



Reasonable standards and exculpating moral ignorance

Nathan Biebel^{1,2}

Accepted: 21 September 2023 / Published online: 2 November 2023
© The Author(s) 2023

Abstract

It is widely agreed that ignorance of fact exculpates, but does moral ignorance exculpate? If so, does it exculpate in the same way as non-moral ignorance? In this paper I will argue that on one family of views explaining exculpating non-moral ignorance also explains exculpating moral ignorance. The view can be loosely stated in the following way: ignorance counts as an excuse only if it is not the result of a failure to meet some applicable reasonable epistemic standard—call this the Reasonable Epistemic Standards Thesis and call views that accept some version of this principle reasonable standards views. I argue that any plausible reasonable standards view ought to allow that moral ignorance exculpates, at least sometimes, and defend such views against the charge that they are susceptible to clear counterexamples.

Keywords Moral responsibility · Ignorance · Excuse · Moral ignorance

1 Introduction

It is widely agreed that ignorance of fact sometimes exculpates, but does *moral* ignorance exculpate? If so, does it exculpate in the same way as non-moral ignorance? The answer to these questions might seem straightforward: False moral beliefs that lead one to commit wrong actions should exculpate under the same conditions that non-moral ignorance exculpates. But this claim is the subject of considerable disagreement. Some theorists claim that theories that seem to get the right answer in cases of non-moral ignorance commit us to excusing paradigm cases of blameworthiness when applied to moral ignorance.¹

¹ Indeed, some theorists have claimed that moral ignorance never exculpates (Harman 2011, 2015, 2017, 2019; Talbert 2013, 2017; Arpaly 2003; Hieronymi 2008; Guerrero 2007).

✉ Nathan Biebel
nbiebel@tulane.edu

¹ Jagiellonian University, Kraków, Poland

² African Center for Epistemology and Philosophy of Science, University of Johannesburg, Johannesburg, South Africa

In this paper I will argue that on one family of views explaining exculpating non-moral ignorance also explains exculpating moral ignorance. The view can be loosely stated in the following way: ignorance counts as an excuse only if it is not the result of a failure to meet some applicable reasonable epistemic standard—call this the Reasonable Epistemic Standards Thesis and views that accept some version of it I shall call reasonable standards views. The paper will proceed in the following way. In Sect. 2 I explain what constitutes moral ignorance and introduce some paradigm cases of wrongdoing from moral ignorance in the literature. In Sect. 3 I explain reasonable standards views as they apply to non-moral ignorance and put forth a theory-neutral general schematic that any plausible reasonable epistemic standard will fit. In Sect. 4 I take that general schematic and test it against our intuitions in the paradigm cases of moral ignorance from Sect. 2. However, those are all cases where we are meant to have the intuition that the agent in each case is blameworthy, so in Sect. 5 I discuss some examples of cases where reasonable epistemic standards yield the plausible verdict that an agent is not blameworthy as a result of acting from moral ignorance. Finally, in Sect. 6 I defend this Account against purported counterexamples.

2 Moral ignorance

Imagine the following set of cases found in the literature.

Jeff is a middle-aged middle manager in a mid-size company located somewhere in the Midwest. To him has fallen the task of alerting “downsized” employees of their new status as job seekers in a gloomy economy. That Jeff has the task is unfortunate for those about to be laid off, not only because they are about to lose their jobs, but—to add insult to injury—because Jeff is a jerk. He is rude and inconsiderate about the feelings of others. And he is unreflective about it. When people react poorly to his behavior...he always writes it off as a shortcoming on the part of others. One afternoon, his superiors tell him that he needs to give notice to a group of long-time employees that they will be laid off. He does tell them, but in an altogether rude and insensitive fashion. (Vargas, 2005, pg. 271).

Consider Mr. Potter, a powerful businessman who holds false moral views. He takes certain business practices—such as liquidating Baily’s Building and Loan and sticking it to the poor families of Bedford Falls—to be “permissibly aggressive,” when in fact they’re “reprehensibly ruthless.” This leads him to do bad things, though he doesn’t understand that he’s acting badly, which means that he’s acting out of a certain kind of ignorance. He’s fully aware of

all the circumstances, but he applies flawed normative principles...and comes up with bad decisions. (Fitzpatrick, 2008, pg. 599–600).²

JoJo is the favorite son of Joe the First, an evil and sadistic dictator of a small, undeveloped country. Because of his father's special feelings for the boy, JoJo is given a special education and is allowed to accompany his father and observe his daily routine. In light of this treatment, it is not surprising that little JoJo takes his father as a role model and develops values very much like Dad's. As an adult, he does many of the same sorts of things his father did, including sending people to prison or to death or to torture chambers on the basis of whim. He is not *coerced* to do these things, he acts according to his own desires. Moreover, these are desires he wholly *wants* to have. When he steps back and asks, "Do I really want to be this sort of person?" his answer is resoundingly "Yes," for this way of life expresses a crazy sort of power that forms part of his deepest ideal. (Wolf, 2003, pg.379, emphasis in the original).

For my purposes here, following Mason (2017), I will understand these cases as instances of *moral ignorance*. In what follows, I shall take moral ignorance to be ignorance that has as its content the moral status of some feature(s) of one's action. That is, moral ignorance happens any time that one fails to believe, either because one believes it is false or because one fails to positively believe that it is true, *that some action or attitude is wrong*. On this understanding of the cases above, the agents in question are all morally ignorant; each of these agents fails to believe their action is wrong. Importantly, none of these agents are ignorant of the facts on the ground. What they are ignorant about is that these facts *make their actions wrong*. Indeed, I shall treat these cases as *paradigm* cases of moral ignorance. With this understanding of moral ignorance in hand, in the next section I will explain reasonable standards views and the general Reasonable Epistemic Standards Thesis to which I think they all subscribe.

3 Reasonable epistemic standards

To understand reasonable standards views, it will be helpful to consider a case of non-moral ignorance. Imagine a doctor is seeing a patient. Despite knowing that it is her responsibility to check every patient's chart thoroughly, she only gives her current patient's chart a cursory glance. This cursory glance causes the doctor to miss a note stating that the patient is allergic to penicillin. As it turns out, penicillin is the typical prescribed medication for what ails the patient. Because the doctor is ignorant of the patient's allergy, the doctor prescribes penicillin. The patient has a severe allergic reaction and nearly dies.

² Fitzpatrick develops this case based on the titular character of Mr. Potter from Frank Capra's famous film *It's a Wonderful Life* (1946), but it is meant to be a development of a more generic case offered by Rosen (2004).

Intuitively, the doctor is blameworthy for the harm caused to the patient despite her ignorance of the patient's allergy. Since ignorance is often an excuse, we need to explain why this doctor's ignorance is not. According to the family of views under consideration, the reason the doctor's ignorance is not an excuse is that it results from her failure to meet an applicable *reasonable epistemic standard*. There are many fine-grained ways of defining what counts as 'reasonable' in this epistemic sense.³ Her ignorance is unjustified (Biebel, 2018; Gibbons, 2013); her beliefs are not formed responsibly (Peels, 2017); she has been negligent in the management of her opinion (Rosen, 2003, 2004, 2008; Zimmerman 1997, 2008; Smith, 1983); her lack of belief is the result of an epistemic vice (Fitzpatrick, 2008, 2017; Mason & Wilson, 2017; Montmarquet, 1995); she has failed to exercise her rational capacities when she should (Clarke, 2014; Sher 2009; Ayars, 2021); she has failed to care adequately about discovering the moral features of her action (Johnson King 2019).⁴ These descriptions (and more) are different ways of understanding what it means to say that the doctor's ignorance was the result of a failure to meet some applicable reasonable epistemic standard.

