

DOES MORAL IGNORANCE EXCULPATE?

Elizabeth Harman

Abstract

Non-moral ignorance can exculpate: if Anne spoons cyanide into Bill's coffee, but thinks she is spooning sugar, then Anne may be blameless for poisoning Bill. Gideon Rosen argues that moral ignorance can also exculpate: if one does not believe that one's action is wrong, and one has not mismanaged one's beliefs, then one is blameless for acting wrongly. On his view, many apparently blameworthy actions are blameless. I discuss several objections to Rosen. I then propose an alternative view on which many agents who act wrongly are blameworthy despite believing they are acting morally permissibly, and despite not having mismanaged their moral beliefs.¹

Gideon Rosen has argued for the following Broad Conclusion:

A person who acts wrongly is blameworthy for so acting only if the action itself is a case of clear-eyed *akrasia* or the action results from a case of clear-eyed *akrasia*.²

While the Broad Conclusion is stated in terms of *akrasia*, its main upshot is regarding *moral knowledge* and *moral ignorance*. An akratic action is one done in the belief that one should act differently. Someone who acts wrongly while ignorant that he should act

¹ For helpful comments on drafts of this paper, I am grateful to Charles Beitz, Peter Graham, Alexander Guerrero, Adam Hosein, and audiences at the CUNY Graduate Conference, the Inland Northwest Philosophy Conference, the New York Society for Women in Philosophy, Princeton University's Program in Law and Public Affairs, Rutgers University, the University of Texas at Austin, and Yale Law School. I also thank the participants in my spring 2009 graduate seminar.

² Gideon Rosen, 'Skepticism about Moral Responsibility,' *Philosophical Perspectives* 18 (2004), 295–313; the paper uses what I am calling the Broad Conclusion to argue for a further claim, that we never have *knowledge* that someone is blameworthy for an action. (This follows from the further premise that we never have knowledge that clear-eyed *akrasia* has occurred.)

While Michael Zimmerman ('Moral Responsibility and Ignorance,' *Ethics* 107 (1997), 410–426) and Susan Wolf ('Moral Saints,' *Journal of Philosophy* 79 (1982), 419–439) make arguments that are in some respects similar to Rosen's, I will focus on Rosen's argument.

differently is not acting *akratically*. The Broad Conclusion implies that someone who acts wrongly without believing that he should act differently is blameworthy only if his action results from a case of clear-eyed *akrasia*: that is, only if it results from some earlier behavior performed in the full belief that it should not be performed. On Rosen's view, all blameworthiness has its root in wrong behavior that was performed while the agent knew she should act differently.

Others have shown that Rosen's argument does not succeed in establishing the Broad Conclusion. However, Rosen's argument may still seem to establish a narrower conclusion which would imply that many ordinary apparently blameworthy actions are in fact blameless. I will argue that Rosen's narrower conclusion is false, and I will develop a view of blameworthiness on which these ordinary apparently blameworthy actions are indeed blameworthy.

Section 1 sketches Rosen's argument for his Broad Conclusion. Section 2 explains some compelling points made by Alexander Guerrero, Michelle Moody-Adams, and William FitzPatrick, which undermine parts of Rosen's argument, and outlines the narrower conclusion that Rosen's argument may seem to establish after these points are taken into account. Along the way, section 2 also offers a new objection to Rosen's argument and raises an objection to a claim of FitzPatrick's. Section 3 argues that the narrower conclusion is nevertheless significantly revisionary about blameworthiness: it implies that many ordinary apparently blameworthy actions are in fact blameless. Section 4 proposes a view on which these actions are indeed blameworthy. Sections 5 clarifies the proposed view and defends it in the face of some objections. Section 6 offers a way of understanding the dispute between the view I propose and Rosen's view, and offers my most important objection to Rosen's paper.

1. Rosen's Argument

Let's begin with the natural thought that ignorance of non-moral matters can exculpate. Suppose Anne spoons some cyanide into Bill's coffee, but she does not know it is cyanide: she believes that she is spooning sugar. It seems that Anne is blameworthy for poisoning Bill only if she is blameworthy for her ignorance. If we stipulate that Anne is not blameworthy for her ignorance, then it seems Anne is not blameworthy for poisoning Bill.

Rosen draws a general lesson from this case:

Lesson:

A wrong action (or omission) done from ignorance is not a locus of *original responsibility*; rather, it is at most a locus of *derivative responsibility*: if the agent is blameworthy for the wrong action (or omission), he is so blameworthy only in virtue of being blameworthy for the ignorance.³

To understand the notions of *original responsibility* and *derivative responsibility*, Rosen offers the example of someone who goes berserk and destroys his hotel room while high on a drug: his actions are not loci of *original responsibility*, but they may be loci of *derivative responsibility*, in that he may be blameworthy for them by being blameworthy for having taken the drug.

Rosen furthermore claims:

An agent is blameworthy for some ignorance only if he has behaved in ways that constitute a violation of our *procedural epistemic obligations* and this behavior has resulted in his ignorance.

Rosen points out that we are obligated to take precautions against harming others (and against other wrongdoing) and that among the precautions we are obligated to take are epistemic precautions. He writes, 'these procedural obligations are always obligations to *do* (or to refrain from doing) certain things: to ask certain questions, to take careful notes, to stop and think, to focus one's attention in a certain direction, etc. The procedural obligation is not itself an obligation to know or believe this or that. It is an obligation to *take steps* to ensure that when the time comes to act, one will know what one ought to know' (301). Rosen stresses that we are not obligated to take all the precautions that we could take, and he suggests that we are obligated to act as a 'person of ordinary prudence' would act. I will use the expression 'mismanagement of one's opinion' to describe any violation of our procedural epistemic obligations.

³ Note that 'from ignorance' here means from a position of ignorance; it does not require that the ignorance is a cause of the action.

Rosen claims blameworthiness for ignorance can only arise from blameworthiness for *behavior* to do with the management of one's beliefs because he claims a belief itself cannot be a locus of original blameworthiness; he claims beliefs are not voluntary in the way that any locus of original blameworthiness must be voluntary; he says that a belief is something that 'simply happens in me or to me' (302).

