

The Irrelevance of Moral Uncertainty

by Elizabeth Harman
eharman@princeton.edu

I. Moral Uncertainty

Consider this case:

B is deciding whether to ϕ . B believes that ϕ ing is morally required, while failing to ϕ would be morally wrong. B thinks that failing to ϕ wouldn't be deeply morally wrong; it would only be minorly morally wrong. Although B believes that ϕ ing is morally required, B isn't 100% sure. B is 90% sure. B has 10% confidence that ϕ ing is actually morally wrong, and indeed is deeply morally wrong.

What should B do? It might seem that B should ϕ ; after all, B believes that ϕ ing is morally required. But consider this line of argument:

B ought to refrain from ϕ ing. Taking a 10% chance is taking a substantial chance of doing a deeply morally wrong thing. One should be very averse to risking doing a deeply wrong thing. It is better to do what is very likely a minorly wrong thing to do than to risk doing what is a deeply wrong thing to do. Suppose that B does ϕ , ϕ ing is in fact morally wrong, and a victim of B's ϕ ing were to later try to hold B responsible for ϕ ing. Could B defend him- or herself by saying "but I believed that ϕ ing was morally required"? No! B knew that B was risking doing a deeply morally wrong thing, only for the sake of what B believed was a minor moral requirement. For this reason, B is blameworthy for ϕ ing, and it is appropriate to hold B responsible for ϕ ing.

This way of thinking about cases like B's is very attractive, and there is a very interesting philosophical literature that takes this way of thinking to be correct and then seeks to answer further questions that arise. I will call a proponent of this line of argument an Uncertaintist, and the view that this line of thought is correct Uncertaintism. According to Uncertaintism, an agent's moral uncertainty (and specific moral credences) are crucially relevant to how the agent should act.¹

¹ Uncertaintist thinking appears in Ross 2006, Sepielli 2008, and Sepielli 2013. (Though these authors do not appeal to considerations of blameworthiness to support their claims.) Related thinking appears in Lockhart 2000, Guerrero 2007, and Moller 2011; all three of these authors claim that an agent's actions should be guided by her moral credences. My argument against Uncertaintism can be adapted to target these three views, as I will explain in footnotes.

Uncertainty begins by considering and rejecting the following view: an agent should be guided by any moral claims she believes. Cases like B's bring out that sometimes an agent takes a big moral risk by being guided by what she believes. If her beliefs hold that very little of moral significance is at stake, but she has a small credence in the claim that a great deal is morally at stake, then she may be taking a big moral risk in simply ignoring her small credence. B's ϕ ing would be like a homeowner's failing to buy fire insurance: one believes that one's home won't burn down, but the small credence one gives to a fire makes it reasonable to buy fire insurance even though there is simply a net loss if no fire occurs.

Uncertainty then consider the following view: one should maximize the expected moral value of one's actions. This is not a consequentialist view. Rather, the view is that some actions are morally worse than others, according to various moral claims. One should maximize the expected moral value *of one's action*. On this view, B gives 90% credence that ϕ ing is morally good, but only minorly morally good, 90% credence that failing to ϕ is minorly morally bad, 10% credence that ϕ ing is very morally bad, and 10% credence that failing to ϕ is very morally good. The overall expected moral value of his action is maximized by failing to ϕ .

The view that one should maximize expected moral value faces a difficult puzzle, which can be illustrated by the following case. Suppose that a person has some credence in Utilitarianism and some credence in Kantianism, and she is trying to decide whether to push one person in front of a train to stop it from hitting five people. Utilitarianism holds that pushing is morally required, while Kantianism holds that it is morally wrong. Does Utilitarianism hold that letting five die is just as morally bad as the Kantian holds that killing the one is morally bad? The answer to this question matters if agents should maximize expected moral value; but it is very hard to answer it. More generally, it is hard to know how to compare the moral values that different moral claims or principles assign to an agent's options.²

While this puzzle is interesting, I will argue that this literature is based on a mistake: the mistake of thinking that Uncertainty is true. Once we see that Uncertainty is false, we will see that we do not really face this puzzle.

To see why Uncertainty is false, let's look more closely at the initial line of argument I laid out concerning B, which the Uncertainty endorses. I will make 3 observations about the Uncertainty's view.³

First: In saying that B should refrain from ϕ ing, the Uncertainty is making a *moral* claim about B, not a claim about how B might best pursue B's goals. (The description of B does not mention whether B cares to do what is moral; though the description may imply that he does.) Suppose we consider Bill, a version of B who

² Lockhart 2000, Ross 2006, Sepielli 2008, and Sepielli (manuscript) have offered solutions or partial solutions to this puzzle.

³ These are observations about Uncertainty as I understand it, and what I take to be the best version of Uncertainty. It is a subtle question whether these observations apply to what has actually been written by Ross and Sepielli, as is whether they apply to related thinking by Lockhart, Guerrero, and Moller. I will take these questions up in footnotes throughout the paper.

doesn't care about morality at all. We might run through that same line of argument about Bill, finding it just as compelling: Bill should be morally cautious, even though we know Bill won't be morally cautious (or if he does the morally cautious thing, it won't be out of concern for morality). The claim that B ought to refrain from ϕ ing is a claim about what B ought to do, morally.

Second: In saying that B should refrain from ϕ ing, the Uncertaintist is making a moral claim that is *subjective* rather than *objective*. To see the distinction, consider the following case. Anne's husband is dying and Anne gives him what she has every reason to believe is the cure to his illness. In fact it is poison. The following normative claims both seem true of Anne:

Anne does something she shouldn't do. (She poisons him.)

Anne does something she should do. (She gives her husband what she has every reason to believe is his cure.)

The first normative statement is true, and it is an *objective* moral truth. The second normative statement is true, and it is a *subjective* moral truth. Loosely speaking, objective moral truths are made true by features of the world independently of anyone's (perhaps mistaken) beliefs about what the world is like, while subjective moral truths are made true by someone's mental states (which may be mistaken), often but not always the agent's mental states.⁴

Whenever a subjective moral claim "A ought to refrain from ϕ ing" is true and is made true by A's beliefs and credences (rather than by an advisor's or an observer's beliefs and credences), one can express the same truth by saying "Given A's whole mental state, A ought to refrain from ϕ ing." That is what the Uncertaintists are saying. They are saying: given the entirety of B's beliefs and credences, B ought to refrain from ϕ ing.

My third observation is that the Uncertaintist is committed to the following claim: if someone is caught in the grip of a false moral view (if she is sure of that view), then she ought to act as her false moral view requires. This is simply the limiting, uninteresting case of one's moral credences determining how one should act. According to Uncertaintism, if one is sure that one is morally required to ϕ , then one should ϕ .

I will now argue that Uncertaintism is false. Uncertaintism is the view that in light of an agent's whole mental state—given all of her credences—she should do the morally cautious thing in cases like B's. Furthermore, someone caught in the grip of a false moral view should do what that moral view holds is morally required. These are subjective moral claims, made true by the agent's whole mental state. But the following principle holds regarding such subjective moral claims:

An agent is blameworthy for her behavior only if she acted as she subjectively should not have acted.

⁴ There is quite a rich literature on subjective and objective normative statements, and on subjective claims made by advisers and remote observers. Subjective statements may also have true readings relative to an agent's *evidence* rather than her *credences* (see section VIII). See Dowell (forthcoming); Jackson 1991; Kolodny and MacFarlane (manuscript); MacFarlane (forthcoming); Smith (manuscript); and others.

Consider Anne, who acts wrongly by poisoning her husband. She is not blameworthy because it is not the case that she subjectively should not have acted that way. Subjectively, she should have given him what she took to be the cure.

Of course, Anne might be blameworthy for *causing* the poisoning if she is blameworthy for having the credences she has. If Anne didn't look at the bottle before giving it to her husband, for example, she may be blameworthy. But Anne would not be blameworthy for *behaving* as she did; she wouldn't be blameworthy for poisoning him. Rather, she would be blameworthy for *causing* the poisoning to occur by failing to look at the bottle.

Uncertaintism makes some subjective moral claims. They have implications for blameworthiness. They imply, for example, that in B's case, if ϕ ing is indeed morally required, as B believes, but B does not ϕ , then B is blameless for failing to ϕ , because B did as B subjectively ought to have done. According to Uncertaintism, B's failure to ϕ would be blameless for the same reason that Anne's poisoning of her husband is blameless: they did the right thing, given their beliefs and credences.

