

Don't Know, Don't Care?

ZOË JOHNSON KING
New York University

Note: This is a draft. Nonetheless, feel free to share it with gleeful abandon – just bear in mind, if you want to cite it, that it might change.

1. Introduction

The title of this paper should really be “ $\neg(\text{Don't Know} \rightarrow \text{Don't Care})$ ”. My thesis is that moral ignorance does not always involve, and so does not imply, a failure to care adequately about that which is in fact morally significant. This entails that, even if we are blameworthy for failing to care adequately about what is in fact morally significant, it does not follow that all moral ignorance is blameworthy.

In §2 I clarify my thesis. In §3 I then prove the thesis; I offer proof of concept in the form of three cases in which agents are morally ignorant despite caring adequately about everything morally significant. I note the generalizable features of the cases, so as to leave readers with a recipe for generating cases of their own. In §4 I discuss the upshot of my thesis for the literature on culpable moral ignorance; I argue that it creates room for a view that respects the concerns of those who hold that moral ignorance is typically blameless, within a framework according to which moral ignorance is blameworthy to the extent that it manifests poor quality of will. I close in §5 by suggesting that quality-of-will theorists should change direction in our thinking on this topic. We should focus on articulating the *standards for adequate caring*, i.e. the standards that specify what it is to care “adequately” about something morally significant. It is only once we have a clearer picture of what constitutes adequate caring, what constitutes inadequate caring, and what (if anything) constitutes supererogatory caring, that quality-of-will theorists can fruitfully address questions of culpability for moral ignorance. I offer four promising avenues for future research on this topic.

2. Moral ignorance

Moral ignorance is ignorance of a moral fact. A moral fact is a fact that includes some evaluative or deontic component. For example, facts about which acts are right and which are wrong are moral facts. So are facts about which things are (dis)valuable, or about the relative degrees of value of two or more valuable things. Facts about the existence of moral reasons or duties, their relative strength, and the normative relationships between them — such as defeat or lexical priority — are also moral facts. As I'll construe things here, facts about the nature or extension of thick properties¹ – like justice or kindness – are also moral facts, since they are facts that include evaluative or deontic components. So, for example, if someone takes another person's

¹ By “thick properties” I mean the properties referred to by thick concepts (on which see e.g. Roberts 2013, Väyrynen 2013). I assume that there are such things. Others deny this. I will not get into that dispute in this paper.

coffee and runs out of the coffee shop while failing to know that this is *wrong* or *unfair* or *rude* or *contrary to duty* or *against the balance of moral reasons* or *conducive to a bad state of affairs*, then she exhibits moral ignorance. But her failure to know that the substance she has stolen is *coffee* would not be moral ignorance, since facts about what is coffee are not moral facts.

In this paper, I focus on moral ignorance of the first kind listed above: ignorance about which acts are right and which are wrong. But this is just to avoid speaking in long disjunctions. None of my points pertain specifically to ignorance of facts about rightness and wrongness. I intend for everything I say to generalize to all the other kinds of moral ignorance, and will touch on some of these others when discussing examples in the next section.

It often happens that someone is ignorant of a moral fact, but only because she is ignorant of one or more related non-moral facts. For example, someone may cheerfully grab what she thinks is her coffee and walk out of the coffee shop, not realizing that the item in her hand is in fact someone else's coffee. She may know full well that it is wrong to steal. Yet she may be ignorant of the moral fact that what she is doing is wrong, simply because she is ignorant of the non-moral fact that what she is doing constitutes stealing. Since this agent knows that stealing is wrong, she would have known that what she is doing is wrong had she but known that it is stealing. In this case, then, only the agent's ignorance about the non-moral nature of her act prevents her from knowing the truth about its moral status.

As far as I am aware, everyone agrees that moral ignorance based solely on non-moral ignorance in this way need involve no failure to care adequately about what is in fact morally significant. In the coffee shop case, the agent's ignorance of the fact that she is doing something wrong need involve no failure to care adequately about the coffee's rightful owner, the value of coffee, or anything else morally significant. That is because her ignorance of the fact that she has grabbed someone else's coffee need involve no failure to care about such things. If the baristas announce that a coffee is ready by calling the customer's name, and she has the same name as a previous customer, then she may innocently mistake another person's coffee for her own while caring adequately about everything morally significant.

My thesis is not about cases like this. In these cases, the agent's ignorance is not fundamentally moral. I am interested in cases in which the agent's ignorance is fundamentally moral: in which she is aware of all morally significant facts about the non-moral nature of her situation, but does not fully appreciate these facts' moral significance. This fundamentally moral ignorance amounts to ignorance of general principles about which kinds of acts are right or wrong (or honest, cruel, respectful, sanctimonious, etc.), or about what is valuable and what degrees of value different things have, or about what reasons or duties we have, their strength, or the normative relationships between them. We might call this "*pure* moral ignorance". I will just call it "moral ignorance".

3. Examples

My thesis is that moral ignorance does not always involve a failure to care adequately about what is in fact morally significant – or, equivalently, that it is possible for someone to be morally ignorant while caring adequately about what is in fact morally significant. As proof of concept, I will now offer three cases. They are all cases of agents who do not know that what they are doing is wrong, despite caring adequately about all morally significant features of their circumstances, including the very features that in fact make their acts wrong. The cases are all based on real events.

GENTRIFICATION: Grace knows that, as gentrification spreads through major cities, poorer communities of people who have lived in certain areas for a long time can come to find it unaffordable to stay there. Grace's neighborhood is beginning to gentrify, and she is committed to protecting marginalized communities in her area from the threat of gentrification. So she participates in anti-gentrification protests and boycotts of new hipster businesses. One of the businesses that Grace targets is an art gallery, which has publicly committed to promoting the work of local Latina women artists. In fact, the gallery's publicity would greatly benefit these women by affording them some much-needed social recognition. And it would do so in a way that more than compensates for the gallery's marginal contribution to rising house prices, with the result that the gallery would have a net positive impact on the interests of marginalized communities in Grace's area — if only it weren't impeded by boycotts and protests. But Grace does not realize this. She doubts that the gallery's publicity does much to further local women's interests, because she under-appreciates the value of social recognition as a component of human well-being. So she includes the gallery in her boycotts and protests. Thus Grace's activity in fact hampers the interests of marginalized communities in her local area, despite the fact that she is trying to promote their interests.