Rosen also claims:

Every instance of morally wrong behavior (whether an action or an omission) is either a case of behavior from ignorance or a case of clear-eyed *akrasia*.

In a case of wrong action (or omission) due to clear-eyed *akrasia*, a person performs a wrong action (or an omission) while aware that she has most reason to act otherwise. In any other case, Rosen claims, the action in question is done from ignorance, because it is done at least from ignorance *that* one has most reason to act otherwise. (Thus, Rosen claims that even someone who knows she is acting wrongly but does not know that morality is overriding – and so, that she has most reason to act otherwise – is acting in ignorance. I will not address this implication of his views.)

What emerges from the three claims I have highlighted is the following. A wrong action (or omission) is either a case of clear-eyed *akrasia* or it is an action (or omission) from ignorance. A wrong action (or omission) from ignorance is not a locus of original responsibility; if it is blameworthy then it is blameworthy only in virtue of blameworthiness for an earlier behavior which constitutes violation of procedural norms in the management of one's beliefs. This earlier behavior is similarly either itself clear-eyed *akrasia* or done from ignorance, in which case it is blameworthy only if an earlier behavior is blameworthy, which is itself blameworthy only if it is clear-eyed *akrasia* or it results from blameworthy behavior.

Rosen illustrates the application of his principles with an example of a surgeon who has type A blood transfused into a patient who has type B blood. The surgeon forgets to double-check the chart, although double-checking at such a time is standard procedure. The surgeon's ignorance that the patient has type B blood means she is not originally responsible for her transfusion order. But she may be derivatively responsible, if she is responsible for her failure to double-check the chart. When she

forgot to double-check the chart, she was not thinking about the chart, and she did not know she should check the chart. So her failure to check the chart was itself in ignorance of the fact that she should do so, and she is blameworthy for that only if she is blameworthy for some earlier behavior which resulted in her ignorance that she should check the chart. If she tends to be forgetful about checking such charts, then she should have asked someone to remind her. But if this is a case of forgetfulness that was not particularly foreseeable, then she is not blameworthy.

I will also mention how the argument applies to a case Rosen discusses in an earlier paper.⁴ Consider an ancient slaveholder, who believes it is permissible to keep slaves. Ancient slavery was not race-based and was not based on the belief that some group of people was inferior to another group; rather, slaves had been captured in battle. Being a slave was seen as an unfortunate state that in principle anyone might have found himself in. So it is plausible that a person could know all the non-moral facts and yet believe that this slavery was permissible. At this time, while everyone thought it was awful to *be* a slave, no one suggested that it was wrong to keep slaves. Is the ancient slaveholder blameworthy for keeping slaves? He is not originally blameworthy, because of his moral ignorance. Is he blameworthy for his ignorance? It does not seem he has been negligent in the management of his opinion. He doesn't ignore arguments against slavery: none are offered to him. He thinks about morality as much as any ordinary person. On Rosen's view, the slaveholder is blameless because he is ignorant that his behavior is morally wrong and he has not violated any procedural norms in the management of his beliefs.

2. Objections to Rosen's Argument

I will discuss four objections to Rosen's argument. The first objection is due to Alexander Guerrero; he raised it as an objection to an earlier paper of Rosen's. The second objection is new; it emerges from consideration of the first objection. The third objection emerges from consideration of a paper by Michelle Moody-Adams; her paper was published before Rosen's papers on

⁴ Gideon Rosen, 'Culpability and Ignorance,' *Proceedings of the Aristotelian Society* CIII (2003), Part 1.

this topic but it addresses the same issues. The fourth objection is a generalization of points made both in Moody-Adams's paper and in a paper of William FitzPatrick's; FitzPatrick's paper is addressed to the argument of Rosen's that I have presented. In the course of discussing FitzPatrick's paper, I will object to a claim he makes about the nature of a certain class of cases of ordinary blameworthiness. At the end of this section, I will describe the narrower conclusion that Rosen's argument might seem to establish even in the face of these four objections.

The first objection is a generalization of an objection made by Guerrero.⁵

Let's consider again the case of Anne, who poisons Bill by spooning cyanide into his coffee. Anne believes she is spooning sugar, and is blameless for her false belief. Anne is blameless for poisoning Bill. Rosen takes the lesson of cases like this to be that an action done from *ignorance* is not a locus of original responsibility. However, that is the wrong lesson. Consider a variant of the case. Alice spoons cyanide into Bob's coffee. Alice is *ignorant* that she is spooning cyanide, in that she lacks a belief that it is cyanide; rather, she has .5 credence that it is sugar and .5 credence that it is cyanide. It is false that Alice is blameworthy for spooning the cyanide only if she is blameworthy for her ignorance. Rather, Alice's spooning of the cyanide is a locus of original responsibility: she is directly blameworthy for poisoning Bob. (In a later paper, Rosen grants that in cases like that of Alice and Bob, ignorance does not in any way exculpate.⁶)

The following lesson might be drawn: it is not *ignorance* that can exculpate, but *false belief*. Only the following weaker lesson seems warranted by reflection on cases of action from non-moral ignorance:

Weaker Lesson:

If an agent acts wrongly while falsely believing that p, and if p were true then her action would be morally permissible, *then* the agent's action is not a locus of original responsibility: if she is blameworthy for acting, then she is so blameworthy only in virtue of being blameworthy for her false belief.

⁵ Alexander Guerrero, 'Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution,' *Philosophical Studies* 136 (2007), 59–97. I omit discussion of Guerrero's complex paper; the example I present makes the point his paper presses, but it is not his example.

⁶ Gideon Rosen, 'Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility,' *Journal of Philosophy* 105 (2008), 591–610, footnote 14.

(Note that I do not endorse this principle; I simply claim that the Weaker Lesson is a better inference to draw from the cases than the Lesson Rosen does draw, which is not plausible in light of the cases.⁷)

The upshot is that the generality of Rosen's conclusion is substantially limited. Note that Rosen claimed:

Every instance of morally wrong behavior (whether an action or an omission) is either a case of behavior from ignorance or a case of clear-eyed *akrasia*.