We can represent a general form of this family of views with the following principle.

Reasonable epistemic standards thesis (REST)

One's ignorance is not exculpating only if (i) there is an applicable epistemic standard that it is reasonable to demand that one meet, and (ii) one's ignorance is the result of a failure to meet that standard.

For my purposes here, we need not commit to any specific explanation of this reasonable epistemic standard. We need only employ the common-sense course-grained schematic notion of reasonable epistemic standards that the above descriptions attempt to capture in more fine-grained ways. This general common-sense schematic includes two criteria. (1) The resultant belief (or more properly, doxastic state) is sufficiently supported by one's accessible evidence and (2) this accessibility is to be understood in terms that include some degree of reflection that avoids obvious reasoning errors and some degree of minimal required effort to gather evidence—especially counterevidence.⁵

³ Plantinga (1993), and Goldman (1999), Alston (1988), Feldman (1988b), and Oliveira (2015, 2018) all argue that deontic notions do not apply to beliefs because beliefs are not voluntary. The responses have been numerous. See, for example, Chrisman (2008), Chuard and Southwood (2009), Huss (2009), McHugh (2013), Nickel (2010), Nettleman (2007), Steup (2008; 2017), Weatherson (2008), Peels (2017), and Hieronymi (2008) among many others.

⁴ In the interests of transparency, Johnson King doesn't put it quite like this. But she explains that "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 the thing." (pg. 421). So, if I understand correctly, for Johnson King, caring adequately (often) manifests itself as adequate epistemic effort to discover the features of one's act that make it wrong.

⁵ It might be helpful to think of this in terms of negligence, where the required effort is based on the risk involved in being wrong. See King (2009, 2014) and Hart (1968).

While this common-sense notion of a reasonable standard is meant to be a theory-neutral schematic, and as such is left purposefully vague, some clarification is necessary. First, the degree of reflection required and the steps one must take to gather more evidence depend on one's context. A baseball umpire is not open to criticism if he calls a close pitch a strike when he would have called it a ball had he seen it from another angle. A home buyer *is* open to criticism if she decides to purchase a home prior to receiving an inspection report that reveals a serious flaw in the house.⁶ Gaus (2010, ch. 13.5) helpfully remarks that our ordinary practices indicate that there are two parameters that determine any plausible, morally relevant, reasonable epistemic standard—a minimum and a maximum. On the one hand, we must not require that ordinary people perform extraordinary epistemic feats to avoid blame. In other words, the standard cannot be extraordinarily difficult to meet because it is meant to apply to all moral agents. On the other hand, we must not have a standard that is so easy to meet that even those making obvious mistakes meet the standard. After all, “It is no excuse for making an obvious mistake that the correct answer is a little less obvious” (pg. 248). Moreover, we expect that people will often fail to meet this standard. “Our very practice of morality presupposes that the relevant reasons are sometimes not very easily accessible. Moral reasons are not always glaring; it often...takes some thought to see what the right thing to do is” (pg. 256). In short, the standard must be set such that everyone to which it applies *can* meet it, but it must not be set such that everyone always (or too often) *does* meet it.

The second point of clarification is that the standard does not merely require that one do *a lot* of thinking about an issue. There are many examples of beliefs that fail to meet a reasonable standard because they are clearly and obviously defeated by the evidence, and *this can be true even if the believer has undergone a great deal of reflection*. Consider beliefs like “the climate is not changing due to human activity,” “vaccinating children is linked to autism,” or “the Earth is flat.” This is important because the reason these beliefs are treated as sub-standard is that we operate under a certain defeasible presumption that has not yet been adequately acknowledged in the literature.⁷ The presumption is the agents whose beliefs we are evaluating are *sufficiently similar* to us in competence, capacities, and circumstances. Some propositions, say, that the Earth is flat, are so obviously ruled out by the evidence that we judge it unreasonable for anyone *sufficiently like us* in rational capacity and epistemic circumstances to believe it—that is, nobody to whom the standard applies lacks access to evidence that undermines a belief that the Earth is flat. It is important to note, however, that this presumption is defeasible. When we discover a person is not sufficiently like us in their ability to access evidence, then we no longer hold them to the same standard. It is unreasonable for people like us, in the twenty-first century

⁶ This example is taken from Gaus (2010 ch. 13).

⁷ Though, see Field (2021) who argues that we ought to take into account how difficult it might be for certain neuroatypical individuals to discover the moral truth when assessing their blameworthiness. Here, I expand this to include factors that apply to neurotypical individuals as well, though I think Field's view can easily accommodate this.

with access to the internet and modern education to believe that the Earth is flat. For the average person born in the fifth century BCE the opposite might be true.⁸

Note that the schematic account of reasonable epistemic standards described above only demands that there are upper and lower boundaries for what we can require. This allows for a great deal of wiggle room inside of those boundaries. Indeed, we allow a great deal of room for differences of opinion; two people with similar intellectual capacities with the same set of evidence may come to different conclusions while neither has violated any reasonable epistemic standard.⁹ Despite considerable room for disagreement within the boundaries of this reasonable standard, and despite the necessarily vague nature of those upper and lower limits, in practice we have a fairly good intuitive grasp of what clearly counts as reasonable (and what clearly doesn't). Consider again our doctor. Her evidence is that the patient's chart reveals that the patient is allergic to penicillin. It is, at minimum, reasonable to expect a competent doctor to read her patient's chart carefully, and this is the sense in which the evidence is accessible to her—if she had done what was minimally required, then presumably she would have seen that the patient has this allergy. It would be odd, though, to say that the evidence is accessible to her if the chart made no mention of an allergy. Perhaps she could run a test if she were suspicious, but if there is no indication of an allergy then she has no reason to be suspicious and hence it would be expecting too much to demand she take steps to investigate.

Some theorists have argued that their individual, more fine-grained versions of REST extend to moral ignorance (Rosen 2004; Zimmerman, 2017; Levy, 2009; Sher, 2017; Fitzpatrick, 2008, 2017). While I find these arguments plausible, we need not commit ourselves to any particular view to establish that REST extends to moral ignorance in its most general form. As such, my claim in what follows is that anyone who holds a reasonable standards view ought to agree that moral ignorance can exculpate. As such, in the next section I will test REST against the three cases of moral ignorance from Sect. 2 to see if it gets the right conclusions. Each of those three cases is meant to be one in which the agent is morally ignorant, and yet they are intuitively blameworthy. So, I will consider it a success if a plausible explanation in each case is that we judge that they remain blameworthy despite their ignorance because the agent fails to meet an applicable reasonable epistemic standard. This is only half the job, however, for we would still need to establish that there are cases in which moral ignorance exculpates. In Sect. 5 I discuss a few such cases from the literature and also offer my own. If any of those cases succeed, then it should establish the prima facie plausibility of the claim that REST applies to moral ignorance.

⁸ This, of course, is an empirical question about what facts regarding the shape of the Earth were available to the ordinary person in the fifth century BCE. Rosen (2003) makes a similar point, though he does not mention the sufficient similarity requirement. See also Weiland (2017) who puts this in terms of accessibility, though Weiland never really explains what "accessible" means.

⁹ There has been much recent debate among epistemologists about whether epistemic peers with the same evidence can rationally disagree. See Feldman & Warfield (2010) for various positions in that debate.