NAMES: McKenzie very much wants to make her academic discipline more welcoming for people of color. (She is herself a woman of color.) To that end, she has created a diversity task force at her academic institution, served in multiple mentoring programs, and developed policies in her department aimed at recruitment and retention of people of color among faculty, graduate students, and undergraduate majors. In fact, one small but easy and cost-free way to make McKenzie's discipline more welcoming for people of color is to choose names for characters in examples that are not all Anglo-American names, thus avoiding the expressive harm that occurs when one acts as if being Anglo-American is seen as "normal" within the discipline. But the fact that this is an expressive harm has never occurred to McKenzie. She is focused on her activism, and has not reflected on this small way in which her policy goals relate to the details of her research. So, when she names the characters in her examples, McKenzie unreflectively follows the local trend of calling them Alice, Bob, Clara, Dave, Edward, and so on.

NON-VIOLENT PROTEST: John and Xavier are two Black Lives Matter protestors. They are both committed to the value of non-violent protest, even in the face of violent responses from police officers; they believe that "two wrongs don't make a right", and thus that violent responses to police violence are rarely morally justified. But John and Xavier don't think that violence is *never* morally justified. They recognize that violence is permitted in egregious cases of self- or other-defense. And they recognize that there are hard cases, in which it is unclear whether police violence is sufficiently severe to defeat the moral presumption in favor of non-violence and call for violent resistance. At a protest, John and Xavier encounter one of these hard cases. In the heat of the moment, they are both deeply unsure what to do, and their emotions are running high. But each of them makes a judgement call. And they make different calls: John fights back, while Xavier restrains himself.

Each case raises slightly different issues. I'll discuss them in turn.

In **GENTRIFICATION**, Grace does something morally wrong: she hampers the interests of marginalized communities in her local area. But Grace is motivated to act by her concern for *promoting* the interests of marginalized communities in her local area. She cares deeply about the issue. And this deep moral concern explains why she acts as she does. So the trouble is not that Grace doesn't care, but that she is mistaken about what the thing that she cares about requires her to do. And she is mistaken about this because she is mistaken about the nature of people's interests; she underappreciates the value of social recognition as an important human good. This mistake leads Grace to inadvertently undermine the very thing she cares about. To emphasize: this is not a case of moral ignorance based solely on non-moral ignorance. It's not that Grace doesn't know that the gallery would bring the local artists some social recognition, but rather that she does not realize how good this would be for them. (If it helps, we can imagine that some of the artists are members of Grace's activist group, that they have told her that the social recognition the gallery would afford them is less important than its contribution to rising house prices, and that she defers to them — people can be mistaken about their own interests, so this is a possible version of the case.) In this case, I claim that Grace's moral ignorance does not show that she fails to care adequately. Instead, it just shows that the case is sufficiently hard that someone with the best will in the world can still get it wrong. Grace has more concern for the interests of marginalized communities than most of us do, but even this unusually high degree of concern is not enough to render her omniscient about those interests.

In **NAMES**, McKenzie does something morally wrong: she enacts a small expressive harm, thus making her academic discipline less welcoming for people of color, which she could easily have avoided at no cost. But McKenzie cares deeply about making her discipline more welcoming for people of color — again, she has more concern for this than most of us do. She displays this concern through her activism, mentoring, and policy initiatives, despite never having stopped to reflect on what her political goals require when it comes to choosing names in examples. In this case, again, McKenzie's ignorance is not based solely on non-moral ignorance. McKenzie knows what names she is using in her examples. If she were to reflect, it would strike her that they are Anglo-American names, and she would realize that what she is doing constitutes an expressive harm and thus a moral wrong. The trouble is just that she does not reflect on the issue. As with Grace, McKenzie's unusually high degree of concern for something morally significant does not render her omniscient about everything that this consideration calls for. I claim that this does not show that McKenzie fails to care adequately. On the contrary, she clearly does care adequately about the very consideration that makes her act wrong; she cares in other ways.

In **NON-VIOLENT PROTEST**, At least one of John and Xavier does the morally wrong thing in their circumstances, either by failing to protect innocent people from egregious violence or by committing an act of violence with insufficient justification. Whoever acts wrongly fails to realize that what he is doing is wrong; he believes that it is morally required. And, once again, this is not a case of moral ignorance based solely on non-moral ignorance. We can stipulate that John and Xavier both know exactly what the effects of violent resistance would be. Their ignorance concerns whether the principle of non-violence admits of an exception in these circumstances, or not. This is a question about the normative relationships that hold between our reasons or duties, and hence a moral question. Nonetheless, I claim that whoever fails to realize that his behavior is wrong does so while caring adequately about all morally significant features of his circumstances, including those that in fact make his act wrong. Both John and Xavier are deeply committed to non-violent protest, and both are deeply moved by the suffering of their fellow protesters. This is why their emotions run high in the heat of the moment. We may imagine that they are in states of intense turmoil, caused by their recognition of what is at stake and by the depth of their concern. John and Xavier clearly care about everything morally significant in their circumstances at least as much as can reasonably be expected of them. Thus, they care adequately.

In these cases, I am inclined to make a stronger moral claim: Grace, McKenzie, John, and Xavier care about everything morally significant in their circumstances to a degree that *surpasses* what could reasonably be expected of them, and is thus supererogatory. I will return to this possibility in §5.

These cases suggest three ways in which someone may fail to realize that what she is doing is wrong while caring adequately about everything that is in fact morally significant. First, she may care adequately about a certain consideration and know that it is at stake in her circumstances, but be ignorant about its precise nature, and thus ignorant about what sort of response it calls for (as in **GENTRIFICATION**). Second, she may care adequately about a certain consideration without recognizing that it is at stake in her circumstances, having failed to reflect on the moral significance of what she is doing (as in **NAMES**). Third, she may care adequately about two or more morally significant considerations that she sees are at stake in her circumstances, but be ignorant about the normative relationships between them (as in **NON-VIOLENT PROTEST**).² Moral ignorance involves no failure to care adequately in any of these types of case, because the agents in these cases *don't* fail to care adequately. Moral ignorance cannot be attributed to a dearth of moral caring in someone who exhibits no dearth of moral caring. So, the agents' moral ignorance in these three types of case is not attributable to their dearth of moral caring.

If the reader is unsatisfied by one or more of my examples, she is welcome to come up with examples of her own. As long as she can think of any cases in which someone cares adequately about a certain morally significant consideration despite being partially ignorant of its precise nature, despite failing to realize that it is at stake in her circumstances, or despite failing to appreciate its relative importance or the normative relationships that it bears to other morally significant considerations, she can think of cases that offer proof of concept for my central thesis.