But in conjunction with the Weaker Lesson, to establish the Broad Conclusion, Rosen would need this claim instead:

Every instance of morally wrong behavior (whether an action or an omission) is either a case of behavior from false belief in *p*, such that if *p* is true then the behavior is permissible, or a case of clear-eyed *akrasia*.

This claim is false. Many instances of morally wrong behavior do not involve clear-eyed *akrasia* because the agent lacks a firm view about whether his behavior is wrong; such behavior involves lack of true moral belief – and hence, *moral ignorance* – without involving moral *false belief*. It now seems that, contra Rosen's original argument, such behavior may well be a locus of original responsibility. Just as Alice's uncertainty regarding whether she was spooning cyanide did not undermine her being originally responsible for so acting, similarly mere uncertainty regarding whether one is acting wrongly does not undermine original responsibility for one's action. If someone acts wrongly while genuinely unsure whether her action is wrong, we need not investigate whether she is blameworthy for being unsure to know whether she is blameworthy: she may well be blameworthy simply for doing what she did, which she believed might well be wrong. (Of course she may not be blameworthy if there are other exculpatory factors, for example if she is under the influence of a powerful drug. Or she may not be blameworthy if she has no option available to her that she is sure is morally permissible.)

⁷ In footnote 21, I discuss another way Rosen might revise the lesson in response to the first objection.

The second objection emerges from what we have learned so far. Rosen takes his argument to apply to cases of wrongful omissions, including cases of *forgetting* to do something; he treats such forgettings as involving ignorance. Consider the following case. A doctor forgets to check a chart and fails to see that a patient is allergic to a common antibiotic; she prescribes the antibiotic, and the patient dies.⁸ On Rosen's view, in this case, the failure to check is not a case of clear-eyed *akrasia*, because the doctor is not thinking about whether to check. Thus the failure to check is not a locus of original responsibility. Thus the doctor is blameworthy only if there is some earlier behavior which is an instance of clear-eyed *akrasia*. Let's stipulate there is not.

Let's apply the lessons of the prior discussion. What first emerged is that, if anything, it is *false beliefs* that exculpate. Does the doctor have a false belief that it is okay not to check the chart? No! Indeed, she believes that, in cases like this, she should check the chart. Applying the Weaker Lesson does not vindicate the claim that the doctor is blameless.

Once we see that it is only plausible that *false belief* exculpates, not that *ignorance* does, we will cease to see Rosen's arguments as applying to cases of omissions of actions due to failure to think about or attend to certain facts, including moral requirements. While such cases of forgetting to do things do involve a failure to attend to certain facts, they do not involve the formation of other beliefs that falsely imply one's actions are permissible.⁹

(Note that the second objection does not require that a doctor in such a case is blameworthy; I merely claim that Rosen's line of thought provides no argument that she is blameless.)

The second objection implies a specific conclusion about Rosen's discussion of blameworthiness for ignorance. Rosen says that one is blameworthy for one's ignorance only if one is blameworthy for violating certain procedural norms. He furthermore points out that often one will not have thought about these norms as one violated them. This, he claims, means that one is violating the norms in ignorance, and thus one is not originally responsible for the violation. But this is to treat the

⁸ Note: this is a different case from the earlier case of the doctor who does not double-check a patient's blood type.

⁹ It might be objected that the doctor who forgets to check the chart has a general implicit belief that she is not currently doing anything wrong, and that this is enough to render her action blameless. I respond to this objection in section 6.

violation of the norms in the same way that he treats the doctor who forgets to check the chart for allergies. It may be true that people who violate procedural epistemic norms are not attending to the norms as they violate them; but it is not true that people in such cases have false beliefs to the effect that failure to perform the behavior the norms require is permissible. Rosen's argument does not support his claim that many cases of failure to observe procedural epistemic norms are not loci of original blameworthiness.

I will now turn to the third and fourth objections to Rosen's argument. Both objections arise out of Michelle Moody-Adams's insights.¹⁰ Moody-Adams is concerned to rebut the view that people are blameless for wrong actions which are seen, in their cultures, as morally permissible (or even required). She points out that many such mistaken moral views licensed practices which were of great benefit to those who committed the moral crimes in question. Just as in an individual case, a person may fail to realize something out of *motivated ignorance*, because she does not want to realize it, similarly each individual within a whole group of privileged people may fail to realize that their practices are wrong out of motivated ignorance. If a cultural practice continues and is accepted in a society not simply because people are *ignorant* of its moral wrongness, but because people *don't want to see* its moral wrongness, the ignorance looks less innocent and the practice is more plausibly blameworthy.

The third objection is this:

An individual may be blameworthy for a false moral belief, although he has not violated any procedural norms to do with the management of his opinion, if his false moral belief is due to *motivated ignorance*.

I think this objection is persuasive, but I will note that it relies on the claim that motivated ignorance is blameworthy, which is no doubt a controversial claim. Rosen would respond to the third objection by saying that the third objection attributes blameworthiness for believing, which is something that happens to one, not something one does.

¹⁰ Michelle Moody-Adams, 'Culture, Responsibility, and Affected Ignorance,' *Ethics* 104 (1994), 291–309.

The third objection is not explicitly made by Moody-Adams. Her discussion inspires the point, but she is not committed to it. What she commits herself to is a different claim, which is the fourth objection I will discuss:

A person may be blameworthy for ignorance which arises out of mismanagement of her opinion, although this mismanagement did not involve any clear-eyed *akrasia* (nor is there an earlier instance of *akrasia*).

Moody-Adams points out that if someone has false moral beliefs which benefit him, he may well avoid information and reflection which would lead his false beliefs to be revealed as false: he may be motivated to *protect* his false beliefs. He may not be aware that this is the correct account of his behavior, and so he may not be aware that he is violating any procedural norms to do with the management of his opinion, and so he may not be engaged in any clear-eyed *akrasia*.

William FitzPatrick¹¹ also presses the fourth objection, though he does not focus only on motivated ignorance.