4 Paradigm cases of moral ignorance

To see how REST explains our ordinary judgments we will need to know what the empirical evidence suggests those judgments are. Recall the case of JoJo. Wolf (1987 pg. 54) takes it for granted that people will share the intuition that JoJo is not blameworthy for his horrific behavior, but recent empirical evidence suggests otherwise. Faraci and Shoemaker (2010) tested subjects who were told the story of JoJo and compared their results to a group of control subjects, which were given a version of the case describing JoJo's father Jo the First.¹⁰ Jo the First is a tyrannical leader who regularly tortures innocents for fun, and he believes it is his right to do so. However, Jo the First was not subject to JoJo's seclusive upbringing, he arrived at his moral ignorance on his own. The purpose of Faraci and Shoemaker's test was to see whether Wolf is correct that people intuit that JoJo is *not blameworthy* in virtue of his bad upbringing. The control group unsurprisingly blamed Jo the First to nearly the highest degree possible. Interestingly, while subjects allotted JoJo slight leniency compared to Jo the First, they still judged JoJo blameworthy to a *significant* degree.¹¹ I want to highlight two things in this result. First, it strongly suggests that subjects *do not believe* that JoJo's moral ignorance is fully exculpating. The second is that, while he is not *fully* excused, the amount of blame we think JoJo deserves is mitigated somewhat compared to his father.¹² This, I think, is just what REST predicts.

First, consider Faraci and Shoemaker's original motivation for conducting their studies. I quote at length.

Wolf takes it as a fundamental datum that our pre-theoretical intuitions converge on the judgment that JoJo is not a responsible agent. But do they? Whenever we introduced the case to students, they always needed considerable coaching to come to the conclusion Wolf wants. They resisted the idea that JoJo is not responsible *primarily because they found it extremely hard to believe that JoJo would not be able to recognize that torture is wrong*. In response, various features of the case would have to be stressed or exaggerated, e.g. the isolation in which JoJo grew up, in a "small, undeveloped country"—an island, it was proffered, with no communication links to the outside world, with heavily propagandized internal media, etc. Eventually the students would reluctantly agree to the "intuition," but at a price: the case now seemed quite precious. JoJo was now taken to live in an airtight vacuum, cut off from the world as we know it, and he was being rescued from responsibility...by a kind of forced and surreal *ignorance*. When the case was brought back into the real world, focused on someone like Uday Hussein, son of Saddam, the intui-

¹⁰ Faraci and Shoemaker (2014, 2017) performed tests of a similar nature that they (and I) take to replicate and confirm these results.

¹¹ Jo the First was rated at a mean of around 6 out of 7 (where 7 was "completely blameworthy" and 1 was "not at all blameworthy") and JoJo was rated at a mean of around 5 out of 7.

¹² See Sliwa (2020) for an argument that moral ignorance never fully exculpates, but it can mitigate blameworthiness. More on Sliwa's view below.

tion that he was responsible seemed to return in full force. (Faraci and Shoemaker 2010, pg. 324, first emphasis added, second emphasis original)

One reason that extensive coaching is needed to drag Wolf's "intuition" out of students comes directly from the way the case is described. The case says, "He is not *coerced* to do these things, he acts according to his own desires... When he steps back and asks, 'Do I really want to be this sort of person?' his answer is resoundingly 'Yes,.'" Wolf's intention for this description no doubt was to make JoJo seem insane because her point is that there is a "sanity" condition on moral responsibility. However, rather than showing JoJo insane, it serves the opposite function. It reinforces the standing presumption mentioned in the previous section, namely, that JoJo is *sufficiently like us*. If JoJo is rational enough to stop and think over what he is doing, then JoJo is subject to the same reasonable standards as everyone else—a standard that we do not impose upon the insane because the insane are not sufficiently like us.¹³ The fact that he gets things wrong is not, as far as our intuitions are concerned, an indication that he's insane so much as an indication that he's just ignoring what is plainly before him.

Additionally, the fact that generating the intuition that he is not blameworthy requires describing JoJo in complete isolation indicates that we think of JoJo (the JoJo described in Wolf's original case) as having access to the same set of evidence as us. What evidence is that? The evidence before him is empirical in that he is aware that he is hurting people and he is aware that they object to his hurting them. If JoJo really is sufficiently like us, then these facts ought to be enough for him to *see* that he has decisive reason not to believe his behavior is permissible. Of course, if he is not sufficiently like us then JoJo is not subject to the same reasonable standards, likely because he really is insane in the way Wolf suggests. However, Faraci and Shoemaker's evidence indicates that, regardless of whether it is true, we *judge* JoJo like he is one of us—a participant in the social-moral community to which a reasonable standard applies. Since we see him as sufficiently like us, we treat JoJo like we treat someone who believes the Earth is flat. In other words, we think JoJo's moral ignorance must be the result of JoJo ignoring evidence that is *so easy to acquire and interpret* that we think that only people in complete isolation, cut off from the outside world lack access to it.¹⁴

There is another interesting feature of these findings. Recall that Faraci and Shoemaker found that people tend to mitigate blame for Wolf's version of JoJo to a small degree. As Faraci and Shoemaker hypothesize (and I agree), our judgements of blameworthiness for JoJo seem to give him some small amount of credit for the *difficulty* of his circumstances. JoJo has exceptionally good, even obvious evidence before him that what he is doing is wrong, but we blame JoJo *less* than Jo the first because it is, by hypothesis, more difficult for JoJo to see and understand his evidence. His belief is well outside the boundaries of what can be considered reasonable for any normal functioning moral agent, but we give JoJo some small leeway because his circumstances make

¹³ The insane are typically regarded as exempt from our usual moral judgments. See Strawson (1962). However, as Field (2021; 2022) rightly points out, other neuroatypical conditions may render one blameless without making one *exempt* from our usual moral judgements.

¹⁴ Indeed, Harman (2019) claims that even people in complete isolation have access to such evidence.

it difficult for him *to function* as a normal moral agent. We do not view him as exempt from moral judgement as Wolf wants us to, but we do recognize that he is in circumstances where exercising his capacities as a moral agent is difficult.¹⁵

So, REST seems consistent with our intuitions about JoJo,¹⁶ what of Mr. Potter? There is no reason to think the same things Faraci and Shoemaker found would not also hold in the case of Mr. Potter. Again, unless we conceive of him as being *incapable* of coming to the right belief, we are not moved by the fact that, even on reflection, he comes to an erroneous conclusion. People in general regularly reflect and draw the wrong conclusions. In many cases those beliefs run counter to clear evidence to the contrary. Such beliefs are unreasonable. Even if one who is sufficiently like us spends a lot of time reflecting carefully on one's belief that the Earth is flat, it is still an unreasonable belief given the accessible evidence. Mere reflection does not make one immune to error, culpably biased reasoning, or culpably bad evidence gathering. Does Mr. Potter have access to clear evidence? Yes. He knows that sticking it to the poor people of Bedford Falls would make their lives unnecessarily difficult, and he further knows that they strenuously object.¹⁷ Our intuitions (and practices) indicate that we think a competent moral agent—the type of agent to which REST applies—would, all else equal, see those facts as decisive reason to believe one's actions are wrong.¹⁸

What of Vargas's case of Jeff the Jerk? As Vargas rightly points out, the strong intuition is that Jeff is blameworthy. Vargas tries to generate doubt about Jeff's blameworthiness by describing the history of Jeff's jerk-like behavior: As a teenager Jeff had a powerful desire to appear attractive to the opposite sex. Jeff noticed that others who behaved like jerks seemed to be attractive to the opposite sex and, being young and full of hormones that addle his judgment, Jeff decided to behave like a jerk. His tactic worked, and this began Jeff on a long road of continual jerk-like behavior until acting like a jerk became deeply embedded in his character. Vargas thinks that Jeff's teenage decision is the real target of reasonable standards views in that it is where his jerk-like character originated. However, Jeff clearly could not

¹⁵ Note, also, that there's nothing about the case that implies JoJo is neuroatypical in Field (2021)'s sense. It is his circumstances that cause his epistemic difficulties, not his neurology.