4. Moral responsibility for moral ignorance

The literature on moral responsibility for moral ignorance has lately been overshadowed by some strong views. One holds that people are rarely, if ever, blameworthy for their moral ignorance. The other holds that people are always blameworthy for their moral ignorance, unless it is based solely on blameless non-moral ignorance (as in the coffee case from §2). Each of these views is a recognizable descendant of one of the main historical traditions in philosophical thinking about moral responsibility. But these traditions, properly understood, may not in fact imply such strong views about moral ignorance. Moreover, we can steer an appealing path between them once we recognize that not all moral ignorance involves a failure to care adequately about what is in fact morally significant. Showing this is my task for the present section.

The first strong view is Gideon Rosen's. Over a series of recent papers, Rosen has argued that culpable moral ignorance is much rarer than one might think, and that, as a result, we should be skeptical about the

² This kind of case is often discussed in the existing literature, in which it is common to use the phrase "hard cases" to refer to cases in which two values are at stake and their relative degrees of value are unclear (see e.g. Rosen 2002, p.70, FitzPatrick 2008, p.600, n.24). However, I think that prior discussion of this phenomenon has been led astray by a focus on examples that are bad illustrations of it, since they are not genuine hard cases. The primary examples in the literature involve slaveholders who do not realize that slavery is morally questionable, men who do not realize that women are their moral equals, a former U.S. president who did not realize that it is wrong to use nuclear weapons against civilians in order to end a war, and a businessman who cannot tell whether his business practices are "reprehensibly ruthless" or "permissibly aggressive" (Rosen 2008, p.305). These all seem like such clear cases of motivated moral ignorance that it is hard to take them seriously. (For a similar complaint and some more realistic examples, see [REDACTED]). So, it is good that real life offers plenty of examples of genuine hard cases. I take myself to have offered one.

culpability of any particular bit of moral ignorance that we encounter. In his (2002), he argues that ignorance — including moral ignorance — is culpable just in case it is “the upshot of recklessness or negligence or deliberate misconduct in the management of [one’s] opinion” (p.**). In his (2004), he then argues that this implies that most moral ignorance is blameless, since the prior acts or omissions that give rise to moral ignorance rarely amount to recklessness, negligence, or deliberate misconduct. They would do so only if they involved a knowing violation of our epistemic obligations to “take steps to ensure that when the time comes to act, one will know what one ought to know” (*ibid.*, p.301). But people rarely *knowingly* violate these obligations — rather, at the time when they are failing to gather or reflect on evidence in a way that amounts to violating an epistemic obligation, they usually do not realize that they are under such an obligation. This is itself a kind of normative ignorance. The culpability of this further bit of normative ignorance is then in question, setting us off on a regress. Rosen assumes that the sort of “clear-eyed epistemic akrasia” (p.**) that could be an original locus of responsibility for moral ignorance is rare. He thus doubts that much moral ignorance is blameworthy. In his (2008), Rosen elaborates on the general way of thinking about blameworthiness that underpins his position on moral ignorance; he considers and rejects the view that we are blameworthy for failing to care adequately about what is morally significant, suggesting that someone who “nursed his indifference or negligently allowed it to fester” (p.608) might be blameworthy for doing so, but that he is not blameworthy if his indifference was out of his control.

This approach bears the hallmarks of a traditional *voluntarist* position on moral responsibility. Voluntarists hold that we are responsible for what we voluntarily choose or cultivate, for what we recklessly or negligently fail to prevent, and for nothing else. Everything else is just something that happens to us, rather than something we do. And it is inappropriate, says the voluntarist, to praise or blame someone for something that just happens to her. This is the approach that leads to traditional worries about whether determinism is compatible with moral responsibility.³

It is by no means obvious that Rosen’s line is the right line for voluntarists to take on moral ignorance. Whether it is the right line depends on some thorny empirical questions. Voluntarists recognize that *motivated moral ignorance* is blameworthy; it is blameworthy for someone to negligently avoid pursuing a line of moral inquiry, or to accept a conclusion that she knows to be poorly thought-through (and perhaps indefensible under further scrutiny), out of a sense that further inquiry may lead to the awkward conclusion that she should substantially revise her behavior and an aversion to doing so. This amounts to an akratic violation of our epistemic obligations. Nonetheless, my impression is that people do this sort of thing a lot. If that is so, then a lot of moral ignorance is blameworthy on voluntarist grounds. But whether I am right about this is a thorny empirical question. I will not settle the matter here; for present purposes we may simply note that voluntarism may have substantially harsher implications regarding moral ignorance than Rosen thinks it does.⁴

The second strong view is Elizabeth Harman’s. In response to Rosen, Harman has argued that moral ignorance is always blameworthy, except in cases in which it is based solely on blameless non-moral ignorance. Harman holds that “we are morally obligated to believe the moral truths relevant to our actions” (2011, p.459), and that we are blameworthy for failing to do so if we understand enough about our actions’ non-moral nature to be “in a position to realize” the moral truth (*ibid.*, pp.462-464). In explaining why this might be, Harman says that “beliefs (and failures to believe) are blameworthy if they involve inadequately caring about something morally significant”, and then that “believing a certain kind of behavior is wrong

³ Incompatibilism has a long and storied history. For a sophisticated contemporary defense, see Levy (2011).

⁴ For excellent discussions of motivated moral ignorance see Moody-Adams (1994), Mills (2007), Alcoff (2007), and Pohlhaus (2012). Cf. the discussions of willful ignorance in Yaffe (2018) and [REDACTED].

on the basis of a certain consideration is a way of caring about that consideration" (*ibid.*, p.460). Harman is not fully explicit about the argument here. But the reasoning seems to be that, if believing an act wrong on the basis of a certain consideration is a way of caring about that consideration, it follows that failing to believe that an act is wrong, when it is in fact wrong, amounts to failing to care adequately about whatever consideration in fact makes it wrong. Since it makes an act wrong, this consideration must be morally significant. So failing to believe that an act is wrong, when it is in fact wrong, amounts to failing to care adequately about something morally significant. Given that failures to believe are blameworthy when they involve inadequately caring about something morally significant, it follows that all failures to believe that an act is wrong, when it is in fact wrong, are blameworthy.