The Uncertaintist is committed to the view that those caught in the grip of false moral views, who do morally wrong things while sure that those things are morally required, are blameless for their behavior. According to Uncertaintism, these agents acted as they subjectively should have acted. Uncertaintism is committed to the view that being caught in the grip of a false moral view exculpates. That view is false, I claim, and so Uncertaintism is false.

My Main Argument:

1. Uncertaintism implies that being caught in the grip of a false moral view is exculpatory.
 2. It is not true that being caught in the grip of a false moral view is exculpatory.
- Therefore:
3. Uncertaintism is false.

This argument has a controversial premise: premise 2. I won't seek to *establish* that premise in this paper, but in section III, I will say some things to defend it and to make it seem plausible. (In section VI, I will discuss whether premise 1 can be denied.)⁵

⁵ My main argument can be adapted to target the views of Lockhart, Guerrero, and Moller. Lockhart claims that an agent should minimize her chances of acting wrongly. Guerrero and Moller claim that, at least for cases of killing, an agent should avoid doing something she believes may well be morally wrong, if she believes her alternative is definitely morally permissible. As I read these authors, all three of their views imply that an agent caught in the grip of a false moral view should act as her moral view dictates; thus, all are committed to the view that being caught in the grip of a false moral view is exculpatory. Guerrero is an interesting case here, because his primary goal in his 2007 is to point out that moral *ignorance* does not exculpate. He considers cases in which people do morally wrong things while having a small credence in the true moral view that their actions are wrong; he argues that they are blameworthy. (As we will see, I agree about those cases, but not because of the agents' moral credences.) But he argues that one

In section II, I will discuss the implications of my main argument by offering and explaining an alternative to Uncertainty. In section III, I will support my main argument by explaining and defending premise 2. In sections, and VI, I will discuss objections to my arguments. In section IV, I also offer another argument against Uncertainty.

II. The Alternative to Uncertainty

In this section, I will discuss the implications of my argument.

Because Uncertainty is false, the puzzle we discussed above, about how to compare moral value between conflicting moral views, is not important. It may be interesting *as a puzzle*; but nothing normatively important hangs on solving it.

If Uncertainty is false, what is true in its place? Here is my proposal:

Actualism: A person's moral beliefs and moral credences are usually irrelevant to how she (subjectively) should act. How a person (subjectively) should act usually depends solely on her non-moral beliefs and credences; her moral beliefs and credences are relevant only insofar as they provide warrant for beliefs and credences about what her non-moral situation may be.⁶

Why is "Actualism" a good name for this view? Because according to Actualism, what a person (subjectively) should do depends crucially on what's *actually* the true moral theory, and not on what people *believe* is the true moral theory.⁷ The contrast is between what is *believed* and what is *actual*. "Actual" here doesn't have anything to do with *contingency*, as our moral requirements (at a sufficient level of generality) are necessary.

What does Actualism say about B? As B's story stands, it is too underdescribed to settle what B should do. Different ways of filling out B's story (or a similar story) will lead to very different results as to how B should act. Consider this way of filling out B's story:

should refrain from doing *what one believes might well* be morally wrong (at least in cases of killing), which implies one should refrain from doing what one is certain is morally wrong.

Weatherson (manuscript) argues against views along the lines of Lockhart's, Moller's and Guerrero's by arguing that such views are implausible in cases of *prudential* uncertainty and then arguing that moral uncertainty is more analogous to prudential uncertainty than to non-moral uncertainty.

⁶ This is my initial statement of Actualism. It implies both that Uncertainty is false and that the related claims of Lockhart, Guerrero, and Moller are false. (See footnote 5.) In section VIII, I offer a revised statement of Actualism which explicitly denies two possible revisions to Uncertainty.

⁷ "Actualism" is used as a name for various philosophical views, including even a view in ethics. I'm not using it in any of these already existing ways. I'm using it as a new, stipulative way.

Barbara is a police officer who has bad marksmanship and has promised her superior that she'll practice today. She is at a shooting range, deciding whether to shoot at the target in front of her. She is 90% sure that this is a normal shooting range environment, in which no one would be in any danger if she shoots. But she has a 10% credence that there is a cleaning person moving around behind the targets, who would very likely be shot if Barbara were to shoot at the target in front of her. If the shooting range were empty, Barbara would be morally required to practice. If the shooting range had a person behind the targets, it would be morally wrong to shoot. Barbara has a 90% credence that she's in a scenario in which it is morally required that she shoot (to practice, and to keep her promise) and she has 10% credence that she's in a scenario in which it is morally required to refrain from shooting. (These are the credences she has regarding the *objective* moral truth about her situation.)

In this case, it is obvious that Barbara should refrain from shooting, for the reasons given in the initial line of argument we considered. Barbara should not take the moral risk of killing someone just to be sure to practice and keep her promise. But what's important, according to Actualism, is not that Barbara believes (as any ordinary person does) that shooting an innocent cleaning woman in this context would be morally wrong, but simply that she knows that the cleaning woman might be there. What makes it the case, according to Actualism, that Barbara should refrain from shooting is simply that she'd be *taking a risk of killing someone*, not that she'd be *doing something that she knows is taking a risk of doing something objectively morally wrong* – though that is also true in this case.

Now consider this story which is very much like B's story:

Bob's daughter Sue has been asking him to teach her to drive and he has finally promised to do so. He already taught her twin brother to drive. Bob has a 90% credence that he is morally required to teach Sue to drive. But Bob has recently been listening to some conservative speeches about the morally appropriate place of women in society. According to the conservative speaker he's been listening to, women should not drive and no one should teach a woman to drive; in fact that is a grave moral wrong. The conservative speaker does not challenge any non-moral facts Bob already believes; the challenge is simply to Bob's normative beliefs. Bob thinks the conservative speaker is probably wrong; he's 90% sure of that. But Bob finds the conservative picture being offered somewhat compelling, so that he is 10% sure it is the correct picture. (He's quite sure that a compromise position is false; so he's simply torn between the liberal ideals he grew up with and the more conservative picture he's learning about.) Bob has a 90% credence that failing to teach Sue would be wrong, but not very seriously wrong; he has a 10% credence that teaching her to drive would be deeply morally wrong. (These are the credences he has regarding the objective moral truth about his situation.)

Uncertainty holds that, if the conservative picture holds that teaching a woman to drive is wrong enough, Bob should not teach Sue to drive. He would be taking a moral risk that he should not take.

But that is false. Bob should teach Sue to drive. Breaking his promise would be treating her badly, just because she is female, in a way that a father should not mistreat his daughter. (Let's not be distracted by the fact that it wouldn't be a *minor* wrong to Sue, so it's not an elaboration of the initial case of B (but rather is a similar story). That's true. But just so long as the conservative picture holds that teaching a woman to drive is wrong enough, then Uncertaintyism will imply that Bob should not teach Sue to drive.)

Bob knows his situation. He knows that he would be treating his daughter differently from his son, simply because she is a woman. He knows that women are not inherently less good drivers, or less intelligent than men. (Remember, the conservative picture does not challenge the non-moral beliefs that Bob already had as a liberal.)

If Bob refuses to teach Sue to drive, breaking his promise, her resenting him for it would be reasonable and appropriate. Even if she knows that he refrains because he does not want to take this moral risk, her resentment is still appropriate. Bob would be blameworthy for refusing to teach Sue to drive.

These are the claims that Actualism makes about the case. These are the claims I make about the case. Whether they seem plausible partly depends on what one thinks about whether being caught in the grip of a false moral view is exculpatory. If it is, then being "partly caught" in the grip of a false view, as Bob is, would also seem to be exculpatory. If it is not, then being "partly caught" is not exculpatory. So, here I note again that the dispute between Actualism and Uncertaintyism hangs at least in part on a dispute over whether being caught in the grip of a false moral view is exculpatory. I'll say more about that in section III.