¹⁶ I want to stress here that I do not mean to criticize Wolf's conclusions; it may be that JoJo really is insane. My aim here has only been to make sense of our intuitive judgements, and it seems that REST accomplishes this.

¹⁷ Many in the disagreement literature argue that the fact that one's peers (e.g. people sufficiently similar in evidence and rational capacity) disagree counts as a reason to suspend judgement. See, among others, Christiansen (2007), Feldman (2006), Elga (2007), and Kornblith (2010).

¹⁸ Fitzpatrick (2008, pg. 605), who originally gave the case (though see fn. 2 above), suggests that, as long as there are no relevant limitations given Potter's social context or physical and mental capabilities and given that Potter's reflection and subsequent behavior is the result of a voluntary exercise of a vice (overconfidence, arrogance, dismissiveness, dogmatism, and so on) then we can "reasonably expect" Mr. Potter to do better. Neil Levy (2009) denies that we can reasonably expect this. For a reply to Levy see Robichaud (2014). See also Talbert (2013, 2017) for attributionist responses to Fitzpatrick. Mason and Wilson (2017) argue for a similar virtue epistemology-based approach to cultural moral ignorance.

have reasonably been expected to foresee that his teenage decision would lead to him becoming a full blown jerk later in life.¹⁹

I am dubious about the relevance of Jeff's teenage decision. It is not relevant to our assessment of Jeff's *current* blameworthiness that Jeff began developing his jerk-like tendencies as a teenager. The key feature of the case is that Jeff is *unreflective* about it—but surely any plausible reasonable epistemic standard will require that one reflect on the moral implications of one's behavior. To explain why Jeff is blameworthy, we need only point out that Jeff is sufficiently like us, namely, an epistemically and morally competent adult in a context sufficiently like ours. Which means that Jeff's failure to believe that his actions are wrong is the result of a failure to meet an applicable reasonable epistemic standard.

So, it seems that REST fits well with our intuitions, at least in these paradigm cases of moral ignorance. Not only does it prescribe the right judgment, but it also provides a coarse grained framework that fits naturally with our practices into which a more fine grained explanation for that judgment will fit. Even so, notice that these paradigm cases are all cases in which the intuition is that the agent is blameworthy. While it speaks in favor of REST that it does not commit us to excusing agents that are intuitively blameworthy in these cases, it does not yet establish that moral ignorance is ever exculpatory. To establish that claim we need cases where an agent's moral ignorance is epistemically reasonable and because of this they do not seem blameworthy for their behavior. In the next section I explore the possibility of such cases.

5 Exculpating moral ignorance

Paulina Sliwa (2020) rightly points out that the literature on the exculpating power of moral ignorance is rife with examples, like our paradigm cases, where the moral conclusions the agents draw are just obviously wrong. This is probably because such cases make good counterexamples—after all, if reasonable standards views are committed to saying JoJo or Rosen (2003)'s Hittite slave owner are excused then that is an important result. However, very often our moral lives are fraught with occasions where the moral truth is not as glaringly obvious as the fact that torturing innocent people or chattel slavery are wrong. If we are going to find cases where moral ignorance provides an excuse then surely we will find it among these much less obvious cases.

Others have offered such cases, and while I think they are not as strong as they could be, it is worth discussing them so that readers are aware that such cases are out there. If readers are not as dissatisfied with the cases as I am, then they will already be convinced and hence the new cases I offer will simply add to pool of evidence.

Sliwa (2020) offers two cases which are similar in that the agents seem less than fully blameworthy in virtue of their moral ignorance. In her cases the agents involved are unsure what the evidence points toward. Their concern for the moral

¹⁹ Fischer and Tognazzini (2009) offer a lengthy response to Vargas.

truth coupled with their lack of clear conclusions about what they ought to do lead them to seek advice from a trustworthy friend. Because they trust the friend, they take the advice and act accordingly. Unfortunately, on this occasion the friend gives them bad advice, and the agents do the wrong thing while falsely believing it is the right thing.

There are two reasons that I find these cases unsatisfying. First, advice is a form of moral testimony. As such, these cases rely heavily on moral testimony as the only deciding factor in the agent's deliberation, and I don't find cases relying heavily on moral testimony to provide very clear intuitions. Second (and relatedly), Sliwa does not intend for these cases to show the agents are not blameworthy, only that their moral ignorance mitigates blame to some degree. This is because her view is that moral ignorance can excuse *to some degree* but it never *fully* exculpates. I, of course, agree that excuses are not always fully exculpatory (see the discussion of JoJo above). But I don't see why, if moral ignorance can lessen blame by degrees just as any other excuse, that it cannot sometimes lessen blame to a negligible degree just as any other excuse. So, I would prefer cases that invoke the intuition that the agents involved are not blameworthy at all.

Claire Field (2021; 2022) offers some convincing cases of exculpatory moral ignorance arising in people with an atypical neurology. A person with autism, for example, may face significant difficulties in discovering the moral features of her behavior. She may not understand why it is sometimes wrong to speak blunt, hurtful truths because she struggles to understand why the truth is sometimes hurtful, but not for lack of trying. Indeed, often such agents make *more* effort because they are aware of their difficulties. Even so, because they are forced to employ methods that are imperfect, they sometimes get things wrong.

These are great cases, and I agree with Field that no good theory would claim such agents are blameworthy; any *reasonable* epistemic standard will need to take into account a person's epistemic capacities, and those capacities can certainly be affected by one's unique neurology. However, cases of moral ignorance due to an atypical neurology are by definition not typical. In my view cases of exculpating moral ignorance are relatively common, and this is because I think that people can also face difficulties in discovering the moral truth that are not due to an atypical neurology that diminishes their epistemic *capacities*.

Zoë A. Johnson King (2019) offers three potential cases that are meant to be fully exculpatory and that do not involve neuroatypical agents. However, I think the cases are not as clear as Johnson King hopes. I will attempt to briefly explain why, but since the reasons are different for each case, let me first give condensed versions of all three. In the first case, GENTRIFICATION, an agent, Grace, recognizes that gentrification is morally problematic because it results in poorer communities being unjustly forced to move from their long-time homes due to rising housing costs. She is committed to stopping the spread of gentrification in her own neighborhood and does so by joining protests that target new, hipster businesses. What Grace doesn't know is that one such business does more good than harm because it is an art gallery committed to providing exposure to artists from marginalized communities, which more than offsets the slight uptick in property values its presence causes. But Grace doubts this because "she doubts that the gallery's publicity would do much

to further local women's interests..." (pg. 416). When Grace protests the art gallery she is doing something overall morally wrong but she believes she is doing the right thing.

In the second case, NAMES, McKenzie is a woman of color who wants to make her academic discipline more welcoming to people of color like herself. She goes to great lengths to accomplish this goal—she creates a task force, serves in mentoring programs, helps develop updated policies, etc. However, she fails to notice that in her own academic work she *could* use names that are not traditionally Anglo-American. This, Johnson King claims, is wrong because it commits an expressive harm toward underrepresented communities by reinforcing the attitude that “being Anglo-American is seen as “normal” within the discipline” (pg. 416). McKenzie fails to realize this because she is too focused on her activism to reflect on the naming conventions she uses.