FitzPatrick describes the following case, which is based on the movie 'It's a Wonderful Life'; it is an expansion of a case Rosen discusses. Mr. Potter is a businessman who decides to close down a business, which will result in many jobs being lost. Mr. Potter has thought about the morality of his action, and has concluded that his action is 'permissively aggressive' although the action is really 'reprehensibly ruthless.' Rosen's view implies that Mr. Potter is blameworthy for acting only if he is blameworthy for his ignorance. FitzPatrick argues that this must be a case in which the agent mismanaged his beliefs, because it is not a 'hard case' in which someone thinking carefully about how he should act might get it wrong. For Rosen, the question becomes: did Mr. Potter's mismanagement of his beliefs involve an instance of clear-eyed *akrasia*? If so, Mr. Potter may be blameworthy. If not, Mr. Potter's mismanagement was out of ignorance and then we must ask whether he is blameworthy for that ignorance. FitzPatrick disagrees with Rosen: he thinks that Mr. Potter can be blameworthy for the mismanagement of his beliefs even in the

¹¹ William FitzPatrick, 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge,' *Ethics* 118 (2008), 589–613.

absence of clear-eyed *akrasia*. FitzPatrick says the right question is whether Mr. Potter can have been *reasonably expected* not to mismanage his beliefs; furthermore, FitzPatrick concludes, there are versions of the case in which Mr. Potter *can* reasonably have been expected not to mismanage his beliefs. Mr. Potter's mismanagement of his beliefs may have been an instance of an exercise of a vice of Mr. Potter's. He may have been overconfident, arrogant, dismissive, lazy, dogmatic, incurious, self-indulgent, or contemptuous (these are FitzPatrick's adjectives) in not taking adequately seriously, and in not adequately exploring, alternative moral viewpoints regarding business ethics. If Mr. Potter's ignorance resulted from the exercise of any of these vices, then, FitzPatrick claims, Mr. Potter is blameworthy despite there being no instance of clear-eyed *akrasia*. Thus, FitzPatrick thinks this case establishes the claim made by the fourth objection:

A person may be blameworthy for ignorance which arises out of mismanagement of her opinion, although this mismanagement did not involve any clear-eyed *akrasia* (nor is there an earlier instance of *akrasia*).

The idea is that Mr. Potter may have been, for example, *overconfident* in coming to his beliefs, though this expression of that vice did not result from some earlier behavior which was an instance of clear-eyed *akrasia*. Nevertheless, FitzPatrick contends, Mr. Potter is blameworthy.¹²

The fourth objection's claim is controversial. As pressed by both Moody-Adams and FitzPatrick, the fourth objection is supported by the description of a case and the claim that the agent is blameworthy for his mismanagement of his beliefs. Moody-Adams says that mismanagement of beliefs which is in fact due to a desire not to know the truth (though this may be ill understood by the agent himself) is blameworthy, and FitzPatrick says that mismanagement of beliefs which constitutes the exercise of a vice is blameworthy. The claims are compelling, I think. But

¹² FitzPatrick makes another objection to Rosen, which is an objection to Rosen's argument *from* what I call the Broad Conclusion to the further claim that we never have *knowledge* that someone is blameworthy for an action. FitzPatrick argues compellingly that we do often have knowledge that *akrasia* has occurred, and so that the Broad Conclusion does not imply the further claim.

Rosen would not be convinced. To FitzPatrick, Rosen would respond that most of the vices in question, properly understood, involve the *false belief* that S is not mismanaging his beliefs. One's time for moral consideration is not limitless, so it is not a good idea (nor morally required) to consider every idle moral worry. Consider the vices of overconfidence, arrogance, dismissiveness, dogmatism, incuriosity and contemptuousness. These vices seem to involve the *belief* that the other views in question aren't worth taking seriously. These are *false beliefs* in this case, but Rosen would say that, in virtue of Mr. Potter's having such beliefs, his mismanagement of his opinion is out of ignorance (indeed, is vindicated by these false beliefs). The disagreement between FitzPatrick and Rosen amounts to a standoff over the case FitzPatrick describes. Rosen will simply deny FitzPatrick's claims about the case. But others may be moved by the case and see it as showing that Rosen's argument fails.

In discussing Mr. Potter, FitzPatrick makes an argument to which I will now object. FitzPatrick argues that an immoral businessman such as Mr. Potter must have mismanaged his beliefs – that is, violated a procedural epistemic norm – because the case is not a 'hard case.' Here FitzPatrick assumes a wholly implausible picture of how moral learning occurs. Many people have false moral beliefs although they have thought carefully about the questions at issue and they have not violated any procedural norms. For example, for many business practices that are in fact reprehensibly ruthless, we can find plenty of businesspeople ready to offer elaborate, sustained, and serious moral defenses of them. We can also find people who have entered the business world, heard arguments on all sides, and become convinced of the wrong view of business morality.

FitzPatrick wants to preserve many of our claims of ordinary blameworthiness. He wants to vindicate the claim that immoral business practices, which we ordinarily take to be blameworthy, are blameworthy; and he wants to vindicate the claim that the actions of the Bush Administration regarding Afghanistan and Iraq are blameworthy. FitzPatrick quotes those who have commented that George W. Bush himself was remarkably incurious; it does appear that Bush is guilty of mismanagement of his beliefs. But presumably FitzPatrick thinks that all the important players in the Bush administration are blameworthy, and it is far from clear that they all are guilty of mismanagement of their beliefs. Some of

them appear to have had a well-thought-out, though extreme, agenda.¹³

It's important that moral knowledge isn't *easy* to come by, and isn't assured if one manages one's beliefs well, not even if one is exposed to good moral arguments for the truth.

My point is this:

Someone may have been perfectly responsible regarding procedural norms for the management of his beliefs while nevertheless having come to deeply false moral views, even where he is perfectly aware of substantial disagreement. This is true regarding most interesting moral claims, including claims about how businesspeople should behave, how people in government should behave, how we should treat our children, etc. It is true of most (if not all) ordinary moral claims.