Harman makes clear that her work draws on the *quality-of-will* approach to moral responsibility. On this approach, we are praise- or blameworthy for that which manifests the quality of our wills: the degree to which we care, or fail to care, about other people and things that are in fact morally significant. This approach is currently popular in part because it dissolves philosophical worries about determinism. It does not matter to quality-of-will theorist how someone came to have the ill will that she displays in her behavior, in her moral beliefs, or in whatever else we blame, since the view is that, in blaming, we react to the ill will itself.⁵

Harman's is the wrong line for quality-of-will theorists to take on moral ignorance. One immediate problem for this view is that it is simply false that believing an act wrong based on a consideration is a way of caring about that consideration. The standard counterexamples to motivational internalism in metaethics show this to be false. For instance, someone who is suffering from depression, someone who has recently undergone a paradigm shift in her moral beliefs and is struggling to bring her affective responses into line, and someone who enjoys abstract discussions of moral topics but has no inclination to act in accord with her conclusions, can all believe acts to be wrong on the basis of considerations without caring about those considerations.⁶

Weaker relationships do obtain between moral belief and moral care. For example, part of what it is to care about something is to be disposed to pay attention to it, notice when it is at stake, and reflect on the moral significance of an act's potential impact on it. This makes someone who cares about something more likely to notice when an act is wrong in virtue of its impact on this thing. So, while believing something wrong based on a consideration is not itself a way of caring about the consideration, the former attitude is often a causal consequence of the latter, and thus a good indicator of it.

But this does not entail that failing to believe that an act is wrong, when it is wrong, amounts to failing to care adequately about the considerations that make it wrong. The cases from §3 illustrate three ways for

⁵ The popularity of this approach is largely due to the enduring influence of Strawson (1962). For more contemporary discussions, see, for instance, Arpaly (2002), Arpaly and Schroeder (2013), and Smith (2005). Harman cites Arpaly as a view that can support her position on culpable moral ignorance.

⁶ See e.g. Stocker (1979) for a discussion of the depressive, Campbell (2007) for a discussion of paradigm shift, and Svavarsdóttir (1999) for a discussion of jerks. Notice that Harman's claim that believing an act wrong on the basis of a consideration is a way of caring about that consideration entails the wrong verdict about all these cases. It entails that the depressive is doing fine when it comes to moral caring; though her affective and motivational states are numbed, she cares about a great many things, since she has moral beliefs. It downplays the work that the paradigm-shifter must do to integrate her moral character, suggesting that she need to nothing to bring her moral caring in line with her moral beliefs since this work is already done. And it suggests that the jerk who is indifferent to morality displays a great deal of moral care, because he believes that certain features ground wrongness. This is all clearly incorrect.

this entailment to fail to go through. Someone can care adequately about the considerations that make an act wrong while being partially ignorant of their nature (as in **GENTRIFICATION**). Or she can care adequately about the considerations that make an act wrong, but fail to notice that they apply in her circumstances (as in **NAMES**). Or she can care adequately about the considerations that make an act wrong, while caring even more strongly about other morally significant considerations that are also at stake, such that she reaches the wrong view about the relative moral weight of these considerations or the normative relationships between them (as in **NON-VIOLENT PROTEST**). In each of these types of case, the agent makes some mistakes in her moral reasoning, but still cares adequately. In short: adequate caring does not require moral omniscience.

My point here is partly logical and partly moral. The logical point is that, for any morally significant consideration, the set of beliefs that an act is wrong based on this consideration is a proper subset of the set of all things that manifest care for the consideration. Plenty of other things manifest caring. This includes courses of action like McKenzie's activism, emotional states like John and Xavier's distress, and other beliefs – even ones that are false, like Grace's belief that boycotting the gallery is morally required. So, even if it were true that believing an act wrong based on a consideration is a way of caring about that consideration, it would not follow that failing to hold such a belief when an act is in fact wrong amounts to failing to care about the consideration that makes it wrong. By analogy: being a square is a way of being a quadrilateral, but it does not follow that failing to be a square amounts to failing to be a quadrilateral. One might be a rectangle.

The moral point is that some of these other things that manifest caring may amount to caring *adequately*. Characters like Grace, McKenzie, and John and Xavier display degrees of moral concern that are intuitively at least adequate, and perhaps more than adequate. This is because the idea of "adequate" caring suggests a measure of leniency afforded to those who fall short in some respects while getting all the basics right. Caring adequately is not caring maximally. So one can care adequately about a consideration without doing absolutely everything that would manifest care for it. When we assess the quality of an agent's will, then, we should not go looking for things that would manifest caring but that she fails to do. Rather, we should look for the things she does that manifest caring. Since adequate care does not require moral omniscience, we should expect to find plenty of people with some false moral beliefs who nonetheless care adequately. They are the rectangles; they care in other ways.

On this way of understanding the quality-of-will approach, it can accommodate many of the concerns that move voluntarists. Voluntarists are often concerned with cases in which it would take an unusual degree of insight to figure out the moral truth, and someone is morally ignorant only because she does not possess this unusual degree of insight. The voluntarist wants to say that such an agent is not blameworthy for this ignorance.⁷ But, properly understood, the quality-of-will approach explains this intuition. The explanation is simple: caring adequately does not require an unusual degree of moral insight. So someone can fail to grasp moral truths that require an unusual degree of insight while still caring adequately. On the quality-of-will approach, what ultimately explains why the agent is blameless is not the fact that the case would take great moral insight to get right. Rather, this fact indicates that someone with the best will in the world could still get the case wrong. What explains why the agent is blameless is the fact that, though she got the case wrong, she nevertheless cares at least adequately about everything that is in fact morally significant – that is, she displays good quality of will.

⁷ See e.g. Rosen (2002, p.66): "Given the intellectual and cultural resources available to a second millennium Hittite lord, it would have taken a moral genius to see through to the wrongness of chattel slavery. The example is meant to show that blameless moral ignorance is a possibility."

The quality-of-will approach also explains cases in which voluntarists do want to say that someone is blameworthy for her moral ignorance. Voluntarists do blame moral ignorance that is “the upshot of recklessness or negligence or deliberate misconduct in the management of [one’s] opinion”. And the quality-of-will theorist has a ready explanation of why this is blameworthy: recklessness, negligence, and deliberate misconduct all manifest ill will. Recklessness and negligence indicate a failure to care adequately about that which the agent thereby endangers. Deliberate misconduct indicates failure to care adequately about the morally significant features that explain why one’s act constitutes misconduct, and also a failure to care adequately about avoiding misconduct (*de dicto*). Similarly, motivated moral ignorance indicates a failure to care adequately about the subjects of the moral questions that one avoids thinking about, and also a failure to care adequately about working to become a morally better person. So, while the voluntarist and the quality-of-will theorist can agree about which instances of moral ignorance are morally blameworthy, the quality-of-will theorist offers a deeper explanation of why this is so: recklessness, negligence, and deliberate misconduct all reveal malice or indifference toward morally significant things in a way that amounts to failing to care adequately.