In the third case, NON-VIOLENT PROTEST, John and Xavier, two black lives matter protestors, are committed to the value of non-violent protest, even in the face of (most) violent retaliation. However, even their strong commitment has limits. They recognize that in some extreme cases violence may be warranted in order to defend oneself or others. At a protest they face a situation that begins to look like one of these extreme cases. John and Xavier must decide: do they respond with violence or not? “In the heat of the moment they are both deeply unsure what to do... but they each make a judgement call. And they make different calls: John fights back while Xavier restrains himself” (pg. 416–417).

These are condensed versions, but I think they capture the important parts of the cases. They are intriguing and insightful, however, I think there is reason to be suspicious of each as an example of moral ignorance. For example, while it is true that Grace in GENTRIFICATION is ignorant that her protest of the art gallery is overall morally wrong, this belief arises because “she doubts that the gallery’s publicity would do much to further local women’s interests...” (pg. 416). But note that, while morally relevant, this is just an empirical fact about how much benefit the exposure will bring. This is unlike the cases of paradigm moral ignorance above. JoJo isn’t ignorant of the empirical facts; he is aware that his actions cause great harm to others. JoJo is ignorant that *causing harm to others is wrong*. Contrast this with Grace, who knows that impeding the local women’s interests is wrong, but doesn’t realize the empirical fact that *what she is doing impedes their interests*.

NAMES is also not obviously a case of moral ignorance, though for a different reason. We can agree that being actively conscientious in the use of names in one’s examples is something someone can do to help toward the goal of making one’s academic discipline more welcoming to people of color, but Johnson King’s claim is that theorists’ use of generic, but stereotypical Anglo-American names like “Bill” and “Jane” is *wrong*. But is failing to be this conscientious really *wrong*? Maybe, but if it is, its wrongness is not so obvious as to be a common-sense judgement. I think most people who are not privy to arguments regarding the nature of expressive harms in moral philosophy would say that being more conscientious about one’s naming conventions is supererogatory—an effort one can make to help a worthy cause. But by definition supererogatory acts are not obligatory. Perhaps this is incorrect, and it really is wrong to use names like “Bill” and “Jane” in one’s examples,

but even if it is wrong, it is not (at present) so clearly wrong that it can be safely stipulated in examples meant to appeal to widespread intuitions. So, McKenzie is not obviously morally ignorant because her action is not obviously wrong.

For NON-VIOLENT PROTEST, I find the detail of two people using the same evidence to draw different but still reasonable conclusions intriguing, but I worry that the case is weakened by the fact that John and Xavier are forced by time and circumstances to make a judgement call. As such, it's not clear whether John and Xavier really *believe* anything about the rightness or wrongness of their behavior. They each 'went with their gut' so to speak, but I doubt that going with one's gut ought to be considered representative of a *belief* that one is doing the right thing. It seems more like the opposite; an admittance that one doesn't really know what to do—if one did, then one wouldn't need to go with one's gut.

What I think all this reveals is that there are at least two ways an agent might be blameless in virtue of moral ignorance. The first is inspired by the structure of NON-VIOLENT protest, and also by the remarks of Weiland (2017): "slavery is of course a special case. From our perspective, it is hard to imagine how one cannot be in a position to see its wrongness. But how about other issues, such as gay marriage, sexism, private property, euthanasia, famine relief, climate change, animal exploitation, and so on?" Unfortunately, Weiland offers these remarks while also using slavery as his main example, and only suggests that such topics are ones where the moral truth may not be readily accessible to everyone. I say this is unfortunate because it seems to me this is the form moral ignorance most commonly takes. So, I suggest combining the structure of NON-VIOLENCE with one of these less-than-obvious topics about which reasonable people disagree despite sharing the same evidence.

Imagine two people Tyrell and Tanya, both of whom are moral philosophers who hold opposing views. Tyrell has considered the reasons both for and against the permissibility of voluntary euthanasia and has been persuaded by the evidence that it would alleviate incredible amounts of human suffering. He is convinced that we could sufficiently regulate its use to avoid the negative social implications or possibility for abuse. And besides, in Tyrell's view, even if regulation failed, the value of the alleviated suffering far outweighs any disvalue that would occur from such abuse. Conversely, Tanya has considered the undeniable positive value of avoiding terrible human suffering, but she doubts that it is possible to avoid the potentially disastrous negative social implications or the potential for abuse. These things, Tanya thinks, are much more important to avoid than suffering caused by terminal illness, especially when avoiding that suffering requires the intentional killing of a human being. Tyrell and Tanya draw opposing conclusions because their differing values cause them to weigh the evidence slightly differently, but, crucially, neither set of values is clearly unreasonable. Even so, on the assumption that there is a correct answer to the question of whether voluntary euthanasia is permissible, one of them is morally ignorant.

Suppose that Tanya is assigned to a committee that is tasked with coming up with a local policy concerning whether to make voluntary euthanasia legal. Tanya votes 'no' because she believes it is wrong. Suppose further that Tanya is incorrect, and voluntary euthanasia is morally permissible. Is Tanya *blameworthy* for voting to

make it illegal? I do not think she is. Her beliefs are well considered and no obvious mistakes are being made. Her conclusions are incorrect, but the accessible evidence does not make either conclusion clearly beyond the limits of what is reasonable. The same is true of Tyrell *mutatis mutandis*. One of them will be acting wrongly from moral ignorance, but if neither of them would be blameworthy for casting their vote, then we have a case where moral ignorance is indeed exculpatory.²⁰

Such cases are not limited to the views of ethicists and moral philosophers. We encounter such situations regularly. Many people who are not philosophers are convinced by the reasons that speak against voluntary euthanasia. Many others do not find those reasons moving. One of these conclusions is incorrect. It is possible, of course, for someone to hold one or the other of these beliefs in an unreasonable way, perhaps they do not spend sufficient time considering the issue, or perhaps they approach the accessible evidence in a biased way. But it seems to me that the accessible evidence does not settle the issue so decisively that anyone sufficiently *like us* could not reasonably hold either belief.

So then, in cases where there can be reasonable disagreement on a moral issue it seems that moral ignorance can be exculpatory. But there is at least one other way that moral ignorance might exculpate, and that is when the agent in question is not sufficiently like us. I do not mean that the agent is not like us with respect to the capacities necessary to recognize and respond to moral reasons (e.g. the “insane” version of JoJo Wolf wants, or the neuroatypical agents that concern Field). What I mean is that the agent is not like us in *their accessibility of the evidence*. My diagnosis of our intuitions about JoJo claimed that in practice we judge JoJo guilty because JoJo’s belief is *so obviously false* that we think that it is not possible for someone to which the standard applies to reasonably draw the conclusions that JoJo does. However, recall also that anecdotal evidence suggests that we can manipulate this intuition by significantly altering the details of the case. When we describe JoJo’s situation so that he grows up in isolation, completely cut off from the outside world, we reluctantly agree that JoJo’s moral ignorance might be reasonable under those circumstances. In other words, we acknowledge that JoJo’s epistemic circumstances would be so different from our own that we begrudgingly admit *that he is not subject to the same standards as us*. Of course, it is true that we have considerable difficulty imagining circumstances where a normal functioning moral agent could fail to see something we consider so obvious, but such cases are possible. In fact, there is a version of such extreme isolation in the real world.

