of care for human life in the inverted case as he does in the original, she might be inclined to offer the same sort of response as the one considered above concerning the inverted Huck Finn case. In particular, she might claim that Immanuel's morally wrong behavior in the inverted case reflects insufficient *doxastic* care, instantiated in a belief that human life doesn't matter very much.

But even in the inverted case, Immanuel needn't have any judgment that human life doesn't matter (or even that it doesn't matter very much). Agents like Immanuel might maintain only that certain considerations (truth-telling, promise-keeping, etc.) always trump other considerations, and not that the competing considerations are unimportant. So, when in the inverted case Immanuel arrives at the alternative decision to tell the truth, his false belief that telling the truth is obligatory may reflect, not a lack of care for human life, but instead excessive care for truth-telling. Furthermore, this is possible even if (as many detractors of Kantianism maintain) it is not even a close call whether one should prioritize truth-telling over saving a life.

More generally, deontologists may act in accordance with beliefs that they ought to keep their promises, tell the truth, and so on, even though refraining from doing so might prevent great harms. But their behavior needn't reflect any underlying judgment that the interests or well-being of others don't matter. They may, in fact, have great conative *and* doxastic care for the well-being of

others. Their behavior reflects care for genuine moral considerations (e.g., promise-keeping), even if it reflects too much care for them.

Parallel reasoning doesn't neatly apply to the compliant Utilitarian. The Utilitarian judges that keeping commitments, truth-telling, promoting justice, and so on matter morally only insofar as these things promote happiness or desire-satisfaction. So, although the Utilitarian might have great conative care for important moral considerations, he will also judge that these considerations are, of themselves, morally unimportant. When the Utilitarian breaks a promise (that he ought to keep) in order to maximize happiness, his action doesn't merely reflect a concern for happiness; it also plausibly reflects the implicit judgment that promise-keeping has no intrinsic value or importance, and this might be so even if the Utilitarian happens to have great conative care for promise-keeping.

More generally, whether the Sufficient Care Excuse applies to a proponent of a false ethical theory depends upon whether that theory entails, not only that some considerations are (perhaps always) trumped by other considerations, but also that certain considerations are morally unimportant or irrelevant. A judgment that some morally important feature of an action is unimportant or irrelevant reflects a deficiency of doxastic care for that feature (even if one happens to have conative care for it). If, however, a false ethical theory can allow that a moral consideration is morally important, then a committed proponent of the theory may still have both sufficient doxastic and conative care for that consideration while also acting out of moral ignorance due to a mistaken judgment that some competing consideration is more important. Ethical theories that acknowledge a wide range of moral considerations may well leave room for their proponents to possess sufficient care for what matters even in cases where they act from moral ignorance. And, again, this is so even if some such cases aren't close calls.

Harman and Arpaly grant that a morally ignorant agent might have sufficient conative care for the morally relevant features of their behavior. Yet they are prevented from affirming that

morally ignorant agents can care enough seemingly because of the assumption that an agent's moral ignorance (and behavior) must at least reflect a deficiency of doxastic care for what matters morally. The above reasoning casts doubt on this assumption.

The scope of the Sufficient Care Excuse for quality of will views depends partly upon how flexibly we should construe moral care or concern. There is conceptual room to maintain that, while a doxastic component isn't required for care, it is required for sufficient care. But even if sufficient care requires both sufficient doxastic and conative care, the reasoning presented in this section offers a plausible way of extending the Sufficient Care Excuse to some cases where proponents of false ethical theories act wrongly from moral ignorance. Since not all of these are close calls, this set of cases is not merely a subset of close call cases. Thus, the scope of the Sufficient Care Excuse is plausibly larger than it would be were it limited to close calls.

5.3 Delimiting the Sufficient Care Excuse

One might object that, if what I have said so far is correct, then there is no reason why one cannot simply stipulate that slaveowners like Cleo care enough about their slaves' welfare, but just care too much about (for example) their own financial security. It might seem, then, that my argument has the counterintuitive implication that slaveowners, racists, and other morally ignorant agents could too easily be excused for their beliefs and resulting behavior.

Recall that, when critiquing Wieland's argument, I pointed out that caring comes cheap. Here one might respond that, if we can simply *stipulate* degrees (or kinds) of moral concern of morally ignorant agents, then caring enough also comes cheap—at least when it comes to hypothetical agents. To illustrate, imagine a slaveowner that cares very much about her slaves, but cares to an even greater degree about her own financial well-being, and judges that it's morally

permissible to keep her slaves on this basis. If the reasoning I've offered here implies that the Sufficient Care Excuse extends to such agents, then this is a counterintuitive result.

