

How Moral Uncertainty Can Be Both True and Interesting

ANDREW SEPIELLI

1. Introduction

Ethics is hard. Many of us are justifiably uncertain about everything from self-defence to partiality to distributive justice. I find myself struck by the question of what to do in the face of this uncertainty—that is, uncertainty not about “the facts,” but about the fundamental moral significance of those facts. In the past few years, I and others have been trying to make some progress towards a theory of “what to do when you don’t know what to do.”¹ It’s been our presumption that there is, indeed, some sense in which what’s right depends on the probabilities of various moral views being true. We’ve focused our attention on *how* rightness, or at least one kind of “subjective” rightness, depends on these probabilities.² For example, while some of us have claimed that one subjectively ought to maximize expected moral value, others have argued that, instead, one subjectively ought to act on the moral theory that is most probably correct.³

But very recently, this presumption has been called into question. Critics have contended that there’s no important sense in which rightness depends directly on the probabilities of moral claims—that the only interesting answer to the question of “what to do when you don’t know

¹ See Sepielli (2009, 2010, 2014a, 2014b); also Lockhart (2000); Ross (2006a, 2006b), MacAskill (2014), and Gustafsson and Torpman (2014).

² Of course we also think there’s some sense in which what we ought to do depends on which moral views are actually true! The idea that there are not merely “objective” and “subjective” rightness, but indeed, many different possible kinds of the latter, may not be familiar to everyone; it is bruited in Parfit (2011) and Sepielli (forthcoming).

³ The latter view is defended in Gustafsson and Torpman (2014).

what to do” is the one that writers like me have found singularly unhelpful: “It depends which moral views are the *correct* ones.”⁴

Some of the most incisive criticisms can be found in Elizabeth Harman’s recent paper “On the Irrelevance of Moral Uncertainty.” In this chapter, I’m going to offer some replies to Harman. I’m also going to take some of Harman’s arguments as invitations to say more about the motivations for this project than I have in the past.

2. Exculpation and epistemic probability

I read the first half of Harman’s paper as containing two arguments for the “irrelevance of moral uncertainty.” The first argument appeals to the supposed exculpatory impotence of moral beliefs. Harman posits a tight connection between subjective rightness and blameworthiness: One is exculpated for doing an action if that action is subjectively right. So if subjective rightness depends on my moral beliefs or degrees of belief, then exculpation depends on my moral beliefs or degrees of belief.⁵ But, Harman argues and has argued elsewhere,⁶ exculpation does not depend on my moral attitudes. So subjective rightness must not depend on my moral beliefs or degrees of belief. The view she calls “moral uncertainty,” according to which an important sort of subjective rightness *does* so depend, must be false.⁷

Now, one might well question this link between subjective rightness and exculpation. Why not say, instead, that some kinds of subjective rightness are exculpatory and others—namely, the kinds that depend on moral beliefs and probabilities—are not? Harman’s answer is that if a kind of subjective rightness is not exculpatory, then it is “uninteresting.” She could then argue that it is uninteresting *just because* it is not exculpatory, or else argue that it is uninteresting because of whatever deeper features *make* it non-exculpatory. Harman gives us no reason to suspect that she’d go the first route, and that’s a good thing, too, for there are important theoretical and practical roles that subjective norms might play other than the exculpatory one. I’ll discuss one such role later. Harman

⁴ See especially Weatherson (2014), Harman (2015), and Hedden (2016).

⁵ I’ll sometimes refer to degrees of belief as “credences” or “subjective probabilities.”

⁶ See Harman (2011).

⁷ Harman (2015), p. 57.

does go the second route, arguing that the only non-exculpatory forms of subjective rightness are those that are objectionably insensitive to the force of the agent's *non-moral* beliefs. I will address this argument later.

The second argument is officially presented as simply an illustration of the difference between moral uncertainty and Harman's own view, "actualism," on which subjective rightness depends upon which moral theories are actually correct, along with the agent's non-moral beliefs or degrees of belief. But this illustration is one that casts uncertainty in an unfavourable light, and so it functions as shorthand for the argument that uncertainty has counterintuitive implications.