Here is another way of filling out B's case:

Bernard is a slaveholder in the early 1800s in the United States. Like every slaveholder in his community, once a week he beats his slaves if they have disobeyed him or failed to do their work as well as he believes they can. He believes that he is morally required to continue this practice, because (as he has been taught) it is a crucial ingredient in sustaining his community's way of life, which he believes he is morally required to do. But lately Bernard has heard some compelling abolitionist speeches. He thinks abolitionism is misguided; but he's not sure. Bernard is 90% sure that he is morally required to beat his slaves if they disobey, though it is not seriously morally wrong to refrain. He is 10% sure that it is deeply morally wrong to beat his slaves. (These are his credences in the objective moral truth about his situation.)

What should Bernard do? As Uncertaintyism says, Bernard should refrain from beating his slaves. That's true. I believe that cases like Bernard's provide crucial and central motivation for Uncertaintyism. It does seem that Bernard should refrain from beating his slaves. Actualism holds that Bernard should refrain from beating his slaves, but not because of his moral credences. He should refrain because he knows it is *beating his slaves*, he knows it is beating people who are kept captive and who work for him without compensation and without freedom. One should not beat a person in this situation. (One should almost never beat another person, of course. Nor should one keep slaves.) Bernard should refrain because he knows what he is doing, Actualism holds, and because what he is doing is actually wrong.

Uncertainty may be motivated in part by the thought that Bernard is blameworthy if he beats his slaves. But Actualism agrees that he is blameworthy and that he should not beat his slaves; the disagreement is over what explains these truths.

Here is another way of filling out B's case:

B is presented with the option of pressing a red button, without knowing what the button does. Then an advisor speaks. The advisor says, "I'm 90% sure that pushing the button is morally required, though it wouldn't be a grave moral wrong to fail to push it. I'm 10% sure that pushing the button is gravely morally wrong." That is all the advisor says.

What should B do in this version of the case? Let's consider two variants.

Betsy is presented with the option of pressing a red button, without knowing what the button does. Then an advisor speaks. The advisor says, "I'm 90% sure that pushing the button is morally required, though it wouldn't be a grave moral wrong to fail to push it. I'm 10% sure that pushing the button is gravely morally wrong." That is all the advisor says. *Betsy knows this advisor well. Betsy in fact has sound moral views, and knows that the advisor does too.*

Betsy should be cautious and refrain from pressing the button, for just the reasons offered in the initial line of thought we considered. Pushing the button is too risky. In particular, the advisor's speech gives Betsy reason to believe that the advisor is uncertain *what pushing the button does*, but that the advisor has a 10% credence that pushing the button does one of a number of things that the advisor and Betsy both think are gravely wrong—and these things are wrong, because they have sound moral views—while the advisor likely has 90% credence that pushing the button does one of a number of things that the advisor and Betsy both think are morally required, but not seriously morally wrong to do—and those things really are morally required but not seriously wrong to do. So, Betsy should have 10% credence that pushing the button is: killing an innocent person, seriously harming an innocent person, or etc. And she should have 90% credence that pushing the button is: breaking a minor promise, dividing a benefit unfairly, etc. When a person's credences are distributed that way, her case is very much like Barbara's case, and she should be morally cautious.

What's crucial to our understanding of Betsy's case is that her moral uncertainty gives rise to non-moral uncertainty. Because she believes the advisor, she comes to be in a state of non-moral uncertainty that makes it reasonable to be cautious and refrain from pushing the button.

But consider this variant, involving an agent who holds a true moral view about helping hurricane victims but a false moral view about gay marriage:

It is the final moments of a U.S. state's legislative session. Unless a bill is delivered to the statehouse by midnight, the bill cannot be signed by the governor. In the statehouse, Ben is presented with the option of pressing a red button. He knows that the button delivers a piece of legislation to the governor, but he does

not know what the legislation is. He knows that two pieces of legislation were before the legislature, one providing aid to towns recently hit by a hurricane and one allowing gay marriage in the state. The governor has pledged to sign each bill if it is delivered to him in time. An advisor says, “I’m 90% sure that pushing the button is morally required, though it wouldn’t be a grave moral wrong to fail to push it. I’m 10% sure that pushing the button is gravely morally wrong.” That is all the advisor says. Ben knows that he and the advisor agree about the morality of both pieces of legislation: they agree that it is good to aid the hurricane victims; and they agree that it is seriously morally wrong for the state to allow gay marriage or for anyone to do anything to aid the state in allowing gay marriage. Ben is an ordinary person with no duties in the legislation, but due to an odd computer set-up, he alone is able to push this button in time to get the legislation, whatever it is, to the governor. If the governor does not get the legislation today, the voting will be re-done in one month by a newly constituted group elected in a recent election, who are expected to decide differently on both bills. In fact, both bills are good bills that should be enacted, though the hurricane bill is not terribly important because federal aid will also be provided.

In this case, Ben has a 90% credence that pushing the button is morally required though failing to push would not be seriously wrong, and a 10% credence that pushing is gravely morally wrong. But he also has credences regarding the non-moral upshot of each choice: he is 90% confident that pushing would provide the hurricane relief, and 90% confident that failing to push would withhold that relief; he is 10% confident that pushing the button would result in the legalization of gay marriage and 10% confident that failing to push would prevent that legalization. Uncertainty holds that Ben should refrain from pushing the button because it would be taking too big a moral risk. Actualism holds that Ben should push the button, because it is morally required to push the button *in either scenario he is considering*.^{8,9}

⁸ I assume that Ben should push the button because of the substance of the bill, not because the legislature’s decision, whatever it is, should be passed on to the governor. When there is a timing deadline, individual citizens are not obligated to help the legislature get their chosen legislation to the governor in time, if they disapprove of that legislation. Individuals are permitted to allow a technicality to defeat the bill.

⁹ The following could seem like a counterexample to Actualism.

The oracle tells Arthur that pushing the green button is morally required. Arthur has lots of bad moral views, so Arthur forms the belief that pushing the button is X1, X2, X3, . . . or Xn, where these are all in fact morally wrong to do. Actualist holds that Arthur should be guided by his *non-moral* views, so Arthur *should refrain* from pushing the button according to Actualism. But, the objector says, pushing the button does a morally good thing, and the Oracle told Arthur that. Surely he should push.

Actualism implies that Arthur should not push the button. In fact, if Actualism is true, it seems that Arthur is blameworthy for pushing the button. This might seem bizarre, though in section III I offer a view on which false moral belief is typically

III. Do False Moral Views Exculpate?

In this section, I will defend premise 2 of my main argument.

It is controversial whether being caught in the grip of a false moral view exculpates. In this section, I will lay out the relevant issues and briefly argue for and defend the view that being caught in the grip of a false moral view does not exculpate. (When I say that someone is “caught in the grip of a false moral view”, I mean that she is *certain* the false moral view is correct.)

Some people have argued for this claim:

Moral Ignorance Exculpates: A person who acts wrongly is blameworthy for so acting only if she *believed* she was acting wrongly.

The claim that Moral Ignorance Exculpates is clearly false. Even if we take seriously that a person’s moral beliefs and credences can sometimes render her blameless for moral wrongdoing, it is implausible that mere ignorance—that is, failure to *believe* the relevant moral truth—is sufficient for blamelessness. Cases such as those that motivate Uncertainty show Moral Ignorance Exculpates to be false. The mere fact that someone does not believe her action is wrong does not rule out that she has some non-trivial credence that it is wrong, and so her moral beliefs and credences may not vindicate her action at all. One cannot defend one’s wrongful action later by saying, “I didn’t know it was wrong, though I had a 30% credence that it was wrong.” That is not a good defense, if one knew that one’s other option was morally permissible.^{10,11}

The more difficult question is whether *being caught in the grip of a false moral view* renders one blameless for wrongful actions. Here is the claim to consider:

blameworthy; that Arthur is blameworthy for pushing the button may seem less strange if he is also blameworthy for his moral beliefs.

If Arthur does not push the button, then Actualism implies he’s not blameworthy; whereas he may seem to be blameworthy. (The oracle told him it was morally required to push! Yet he didn’t push!) But at least the Actualist can say, in that case, that Arthur is blameworthy for his false beliefs, though not for failing to push.

¹⁰ The point that moral ignorance does not exculpate, because one might be *ignorant* while being *uncertain* as to whether one’s action is morally wrong, is made by Guerrero 2007. Guerrero advocates the principle that if one knows an action might be a morally wrong killing, while the alternative is definitely morally permissible, then one should refrain from acting.