Consider a culture like that of the Sentinelese people (Sasikumar, 2018). The Sentinelese are an Indigenous tribal people that occupy North Sentinel Island, a

²⁰ It might be thought that this case rests on a mistake in that it conflates the permissibility of making something legal with the permissibility of the act itself. It is clearly possible for something permissibly legal to also be morally wrong (consider gossip, for example). However, I think this is a case where we would want the law to match up with morality—if voluntary euthanasia is morally wrong, then legalizing it is tantamount to legalizing certain cases of murder. If it is morally permissible, then making it illegal is the state forcing its citizens to endure cruel and unnecessary painful suffering. Though, see Velleman (1992) for an argument that there might be good reason not to legalize voluntary euthanasia even if it is morally permissible.

small island off the coast of India. They are one of the last cultures to have remained completely isolated from modern civilization and characteristically attack people who approach the island. The Indian government has strictly enforced laws forbidding anyone from accessing the island, leaving the tiny culture to their literal near-complete isolation. Suppose that we discover the Sentinelese people have a widespread cultural belief that it is permissible to kill people who venture onto their beach—something not difficult to imagine, given that they have done that very thing more than once.²¹ For the Sentinelese people, such a belief is probably not violating any applicable reasonable standard. To put the point succinctly, we (probably) cannot say for certain of a Sentinelese person that she *should not be ignorant* that killing innocent people who venture onto the island is wrong. We certainly can say that about everyone we interact with on a regular basis.

In the previous section we saw that reasonable standards views have a good explanation for when our intuitions indicate that moral ignorance clearly does not exculpate in the paradigm cases from the literature. We have now seen how reasonable standards views can explain when moral ignorance does seem to exculpate. This gives us at least good prima facie reason to think that REST applies in cases of moral ignorance. We are now in a position to see how this explanation stacks up against counterexamples in the literature.

6 Responding to counterexamples

Elizabeth Harman claims that, “While it might have seemed that the fact that false non-moral views exculpate can support the claim that false moral views exculpate, there is no support from the former to the latter.” (Harman 2015, pg. 65). To make this claim, Harman argues that any reasonable standards view is open to obvious counterexamples. Consider the following cases she offers.²²

Max works for the Mafia “family” and believes he has a moral obligation of loyalty to the family that requires him to kill innocents when it is necessary to protect the financial interests of the family. This is his genuine moral conviction, of which he is deeply convinced. If Max failed to “take care of his own” he would think of himself as disloyal and he would be ashamed.

Gail is a gang member who believes that she has a moral obligation to kill a member of a neighboring gang as revenge after a member of her own gang is killed, although her victim was not responsible for the killing. This is her genuine moral conviction, of which she is deeply convinced. If Gail failed to “take care of her own” she would think of herself as disloyal, and she would be ashamed. (Harman 2015, pg. 65).

²¹ McDougal (2006), Roy (2018)

²² Harman (2019) offers a few other cases, but I think everything I say regarding Max and Gail will apply equally to those cases as well.

Harman thinks that these are paradigm cases of blameworthiness, but the parties involved are also described as having false moral views. It is morally wrong to kill the innocent for the reasons set out by each case. Of course, one might argue for something consistent with the intuitions we seem to have about JoJo and Mr. Potter above, namely, that we treat them like someone who believes the Earth is flat: it is simply not possible for anyone sufficiently like us to *reasonably* believing something so obviously false. Harman anticipates this, and responds in the following way.

It is a grave mistake to think that people cannot become convinced of deeply false moral views...without violating this kind of procedural obligation. Consider versions of Max and Gail who have thought an ordinary amount about morality and have taken seriously the moral arguments that have been presented to them. They are aware that many people think their moral views are false, but they believe that they understand where others have gone wrong: others have been “suckered” into a “wimpy” morality, when what is really important is taking care of one’s own. These versions of Max and Gail are certainly possible; they would be blameworthy for their wrongful behavior. (Harman 2015, pg. 66)

The kind of procedural obligations Harman is speaking about, “include the obligation to think a reasonable amount about morality in general, the obligation to take seriously moral arguments one hears, etc.” (ibid). Harman gets the term ‘procedural obligations’ from, Rosen (2003), but I think it is clear that she means to include any reasonable epistemic standard that aims to explain how ordinary (non-moral) factual ignorance can exculpate. This is evidenced by the above quote (from pg. 65) asserting that the claim that false non-moral beliefs can exculpate does not provide support for the claim that false *moral* beliefs exculpate. So then, according to Harman, it is easy to imagine instances of the cases where Max and Gail’s moral beliefs meet the applicable reasonable standards. REST, therefore, is allegedly committed to saying Max and Gail are excused under those imagined instances.²³ In other words Harman claims, “The stipulations that I have made are meant to establish that [agents like Max and Gail] have not been irresponsible in the course of forming their beliefs” (2019, pg. 172). My response to Harman is that the stipulations she makes do not entitle her to claim that Max and Gail’s beliefs meet the applicable reasonable standard.

²³ Harman also offers the positive view that moral ignorance never exculpates, and this is because, she claims, that moral ignorance is always the result of a failure to care sufficiently about what morally matters (Harman 2011, 2015, 2017, 2019). This view is one of a family of views, often referred to as *attributionist*, that suggest that blameworthiness is primarily about a failure in the judgement sensitive attitudes (or character, or quality of will) of an agent. On these views, since having a bad character is just what it means to be blameworthy, moral ignorance is never exculpatory. Other attributionists include Heironymi (2004, 2006, 2008), Talbert (2013, 2017), Arpaly (2003), Arpaly and Schroeder (2014), Smith (2005), and I think Mason (2015, 2017) and Mason and Wilson (2017). Weiland (2017), Johnson King (2019), and Field (2021; 2022) all argue against these theorists’ positive claims by pointing out flaws in that view. My task here is not to point out a flaw in attributionist views, but to respond to alleged flaws Harman claims are common among reasonable standards views.

There are two things to note. First, recall from Sect. 3 that part of what generates our common-sense schematic notion of “reasonable standard” is that, when we judge whether a person’s beliefs are reasonable, we presume that person is a normal functioning adult with access to roughly the same evidence as any other normal functioning moral agent would have—that person is sufficiently *like us*. With the information Harman provides, we are left to presume that Max and Gail reason in the same way and have access to mostly the same evidence as the rest of us (e.g. they are not on an isolated island, excessively young, neuro-divergent, etc.). These details strongly indicate that Max and Gail have not met the applicable reasonable standard, even in the stipulated version of the case. A good indicator that they have an unreasonable belief (moral or otherwise) is just that so many of the rest of us do not *just* disagree (there is room for reasonable disagreement as we saw above) we positively *condemn* their beliefs.²⁴ Of course, it is possible for Max and Gail to have a reasonable belief that very few others share. Perhaps Max and Gail discover some evidence that the rest of us do not have, or perhaps there is some nuance to their reasoning that the rest of us have missed. This is not what we see in the details Harman provides, however. The reason Max and Gail think that they get things right where the rest of us have gone wrong is not that they have some extra evidence or have reasoned better. Instead, they conclude that the rest of us are “suckers” and “wimps”, which brings us to the second point.

Considering alternative views “wimpy” and thinking of others as being “suckered” indicates that Max and Gail view the evidence before them in an objectionably biased way (only a sucker would be convinced by that evidence!). Furthermore, it reveals that they see themselves as intellectually superior. If *most people* are convinced by evidence that only a “sucker” would believe, then they are, by their own lights, intellectually superior to the rest of us suckers and wimps.

The trouble with this point of view is that it is, frankly, unreasonable. Again, it is possible that Max and Gail do see the truth where everyone else gets it wrong, but surely there must be principled reasons they can point to, some *evidence* they can offer that other moral agents of similar capacity and context would find lucid. As far we are told, there are no such reasons, we are only told what derisive and dismissive attitudes Max and Gail have toward the conclusions of others. As such, we are once again left to assume their moral beliefs are based on the same evidence and are to be held to the same standards as the rest of us. The stipulations Harman makes, therefore, are not enough to support the claim that Max and Gail’s beliefs are reasonable.