Harman describes a case in which a father starts to suspect that driving a car is inconsistent with the proper role of women in society, and refuses on these grounds to teach his teenage daughter how to drive. If subjective rightness of an interesting sort depends not only on the agent's non-moral beliefs, but on his credences in moral claims as well, then we may have to say that there is some interesting sense in which the father's refusal to teach his daughter to drive is right. But this, Harman tells us, is a mistake.⁸

It will help me to respond to this argument and the previous one to spell out more clearly the kind of uncertainty I'm concerned to defend. It is a version according to which subjective rightness depends on the *epistemic* probabilities, rather than on the *subjective* probabilities, of objective moral norms. This is the version of uncertainty that plays the practical and theoretical roles that I want a moral uncertainty theory to play.

Let me explain what those roles are. At least on my better days, I aim to do actions just insofar as there are objective reasons to do them—*not because* there are objective reasons to do them; that's not my *motivation*—but just *insofar* as there are. But again, I'm not always sure how the balance of reasons tilts. As a general matter, when I want to do an action with a feature, F, and I'm not sure which actions have F, or to what degree they have F, I will often not be able to pursue my aim by acting under the description "Doing an action with F." Of course I can end up doing an action with F; I can bumble into success. But I can't choose it, in the sense that "Doing an action with F" doesn't show up as an object

⁸ *Ibid.*, p. 59.

of choice for me, as a path-tine in the fork. Rather, when my aim is to do an action insofar as it has F, but I am uncertain about F, then at least from my perspective, the only way to faithfully pursue my aim is by choosing an action under another description: “The *best try* at doing an action with F.”

I’ve argued that, basically, the subjectively right act just *is* the act that counts as the best try at doing what one has objective reason to do.⁹ Indeed, this is how I *define* subjective normativity in terms of objective normativity.¹⁰ For me, then, subjective norms are important because it’s by acting on them that I give my best try at doing what I have objective reason to do. This connection between subjective normativity and trying will matter a great deal later in this chapter, but here’s why it matters now: I’ve also argued that the *kinds* of probabilities that I guide my actions by in giving what I see as the best try are *not* credences/degrees of belief/subjective probabilities. They’re the kinds of probabilities that are mentioned in statements that *express* credences¹¹—what writers like Seth Yalcin, Sarah Moss, and myself have called “epistemic probabilities.”¹² And so the kinds of subjective norms I’m concerned to posit are those that make rightness a function of the epistemic, rather than the subjective, probabilities of objective moral claims.

To get a feel for what epistemic probabilities are, and for their role in the guidance of tries/attempts, consider: It is odd to suppose that someone who believes that P will fundamentally guide her behaviour by a norm that says “If you believe that P, do A.” Rather, she will guide it by a norm like “If P, do A.” In other words, she will employ a norm that includes the expression of the belief (“P”) rather than the report of the belief (“I believe P”). Similarly, then, I will want to say that a normatively conscientious agent with a slight credence in P will not fundamentally guide her behaviour by the norm “If you have a slight credence in P, do A,” but rather by a norm like “If there’s a slight chance

⁹ See Sepielli (2012), where I also deal with the predictable objection from “Jackson” or “mineshaft” cases.

¹⁰ Ibid. Mason (2003 and this volume) also draws a link between subjective rightness and trying.

¹¹ Sepielli (2012).

¹² See Yalcin (2007) and Moss (2013). A difference: Some of these writers seem to have an independent notion of what epistemic probability statements are, and are offering expressivist theories of their semantics. I’m using “epistemic probability statement” to simply mean whatever statements are fundamentally used to express one’s credences.

(read: ‘epistemic probability’) that P, do A”—in other words, by a norm that includes the probabilities mentioned in the credence’s *expression*.