5. The standards for adequate caring

On the quality-of-will approach, moral ignorance is blameworthy when it manifests a failure to care adequately about what is in fact morally significant. So, figuring out when moral ignorance is blameworthy is a matter of figuring out what moral ignorance is ruled out by the moral standards specifying what it is to care “adequately”. Call these “the standards for adequate caring”.

There has been a tremendous amount of work done within the quality-of-will tradition in recent years. But, to my knowledge, nobody has attempted to describe the nature and structure of the standards for adequate caring. This is a major lacuna. For, if being blameworthy is a matter of failing to care adequately, then we will have a hard time figuring out what is blameworthy until we figure out what adequate caring amounts to. So I suggest that quality-of-will theorists reorient our inquiry. Before asking about the conditions under which a particular type of state – e.g., false moral belief – is blameworthy, we should first ask about the nature and structure of the standards for adequate caring.

In this final section, I will put my money where my mouth is and make an initial proposal for how to think about the standards for adequate caring. I will show that this proposal can accommodate attractive verdicts on the cases from §2, and will close by identifying four promising avenues for future research on this topic.

Here is how I propose that we conceive of the standards for adequate caring. For each morally significant thing, there is a threshold specifying the minimum level of concern for the thing that can reasonably be expected of people.⁸ This threshold is the standard for adequate caring with respect to this thing. So, for

⁸ Here I follow a trend in the literature on moral responsibility, wherein questions about blameworthiness are phrased in terms of what “could reasonably be expected” of someone. For example, Rosen argues that to blame wrongdoing that stems from blameless moral ignorance “is to expect more than it is reasonable to expect”, and is therefore unfair (2002, p.82). FitzPatrick argues that it is fair to blame moral ignorance that results from the exercise of a vice “where it would have been reasonable to expect people not to exercise those vices” (2008, p.611). And [REDACTED] argues that moral responsibility is in general a matter of meeting, exceeding, or failing to meet reasonable expectations. Quality-of-will theorists sometimes employ the same locution; for instance, Björnsson (2017) argues that blame for ignorance is reasonable if the ignorance results from a failure to care “as can reasonably be demanded”. On this sort of

any pair of an agent and a morally significant thing, there is a fact about whether the agent meets, falls short of, or exceeds the standard for adequate caring with respect to the thing. She falls short if the degree to which she actually cares about the thing is lower than the minimum threshold set by the standard. She meets the standard if the degree to which she cares about the thing is greater than or equal to the threshold. And she exceeds the standard if her degree of concern is above the threshold. In this third case, she cares about the thing *more than adequately* – that, is, to a higher degree than the minimum that could reasonably be expected of her.

Caring about something morally significant to a degree that surpasses the associated minimum threshold is *supererogatory caring*. It is not wrong to care to such a high degree. On the contrary, such a high degree of moral concern is intuitively quite praiseworthy. But nor is it wrong to fail to exceed the threshold, so long as one meets it; falling below the minimum threshold would be blameworthy, but it is not blameworthy to care about something to the exact degree required by the corresponding standard. Hence, caring to a degree that surpasses the threshold is supererogatory – as with other forms of supererogation, it goes “above and beyond” what is required, but is still good to do.

I hold that Grace, McKenzie, John, and Xavier all care supererogatorily about the very features that in fact make their acts wrong. We all should care about protecting the interests of marginalized groups in our communities, about making historically-white environments more welcoming for people of color, and about both non-violence and the prevention of harm to others. But Grace, McKenzie, John and Xavier care about these things to a degree that far surpasses that which could reasonably be expected of us. Thus, they go “above and beyond” in the manner characteristic of supererogation; their high degrees of moral concern are optional, but praiseworthy when exhibited.

This diagnosis makes sense of a natural set of reactions to have to these agents. It seems unfair – churlish, even – to reproach them for their moral ignorance. And if they were to learn that their actions are in fact morally wrong, the fitting response would not be guilt but rather something more like the “Oh no!” attitude one feels when one learns that one’s action was objectively wrong, although it was subjectively right. We can make sense of these intuitions by saying that, notwithstanding their moral ignorance and wrongful action, Grace, McKenzie, John and Xavier care supererogatorily about the very considerations that make their acts wrong. They are not blameworthy, may not be reproached, and need not feel guilty for the ignorance that remains despite their supererogatory care, because they are already surpassing the standards that specify what can reasonably be expected of them.

If I am right about this, then people can make moral mistakes because of supererogatory caring, just as they can make moral mistakes owing to a dearth of caring. To repeat (from §3): caring about something disposes us to pay attention to it, think about its moral significance, and think about the potential moral significance of an act’s having an impact on it. So if someone cares adequately about every morally significant thing at stake in her circumstances, but she cares *supererogatorily* about some of them, then these features may loom large in her moral thinking in a way that misleads her. For example, in **NON-VIOLENT PROTEST**, John concludes that he ought to fight back *because of how much he cares* about his fellow protesters, whereas Xavier concludes that he ought to restrain himself *because of how committed he is* to the principle of non-violence. Each of these agents cares about one of the morally significant features of their circumstances to a greater degree than the already-supererogatory degree to which the other cares about it. So, whomever reaches a

approach, the standards for adequate caring are standards describing the degrees and varieties of moral care that can reasonably be expected of people. I adopt this approach.

false view about how he should act does so not because he cares inadequately, but because he cares about morally significant features of his situation to a degree that more than adequate.

This approach raises a lot of outstanding questions. I will consider four.

One question concerns whether, in addition to standards specifying minimum *absolute levels* of care for each morally significant thing, there are also standards specifying the *relative degrees* of care for two or more morally significant things that someone must display in order to count as caring adequately. Perhaps, for some x and y where x 's moral significance is in fact greater than y 's, to care about y more than x is to care inadequately.⁹ More strongly, perhaps, for some x and y where x 's moral significance is precisely z times that of y , one cares inadequately unless one's degree of concern for x is precisely z times one's degree of concern for y . We can call standards of this form "relative standards", and we can call standards specifying minimum thresholds of adequate care for a single thing "absolute standards". So, one outstanding question is whether there are some relative standards in addition to my proposed absolute standards.