¹¹ Note that, on my view, this is definitely not a good defense because moral belief and credence is not exculpatory (except when it warrants non-moral credences that would be exculpatory). But my point is that *even if one thinks that moral belief and credence can be exculpatory*, one should not think that mere ignorance can be exculpatory.

False Moral Views Exculpate: If a person behaves in a morally wrong way, while certain of a false moral view according to which that behavior is morally required, then she is not blameworthy for her behavior.¹²

My Main Argument in this paper involves claiming that Uncertaintist is committed to the claim False Moral Views Exculpate; and I claim that False Moral Views Exculpate is false.¹³

How might we figure out whether false moral views exculpate? Some writers have seemed to suggest that we can learn that false moral views exculpate from the fact that false *non-moral* views exculpate.¹⁴ Consider Anne, who poisons her husband, certain that she is curing him. Anne is not blameworthy. What is the explanation of Anne's blamelessness? It does seem that Anne's *ignorance* renders her blameless. But how can we characterize the ignorance that renders her blameless? There are two competing principles that would explain Anne's blamelessness. Here is the first:

A person is blameworthy for behaving in a particular way only if her behavior has certain features that make it morally wrong and she believes that her behavior has those features.

This principle implies that Anne is not blameworthy. Anne's behavior was a poisoning of her husband, that makes it morally wrong, but she did not believe her behavior had that feature. Anne believed her behavior was an attempt to save her husband's life, but that feature does not make her behavior morally wrong. There is no feature of her behavior that both makes it morally wrong and is such that she was aware of it.

The second principle that would explain Anne's blamelessness is this one:

A person is blameworthy for behaving in a particular way only if her behavior has certain features that make it morally wrong and she believes *both* that her behavior has those features *and that they make it morally wrong*.

The second principle implies that false moral views exculpate. So, if cases of exculpatory non-moral false views like Anne's are explained by this second principle, then these cases can directly support the claim that false moral views exculpate.

But cases like Anne's are not explained by the second principle. These cases give us no reason to believe the second principle. Anne's blamelessness follows simply from her failure to know that what she is doing *is a poisoning*. (Importantly, she does not even know that it might be a poisoning.) Anne's blamelessness is fully explained by the first

¹² In my 2011, I discuss an argument that moral ignorance exculpates; I argue that the real issue is not whether moral ignorance exculpates but rather whether false belief exculpates, and I argue that it does not. Rosen 2003, Rosen 2004, Wolf 1982, and Zimmerman 1987 offer arguments that moral ignorance or false moral belief exculpates.

¹³ A stronger claim would replace "morally required" with "morally permissible" in the statement of False Moral Views Exculpate. That stronger claim is advocated by Rosen 2004 and others; I argue against it in my 2011.

¹⁴ Rosen 2004 seems to suggest this, though this may not be intended.

principle, which is simply a weaker claim than the second principle. We do not need to make the stronger claim in the second principle to explain Anne's blamelessness.

Thus, 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.

To know whether false moral views exculpate, we must confront cases involving false moral views head on, and ask whether they involve blameworthiness. When we do, we see that many cases of wrongful behavior by agents caught in the grip of false moral views are *paradigm* cases of blameworthiness.¹⁵ Consider these two cases:

Max works for a 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 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.

I claim that Max and Gail are paradigm cases of agents blameworthy for their wrongful actions. They know that they are killing innocent people; this is sufficient for the agents to be blameworthy. Their knowledge that they are killing innocent people is the ground for their blameworthiness.

A proponent of the claim that false moral views exculpate might respond as follows:

While these two agents are not blameworthy for *killing* the innocent people, they are blameworthy for *causing these killings to occur* because they must be blameworthy for having come to have these beliefs. Agents are blameworthy for their moral views if and only if they have violated *procedural moral obligations regarding learning the moral truth* (and these violations were themselves blameworthy).¹⁶ These include the obligation to think a reasonable amount about morality in general, the obligation to take seriously moral arguments one hears, etc., but do not include any substantive obligations to believe certain things. These two agents cannot possibly have fulfilled their procedural moral obligations regarding learning the moral truth.

I have two responses. First, it underplays the extent and gravity of these agents' blameworthiness to treat them as *merely* blameworthy for having come to have certain beliefs, and having thereby *caused* some deaths, while not seeing them as blameworthy

¹⁵ See my 2011.

¹⁶ Rosen 2004 offers this view of blameworthiness for moral beliefs.

for their *killings*, which occurred while they knew full well that they were killing innocent people (in one case, for the financial gain of the Mafia family; in the other case, out of revenge at another person's actions). Second, it is a grave mistake to think that people cannot become convinced of deeply false moral views, such as these, 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 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; and, I claim, they would be blameworthy for their wrongful behavior.

The question of whether false moral belief exculpates turns out to be intimately connected to the question: is ethics hard? Ethics is indeed quite hard, and this is why we see so many cases of false moral belief even among those who have fulfilled their procedural moral obligations regarding learning the moral truth. Sure, 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 think hard about morality, take the arguments they hear seriously, and still *get it wrong*.¹⁷ A failure to appreciate how hard ethics is can make the claim that false moral belief exculpates seem less radical than it is.¹⁸ But that claim is very radical.¹⁹

According to the view that false moral views do not exculpate, people who act wrongfully are blameworthy not in virtue of what their moral beliefs and credences were, but in virtue of what their *non-moral* beliefs and credences were, and how these influenced their choices. Someone who knows she is killing an innocent person, and does so anyway, does not care adequately *to avoid killing the innocent*.²⁰ A view of blameworthiness that can undergird the claim that false moral views do not exculpate is this:

A person is blameworthy for her wrongful behavior just in case it resulted from her failure to care *de re* about what is morally important—that is, from her failure to care adequately about the non-moral features of the world that in fact matter morally.²¹

¹⁷ I make this point in my 2011.

¹⁸ FitzPatrick 2008, in discussing whether false moral belief exculpates, assumes that ethics is only hard in special cases.

¹⁹ My paper "Ethics is Hard. What Follows?" (manuscript) expands on some of the arguments in this section.

²⁰ Unless this is a case in which it is morally permissible to kill an innocent person.

²¹ Here I am adopting a view along the lines of Arpaly 2003.

A person cares *de dicto* about morality if she wants *to be moral*. A person cares *de re* about morality if she wants e.g. *to keep her promises, to help the needy*, etc., and if keeping one's promises and helping the needy are in fact morally important.^{22,23}

A proponent of the view that false moral views exculpate holds the following:

(*) Someone who behaves morally wrongly while caught in the grip of a false moral view (according to which what she is doing is morally required) is blameworthy for this behavior *only if* and *only because* she is blameworthy for her false moral view.

I deny claim (*) but I grant that there is something intuitively compelling about it. I grant that the following weaker claim is true:

(**) Someone who behaves morally wrongly while caught in the grip of a false moral view (according to which what she is doing is morally required) is blameworthy for this behavior *only if* she is blameworthy for her false moral view.

I grant that there is something odd about holding someone blameworthy for her morally wrong behavior while acknowledging that she is blameless for a false moral view according to which that behavior is morally required. But I deny the “only because” part of (*). I endorse only the “only if” part of the claim. I hold that people who do morally wrong things while caught in the grip of false moral views are blameworthy for their actions and are *also* blameworthy for their beliefs. But they are not blameworthy for their actions *merely because* they are blameworthy for their beliefs. Rather, they are blameworthy for both for related reasons – because both their actions and their beliefs involve their failing to care adequately about what matters morally:

Believing that one's wrong action is morally required involves caring inadequately about the features of one's action that make it morally wrong, because believing that an action is morally wrong on the basis of the features that make it wrong is a way of caring about those features.²⁴

²² A consequentialist would deny that keeping one's promises is in itself morally important.

²³ One might object to the view I have proposed that someone might care inadequately about things that matter morally because of a psychological problem such as depression. I endorse the development of this view in Arpaly 2003 which holds that when a person fails to act as she should because of depression, this is not a case in which she fails to care adequately about what matters morally. We can distinguish the psychological mechanism of depression from the psychological mechanism of inadequate caring; depression may interfere with caring rather than eliminating it. So the view I endorse can accommodate the thought that psychological problems such as depression mitigate blame.