Now, just as the truth conditions for “P” can diverge from those for “S believes that P,” those of “There’s a 0.2 epistemic probability that P” can diverge from those for “S has a 0.2 credence that P.” Meta-ethical expressivists often take pains to make a similar point (and *are* pained by others’ failures to fully grasp it): Holding the view that ethical claims express, say, pro-attitudes, in no way commits one to the view that the truth conditions for ethical claims bear some relation to the truth-conditions for claims *about* those attitudes. To think otherwise is to conflate expressivism and subjectivism.

So I can grant, for my purposes, that there’s no interesting sense in which subjective rightness depends on the agent’s moral credences. I claim only that there’s a sense in which it depends on the epistemic probabilities of moral claims. But this doesn’t give us a counterintuitive result in the “father–daughter” case, for the epistemic probabilities here are independent of the father’s credences. The father’s credence in the claim about the “proper role” of women is substantial, but the epistemic probability of this claim’s being true is zero or thereabouts (I say, expressing my own credences). I think it’s zero because I think morality is in some way about things like pleasure, pain, autonomy, desire, and so on—about beings in their capacities as agents or as patients—and not fundamentally about “proper roles” of anything, certainly not of people, and most certainly not of women as such. Others, of course, may think it’s zero for entirely different reasons, or may be more tempted by the father’s position.

And while I won’t go so far as to offer a worked-out view of praise and blame, it’s plausible that moral credences are exculpatory insofar as they line up with the epistemic probabilities of moral propositions. This would have to be developed in some detail to avoid absurdity. For suppose you accept a highly complicated moral theory that assigns value to each of pleasure, pain, autonomy, etc., but that is inconsistent in a subtle and difficult-to-locate way. Having located the inconsistency, I assign an epistemic probability of zero to the theory. Still, I think that your acceptance of that theory may be exculpatory, while our father’s acceptance of his theory is not. Your failure is one of cleverness; his is one of character. So we might instead say something like: If a theory assigns value to A, B, C, and D, the exculpatory

potential of a belief in the theory is a function not (only) of the theory's epistemic probability, but (also) of the epistemic probabilities of A, B, C, and D each being valuable.

Now, Harman may well want to reject norms that are relative to the epistemic probabilities of moral claims, too. And indeed, her arguments that I've yet to consider arguably apply to such norms just as much as they do to those norms that are relative to the agent's credences in moral claims. My claim here is just that my version of uncertainty, which posits only epistemic-probability-relative norms, is under no pressure to ratify the father's actions as right, or else let him off the hook for his actions, on the grounds that they're appropriate in light of his moral credences. My uncertainty affords no right-making role to the agent's credences.

3. Harman's Scylla and Charybdis

The arguments of Harman that especially concern me, though, are a response to an anticipated objection to her arguments discussed above, and a response *to that response*. The objection she anticipates is something you often hear from uncertaintyists like me—namely, that our critics seem to want to pick a fight where there's none to be picked; when an anti-uncertaintyist argues that rightness can't depend on probabilities of moral propositions, she can't be talking about the particular kind of subjective rightness that the uncertaintyist is positing, for that kind is relative on such probabilities *by definition*. The anti-uncertaintyist like Harman is talking past us, we want to say, their claim to the contrary notwithstanding.

Harman does not deny that there is an interesting sort of rightness that is relative to all of the agent's beliefs, including his moral beliefs. But she claims that the uncertaintyist is substantively wrong about which actions are right in this sense.¹³ Specifically, she argues that even though this sort of rightness is *relative* in a certain sense to a set of attitudes that includes moral beliefs, it doesn't *depend* on moral beliefs; these are idle wheels, normatively speaking. Harman anticipates that her argument may prompt an uncertaintyist to retreat to a kind of rightness that is

¹³ Harman (2015), p. 70.

relative to an agent's moral beliefs only, rather than to all of his beliefs. But this kind of rightness, she claims, is *uninteresting*.¹⁴

So Harman thinks that there are two interpretations of the uncertain-tist position. On one interpretation, uncertainty is false. On the other, it is uninteresting.