To sum up my points let me refer to something Harman says:

Ethics is indeed quite hard, and this is why we see so many cases of false moral belief even among those who have [met the applicable reasonable standard]. Some people do not think hard enough about morality in general, or they dismiss moral arguments that they ought to take seriously. But many people

²⁴ Talbert (2013), an attributionist, seems to agree with this point. He argues that the degree to which we disagree with the moral views of others matches the degree to which we judge they are blameworthy. I would add to this assessment that we often do not blame someone at all, despite disagreeing, when we think that their conclusions are not unreasonable.

think hard about morality, take the arguments they hear seriously, and *still get things wrong*. A failure to appreciate how hard ethics is can make the claim that false moral belief exculpates seem less radical than it is. But that claim is very radical. (Harman 2015, pg. 66, emphasis in the original)

Harman is right that ethics is hard, and that we often get things wrong despite caring about, and taking reasonable steps to discover the moral truth. Harman is also right that it is a grave mistake to think that people cannot be convinced of false moral views without violating any applicable reasonable epistemic standard. People can, and do, hold perfectly reasonable, and yet false, moral views, as we saw in Sect. 4. However, it is an equally grave mistake to think that such standards are so easily met that Max and Gail's beliefs would count as reasonable.²⁵ Indeed, it seems to me that Max and Gail do not even meet the criteria that Harman sets out. I cannot see how the attitude that others have been "suckered into a wimpy morality" is evidence that Max and Gail are "taking the arguments they hear seriously." Instead, Max and Gail's attitudes reveal that they "dismiss moral arguments that they ought to take seriously." So, even by Harman's own lights Max and Gail's beliefs are not reasonable.

I want to urge caution, however. My claim here is only that the details Harman offers do not entitle her to the stipulation that Max and Gail hold reasonable (but false) moral beliefs.²⁶ It must be stressed, however, that we certainly can fill in the details in such a way that Max and Gail's false moral views meet a minimum reasonable standard. This will not help Harman though, for when we fill in the details in this way, they are no longer paradigm cases of blameworthiness. Here is why.

Recall from Sect. 3 that when we describe JoJo's situation so that he grows up in isolation, completely cut off from the outside world, we reluctantly agree that JoJo's moral ignorance might be reasonable under those circumstances, and there is no reason to think that Max and Gail are different. We can redescribe the case so that they are in such unusual circumstances that the typical standards that apply to the rest of us do not apply to them. Note, though, that when we describe the context in this way the case becomes, as Faraci and Shoemaker put it, "quite precious." Moreover, just as we begrudgingly give up the intuition that JoJo is blameworthy under such circumstances, so too would we give up the intuition that Max and Gail are blameworthy. In other words, the more we make Max and Gail's epistemic circumstances foreign to our own, the less they can be held to the same standards.

To sum up, I grant that claiming that moral ignorance exculpates according to REST would be radical if it entailed that we should excuse Max and Gail, but it does not. Indeed, the opposite seems true. Intuitively, Max and Gail are paradigm cases of blameworthiness *precisely because* their moral ignorance is the result of a failure to meet an applicable reasonable epistemic standard. Furthermore, when we describe

²⁵ Indeed, in moral contexts we typically raise our standards because we think that moral beliefs are much more important to get right—Harman's cases of Max and Gail are not sensitive to this fact.

²⁶ Interestingly, Harman (2019) makes an even stronger claim. False moral beliefs are always unreasonable because one always has access to sufficient counter evidence. Weiland (2017) argues convincingly that this is false.

Max and Gail in such a way that their beliefs seem reasonable, just like JoJo or the Sentinelese, I think we are far less sanguine about their blameworthiness.

7 Conclusion

In this paper I have argued that the Reasonable Epistemic Standards Thesis can extend its application to moral ignorance. I argued that the Reasonable Epistemic Standards Thesis fits very well with our intuitions in three paradigm cases of moral ignorance found in the literature. However, all of those cases are such that the agents in them are blameworthy, so I then offered cases where moral ignorance plausibly does count as an excuse. I then considered and responded to Harman's purported counterexamples. It seems then, that without some reason to think otherwise, anyone who has some version of a reasonable standards view ought to claim moral ignorance can sometimes exculpate.

Acknowledgements This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 740922). I would like to thank the members of the Normative and Moral Foundations of Group Agency Research Group at the University of Vienna, especially Herlinde Pauer-Struda, Olof Leffler, Grace Paterson and Niels de Haan for their helpful discussion. And an especially large debt of gratitude is owed to Veli Mitova for extensive discussion and comments on an earlier draft of his paper, which greatly helped in clarify my thoughts. I would also like to thank David Shoemaker, Daniel Tigard, Jesse Hill, Eric Brown, Nate Stout, and Nick Sars for helpful discussion. Finally, to the members of the Early Career Researchers Workshop where this paper was presented in August 2022.

Declarations

Conflict of interest There are no conflicts of interest to report.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alston, W. (1988). The deontological conception of epistemic justification. *Philosophical Prospectives*, 2, 257–299.
- Arpaly, N. (2003). *Unprincipled virtue*. Oxford University Press.
- Arpaly, N and Tim A.S. (2014) *In Praise of Desire*. Oxford University Press: New York
- Ayars, A. (2021). Blaming for unreasonableness. *Journal of Ethics and Social Philosophy*, 19, 56–79.
- Biebel, N. (2018). Epistemic justification and the ignorance excuse. *Philosophical Studies*, 175(12), 3005–3028.
- Chrisman, M. (2008). Ought to believe. *Journal of Philosophy*, 105, 346–370.