FitzPatrick supposes that it is only in 'hard cases' that someone may think seriously about morality, in the face of views that oppose his own, and come to the wrong view. And he claims that business ethics and the ethics of government are not hard cases. Unfortunately, ethics is not that easy: if we use 'hard cases' in FitzPatrick's sense, then most cases are hard cases.

I will draw three lessons from my discussion of these four objections to Rosen. First, it is not at all plausible that mere ignorance exculpates; at best, it is false belief that exculpates. Second, motivated ignorance may be blameworthy although it is not the result of an earlier instance of clear-eyed *akrasia*. Third, some mismanagement of belief may be blameworthy without involving clear-eyed *akrasia*. Taking these lessons into account, it may seem that Rosen's argument still supports the following:

Narrower Conclusion:

If

- (a) a person acts wrongly while believing a false claim, p,
- (b) if p is true then the action is permissible,
- (c) the false belief did not result from mismanagement of belief, and

¹³ We should not be misled by the fact that they were incurious about some things, such as the truth about global warming: there were some things they wanted little information about because they were determined to act in a particular way regardless of that information.

(d) the false belief is not a case of motivated ignorance,
then
the person is blameless for so acting.

Clauses (a) and (b) place a restriction to cases of false beliefs that license the action, taking the first objection seriously. Clause (c) restricts us to cases that do not involve mismanagement of beliefs. Note that this restriction is more restrictive than is necessary in light of the fourth objection. I am restricting in this way to set the fourth objection aside. Clause (d) restricts us to cases that do not involve motivated ignorance, taking the third objection seriously.

3. The Significance of the Narrower Conclusion

The Narrower Conclusion is still quite strong; it implies that in many ordinary cases of apparent blameworthiness, the agents are in fact blameless. I will now describe a number of cases in which we would ordinarily take the agents to be blameworthy, but the Narrower Conclusion implies they are not. I stipulate for all of these cases that clauses (c) and (d) apply. It is important to emphasize this stipulation because it is easy to imagine variants of the cases in which either motivated ignorance or mismanagement of beliefs does play a role. Nevertheless, there are many actual cases of the types I will outline in which ordinary people come to false moral views without having mismanaged their beliefs. In many of the cases below, I stipulate that the agents have thought long and seriously about the relevant moral questions, have considered opposing views, and have come to false moral views. In some of the cases, the agents are not aware that people might disagree with them; these agents have not thought particularly hard about the moral views in question, but they have thought an ordinary amount about morality in general, and so they have not violated any procedural norms.

As I've already suggested, one case is that of a businessman who has thought seriously about the ethics of business and who has concluded that what are in fact 'reprehensibly ruthless' practices are really 'permissibly aggressive.' This businessman is bothered by the fact that some people criticize his practices. He isn't surprised that those who are hurt by the practices say they are

ruthless, but it bothers him that his brother says so. He discusses it at length with his brother and also with others. He comes to the view that it is better for the economy overall if companies operate to maximize profits, and that economic progress always has casualties. He believes that by pursuing what is in the best interests of his shareholders, he is playing a valuable role in the economy. This is not a great view. But someone may well have thought seriously about the morality of business practices and come to such a view. Indeed, many actual businesspeople have thought about the morality of business, have not been irresponsible in the management of their opinions, and have false views that license their behavior.

Two more examples are variants of cases Rosen discusses. One is a case of a 1950s American father who doesn't save to send his daughter to college but does save to send his sons to college: he is a sexist in his treatment of his children. The other is an ancient Hittite slaveholder who keeps slaves. Rosen stipulates that these two agents do not believe their actions are wrong. He says they do not think about whether the actions are wrong at all. In the versions of the cases I will focus on, the agents have *implicit* moral beliefs that their actions are permissible. (It's not clear if this is what Rosen has in mind.¹⁴) I stipulate that these agents care about being moral and they think about morality to a reasonable degree.

Another kind of case involves people who are raised to believe in an ethics of 'everyone should take care of his own.' Consider three cases along these lines. One is someone raised in a Mafia family who goes into the family business and believes in an ethics of deep loyalty to the family business group and no moral obligations to those beyond it; this man kills a store owner who won't 'pay for protection'. Another is someone raised in a tribal group that is in conflict with other nearby tribal groups; he kills a member of a rival group in order to assert his own group's dominance. Finally, consider someone raised in a poor neighborhood of an American inner city. He joins a gang and believes that one should show deep loyalty to one's own and one should be willing to do anything to

¹⁴ Guerrero 'Don't Know, Don't Kill' discusses several variants of Rosen's cases, in some of which the agents simply *lack beliefs* about whether their actions are wrong, rather than having beliefs that their actions are permissible. As Guerrero points out – this is an instance of the first objection I discussed – mere lack of belief that one's action is wrong does not exculpate.

others for the sake of the gang. These three agents have not been irresponsible in the management of their beliefs. Let's suppose that the tribe member is not aware of other moral views. The Mafia member and the gang member are aware of people who disagree with the moral views they hold. But they believe others have been suckered into a false sense of their duties.

Another kind of case is given by actions which are currently deeply morally contested; I will give three such examples. In giving these examples, I make controversial moral claims; but these claims are not crucial to my points in this paper. Consider someone who believes abortion is wrong and who yells at women outside abortion clinics. It is wrong to yell at women outside abortion clinics: these women are already having a hard time and making their difficult decision more psychologically painful is wrong. But this person acts in a way that would be permissible if her moral views were true. Another example is someone who believes abortion is wrong and who kills an abortion doctor, in a part of the country where there is good reason to think that this doctor's death will reduce the number of abortions. This person believes that he ought to kill abortion doctors if doing so would reduce the number of abortions that would be performed. A third example is someone who believes homosexuality is wrong who organizes a campaign against the legalization of gay marriage. He believes he is doing something morally good in organizing the campaign; in fact, in working to further oppression, he is acting wrongly.¹⁵

I have stipulated that in all of the above cases, the agents have not violated any procedural norms to do with the management of their beliefs. Someone might object that some of the moral views in these cases are *too crazy* for anyone to have come to them without having behaved poorly in the management of her beliefs. But that seems manifestly false.