If there are relative standards, then **NON-VIOLENT PROTEST** as I have described it may in fact be impossible. For it might be that, whichever of non-violence and preventing harm is in fact more important in John and Xavier's circumstances, he who cares more about the other consideration thereby fails to care adequately by violating the relative standards. The same may hold in **GENTRIFICATION**; Grace may fail to care adequately in virtue of failing to care about social recognition for local Latina women artists *more* than she cares about their ability to continue living in her community, though she cares about each of these things to a more-than-adequate degree.

There is some reason to posit relative standards in addition to absolute standards, for they can explain what is wrong with someone's caring *selectively*. For example, suppose that someone cares a very great deal about promoting educational opportunities for children from their own ethnic group, but cares much less about promoting educational opportunities for other children. Even if the absolute degree to which he cares about promoting educational opportunities for other children is above the minimum threshold, the discrepancy between the degrees to which he cares about these two (roughly) equally important things still indicates that something is wrong with his degrees of moral care. Positing relative standards allows us to say what is wrong in an intuitively appealing way: the degrees to which he cares about multiple morally significant things do not "match" the facts about the relative importance of these things.

However, positing relative standards in addition to absolute standards would make the standards for adequate caring very precise. I doubt that they are quite so precise. The idea of "adequate" caring suggests a measure of sympathetic understanding afforded to those who make some mistakes while getting all the basics right. This is better captured by standards specifying minimum absolute levels of moral concern that one must show for each morally significant thing than by standards specifying the relative degrees of care that one must display toward any two or more morally significant things. The latter sort of standards are much harder to meet, raising the bar for adequate caring implausibly high. If there are no relative standards then we can say that, even if someone meets the standard for adequate caring with respect to some morally significant thing, it would still be better for her to care about it even more (though this is not required – it is supererogatory). But this does not hold if there are relative standards. If there are, then someone can

⁹ This way of thinking about the standards for adequate caring is suggested by Holly Smith's classic paper on culpable ignorance (1983), which construes ill will in terms of blameworthy "configurations" of attitudes. For Smith, what are blameworthy are things like someone's not caring enough about x to override her concern for y and cause her to act so as to promote x , rather than someone's not caring enough about x *simpliciter*. I argue against this view in what follows.

violate the standards for adequate caring just by increasing her degree of concern for one thing, caring about everything else to the same (perhaps quite high) degree as before, as she may thereby upset the ratio between two or more degrees of care in a way that violates the relative standards. This seems a little harsh.

More worryingly, if there are relative standards then someone can improve her moral care just by “leveling down”. Suppose that someone initially cares to an adequate or more-than-adequate degree about each morally significant thing, and that she then suddenly reduces the degree to which she cares about everything — her degrees of moral care go down across the board. But suppose that her degrees of care for different things reduce by different amounts, such that this reduction brings the ratios between her degrees of care closer to those specified by one of the relative standards. This is a strict deterioration in her moral caring; she cares about everything less than before. Yet relative standards would entail that such a deterioration can be praiseworthy. I find this implausible. Hence, I favor the simpler view that it is good to care about what is in fact morally significant, and better to care about these things more rather than less.¹⁰

Absolute standards can still incorporate information about the relative moral significance of things. If x is more morally significant than y , then the minimum threshold required for caring adequately about x will be higher than it is for y ; it takes more to care adequately about that which is in fact more important. This means that when people say things like “You focus too much on y , and neglect x ”, we need not hear this as an allegation that their addressee has violated relative standards. Instead, we can hear them as saying that their addressee falls short of the standards for adequate caring with respect to x , in a way that is causally connected to their degree of concern for y .¹¹ The real problem in such a case is not that the agent cares about x more than y , but that they don’t care enough about x .

Even if there are no relative standards, more caring may not always be better. For there may be *upper limits* on the degree to which someone may care about certain morally significant things and still count as caring adequately — in other words, there may be some morally significant things that one can care about too much as well as too little. This would mean that the standards for adequate caring have a structure such that there is a threshold below which one cares inadequately, a range of permissible degrees of care, and then an upper limit above which one also cares inadequately. (Alternatively, there may be three thresholds for each morally significant thing: one below which one cares inadequately, one above which one cares supererogatorily, and one above which one cares altogether too much.) So, a second outstanding question about the standards for adequate caring is whether there are any such upper limits.

I grant that we often speak of people caring “too much” about something. But I am not sure that this is good evidence of upper limits. Talk of caring too much is not always best understood as an objection to someone’s absolute degree of concern. Rather, it is often an objection to the *manner* in which the person

¹⁰ There is a way of thinking about what caring is according to which my view makes no sense, and relative standards are the only kind of standards for adequate caring that there can be. On one way of thinking, all the facts about how much someone cares about something are relative; we determine how much *more* someone cares about x than y by seeing how much y she would forego in order to get x , but we don’t have a way of determining the absolute degree to which she cares about x independently of these comparisons. This is suggested by a behaviorist approach, or one that understands caring in terms of utilities (since utility-functions have an arbitrary zero point). But I hold that this way of thinking is mistaken. There are facts about the absolute degree to which each person cares about each thing. This means that two people could care about everything to precisely the same *relative* degrees, thus having the same behavioral dispositions and the same utility-function, while one cares about everything to a muted degree while the other cares about everything with passionate intensity. In short: some people just care more than others.

¹¹ This might be because, given the agent’s limited cognitive capacities, their degree of concern for x “crowds out” their degree of concern for y . I am not sure of the extent to which this is a genuine cognitive phenomenon.

cares; they care obsessively, pathologically, or in some other morally inappropriate way. Degree and manner come apart, such that even a low degree of care can manifest pathologically, leading to anxiety, intrusive thoughts and the like. For example, someone suffering from obsessive-compulsive disorder may feel compelled to turn light switches on and off a certain number of times before leaving a room, but it would be a stretch to say that she thereby manifests a high degree of care for the flipping of light switches. And, when it is not an objection to the way someone cares, the accusation of caring “too much” is often best understood along the lines suggested above: there is some other thing about which the person cares inadequately, and their failure to care adequately about this other thing is causally connected to their high degree of concern for that which we say they care about too much. In this case, the agent’s high degree of concern is not the real problem. Moreover, I am inclined to think that a high degree of concern for some morally significant thing is still praiseworthy in and of itself, even if it is causally related to a blameworthy lacuna in the agent’s degree of moral concern for another thing. In other words: it is bad that she doesn’t care enough about y , but her caring a lot about x is not thereby impugned — that’s the good part.