4. On the charge that uncertainty may be uninteresting

Let me reply first to Harman's "uninterestingness" claim, which she motivates with this example:

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?¹⁵

Harman allows that there is a sense in which Nora ought to believe that the earth is no more than 6,000 years old. It is the sense of 'ought' that is relative only to the evidence constituted by Tom's assertion. But Harman says that that sense is uninteresting, and that so is the similar sense of 'ought' that is relative to moral beliefs alone.¹⁶ (I mentioned above that I really care more about an epistemic-probability-relative subjective 'ought' than a belief-relative one. But in replying to Harman, I'll switch back-and-forth between them, hopefully in a way that's not too confusing.)

I agree that the sense of 'ought' at issue in Nora's case is uninteresting. It plays no important practical or theoretical role. But I deny that the moral uncertain-tist 'ought' under discussion is similarly useless or uninteresting. There are some subsets of the probabilities relative to which there is an interesting 'ought', and other subsets relative to which there is only an uninteresting 'ought'. In other words, there are some 'oughts' that "carve" probability-sets at "the joints," and others that don't.

¹⁴ Ibid., p. 72.

¹⁵ Ibid.

¹⁶ Ibid.

Let's have a look at some uninteresting, joint-missing 'oughts', starting with the one above. Harman does not tell us whether Nora has much lower-order evidence regarding the earth's age—e.g., geological samples, fossils, and so forth. What is clear is that Nora has lots of higher-order evidence, regarding the evidentiary import of the samples, fossils, and so forth. Specifically, she has: (i) what nearly all the scientists say; (ii) what Tom says; and indeed, (iii) what she herself thinks.¹⁷ There is something peculiar about a kind of epistemic 'ought' that is relative only to a bit of higher-order evidence, rather than to Nora's higher-order evidence generally, or indeed, to all of Nora's evidence.

Here is an analogous case involving practical rather than epistemic reasons: Nora's friend has dared her to get a tattoo of The Great Cthulhu. There are reasons to get the tattoo: it would be funny, she'd feel cool about accepting a dare, and the tattoo would be a good way to start up conversations. But of course there are very strong reasons not to get it. We might say that, relative only to facts that constitute the reasons in favour of getting the tattoo, Nora ought to get the tattoo. But that 'ought'—call it the "blinker objective 'ought'"—is indeed uninteresting.

That was a case involving *objective* practical reasons. But one might think that a case involving *subjective* practical reasons would be more apposite, since of course those are the kinds of reasons that the uncertainty is concerned to posit, and because those reasons are more like the kind of higher-order evidence at issue in Harman's example. It is not hard to cook up a case. Nora recognizes that there is some chance that if she presses a button it will kill a person, and some chance that if she presses it it will save a person's life. We might say that she ought, relative only to the chance of saving someone, to push the button. But again: uninteresting.

What's so uninteresting about the 'oughts' in these three cases? It's not, I want to suggest, that they're relative only to a subset of beliefs or probabilities. It's that they're relative to unimportant subsets. But norms relative to the probabilities of moral propositions are important, since these specify for the agent who is uncertain about objective reasons what would count as the best try at acting in accordance with them. And it's

¹⁷ On higher-order evidence, see, e.g., Christensen (2010).

by acting on norm of this sort that she sees as correct that she would give her best try.¹⁸

A norm according to which rightness did not depend on the probabilities of objective moral claims, but rather on which such claims were actually correct—e.g. Harman’s “actualism”—would not, by anyone’s lights, correctly specify the morally uncertain agent’s best try. It would not do so because it would leave the agent uncertain what to do. But an action cannot count as the best try for me unless I am sure that it is the best try; that’s part of what it is for it to be *my* best try, or best shot. Therefore, a theory that correctly specifies the best try for an agent is one that, necessarily, the agent cannot be consciously uncertain of. The action-description “The best try at doing the action that has F” is one that I can only regard as holding true of an action if that description is in my “choice box”—only if it’s a description I can act under, such that my action counts as norm-guided.