This objection provides a reason to delimit the Sufficient Care Excuse. Accordingly, one might maintain that moral ignorance excuses only when it involves a mistaken judgment about the relative importance of (a) genuine moral considerations and when (b) there is not a vast difference in the actual importance of those considerations.²⁸ It's certainly not clear that the slaveowner in the above case satisfies both (a) and (b), since there is no obvious candidate that counts as a genuine moral consideration *and* that isn't vastly outweighed by the considerations concerning the slaves' well-being and right to freedom. In contrast, Clark's case arguably satisfies each of the requirements. Telling the truth about the machine malfunction is something Clark is implicitly committed to by remaining in his job. A failure to tell the truth, then, would constitute a violation of that commitment and a breach of trust with his manager. Furthermore, honesty (and refraining from deception) is deeply important to our understanding of what we owe to others, morally speaking. Thus, the considerations that Clark prioritizes are genuine and (at least moderately) important moral considerations that, while plausibly outweighed by his children's dental needs, aren't *vastly* outweighed.²⁹

However, it remains to be seen whether there is any principled reason to place this (or a similar) restriction on the Sufficient Care Excuse beyond the fact that doing so coheres better with our intuitions. Toward the end of Section 4 I suggested that, while it may not be true *in general* that caring enough about one thing requires caring more about it than other competing factors (*CER* is false), this may be so with respect to certain considerations. Caring sufficiently (and perhaps in the right *kind* of way) for human well-being or freedom, for example, seems to require caring about these things more than about one's own financial success. Part of caring sufficiently about human life, it might be argued, is prioritizing it above (at least some) other concerns. Additionally, it's

plausible that sufficient care for the well-being of one's own child may require more than just a certain *degree* of care. It may also require being willing, at least to some extent, to sacrifice one's personal interests for the child's well-being. The mother who prioritizes her wardrobe over her son's medical checkup reflects this sort of failure to care sufficiently.

The foregoing considerations do not, however, imply that agents like Clark don't care enough. It's not generally part of caring sufficiently for the well-being of others (even one's own children) that one prioritizes each of their interests above all other moral considerations. In at least some cases we may have sufficient care for what matters morally without our cares being perfectly calibrated to the relative importance of each particular consideration. So, while there are plausibly certain cases where caring enough about some feature requires caring about it more than competing factors, this isn't true generally.³⁰ And if so, then the Sufficient Care Excuse can drive a wedge between quality of will views discussed here and *No Excuse*. Even if these views are correct, moral ignorance sometimes excuses.

Notes

¹ Fitzpatrick (2008) and Talbert (2013) use this terminology. Rosen (2003) calls this type of ignorance 'factual' or 'non-moral' ignorance.

² By *wrong-making features*, I mean *pro tanto* considerations that contribute to the moral wrongness of some behavior. These same considerations may be present in cases of morally permissible or right behavior, and yet be outweighed by competing considerations in favor of the behavior's permissibility or rightness. For instance, the fact that some medicine might cause an allergic reaction counts against prescribing it, but this might be outweighed by some competing consideration (e.g., the medicine is required to save a life).

³ Wieland (2017) calls this 'pure moral ignorance.'

⁴ Theorists that maintain some version of *NQW* include Angela Smith (2005), Pamela Hieronymi (2008), Elizabeth Harman (2011), Matthew Talbert (2013), and Nomy Arpaly (2003, 2015).

⁵ *NQW* is an account of the *conditions* on blameworthiness. It is neutral on the nature of blame, and thus on what it is *to be* blameworthy. Perhaps blameworthiness concerns the appropriateness of certain reactive attitudes (Wallace 1994), or moral protest (Talbert 2012, Smith 2013), or relationship alteration (Scanlon 2008). *NQW* might conceivably be conjoined with any such account.

⁶ While Harman allows for exceptions, she holds that moral ignorance itself never excuses. I explain Harman's justification of such exceptions in footnote 13.

⁷ As Wieland notes, defenses of *No Excuse* often involve the conjunction of *NQW* with the view that moral ignorance always reflects a negative quality of will (2017: 150).

⁸ I do not espouse any particular ethical theory here, nor should the verdicts I offer about certain cases depend upon whether any particular ethical theory is true. I am grateful to an anonymous referee for *Philosophical Explorations* for encouraging me to clarify this.

⁹ As I argue in Section 5, whether these implications really do hold for a particular ethical theory depends upon the version of consequentialism or deontology in question.