A third outstanding question concerns *trade-offs* between different ways to meet the standards for adequate caring. The notion of “trade-offs” finds purchase here because there are many ways to manifest care. For example, forming moral beliefs about the object of one’s care, acting to promote the object of one’s care, and showing distress when the object of one’s care is threatened are all ways of manifesting care – as we saw in §3. So, the total extent to which someone manifests care for something morally significant depends on the extent to which she does each of these three things, and maybe more. Since caring adequately is not caring maximally, one need not do absolutely everything that would manifest care for a morally significant thing in order to count as caring adequately about it; one can fail to do some things that would manifest care but “make up for it” by doing others. One must reach the minimum threshold. But there may be many ways to reach it. The question is what they are.

This question is too large to address in full here. But we can begin to think about trade-offs by reflecting on the cases from §3. In **NAMES**, there are different ways for McKenzie to show that she cares about making her discipline more welcoming for people of color: developing and advocating for policy initiatives on the one hand, and using non-white names in examples on the other. It does not seem to constitute inadequate caring for her not to do the latter, given that she does the former. In this case, then, it seems that McKenzie fails to manifest care in one respect but makes up for it by manifesting care for the same thing in other more important respects. Something similar holds in **GENTRIFICATION**; Grace could have shown that she cares about the interests of marginalized groups in her community by promoting the gallery on the grounds that it furthers their interests, but she instead shows that she cares by boycotting the gallery, since she acts as she does because she thinks that the gallery will hamper the interests of marginalized groups. This does not seem to constitute inadequate caring, though it does constitute a failure to do what would in fact have promoted people’s interests. So, it looks as though someone can make up for failures to do certain things that would manifest concern for a cause even by doing other things that in fact inhibit the cause, so long as this activity nonetheless still shows how much she cares about the cause. In other words, the standards for adequate caring seem to license trade-offs whether or not the agent actually succeeds in promoting the object of her care through the activity that manifests it. Thus Grace can manifest care by forming and acting on false moral beliefs.

That said, I suspect that there are some limits on permissible trade-offs between ways of showing care. For example, I doubt that someone counts as caring adequately just in virtue of forming lots of moral beliefs about something, if she has ample opportunities to act so as to promote the object of her care but never takes any of them. (One may lack opportunities to show care through action due to distance, inadequate resources, disabling mental health difficulties, etc.) These limits on permissible trade-offs might turn out to

entail that certain moral beliefs are necessary for caring adequately. But I am unsure about this. Someone who cares about something should want her understanding of its nature to be as complete as possible, to avoid inadvertently undermining it while trying to promote it. So she should pay attention to the object of her care. And perhaps certain moral facts are so obvious that anyone paying the amount of attention involved in adequate caring would figure them out; for example, perhaps anyone paying an adequate amount of attention to a happy, healthy, adult human being would realize that it is *pro tanto* wrong to kill them. But it is hard to say what proportion of moral facts are like this, if any are. Whether an agent paying adequate attention is guaranteed to form a moral belief depends on how obvious the moral fact is, and on how good the agent is at moral reasoning.¹² I leave this topic as a promising avenue for future research.

The third question concerned trade-offs between ways to show care for one morally significant thing — trade-offs *within* values, so to speak. A fourth question concerns trade-offs between ways of showing care for different morally significant things — trade-offs *between* values, so to speak. The question is: if someone falls below the standards for adequate caring with respect to some morally significant thing *x*, can she make up for it by caring supererogatorily about a distinct thing *y*, such that she still counts as caring adequately about everything morally significant overall?

Everything that I have said so far is neutral on this question. To clarify: one may think that **NON-VIOLENT PROTEST** is a case in which the standards for adequate caring license trade-offs between values. But this is not so, on my diagnosis of the case. I hold that John and Xavier care supererogatorily about all morally significant things in their circumstances. So this is not a case of meeting the standards for adequate caring by making up for inadequate caring with supererogatory caring. It is rather a case of meeting the standards for adequate caring with supererogatory caring across the board.

There are cases of someone devoting her entire life to promoting one morally significant cause, at the expense of other morally significant things. For example, someone might devote her life to the prevention of long-term environmental damage at the expense of promoting the well-being of her immediate family members. Or she might devote her life to promoting the well-being of her immediate family members at the expense of preventing long-term environmental damage. For some ways of filling out the details of such lives, it would seem inappropriate to reproach the agents for not caring more about the morally significant thing that they have neglected, and it would seem churlish to say that they do not count as caring adequately overall. So perhaps these are instances of trade-offs between values. But, again, there are significant limits. Notably, it does not seem as though someone can make up for active hostility toward one morally significant thing by caring about other, entirely separate, morally significant things. For instance, if we learned that a brilliant climate scientist had horribly abused her wife and children while working to prevent long-term environmental damage, I suspect we would be reluctant to say that she cared adequately about everything morally significant. This is why I suspect that the standards for adequate caring have the threshold structure that I have described; there are baseline levels of minimal requisite care for each morally significant thing, and someone does not count as caring adequately about *everything* morally significant if she falls below the baseline level with respect to *anything* morally significant.

It might still be true that the brilliant climate scientist does not have to do as much in order to count as caring adequately about her wife and kids than someone with fewer morally momentous commitments must do in order to count as caring adequately about his wife and kids. This might hold if the standards for adequate caring are *agent-relative*, demanding different amounts of care of different people toward the same things. Agent-relativity would allow for something akin to trade-offs between values: a choice to

¹² See the related remarks on “obvious” moral facts in Groll and Decker (2014).

“specialize” in promoting a certain cause may create special obligations, which then make it the case that the agent’s thresholds for caring adequately about other morally significant things are lower than they otherwise would be. On this picture the standards for adequate caring still have the threshold structure that I have described: it is still true that, for each pair of an agent and a morally significant thing, there is a fact about whether she meets, falls short of, or exceeds the standard for adequate caring with respect to this thing, which determines whether she cares inadequately, adequately, or supererogatorily about the thing. The only difference is that the standards for adequate caring may place the threshold for caring adequately about a certain thing in different places for different people. And, indeed, there are reasons to posit some agent-relativity independently of what we think about trade-offs between values; for example, it is very natural to think that the degree to which someone should care about her own family is higher than the degree to which someone else should care about her family. I leave the topic of agent-relativity in the standards for adequate caring as another fruitful topic for future research.