²⁴ As I understand it, this part of my view goes beyond anything in Arpaly 2003.

False moral belief is blameworthy. Actions done on the basis of false moral belief is often blameworthy. On my view, they are blameworthy for similar reasons.

In this section, I have explained and defended premise 2 of my main argument, my claim that false moral views do not exculpate. I've argued that consideration of the way that false non-moral views exculpate in no way supports the claim that false moral views exculpate. I claim that if we look directly at cases of wrongdoing due to false moral views, we see that these are paradigm cases of blameworthy behavior. And I've offered a view of blameworthiness on which false moral views do not exculpate, which accommodates the thought that if the behavior is blameworthy, then the false moral belief is blameworthy as well.²⁵

So far in this paper, I have offered my main argument that Uncertainty is false: If Uncertainty is true, then false moral views exculpate. But false moral views do not exculpate. So, Uncertainty is false. And I have articulated and explained an alternative to Uncertainty: Actualism. In the remaining sections of the paper, I defend this argument in the face of some objections. In section IV, I discuss the objection that Uncertainty and Actualism do not really disagree. In that section, I also offer a further consideration against Uncertainty, based on an analogy with epistemology. In section V, I discuss an objection that the analogy with epistemology can be used to support Uncertainty. In section VI, I discuss an objection to the first premise of my main argument, which holds that if Uncertainty is true, then false moral views exculpate. In section VII, I discuss another objection holding that Uncertainty and Actualism do not really disagree; this objection holds that Uncertainty is trying to answer a first-person question while Actualism is answering a third-person question. In section VIII, I discuss whether either of two revisions of Uncertainty can resist my arguments; these revisions hold that a person should be guided either by her epistemically justified moral beliefs or by those moral beliefs that she would be epistemically justified in holding.

IV. Is there really disagreement between the Uncertainty and the Actualist?

Let's focus on Bob, who must decide whether to teach his daughter Sue to drive. I have said that the Uncertainty and the Actualist disagree about Bob. The Uncertainty says:

- (1) Bob should not teach Sue to drive.

because teaching Sue to drive would be taking a serious moral risk.

The Actualist says:

- (2) Bob should teach Sue to drive.

because refusing to teach Sue to drive would be limiting her options because she is a woman.

The dialectic here is complicated. The Uncertainty will grant that (2) is true, or has a true reading. To see this, remember the case of Anne, who poisons her husband thinking she is curing him. The following claim can be truly made of Anne:

- (3) Anne should not give the drink to her husband

²⁵ I respond to some objections to this view in my 2011 and in my "Ethics is Hard. What Follows?" (manuscript).

because it is poison. We can also truly say:

(4) Anne should give the drink to her husband
because one should do what one believes will save one's husband's life.

Claim (3) is true of Anne, *ignoring her beliefs and credences about her situation*, we might say; it is made true by what her situation really is. Claim (4) is true of Anne, *given her beliefs and credences about her situation*, we might say.

Similarly, the Uncertaintist can claim that there are three ways for claims about Bob to be true. First:

(2) Bob should teach Sue to drive.
This is true objectively. Ignoring Bob's beliefs and evidence about his situation, given what his situation really is, he should teach his daughter to drive.

But also:

(2) Bob should teach Sue to drive.
This is true for one kind of subjectivity. Ignoring Bob's *moral beliefs and credences*, just focusing on his non-moral beliefs and evidence, Bob should teach Sue to drive. (That is, the same claim is true on two different readings in this case. If Bob had false non-moral beliefs, then it might be different claims that were true on the two different readings.)

Finally,

(1) Bob should not teach Sue to drive.
According to the Uncertaintist, this claim is true *taking into account all of Bob's beliefs and credences about his situation (thus, including his moral beliefs and credences)*. This claim is true for a second kind of subjectivity.^{26,27}

²⁶ To see an example in which three different moral claims might be true, one in each of these three ways, according to the Uncertaintist, consider this case:

John's sixteen-year-old daughter wants to take the morning-after pill. John has two pills in front of him, A and B. John believes that A is an aspirin and B is the morning-after pill. In fact, it is the reverse: A is the morning-after pill and B is an aspirin. John is sure that taking the morning-after pill is wrong and that it is wrong to give it to one's daughter. Consider these three claims:

- (1) John should give his daughter pill A.
- (2) John should give his daughter pill B.
- (3) John should give his daughter neither pill.

(1) is true as an objective moral claim. John should give his daughter what is actually the morning-after pill: pill A. (2) is true as a subjective moral claim relative to John's non-moral beliefs but ignoring his moral beliefs: a person should give his daughter what he takes to be the morning-after pill if she wants to take the morning-after pill. (3) is true relative to John's entire mental state, according to the Uncertaintist.

²⁷ Now that we have distinguished these ways in which moral claims may be true, we can clarify the best interpretation of certain general claims that Uncertaintism makes, such as: "Someone who is 90% sure that ϕ ing is minorly morally wrong, but 10% sure that failing to ϕ is deeply morally wrong, should refrain from ϕ ing." Here the agent's two mentioned credences are best understood as credences in *subjective moral claims that are relative to the agent's non-moral credences, ignoring the agent's moral credences*; the final "should" claim, which the Uncertaintist makes, is best understood as a subjective claim relative to the agent's entire mental state.

The Uncertaintist may then claim that the Actualist and the Uncertaintist do not disagree. Rather, both agree that claim (2) is true. The Uncertaintist simply raises a third question that the Actualist does not appear to be interested in: suppose we do not ignore a person's moral beliefs and credences, but take into account her *whole mental state*. What should she do, on the basis of this whole mental state?

I believe that some Uncertaintists would see the dialectic this way.²⁸ But they are wrong. Actualism is a proposed answer to the very same question the Uncertaintists are interested in, namely: how should a person act, taking into account her beliefs and credences (including her moral beliefs and credences), given that one sometimes must act while experiencing moral uncertainty?

I will now offer an analogy with epistemology to show that the following general principle is false (letting X range over either believing in a certain way or behaving in a certain way):

If an agent is convinced of principles that imply she should X, and she sees that these principles imply she should X, then all things considered, in light of her whole mental state, she should X.

But be forewarned – the moral and the epistemological are analogous in only certain ways, as I will explain.

Uncertaintism is a view about how a person should behave; consider the analogue of the Uncertaintist regarding how a person should believe. The Epistemic Uncertaintist holds that how a person should believe is determined by what she *believes* about how she should believe. If a person is sure that she should form her beliefs in accord with a certain rule, then indeed she should do so. If a person is unsure between two different ways of forming beliefs, then her beliefs should be formed in a compromise way between those two methods, etc. To see why this view fails, consider this case:

Mary believes that a particular way of reasoning is a good way of reasoning. In fact it's not. It involves coming to beliefs on the basis of claims that don't really support those beliefs. This way of reasoning is not a way of becoming justified in the newly formed beliefs. But Mary does believe it is a good way of reasoning. Now Mary considers following a particular line of reasoning. She correctly sees that it is an instance of that way of reasoning, the one she believes to be a good way of reasoning. She correctly sees that this line of reasoning would lead her to believe P. In fact Mary's evidence and the reasoning available to her (ignoring her false belief that it is a good way of reasoning) do not support a belief in P. Should Mary believe P?

²⁸ Seppieli 2008 distinguishes the “Non-Normative Belief-Relative ‘Should’” from the “most belief-relative ‘should’ of all—relative to the agent's beliefs about both the normative and the non-normative” which seem to correspond to the two kinds of subjectivity I mention. He clarifies that his claims are in terms of the latter (as I say the Uncertaintist's claims are).

Mary should not believe P. The mere fact that she believes a bad way of reasoning to be a good way of reasoning does not make that way of reasoning *a good way for her to reason* or a way that she now *should* reason. We can say this while *taking into account* that she believes that it is a good way of reasoning. The true epistemological view, whatever it is, will hold that Mary should not believe P. The claim that Mary should not believe P is an analogue of Actualism's claim that people caught in the grip of false moral views nevertheless should not do the particular deeply morally wrong things they believe to be morally required; and it is an analogue of Actualism's claim that Bob should teach Sue to drive, even taking into account his whole epistemic state.