¹⁰ This phrase intentionally reflects the title of Harman's paper, "Ethics is Hard. What Follows?"

¹¹ My argument should work so long as at least some of the cases I offer are counterexamples to *No Excuse*. The different cases (and their possible variations) discussed in Sections 4 and 5 should safeguard against misgivings that some theorists may have concerning particular cases.

¹² I set aside cases in which an agent is uncertain about the moral status of his behavior, which are addressed by Guerrero (2007) and Harman (2015).

¹³ I include 'all else equal' because Harman (2014) maintains that psychological conditions (e.g., depression) can prevent an agent's behavior from manifesting what she cares about. Harman's view seems to be that, in such cases, an agent might (i) behave wrongly while (ii) believing that her behavior is morally permissible and nevertheless (iii) care sufficiently about the wrong-making features of her behavior (19). Thus, Harman allows that some morally ignorant agents may be blameless for their wrong behavior. Regardless, Harman maintains that it is not the moral ignorance that excuses, but the relevant psychological condition.

¹⁴ In what follows I use the term 'good will' to include to what is sometimes called 'due regard.'

¹⁵ Other positive duties include *pro tanto* duties such as duties of fidelity (e.g., promise-keeping).

¹⁶ This way of parsing the distinction between types of negative quality of will is similar to [Author 2019], and is, as far as I can tell, consistent with Wieland's characterization of negative quality of will along the lines laid out by Arpaly (2003) and Harman (2011).

¹⁷ I frame negative quality of will in terms of an agent's attitudes and beliefs with respect to *wrong-making* features of the agent's morally wrong behavior. Subsequent discussion of an agent's care or concern for the morally relevant features of his behavior (or considerations bearing on the moral status of his behavior) may also refer to an agent's attitudes or beliefs with respect to something that counts *in favor* of his behavior (whether or not that consideration is outweighed by moral considerations against it).

¹⁸ Wieland also considers another interpretation of quality of will, according to which 'S does X with good will (despite the fact that she is ignorant that X is wrong) if S has made a serious attempt to figure out whether X is wrong' (154). He argues that this can be true of morally ignorant agents if the moral truth is not accessible enough to them, and offers an account of accessibility (158ff).

¹⁹ It's worth noting that a defender of *No Excuse* needn't maintain that acting from good will requires acting permissibly, since an agent might be circumstantially ignorant. In order to refute *No Excuse*, Wieland needs to illustrate cases where good will and behaving wrongly *from moral ignorance* go together. This is presumably the purpose of 3**.

²⁰ Those who are unconvinced that his children's interests in this case are strong enough to outweigh the considerations in favor of truth-telling are welcome to substitute his children's need for dental work with something more important (e.g., the need for heating during winter months).

²¹ The judgement that Clark's action is morally wrong does not depend upon the truth of any particular ethical theory. This judgment is perhaps most straightforwardly accommodated by consequentialist theories, according to which Clark's action is justified by the value of its consequences. Deontological theories, as well, can allow for withholding the truth so long as they don't include exceptionless constraints against doing so (e.g., Ross 1930). Virtue-based theories may also allow that the virtue of benevolence toward one's family is to be prioritized over the virtue of truthfulness in this case. While Kantianism, in particular, may imply that Clark's judgment is correct and his behavior morally right, the authors I'm engaged with are not (to my knowledge) committed to Kantianism. Thus, the verdict offered about this case should not be question-begging. I am grateful to an anonymous referee from *Philosophical Explorations* for encouraging me to clarify this point.

²² For a sustained argument that fulfilling one's procedural epistemic obligations does not entail an improved epistemic position with respect to some relevant fact, see [Author 2017].

²³ Harman (2011) explicitly rejects the idea that culpability for one's ignorance depends upon reasonable expectation, and Talbert (2013) explicitly rejects the view that morally ignorant wrongdoers are excused when it was unreasonable to expect them to know better.

²⁴ One might think that the plausibility of Clark's blamelessness hangs on which theory of blame or blameworthiness is assumed here. But this misunderstands how the argument works. If Clark cares enough, then *NQW* straightforwardly implies that Clark isn't blameworthy (and that *No Excuse* is false), *whichever* account of blame or blameworthiness one conjoins with *NQW* (see footnote 5). For example, take the view according to which blameworthiness is a matter of being an appropriate target of moral protest (Talbert 2012, Smith 2013). A theorist who holds this view might claim that

Clark is blameworthy, arguing that it would be reasonable for Clark's family to register moral protest against him in response to his behavior. But this theorist would have to part ways with *NQW* as an account of the *conditions on* blameworthiness, since Clark doesn't meet them.