6. Conclusion

Here are some practical suggestions for how to move forward.

I have suggested that quality-of-will theorists can figure out what our approach implies about culpability for moral ignorance only by rolling up our sleeves and doing the hard work of figuring out what it is to care adequately. We must work toward articulating the nature and structure of the standards for adequate caring.

In so doing, I suggest that we focus on examples that reflect the complexity of everyday moral life. It is easy to figure out that stealing is (typically) wrong. So, it is natural to think that someone is blameworthy if she knows that she is stealing and yet fails to realize that what she is doing is wrong. But it is much harder to maintain that someone is blameworthy if she knows that she is protesting an art gallery that will contribute to rising house prices in a formerly-low-income neighborhood, perhaps forcing some people out of their homes, but that also offers valuable social recognition to local Latina women artists, and knows that her activist group (including some of the artists) has concluded that the gallery has a net negative impact on the community overall, yet fails to realize that what she is doing is wrong. With realistically complex cases in view, it seems seriously mistaken to think that all moral ignorance involves a failure to care adequately. It is thus mistaken to take the quality-of-will approach to imply that all moral ignorance is blameworthy.

Readers may have qualms about some of the particular examples I have offered. If so, they are welcome to devise cases of their own with the same structure. For a case like **GENTRIFICATION**, the important thing is that the agent cares passionately about something morally significant, but is partially mistaken about its nature, such that her attempt to promote it ends up undermining it. For a case like **NAMES**, the important thing is that someone is devoted to a morally significant cause, but fails to notice that something she does undermines it in a small way that could easily have been avoided at no cost. For a case like **NON-VIOLENT PROTEST**, the important thing is that someone faces a clear conflict between two things of genuine moral significance, and she recognizes this, but reaches the wrong all-things-considered judgment because of how much she cares about the thing that is in fact less important in her circumstances. It is easy to come up with cases like these. And, I maintain, they are all cases in which someone acts wrongly despite caring at least adequately about everything morally significant, including the very considerations that in fact make her act wrong. As I have said, the reader should be able to construct parallel cases of her own just as long as she thinks that there are any moral facts that it is possible to fail to know while caring adequately – i.e., as long as she thinks that the standards for adequate caring do not require moral omniscience.

My view has a further practical upshot. On my view, people can be somewhat morally ignorant even while caring adequately – and perhaps supererogatorily – about everything that is in fact morally significant. But there are clearly also a lot of people who care *inadequately* about some things that are morally significant, and some people are morally ignorant because of their dearth of caring. There is a lot of motivated moral ignorance out there, and there is a dispiriting amount of indifference to considerations that do in fact matter morally. On my view, these things are clearly blameworthy. So the practical upshot of my view is that we should stop nitpicking and in-fighting; we should stop blaming people who care at least adequately about everything morally significant, but make some moral mistakes even so, and we should redirect our blame toward those who fail even to care adequately. This has considerable practical bite. It suggests, for instance, that we should not waste time throwing blame back and forth with friends who all care adequately and are sincerely trying to promote that which genuinely matters, even if they do so somewhat ineptly, when we all recognize that there are common enemies to fight.

REFERENCES

- Alcoff, Linda Marie (2007). "Epistemologies of ignorance: three types". In Shannon Sullivan and Nancy Tuana (eds.), *Race and Epistemologies of Ignorance*. Albany, NY: State University of New York Press.
- Arpaly, Nomy (2002). *Unprincipled Virtue: An Inquiry Into Moral Agency*. Oxford: OUP.
- Arpaly, Nomy and Schroeder, Timothy (2013). *In Praise of Desire*. New York, NY: Oxford University Press.
- Björnsson, Gunnar (2017). "Explaining away epistemic skepticism about culpability". In Shoemaker, David (ed.), *Oxford Studies in Agency and Responsibility*. Oxford: OUP. pp. 141–164.
- Campbell, Richmond (2007). "What is moral judgment?", *The Journal of Philosophy* 104(7), 321-349.
- FitzPatrick, William (2008). "Moral responsibility and normative ignorance: answering a new skeptical challenge", *Ethics* 118, 589-613.
- Groll, Daniel, and Decker, Jason (2014). "Moral testimony: One of these things is just like the others". *Analytic Philosophy* 55(1), 54-74.
- Harman, Elizabeth (2011). "Does moral ignorance exculpate?", *Ratio* 24, 443-468.
- Levy, Neil (2011). "Hard luck: How luck undermines free will and moral responsibility". Oxford: OUP.

- Mills, Charles (2007). "White ignorance". In Shannon Sullivan and Nancy Tuana (eds.), *Race and Epistemologies of Ignorance*. Albany, NY: State University of New York Press.
- Moody-Adams, Michele (1994). "Culture, responsibility, and affected ignorance", *Ethics* 104, 291-309.
- Pohlhaus Jr., Gail (2012). "Relational knowing and epistemic injustice: toward a theory of willful hermeneutical ignorance". *Hypatia* 27(4), 715-735.
- Roberts, Debbie (2013). "Thick concepts". *Philosophy Compass* 8(8), 677-688.
- Rosen, Gideon (2002). "Culpability and Ignorance" *Proceedings of the Aristotelian Society* 103(1), 61-84.
- Rosen, Gideon (2004). "Skepticism about moral responsibility", *Philosophical Perspectives* 18, 295-313.
- Rosen, Gideon (2008). "Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility". *Journal of Philosophy* 105(10), 591-610.
- Smith, Angela (2005). "Responsibility for attitudes: activity and passivity in mental life", *Ethics* 115, 236-271.
- Smith, Holly (1983). "Culpable ignorance". *The Philosophical Review* 92(4), 543-571.
- Stocker, Michael (1979). "Desiring the bad: An essay in moral psychology". *Journal of Philosophy* 76(12), 738-753.
- Strawson, Peter (1962). "Freedom and resentment", *Proceedings of the British Academy* 48, 178-211.
- Svavarsdóttir, Sigrún (1999). "Moral cognitivism and motivation". *The Philosophical Review* 108(2), 161-219.
- Vayrynen, Pekka (2013). *The Lewd, the Rude, and the Nasty: a Study of Thick Concepts in Ethics*. Oxford: OUP.
- Yaffe, Gideon (2018). "The Point of Mens Rea: The Case of Willful Ignorance." *Criminal Law and Philosophy* 12, 19-44.

Note: Three further references have been redacted for blind review.