These examples show that, if the Narrower Conclusion is true, then in many cases in which we would ordinarily take people to be blameworthy, they are in fact blameless. The Narrower Conclusion is thus significant and deeply implausible.

¹⁵ If my controversial claims are false, then different examples should be used to make my point. For example, if abortion is morally wrong, then a doctor who believes abortion is permissible is nevertheless blameworthy for performing abortions.

4. My Proposed View

I will now propose a view on which the actions I have just described are indeed blameworthy, as they appear to be. The view I propose denies Rosen's claim that actions done by agents with false moral beliefs cannot be loci of original responsibility. However, the view grants that there is something right about Rosen's position. In particular, it cannot be that a person is blameworthy for a wrong action he believed to be permissible but is blameless for his false belief. While denying that the blameworthiness for the wrong action is *derivative* from the blameworthiness for the belief, the view holds that in such cases the action is blameworthy *only if* the belief is blameworthy as well.

On the view I propose, one can be blameworthy for having false moral beliefs although one has not been in any way irresponsible in the management of one's opinion: one may not have acted in any way procedurally badly. Rather, one has violated some moral norms that apply to *beliefs* themselves, not to the management of one's beliefs. The view holds that we morally ought to believe the moral truth.

Here is the proposed view:

We are morally obligated to believe the moral truths relevant to our actions (and thus not to believe false moral claims relevant to our actions), and we are often blameworthy for failing to meet these moral obligations, even if we have not been guilty of mismanagement of our beliefs, and even if our ignorance is not motivated.

Wrong actions that result from false moral beliefs are not thereby blameless; indeed, they may be loci of *original* responsibility. While both the beliefs and the actions are blameworthy, the actions are not blameworthy *because* the beliefs are blameworthy. Rather, the actions and the beliefs are blameworthy for similar reasons.

These are the essential components of the view. I propose the view to show what an alternative picture to Rosen's could look like, while capturing some of what is plausible about his view (namely, that if the actions in these cases are blameworthy, then the beliefs are too) without resulting in his implausible conclusions.

So far the view leaves open what ultimately explains blameworthiness. One way of spelling out the view further would be to appeal to Nomy Arpaly's view of blameworthiness:¹⁶

An action is blameworthy just in case the action resulted from the agent's caring inadequately about what is morally significant – where this is not a matter of *de dicto* caring about morality but *de re* caring about what is in fact morally significant.

(One *de dicto* cares about morality if one desires *to act morally*. One *de re* cares about what is in fact morally significant if, for example, one desires *not to cause others' suffering* (and this is in fact morally significant).¹⁷)

We might furthermore expand the view beyond Arpaly's own development of it:¹⁸

Beliefs (and failures to believe) are blameworthy if they involve inadequately caring about what is morally significant. Believing a certain kind of behavior is wrong on the basis of a certain consideration is a way of caring about that consideration.

Some failures to believe moral truths relevant to one's actions are not blameworthy. For example, if one blamelessly falsely believes a non-moral claim, and this leads to one's false moral belief, then one's false moral belief does not involve inadequately caring about what is morally significant.

5. Objections to the Proposed View

I will now discuss some objections to the proposed view. This discussion will bring out some further aspects of the view.

The first objection:

In the cases described in section 3, the agents who act wrongly have *epistemically justified* false moral beliefs that their actions

¹⁶ Nomy Arpaly, *Unprincipled Virtue* (Oxford: Oxford University Press, 2003).

¹⁷ Exactly what is involved in *de re* caring about morality depends on the true moral view. On some views, promise-breaking is morally significant; on these views, desiring *not to break promises* would be an instance of caring about what is in fact morally significant. But if utilitarianism is true, then promise-breaking is not morally significant, and this desire would not be an instance of caring about what is in fact morally significant.

¹⁸ I believe this is an expansion of Arpaly's view, but perhaps she intended this to be part of her view.

are permissible. Surely they cannot be blameworthy for having beliefs that are epistemically justified (since they haven't mis-managed their beliefs). But then they cannot be blameworthy for acting *on* their reasonable beliefs.

In favor of the claim that the agents' beliefs are epistemically justified, the objector might offer the following considerations. First, some of the agents' evidence is *testimonial*: people they reasonably take to be authorities about morality have made the claims they believe. Second, many of them have considered arguments for and against their views and they have been convinced by seemingly compelling arguments. Third, some of the considerations they rely on to support their views do genuinely *tell in favor* of their beliefs. Fourth, they have genuinely tried hard to think about the issue and upon reflection have arrived at their views. When we focus on the way that they came to their beliefs and their judgments that their beliefs are the right response to their evidence, it may appear that their beliefs are justified. (This applies to those agents who've considered the moral issue carefully, not to those who haven't considered it at all, such as the slaveholder.)

My response to the objection is that the agents' beliefs are not epistemically justified. Consideration of fallacious, though seductive, lines of thought is not a way of coming to have justified beliefs. Someone who is suckered in by a sneaky instance of affirming the consequent is not justified merely by its having seemed to him that the considerations he relies on do support the conclusion he draws. While testimony can lead to justified belief, we have stipulated that in many of these cases the agents are not *merely* relying on testimony. Consideration of evidence which shows a claim to be false can undermine testimonial justification for the claim.

The objector's claim that the agents' false moral beliefs are epistemically justified is more compelling regarding those agents who are not aware that there is disagreement about their ethical beliefs; these agents hold these beliefs on the basis of testimony, and have not thought the moral questions through for themselves. Nevertheless, this does not show that their beliefs are epistemically justified. It is a hard question what constitutes evidence for moral claims, and in what circumstances a person's evidence is such that the right response to the evidence – the epistemically responsible response – involves believing the moral truth about a certain matter. But I claim that ordinary people who

know the non-moral facts of what they are doing, when they do wrong things, often do have *sufficient evidence* that their actions are wrong. For this reason, their beliefs that their actions are permissible are unjustified. (This is a substantive and controversial claim I will not defend here.)