²⁵ I am grateful to [Blinded] for suggesting this way of parsing this distinction.

²⁶ Some may object that my judgment that agents in close-call cases act wrongly presupposes a specific "model answer" approach to ethics. Such theorists are invited to consider in Section 5.2 another category of sufficient care excuses that extends beyond close calls. I am grateful to an anonymous referee from *Philosophical Explorations* for raising this concern.

²⁷ Arpaly refers to the actions of Huck and Immanuel as instances of 'inverse akrasia.'

²⁸ I am grateful to [Blinded] for suggesting (though not endorsing) this as a possible limiting principle. Though I am not committed to it, something like it may be a promising way to delimit the Sufficient Care Excuse.

²⁹ Some may object that the introduction of this limiting principle places the Sufficient Care Excuse in a tight spot: on one hand, some won't be convinced that close-calls are clear cases of moral ignorance. On the other hand, the further apart competing moral considerations are in terms of their relative importance, the more likely it is that a case is excluded by the suggested limiting principle (perhaps some will, on these grounds, think that the Kantian case of prioritizing truth-telling over saving a life won't meet the suggested limitation). There are two responses to this concern: first, some theorists will accept close calls as cases of moral ignorance. Second, there is considerable room for cases where the difference in competing moral considerations is neither slight nor vast, but moderate. The reasons I've presented here make it plausible that Clark's case is one such case. Some cases of good faith proponents of false ethical theories will also fall in this range. I am grateful to an anonymous referee from *Philosophical Explorations* for encouraging me to address this concern.

³⁰ What morality demands with respect to *de re* care may be context-dependent, and thus sensitive to a host of considerations. If so, then formulating a general principle expressing what morality demands in this respect may not be possible.

References

[Author]. 2017.

[Author]. 2019.

Arpaly, Nomy. 2002. "Moral Worth." *The Journal of Philosophy* 99 (5): 223-245.

Arpaly, Nomy. 2003. *Unprincipled Virtue: An Inquiry Into Moral Agency*. New York: Oxford University Press.

Arpaly, Nomy. 2015. "Huckleberry Finn Revisited: Inverse Akrasia and Moral Ignorance." In *The Nature of Moral Responsibility: New Essays*, edited by Randy Clarke, Michael McKenna, and Angela Smith, 141–56. New York: Oxford University Press.

FitzPatrick, William J. 2008. "Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge." *Ethics* 118 (4): 589-613.

FitzPatrick, William J. 2017. "Unwitting Wrongdoing, Reasonable Expectations, and Blameworthiness." In *Responsibility: The Epistemic Condition*, edited by Philip Robichaud and Jan Willem Wieland, 29-46. New York: Oxford University Press.

Guerrero, Alexander A. 2007. "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution." *Philosophical Studies* 136: 59-97.

Harman, Elizabeth. 2011. "Does Moral Ignorance Exculpate?" *Ratio* 24 (4): 443–468.

Harman, Elizabeth. 2014. "Ethics is Hard. What Follows?" Manuscript.

Harman, Elisabeth. 2015. "The Irrelevance of Moral Uncertainty." *Oxford Studies in Metaethics* 10: 53–79.

- Hieronymi, Pamela. 2008. "Responsibility for Believing." *Synthese* 161 (3): 357-373.
- Levy, Neil. 2009. "Culpable Ignorance and Moral Responsibility: A Reply to FitzPatrick." *Ethics* 119 (4): 729-741.
- Rosen, Gideon. 2003. "Culpability and Ignorance." *Proceedings of the Aristotelian Society* 103: 61-84.
- Ross, W.D. 1930. *The Right and the Good*. New York: Oxford University Press.
- Scanlon, T. M. 2008. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.
- Smith, Angela. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life." *Ethics* 115: 236-71.
- Smith, Angela. 2013. "Moral Protest and Moral Blame." In *Blame: Its Nature and Norms*, edited by D. Justin Coates and Neal A. Tognazzini, 27–48. New York: Oxford University Press.
- Talbert, Matthew. 2012. "Moral Competence, Moral Blame, and Protest." *The Journal of Ethics* 16 (1): 89-109.
- Talbert, Matthew. 2013. "Unwitting Wrongdoers and the Role of Moral Disagreement in Blame." In *Oxford Studies in Agency and Responsibility: Volume 1*, edited by David Shoemaker, 225–245. Oxford: Oxford University Press.
- Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Wieland, Jan Willem. 2017. "What's Special About Moral Ignorance?" *Ratio*, 30 (2): 149-164.