- Christiansen, D. (2007). Epistemology of disagreement: the good news. *The Philosophical Review*, 116(2), 187–217.
- Chuard, P., & Southwood, N. (2009). Epistemic norms without voluntary control. *Nous*, 43, 599–632.
- Clarke, R. (2014). *Omissions*. Oxford University Press.
- Elga, A. (2007). Reflections on disagreement. *Noûs*, 41, 478–502.
- Faraci D. and Shoemaker D. (2014) "Huck vs JoJo Moral Ignorance and the (A)symmetry of Praise and Blame." *Studies in Experimental Philosophy* pp: 7–27.
- Faraci, D. and David S. (2017) "Good selves, true selves: moral ignorance, responsibility, and the presumption of goodness." *Philosophy and Phenomenological Research* pp 1–17.
- Faraci, D., & Shoemaker, D. (2010). Insanity, deep selves, and moral responsibility. *Review of Philosophy and Psychology*, 1, 319–332.
- Feldman, R. (2006). Epistemological puzzles about disagreement. In: *Epistemology Futures* 223. Oxford University Press: Oxford. pp 216
- Feldman, R. (1988). Subjective and objective justification in ethics and epistemology. *The Monist*, 71(3), 405–419.
- Field, C. (2021). Moral appraisal for everyone: neurodiversity, epistemic limitations, and responding to the right reasons. *Ethical Theory and Moral Practice*, 24, 733–752.
- Field, C. (2022). At least you tried: the value of de dicto concern to do the right thing. *Philosophical Studies*, 179, 2707–2730.
- Fischer, J. M., & Neal, T. (2009). The truth about tracing. *Nous*, 43, 531–556.
- Fitzpatrick, W. J. (2008). Moral responsibility and normative ignorance: answering a new skeptical challenge. *Ethics*, 118, 589–613.
- FitzPatrick, William. (2017). Unwitting wrongdoing, reasonable expectations, and blameworthiness. In R. Philip & W. W. Jan (Eds.), *Responsibility the epistemic condition* (pp. 29–46). New York: Oxford University Press.
- Gaus, G. (2010). *The order of public reason*. Cambridge University Press.
- Gibbons, J. (2013). *The norm of belief*. Oxford University Press.
- Goldman, A. (1999). Internalism exposed. *Journal of Philosophy*, 96, 271–293.
- Guerrero, A. A. (2007). Don't know, don't kill: moral ignorance, culpability, and caution. *Philosophical Studies*, 136, 59–97.
- Harman, Elizabeth. "Moral Testimony Goes Only So Far." *Oxford Studies in Agency and Responsibility*. Ed. David Shoemaker. Vol. 6. New York: Oxford University Press, 2019. 165–185.
- Harman, E. (2011). Does moral ignorance exculpate? *Ratio*, 24, 443–468.
- Harman, E. (2015). The irrelevance of moral uncertainty. *Oxford Studies in Metaethics*, 10, 53–79.
- Harman, E. (2017). *When is failure to realize something exculpatory? responsibility the epistemic condition* (pp. 117–126). Oxford University Press.
- Hart, H. L. A. (1968). *Punishment and responsibility*. Oxford University Press.
- Hieronymi, P. (2004). The force and fairness of blame. *Philosophical Perspectives*, 18, 115–148.
- Hieronymi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87(1), 45–74.
- Hieronymi, P. (2008). Responsibility for believing. *Synthese*, 161, 357–373.
- Huss, B. (2009). Three challenges (and three replies) to the ethics of belief. *Synthese*, 168, 249–271.
- King, Johnson, & Zoë, A. (2019). Don't know, don't care? *Philosophical Studies*, 177, 413–431.
- King, M. (2009). The problem with negligence. *Social Theory and Practice*, 35(4), 577–595.
- King, M. (2014). Getting traction without tracing: a (partial) solution for control-based accounts of moral responsibility. *European Journal of Philosophy*, 22, 463–482.
- Korblith, H. (1983). Justified Belief and epistemically responsible action. *The Philosophical Review*, 92(1), 33–48.
- Kornblith, Hillary. (2010). Belief in the face of controversy. In F. Richard & A. Ted (Eds.), *Disagreement warfield* (pp. 29–52). Oxford: Oxford University Press.
- Levy, N. (2009). Culpable ignorance and moral responsibility: a reply to Fitzpatrick. *Ethics*, 119, 729–741.
- Mason, E. (2015). Moral ignorance and blameworthiness. *Philosophical Studies*, 172, 3037–3057.
- Mason, E. (2017). Moral incapacity and moral ignorance. In Rik Peels (Ed.), *Perspectives on ignorance from moral and social philosophy* (pp. 30–52). Routledge: Routledge Studies in Contemporary Philosophy.
- Mason, Elinor, & Wilson, Alan T. (2017). Vice, blameworthiness, and cultural ignorance. In R. Philip & W. W. Jan (Eds.), *Responsibility: the epistemic condition* (pp. 82–100). Oxford: Oxford University Press.

- McDougal, Dan. "Survival comes first for Sentinel islanders—the world's last 'stone-age' tribe." *The Guardian* 12 February 2006.
- McHugh, C. (2013). Epistemic responsibility and doxastic agency. *Philosophical Issues*, 23, 132–157.
- Montmarquet, J. (1995). Culpable ignorance and excuses. *Philosophical Studies*, 80(1), 41–49.
- Nickel, Philip J. (2010). Voluntary belief on a reasonable basis. *Philosophy and Phenomenological Research*, 81, 312–334.
- Nettleman, N. (2007). *Blameworthy belief: A study in epistemic deontologism*. Springer.
- Oliveira, L. R. (2015). Non-agential permissibility in epistemology. *Australasian Journal of Philosophy*, 93(2), 289–294.
- Oliveira, L. (2018). Deontological evidentialism and ought implies can. *Philosophical Studies*, 175, 2567–2582.
- Peels, R. (2017). *Responsible belief*. Oxford University Press.
- Plantinga, A. (1993). *Warrant: the current debate*. Oxford University Press.
- Robichaud, P. (2014). On culpable ignorance and akrasia. *Ethics*, 125, 137–151.
- Rosen, G. (2003). Culpability and ignorance. *Proceedings of the Aristotelian Society*, 103, 61–84.
- Rosen, G. (2004). Skepticism about moral responsibility. *Nous*, 18, 295–313.
- Rosen, G. (2008). Kleinbart the oblivious and other tales of ignorance and responsibility. *The Journal of Philosophy*, 105(10), 591–610.
- Roy, S.K. (2018). American killed on remote Indian island off-limits to visitors. *Reuters* 21 November
- Sasikumar, M. (2018). The Sentinelese of the north sentinel island: concerns and perceptions. *Journal of the Anthropological Survey of India*, 67(1), 37–44.
- Sher, G. (2009). *Who knew?* New York: Oxford University Press.
- Sher, George. (2017). Blame and moral ignorance. In R. Philip & W. W. Jan (Eds.), *Responsibility the epistemic condition* (pp. 101–116). New York: Oxford University Press.
- Sliwa, P. (2020). Excuse without exculpation. *Oxford Studies in Metaethics*, 15, 72–95.
- Smith, A. (2005). Responsibility for attitudes: activity and passivity in mental life. *Ethics*, 115(2), 236–271.
- Smith, H. (1983). Culpable ignorance. *The Philosophical Review*, 92(4), 543–571.
- Steup, M. (2008). "Doxastic Freedom." *Synthese* 161: 375–392.
- Steup, M. (2017). Believing intentionally. *Synthese*, 194, 2673–2694.
- Strawson, P. F. (1962). Freedom and resentment. *Proceedings of the British Academy*, 48, 1–25.
- Talbert, M. (2013). Unwitting wrongdoers and the role of moral disagreement in blame. *Oxford Studies in Agency And Responsibility*, 1, 225–245.
- Talbert, M. (2017). Akrasia, awareness, and blameworthiness. In R. Philip & W. W. Jan (Eds.), *Responsibility the epistemic condition* (pp. 47–63). New York: Oxford University Press.
- Vargas, M. (2005). The trouble with tracing. *Midwest Studies in Philosophy*, 29, 269–291.
- Velleman, J. D. (1992). Against the right to die. *The Journal of Medicine and Philosophy*, 17, 665–681.
- Weatherston, B. (2008). Deontology and Descartes' demon. *Journal of Philosophy*, 105, 540–569.
- Wieland, J. W. (2017). What's special about moral ignorance? *Ratio*, 30, 149–164.
- Wolf, S. (1987). Sanity and the metaphysics of responsibility. *Responsibility, character, and the emotions: new essays in moral psychology ferdinand david schoeman* (pp. 46–62). Cambridge: Cambridge University Press.
- Wolf, Susan. (2003). Sanity and the metaphysics of responsibility. In W. Gary (Ed.), *Free Will 2nd* (pp. 372–387). Oxford: Oxford University Press.
- Zimmerman, M. J. (1997). Moral Responsibility and Ignorance. *Ethics*, 107, 410–426.
- Zimmerman, M. (2008). *Living with uncertainty*. New York: Cambridge University Press.
- Zimmerman, M. (2017). Ignorance as a moral excuse. In Rik Peels (Ed.), *Perspectives on ignorance from moral and social philosophy* (pp. 77–94). New York: Routledge.