My considered view is that the agents in all of the cases I described in section 3 do not have epistemically justified beliefs. But if they do have epistemically justified beliefs, I accept the implication of my proposed view that they are blameworthy for having beliefs that are epistemically justified, and they are blameworthy for acting on those beliefs.¹⁹ I see their cases as instances of constitutive moral bad luck: they are unlucky to have found themselves in circumstances that have caused them to be (in some respects) morally bad people.

The second objection to the proposed view:

The beliefs in question are epistemically justified, so believing otherwise would be epistemically unjustified. We cannot be morally required to believe in a way that would be epistemically unjustified.

Like the first objection, this objection claims that the beliefs in question are epistemically justified. I will grant for the sake of argument the second objection's claim that the agents' beliefs are epistemically justified.

My response is that there can be cases in which a person is epistemically justified in believing *p*, but *p* is false, and she is *in a position to realize* that *p* is false, such that she would become epistemically justified in believing not-*p*: she is in a position to become justified in believing not-*p*. For example, consider the following case. I studied some advanced math in college, but that was years ago. My friend Moon, who has a PhD in math, might tell me that a particular mathematical claim, *p*, is true. I may

¹⁹ There is independent reason to think that people can be blameworthy for having epistemically justified beliefs, and for acting on these beliefs (even when the blameworthiness cannot be traced to earlier wrong behavior). Simon Keller ('Friendship and Belief,' *Philosophical Papers* 33 (2004), 329–351) and Sarah Stroud ('Epistemic Partiality in Friendship' *Ethics* 116 (2006), 498–524) have argued that we owe our friends the benefit of the doubt, such that we ought to refrain from believing ill of them in some cases in which such beliefs would be epistemically justified. On this view, someone who thinks ill of a friend when she owes it to her friend to withhold belief is blameworthy for her belief; and if she acts wrongly on the basis of the belief, it seems she is blameworthy for so acting.

remember enough math to fully understand what *p* says. I come to believe *p* on the basis of Moon's testimony, and I am justified in so believing. However, *p* is false, and I still remember enough of the relevant math that I *could* realize *p* is false and I could come up with a proof of not-*p*. It is not assured that if I try to think about whether *p* is true, I will have this realization; it may not come to me. But I could have this realization; I am *capable* of having it. Thus, while my belief in *p* is epistemically justified, I am in a position to become justified in believing not-*p*: I have the mathematical tools to go through a thought process that would undermine my justification for believing *p* and render me justified in believing not-*p*. This case is analogous to the case of people who have false moral beliefs (if, as I am granting to the objector, those beliefs are epistemically justified): they *could* realize the moral truth, though for many of them there may be no guarantee that they will if they try.

The third objection:

One can't be blameworthy for something's not *occurring* to one, even if it's *possible* that it occur to one, and particularly if there's no particular step one might have taken that would have ensured that it occur to one.

This objection is a natural response to what I have said in response to the second and third objections. However, the claim the objector makes is simply false. One can be blameworthy for something's not occurring to one. A parent who forgets to pick her child up at school is blameworthy for forgetting, even in a case in which she took what everyone would agree are reasonable precautions against forgetting. A doctor who forgets to check a chart for an allergy is blameworthy for forgetting, even in a case in which she took what everyone would agree are reasonable precautions about forgetting. Of course *forgetting* is different from *failing to realize*, though both are cases of something's not occurring to one. The objector might grant that one can be blameworthy for *forgetting* while denying that one can be blameworthy for not *realizing* something. But there are counterexamples to that claim too. Suppose that we simultaneously experience a blackout and smell a gas leak, and I decide to light a match to navigate my way through the building. I do not *realize* that this may cause an explosion. I *could* realize it, but there's no particular step I could have taken that would have ensured that I realize it. I may still be blameworthy for

my failure to realize, and for risking an explosion. By its nature, *realizing* something, like *remembering* something, is such that when one is capable of doing it, it's also true that one might fail to do it; and is such that there is no particular step one can take to ensure that one do it, short of already doing it.

6. Understanding My Disagreement with Rosen

In this section, I will step back and make some general comments aimed at illuminating the dispute over whether this claim, to which Rosen is committed, is true:

Weaker Lesson:

If an agent acts wrongly while falsely believing that p, and if p were true then her action would be morally permissible, then the agent's action is not a locus of original responsibility: if she is blameworthy for acting, then she is so blameworthy only in virtue of being blameworthy for her false belief.

My proposed view is just one instance of a family of possible views on which Weaker Lesson is false. In this section, I will set aside the further features of my proposed view and focus on my dispute with Rosen regarding Weaker Lesson.

Weaker Lesson is supposed to be supported by the case of Anne, who spoons cyanide into Bill's coffee while thinking it is sugar (and who is blameless for her false belief). Anne is blameless for poisoning Bill; Rosen would say that Weaker Lesson tells the right story about why she is blameless. Anne's belief that she is spooning sugar is such that, if it were true, her action would be permissible, and she is blameless for having this belief. These two facts, according to Rosen, are why she is blameless.

I think Anne's blamelessness stems not from the *presence* of a particular false belief, but from the *absence* of any beliefs that make her blameworthy. I think the right lesson to draw is this:

The Right Lesson:

If an agent acts wrongly and nothing in her epistemic state can ground her blameworthiness for acting as she does, then the agent's action is not a locus of original responsibility: if she is blameworthy for so acting, then she is so blameworthy only in virtue of being blameworthy for her epistemic state.

There is nothing at all in Anne's epistemic state to which we could appeal to ground her blameworthiness. She does not knowingly do something morally wrong; nor does she knowingly poison someone; nor does she knowingly risk poisoning someone, etc.²⁰

Rosen should agree with me that the Right Lesson makes a true claim. My disagreement with Rosen is not over whether the Right Lesson is true, but over what in an agent's epistemic state can ground her blameworthiness for acting as she does. Rosen thinks that an agent who knowingly enslaves another, who does not know that doing so is wrong, *lacks* anything in his epistemic state which grounds blameworthiness for his action. I disagree. I think that knowingly enslaving someone is itself blameworthy: the feature of the slaveholder's epistemic state that renders him blameworthy is his knowledge that he is enslaving someone. (While Rosen should agree with me that Right Lesson is true, he will hold that Weaker Lesson (or a claim along its lines) is also true and is explained by Right Lesson; I deny this.)