Mary's case shows that the following general claim is false, where X ranges over either behaving in a certain way or believing in a certain way:

If an agent is convinced of principles that imply she should X, and she sees that these principles imply she should X, then all things considered, in light of her whole mental state, she should X.

While this principle would indeed imply that Uncertaintyism is right about cases involving certainty, consideration of Mary's case shows that this principle is false. Seeing how it is false in the epistemological case helps to see how it may be false in the moral case, I claim.

Mary's case also provides a counterexample the following more general claim (again letting X range over either behaving in a certain way or believing in a certain way):

An agent's beliefs about how she should X alone determine how she should X.

This general claim would imply that Uncertaintyism (or something like it) is true, and that Actualism is false; but this general claim is false.

For all I have said so far in this section, it might seem that we could state the disagreement between Actualism and Uncertaintyism as follows. Uncertaintyists believe that claims like (1) and (2) both have true readings. Actualists believe, rather, that claims like (1) are false on any reading. But the dialectic is not that simple. The Actualist need not deny that there is *any* true reading of (1), but simply that the Uncertaintyist's intended reading of (1) is false. The Actualist can grant that there is a true reading of (1), while denying both that it is what the Uncertaintyist intends and that it is interesting.²⁹

How can the Actualist see (1) as having a true reading? The true reading occurs in a speech like this:

“If Bob were to decide what to do by being guided by his moral beliefs and credences, what should he do? He should refrain from teaching Sue to drive.”

This isn't a crazy way to talk. But it doesn't tell us anything interesting about Bob, because he shouldn't be guided in that way. This speech is an analogue of the following.

²⁹ What's really crucial is that, while there may be a true reading of (1), it isn't an interesting claim.

“What should Mary believe, if she were to be guided by her beliefs about how she should believe? She should believe P.”

The claim that Mary should believe P may be true in this context. But it’s not an interesting truth. It is no more interesting that the last claim in this speech:

“Nora has a lot of evidence about whether the earth is more than 6,000 years old. She knows that Tom says the earth is only 6,000 years old. She knows that there is a scientific consensus that the earth is more than 6,000 years old. But let’s ignore her information about the scientists. Focus on what she knows about what Tom says. If Nora’s belief is to be guided by her knowledge of what Tom says, what should she believe? She should believe that the earth is not more than 6,000 years old.”

The final claim in this speech is true. But it is not an interesting truth.

Actualism holds that there are contexts in which (1) expresses a truth, but it is as uninteresting as the claim that Nora should believe the earth is not more than 6,000 years old. Most importantly, Actualism holds that Uncertaintists put forward claims like (1) in contexts in which those claims are not true. Uncertaintists claim that a person who faces moral uncertainty, such as Bob, is doing the best he can *in light of his full epistemic situation*, including both his moral and non-moral beliefs and credences, when he chooses not to teach Sue to drive. That claim is false.³⁰

V. Objection: the Analogy with Epistemology can be used to establish Uncertaintism

At this point, someone might object as follows on behalf of Uncertaintism. Consider the following case.

George follows a particular line of reasoning to a conclusion. It leaves him with a high credence in P. In fact, this is good reasoning. But George has recently been told, by someone he respects, that this kind of reasoning is bad reasoning. George was not convinced, but he has a small credence that this person was correct; his credence is 10%.

The objector’s claim: George should be less confident in P than if he did not have this small doubt about this form of reasoning: his credence in P should be *lower* than it should be in the alternative.

³⁰ It’s a bit unclear how to read Sepielli 2008 on this question. On the one hand, he does say that his claims are in terms of the “most belief-relative ‘should’ of all—relative to the agent’s beliefs about both the normative and the non-normative”. That’s how I read him. On the other hand, he comments that he is making claims about “local rationality” rather than “global rationality”, explaining what an agent should do “relative to” her credences in claims of the form “action A is better than action B” etc. I take this to stipulate that an agent might have other *moral* beliefs that Sepielli’s view doesn’t take into account; I do *not* take him to be saying that his claims ignore the agent’s *non-moral* beliefs.

If the objector's claim about George is right, then in the epistemic case, a person's beliefs and credences about how she should believe are relevant to what she should believe. This might seem to support Epistemic Uncertainty.

I will respond to this worry in two ways.

First, I will make a dialectical point. I brought up the question of whether Epistemic Uncertainty is true to show that it is not in general true that a person's beliefs about how she should X alone determine how she should X, when we allow X to range over both behaviors and beliefs. Rather, other aspects of a person's total mental state are relevant to whether she should X. If she believes that her total mental state warrants her in Xing, or she believes a principle that implies that her total mental state warrants her in Xing, it may nevertheless be the case that she should not X and that her Xing would not be at all warranted. Mary's case shows all of this. So the epistemic analogy is initially brought up as a *counterexample* to a general principle that would support Uncertainty.

So, Mary's case shows that one's beliefs about whether one should X do not *alone determine* whether one should X. If the objector's claim about George is correct, then George's case shows that, at least sometimes, one's beliefs about whether one should X are relevant to whether one should X.

The Actualist acknowledges that sometimes one's beliefs about how one should behave are relevant to how one should behave – in particular, in cases in which one's beliefs about how one should behave are themselves the warrant for further beliefs about what one's non-moral situation is.

I think the lesson of Mary's case and George's case is that there is a substantive *epistemological question* "how should this agent believe?" that is answered by consideration of the agent's whole mental state. The answer cannot be read off of the agent's beliefs about how she should believe. Similarly, I claim, there is a substantive *moral question* "how should this agent behave?" that is answered by consideration of the agent's whole mental state. The answer cannot be read off of the agent's beliefs and credences about how she should behave. But epistemology and morality are different, and the way that these substantive questions get answered may be different.³¹

Now for my second response to the objection. It is not obvious to me that the objector's claim about George must be true.³²

³¹ Actualism holds that moral beliefs are relevant to how one should act only if they are relevant to how the world may be (non-morally). Is there any kind of analogous epistemological claim that we could consider? Perhaps this claim: beliefs about how one should reason are relevant to how one should believe only if they are relevant to what may be the case. The problem is that this may not rule out the relevance of any beliefs about how one should believe, unlike the moral claim. Any belief about how one should believe that one is considering relying on is thereby potentially relevant to how the world may be. This shows that the distinction between moral beliefs and non-moral beliefs is importantly different from the distinction between beliefs about how one should believe and other beliefs.

³² The Objector makes his claim; my suggestion is that the Objector's claim may be false; a third option is that there is more than one sense of "epistemically justified" or

Nevertheless, I do find the objector's claim about George plausible and I am not going to deny it. I want to be very clear about this. *An Actualist does not need to deny the claim about George.*

VI. Objection: The first premise of my main argument is false

In section IV, I distinguished several types of subjective moral claims. This provides a ways for the Uncertaintist to object to the first premise of my main argument. My main argument is:

1. Uncertaintism implies that being caught in the grip of a false moral view is exculpatory.
 2. It is not true that being caught in the grip of a false moral view is exculpatory.
- Therefore:
3. Uncertaintism is false.

My argument for premise 1 can be unpacked as follows:

- (i) Uncertaintism claims that agents should act as their moral credences dictate. (This is a subjective "should".)
 - (ii) A person is blameworthy for some behavior only if she should not have behaved in that way. (This is a subjective "should".)
 - (iii) Persons caught in the grip of false moral views, who behave in ways those views morally require, *should* behave as they do, according to Uncertaintism.
- Therefore:
- (iv) Persons who behave morally wrongly while caught in the grip of false moral views (that require their behavior) are not blameworthy for their behavior, if Uncertaintism is true.

An objector might hold that this argument equivocates between two different kinds of subjective claims: subjective claims that are true relative to an agent's moral beliefs *but ignoring her non-moral beliefs and credences*; and subjective claims that are true relative to an agent's entire mental state, including all her beliefs and credences. Claim (ii) uses a subjective "should" that is relative to an agent's entire mental state. But claim (i), the objector holds, uses a subjective "should" that is relative only to the agent's moral beliefs and credences, ignoring her non-moral beliefs and credences. The objector holds that my argument for premise 1 equivocates.

"rational" and that the claim is true on one reading and false on another. Miriam Schoenfield (manuscript) develops a view according to which there are two senses of "rational". One credence in P may be rational for George, in light of his evidence. Another credence for P may be rational for George, in light of his doubt about his reasoning process. Sepielli 2013 similarly offers a view on which there are various kinds of local rationality (rationality relative to some of an agent's evidence and other credences), but there is no such thing as a credence being globally rational, or rational in light of the agent's entire epistemic situation.

My response to this objection is that if this objection is correct, then Uncertainty is an uninteresting claim. It is no more interesting than the claim that Nora should believe the earth is not more than 6,000 years old, discussed above.³³

VII. Objection: the Uncertainty is trying to answer a First-Person Question, while the Actualist is answering a Third-Person Question

Someone might respond to what I have said so far by saying that I have missed what really motivates Uncertainty, what makes it seem plausible and what question it is really answering. To see the appeal of Uncertainty, don't describe cases involving other people, in which *we know* the moral truths in question. Rather, think about a case in which *you* are the agent, and think about a moral question on which *you yourself* experience genuine moral uncertainty. You are trying to decide whether to reveal to your friend that his partner is cheating on him; on the one hand, you owe him loyalty and perhaps honesty; on the other hand, you might ruin his life. Should you tell him? You are trying to decide whether to fail a student who has submitted failing work but will be kicked out of school if she fails. Should you fail her? You are trying to decide whether to put your child into a private school, depriving her of peers from a range of socio-economic classes but ensuring that she'll get a better education. What do you owe her? These situations involve difficult moral questions, and they involve moral risk, when considered from the first person.

The Uncertainty will say: From the first person, we cannot do better than use the moral beliefs and credences we have. It does us no good to say that we should act in light of our *non-moral* beliefs and credences, in combination with the *actual moral truth*, as Actualism says we should. Our best access to the actual moral truth is through our moral beliefs and credences.³⁴

To see that this speech is misguided, simply consider the analogous speech from the Epistemic Uncertainty: It does no good to tell someone to be guided by the *actual* principles of reasoning; all she can do is be guided by her *beliefs and credences* about how to reason.

It's true that the Actualist might be seen to give "advice" that is sometimes hard to follow. But Actualism is not in the business of giving advice; nor is Uncertainty. The question both answer is: what *should* agents in cases of moral uncertainty do?

Similarly, epistemology offers an answer to the question: what should agents believe? Some agents are poorly placed to believe as they should. Some do poorly, and cannot see that they are doing poorly. But they are still doing poorly.

³³ While my reading of Sepielli 2008 is that he does not make what I call the uninteresting version of his claim (but rather makes the interesting claim which I say implies that being caught in the grip of a false moral view is exculpatory), my reading of Sepielli 2013 suggests he would prefer to embrace what I call the uninteresting version of his claim, while disputing that it is uninteresting.

³⁴ Sepielli 2008 writes "we cannot base our actions on the correct normative standards . . . we cannot guide ourselves by the way the world is, but only by our representations of the world" (p 8).

Agents caught in the grip of false moral views, or swayed by some level of credence in false moral views, who thereby act as they should not act, are doing poorly. When they know what actions they are performing (they know the non-moral features of their actions), they are acting as they should not act.

It is true that they are doing well at something else, but something that is not worth doing well at, the Actualist claims. They are doing well at being guided by their moral beliefs and credences. Does this have moral value? Is that what a person who faces moral uncertainty should do?

Whether it has moral value to be guided by one's moral beliefs and credences hinges on whether caring about morality is a virtue. (That is, whether it is a virtue to want *to be moral*.) As may be surprising, Actualism holds that caring about morality is often not a virtue, and being guided by one's concern to be moral is often not a virtue. Caring about morality often *is* a virtue, but only because that care is often found in people with a great deal of moral knowledge, whose care about morality is care for the things that actually matter morally.^{35,36}

VIII. Two Revisions of Uncertainty

In this section, I will consider an objection that offers two revisions to Uncertainty that may seem to avoid some of the worries I have raised and that provide some of what Uncertainty offers. They are:

Revision #1 of Uncertainty: An agent should be guided not by her actual moral beliefs and credences but only by her epistemically justified beliefs and credences.

Revision #2 of Uncertainty: An agent should be guided not by her actual moral beliefs and credences but by the moral beliefs and credences that she would be justified in holding.

The objector claims that one or both of these views does better than Uncertainty at addressing the worries I have raised.

The following questions are relevant:

- (a) Is it possible to be epistemically justified in being certain of a false moral view?
- (b) Is it possible to be epistemically justified in believing a false moral view?
- (c) Is it possible to be epistemically justified in having some credence in a false moral view?

³⁵ Actualism could instead adopt the alternative view that caring about morality is a virtue but is a virtue that is swamped by the vice of behaving subjectively wrongly.

³⁶ Michael Smith has argued that caring about morality, and acting in a certain way *because it is morally required* rather than because of the reasons that make it morally required, is being a moral fetishist, and is thereby objectionable. (See Smith 1994.) This is a different thought than the Actualist's thought that some concern for morality lacks moral value because it is in fact concern for things that do not matter morally.

In other work, I have entertained an answer of “no” to question (b).³⁷ This would imply an answer of “no” to question (a) as well. I assume that those interested in rescuing Uncertainty by offering Revision #1 or #2 believe that the answer to all three questions (a), (b), and (c) is “yes”. (I discuss the combination of either revision with a “no” answer to some or all questions below.) So, let’s assume for now that the answer to all three questions is “yes”; and let’s assume that *testimony* provides a way of being epistemically justified in believing false moral views.

We should note that these revisions of Uncertainty abandon some of what might have motivated Uncertainty. The idea that a person cannot do better than use her own beliefs (mentioned in section VII) is given up as motivation for Uncertainty. One Uncertainty has suggested that we should separate the question of what credences a person should have from the question of how she should act in light of the credences she actually has and has said that Uncertainty answers the latter question only.³⁸ (But he has since changed his mind.³⁹) This is also abandoned by the revisions.

Can my main argument be revised to address these new views? It can.

My Main Argument, Revised to Target Revisions #1 and #2 of Uncertainty:

- 1’. If Revision #1 or Revision #2 of Uncertainty is true, then being epistemically justified in being certain of a false moral view is exculpatory.
 - 2’. It is false that being epistemically justified in being certain of a false moral view is exculpatory.
- Therefore:
- 3’. Neither Revision #1 nor Revision #2 of Uncertainty is true.

I endorse premise 2’ of this revised argument for the same reasons that I endorse premise 2 of the original argument. Consideration of the way that justified non-moral certainty exculpates gives us no reason to think that justified moral certainty exculpates. The very same kinds of cases I discussed, the Mafia family member and the gang member who kill innocent people, provide counterexamples to the claim that justified moral certainty exculpates. If we take *testimony* seriously as a source of justification for false moral belief, then these agents may be epistemically justified while they are nevertheless paradigm cases of blameworthy agents. Finally, we can grant that it would be weird for their actions to be blameworthy while their beliefs are blameless, claiming that their beliefs are indeed blameworthy.

One might object that epistemically justified beliefs cannot be morally blameworthy. But why not?

One (admittedly controversial) example of an epistemically justified but morally blameworthy belief is a sister’s belief in her brother’s guilt of a serious crime given *just enough* evidence to make that belief epistemically justified; arguably, she should give him the benefit of the doubt and withhold belief. We can make such a claim even while

³⁷ In my 2011.

³⁸ Sepielli 2008.

³⁹ Sepielli [year].

acknowledging that *more than enough* evidence to make the belief epistemically justified would also make the belief blameless; the moral duty to give one's sibling the benefit of the doubt just requires being somewhat reluctant to believe badly of him.⁴⁰

One might think that epistemically justified beliefs cannot be morally blameworthy because a person who is epistemically justified *could not have believed differently*, and a person cannot be blameworthy for something if she could not have done otherwise. But it is not in general true that epistemically justified beliefs are such that one could not have believed differently. I might take your word for something, and come to be epistemically justified in believing it, even though I was capable of thinking it through for myself, and had I done so, I would have realized it was false.⁴¹

I claim that this revised argument is effective against Revisions #1 and #2 of Uncertaintyism.