I am now in a position to tie up a loose end from my earlier discussion.

In section 2, I discussed the example of a doctor who forgets to check the chart to see whether her patient is allergic to the drug she is going to administer; I said the doctor, who does not think of whether to check the chart, has no false belief such that if it were true her action would be permissible. On the contrary, she knows that she is administering a drug, she knows that she is not currently checking the chart, and she knows that in such cases she ought to check the chart; she is not explicitly thinking about these facts, but she does know all of them. In footnote 9 I mentioned the objection that the doctor must have an implicit belief that she is not currently doing anything wrong. My response to this objection is that *even if* the doctor has that implicit belief, that is not enough to show that her action is not a source of original responsibility: while she may have that implicit belief, her *whole epistemic state*

²⁰ Note that the Right Lesson is compatible with a view on which it is a *global* question whether an agent's epistemic state renders her action blameworthy: some beliefs may ground blameworthiness for Aing in one case, but the same beliefs would not ground blameworthiness for Aing if other beliefs are present. (For example, Sam may be blameworthy for knowingly pushing someone to the ground. But Bob may not be blameworthy though he knowingly pushed someone to the ground, since Bob had the further belief that this was the only way to save that person's life.)

Indeed my objection to Weaker Lesson is that it does not allow the *whole* of a person's epistemic state to determine whether she is blameworthy.

includes knowledge of facts that render her action blameworthy. When deciding whether a person's epistemic state gets her off the hook – such that if she is blameworthy at all, it is only for being in that epistemic state – we must consider the *whole* of her epistemic state, not merely a single belief within it.

(The example of the doctor who forgets to check the chart highlights a more general phenomenon: that a person holds a particular belief such that, if true, her action would be permissible, does not ensure that she does not hold other beliefs which straightforwardly imply her action is impermissible. Rosen thinks that whether a person is blameworthy crucially depends on the person's moral beliefs. But given cases in which a person's moral beliefs are not fully coherent, even he should deny Weaker Lesson. I am not sure how to state the best version of his view.²¹)

The dispute between Rosen and me regarding Weaker Lesson (suitably fixed up) should be understood as a dispute about whether the Right Lesson implies Weaker Lesson. Rosen and I agree that Anne's blamelessness is explained by her epistemic state; but we disagree about what features of her epistemic state make her blameless, and what the lesson is for cases of moral false belief. We also disagree about how much *help* consideration of Anne's case gives us when we turn to consider cases of agents who act wrongly, with full factual information but with false moral beliefs. I think all we learn from Anne's case is that a person's epistemic situation, considered as a whole, may contain no features which ground her blameworthiness, even though she is acting wrongly.

My claim that Anne's case (or any case of mere non-moral ignorance) does not provide the lesson Rosen draws is further bolstered by reflection on the stated version of Rosen's view, which is that Anne's *ignorance* renders her blameless. Rosen does not characterize the ignorance Anne has, and so when he says that

²¹ One possible revision of the lesson is the following, which Rosen has suggested to me in conversation about cases of uncertainty: an agent who is ignorant that her action is wrong *in any sense* (either objectively wrong or subjectively wrong) is blameworthy for acting only if (and only because) she is blameworthy for her ignorance. With this revision, Rosen can acknowledge that Alice is blameworthy when she spoons what she has a .5 credence is cyanide into Bob's coffee: Alice knows that taking a 50% risk of poisoning someone is wrong. But this revision will not correctly handle cases that involve uncertainty whether an action is subjectively wrong: an agent who acts wrongly with 50% credence her action is subjectively wrong is blameworthy for so acting (if she knows the alternative is subjectively permissible). More seriously, this revision runs into the problem I lay out in the final paragraph of this section.

his principles regarding ignorance apply *in full generality* (that is, to moral ignorance as well as factual ignorance), he faces a problem. The relevant principle is that an agent who acts in ignorance is blameworthy for acting only if (and only because) she is blameworthy for her ignorance. This principle is elliptical; not any ignorance will do, so we must characterize the ignorance. Rosen's principle could be read as either of the following claims. (a) An agent who acts wrongly in ignorance of every fact that is sufficient for the wrongness of her action is blameworthy for acting only if (and only because) she is blameworthy for this ignorance.^{22,23} Or (b) an agent who acts wrongly in ignorance *that* her action is wrong is blameworthy for acting only if (and only because) she is blameworthy for this ignorance. The problem Rosen faces is that he needs to be appealing to principle (b); principle (a) does not have the implications he draws for cases of moral ignorance.²⁴ However, the case of Anne does not support (b); while Anne's case supports (a), reflection on Anne's case need not lead anyone to endorse (b). Whether (b) is plausible depends on whether one thinks moral ignorance is exculpating; (b) gains no independent plausibility from consideration of Anne's blamelessness.

7. Conclusion

I have discussed some objections to Rosen's Broad Conclusion and argued that, even in the face of these objections, Rosen's arguments still may appear to support the Narrower Conclusion, which has sweeping and implausible implications. I have described some cases of actions which are blameworthy but which the Narrower Conclusion implies are blameless. I have proposed a view on which these actions are blameworthy, contrary to Rosen's view. I have defended the view against some objections. I have not independently argued for the proposed view, though it is supported by the claim that the actions I described are indeed

²² I state the principle in this way to cover cases in which an action is wrong for two independent reasons; an agent ignorant of the fact that makes his action wrong in one way is not blameless if he is aware of the other fact that makes his action wrong.

²³ Note that each fact sufficient for the wrongness of an action will be a complex fact.

²⁴ An agent who is morally ignorant but not factually ignorant knows her factual situation, and so knows a fact sufficient for the wrongness of her action.

blameworthy. I have not aimed to establish the proposed view, but to set it out as an alternative to the picture Rosen offers. Finally, I have argued that consideration of cases in which false non-moral beliefs exculpate does not support Rosen's view that false moral beliefs exculpate.

*Department of Philosophy
Princeton University
Princeton, NJ 08544
eharman@princeton.edu*