Next we should discuss whether some other parts of my discussion would be less effective aimed at Revision #1 or Revision #2. Though I was not explicit about this, in section 2, my discussion of Bob—who should teach his daughter Sue to drive—provides another argument against Uncertaintyism: the Uncertaintyist holds that Bob should refrain from teaching Sue to drive; but that's false; he should teach her to drive. An advocate of Revision #1 or #2 might say: Bob's 10% credence in a sexist morality is clearly unjustified, so this case is unpersuasive against either revision of Uncertaintyism. (Indeed, an Uncertaintyist did offer me this response.) But Bob's 10% credence might well be justified. If we take seriously that people can be justified in false moral views, then we should take seriously that testimony can be a source of justification—indeed, can *make one* justified in having false moral views. Deeply sexist moralities have not only been held by many past societies but are currently held by many people, all over the world. In the United States, the sexists don't say that women shouldn't drive; but some are still deeply sexist.

I also argued against Uncertaintyism by using an analogy with epistemology. An advocate of Revision #1 or Revision #2 might say the following:

The case of Mary was intuitively compelling because we imagine that she is unjustified in believing that a bad way of reasoning is a good way of reasoning. If we suppose that she is justified in believing that a bad way of reasoning is a good way of reasoning, then it is actually the case that she *should* believe P, which she arrives at via some bad reasoning she justifiably believes is good reasoning. This supports Revision #1 or Revision #2 of Uncertaintyism: one *should* rely on one's *justified* beliefs about whether one should X, where Xing is believing or behaving. Here I am inclined to disagree about Mary. If she is justified in believing this is a good way of reasoning simply because she is told by someone that it's a good way of reasoning, she nevertheless isn't justified in believing P, and she should not believe P. Simply being told that something is a good way of reasoning can't *make it* a good way for you to reason, if it isn't a good way to reason.

Finally, let's turn back to questions (a), (b), and (c).

⁴⁰ Keller 2004 and Stroud 2006 argue that we owe our loved ones the benefit of the doubt, and that we may owe it to them to refrain from holding beliefs that would be epistemically justified. Lackey (manuscript) disagrees.

⁴¹ See my 2011.

- (a) Is it possible to be epistemically justified in being certain of a false moral view?
- (b) Is it possible to be epistemically justified in believing a false moral view?
- (c) Is it possible to be epistemically justified in having some credence in a false moral view?

As I mentioned, I have entertained elsewhere that the answer to (b) may be “no”, which would make the answer to (a) “no” as well. Consider that ordinary life experience gives people a lot of evidence that is relevant to our moral requirements, and consider that listening to bad arguments, even if intuitively compelling, is not in general a way to become justified in believing something. It is not at all clear to me that, given the total of a person’s epistemic situation, including her life experience evidence, it is ever epistemically justified to hold a false moral belief. If the answers to (a) and (b) are “no”, but the answer to (c) is “yes”, then Revisions #1 and #2 of Uncertainty disagree with Actualism only in cases in which people have some justified credence in a false moral claim but do not believe it (or in which people would be justified in a credence that falls short of belief). In this situation, my Main Argument cannot be revised to target Revisions #1 and #2 because there are no cases in which people are epistemically justified in being certain of false moral views, so it is vacuously true that epistemic certainty in a false moral view exculpates. But a similar argument can be made, using cases such as Bob’s. If Uncertainty is true, then Bob is not blameworthy for refusing to teach his daughter to drive; but Bob is blameworthy for refusing to teach his daughter to drive; therefore, Uncertainty is false.⁴²

Finally, if the answers to all of (a), (b), and (c) are “no”, as in my view they may be,⁴³ then Revision #2 of Uncertainty does not disagree with Actualism regarding how agents caught in the grip of moral uncertainty should behave.⁴⁴

In response to Revisions #1 and #2 of Uncertainty, here is a revised statement of the alternative view:

Actualism (Revised): A person’s moral beliefs and moral credences (*whether justified or not*) are usually irrelevant to how she (subjectively) should act (*as are the moral beliefs and moral credences she would be justified in holding*). How a person (subjectively) should act usually depends solely on her non-moral beliefs and credences; her moral beliefs and credences (*or those she should hold*) are relevant only insofar as they provide warrant for beliefs and credences about what her non-moral situation may be.

Conclusion

⁴² Note that it is hard to understand why (a) and (b) would have the answer “no” while (c) has the answer “yes”.

⁴³ Though I’m certainly not *claiming* that “no” is the correct answer to (a), (b), and (c).

⁴⁴ Revision #1 may disagree with Actualism because it would not imply that an agent should be guided by the moral truth, but just by the subset of moral truths that she actually believes. I postpone further discussion of this combination of views (Revision #1 plus answers of “no” to (a), (b), and (c)) to a later date.

I have argued that Uncertainty is false. If Uncertainty is true, then false moral views exculpate. But false moral views do not exculpate. So, Uncertainty is false.

There are two main ways of rejecting my argument. First, one might hold that false moral views do exculpate. If the Uncertainty goes that route, she takes on a significant commitment that Uncertainty has so far not acknowledged (as far as I know). Second, one might hold that Uncertainty is not committed to the view that moral false views exculpate. If the Uncertainty goes that route, it turns out that her claims are not interesting. The Uncertainty is not telling us how an agent should act, in light of her whole mental state (including both moral and non-moral beliefs and credences); rather, she is simply telling us how agents should act, ignoring their non-moral beliefs and credences.⁴⁵

Works Cited:

- Arpaly, Nomy 2003 *Unprincipled Virtue* Oxford University Press.
- Dowell, Janice (forthcoming) "Flexible Contextualism about Deontic Modals: A Puzzle about Information-Sensitivity" *Inquiry*
- FitzPatrick, William 2008 "Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge" *Ethics* 118: 589-613.
- Guerrero, Alexander A. 2007 "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution" *Philosophical Studies* 136: 59-97.
- Harman, Elizabeth 2011 "Does Moral Ignorance Exculpate?" *Ratio*
- Harman, Elizabeth (manuscript) "Ethics is Hard. What Follows?"
- Jackson 1991. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection" *Ethics*
- Keller, Simon 2004 "Friendship and Belief" *Philosophical Papers* 33: 329-351.
- Kolodny, Niko and McaFarlane, John. Manuscript. "Ought: Between Subjective and Objective"
- Lackey, Jennifer (manuscript) "Why There is No Epistemic Partiality in Friendship"
- Lockhart, Ted 2000 *Moral Uncertainty and Its Consequences* Oxford University Press.
- MacFarlane, John *forthcoming* *Assessment Sensitivity: Relative Truth and its Applications*. <http://johnmacfarlane.net/retruth.web.pdf>
- Moller, Dan 2011 "Abortion and Moral Risk" *Philosophy* 86: 425-443.
- Rosen, Gideon 2003 "Culpability and Ignorance" *Proceedings of the Aristotelian Society*
- Rosen, Gideon 2004 "Skepticism about Moral Responsibility" *Philosophical Perspectives*
- Ross, Jacob 2006 "Rejecting Ethical Deflationism" *Ethics* 116: 742-768.
- Schoenfield, Miriam (manuscript) "Expecting Too Much From Rationality: Why We Need Two Notions Instead of One"

⁴⁵ I owe much thanks to Tyler Doggett, Peter Graham, Alex Guerrero, Sarah McGrath, Miriam Schoenfield, Andrew Sepielli, and Brian Weatherson for comments on drafts of this paper.

- Sepielli, Andrew 2008 “What to Do When You Don’t Know What to Do” *Oxford Studies in Metaethicss*
- Sepielli, Andrew 2013 “What to Do When You Don’t Know What to Do When You Don’t Know What to Do . . .” *Nous*
- Sepielli, Andrew (manuscript) “Normative Uncertainty and Intertheoretic Comparisons of Value”
- Smith, Holly (manuscript) Making Morality Work
- Smith, Michael 1994 The Moral Problem Oxford: Blackwell.
- Stroud, Sarah 2006 “Epistemic Partiality in Friendship” *Ethics* 116: 498-524.
- Weatherson, Brian (manuscript) “Running Risks Morally”
- Wolf, Susan 1982 “Moral Saints” *Journal of Philosophy* 419-439.
- Zimmerman, Michael 1997 “Moral Responsibility and Ignorance” *Ethics* 107: 410-426.