Now I suppose that an objector might deny that a norm that correctly specifies the best try at doing what the objective reasons support is thereby interesting or important. But this would sit uneasily with what I imagine any objector does in the face of conscious moral uncertainty about a significant-enough question. I assume that she would do just what I would do—namely, take into account the chance that *this* moral view is right and the chance that *that* moral view is right, regard there being better and worse ways to take these into account, and strive to take it into account in one of the better ways. In other words, she would strive to try her best to do what she has objective reason to do, and thus treat the notion of the best try as an important one in practice. As Peirce might have said, let us not pretend to doubt in philosophy what we do not doubt in our wills.

The parallel epistemic ‘ought’ is not one Harman uses her “Nora” case to cast as uninteresting. Rather, it is an ‘ought’ that depends for its

¹⁸ In saying this, I’m imagining that the objective normative claims among which I’m uncertain refer to actions using descriptions under which it is possible to act—close to what some might call “basic” action-descriptions; descriptions like “Opening that jar of peanut butter” and “Dialling 911 on my cell phone.” If they employ descriptions like “Maximizing aggregate well-being,” then I will not be able to guide my trying by subjective norms that are relative to their probabilities. Instead, I will need to employ subjective norms that depend on probabilities of objective norms *and* of non-normative facts in a particular way—e.g., the probability of “Doing A maximizes aggregate well-being” and the probability of “One objectively ought to maximize aggregate well-being.”

proper application on the way the agent's first-order evidence stacks up by the lights of *all the agent's higher-order evidence*. Such an 'ought' is interesting, for it attaches to the best try at believing in accordance with the evidence (and at believing the truth), just as the moral-uncertainty 'ought' under discussion attaches to the best try at doing what the objective reasons support. For suppose I believe P, Q, and R, which together constitute a body of lower-order evidence, and then have further beliefs about the higher-order evidence regarding the evidentiary value of P, Q, and R. From the third-personal point of view, we may want to evaluate the thinker's belief-formation in light of all of the evidence, or indeed perhaps just P, Q, and R alone (on the thought that higher-order evidence that misleads with regard to the evidentiary value of lower-order evidence ought to be ignored). But from my point of view as the thinker, it seems as though an 'ought' that depends on the way the evidentiary value of P, Q, and R looks *according to the higher-order evidence* is the one that I'll use—that I'm condemned to use, you might say—to guide my reasoning. I may, of course, decide that the opinions of, e.g., the scientific community, have little-to-no higher-order evidentiary value. Indeed, I may decide this on the basis of my own evaluation of the lower-order evidence. But whatever the pool of higher-order evidence says about the evidentiary value of P, Q, and R, it's an 'ought' relative to this pool that I must use in forming my judgment based on P, Q, and R if I wish to aim at the truth.

5. On the charge that uncertainty may be false

Now let's move to consider the interpretation of uncertainty on which Harman believes it is false. So construed, uncertainty posits reasons and rightness relative to all of an agent's beliefs, or to all the probabilities. Harman accepts that this sort of rightness is interesting, but claims that the uncertainty is substantively mistaken about which actions are right in this sense. Harman is mainly concerned to defend her "actualist" view—that what's right in this very sense depends only on the agent's non-moral beliefs (or the non-moral probabilities), plus whichever moral views are actually correct.

That Harman's view is at least a conceptual possibility helps to highlight a distinction that is sometimes overlooked by uncertainty. When we say that X is right "relative to" a belief or a credence, we could

mean either: (a) the weaker claim that it's not conceptually ruled out that its rightness depends on that belief or credence, or (b) the stronger claim that its rightness indeed does depend on that belief or credence.

Similarly, it's often said that objective rightness is relative to "all of the facts" or "the way the world is generally" in the sense that no fact is conceptually ruled out as relevant. But it would be absurd for, say, a racist to insist, "Well, whether it's objectively right for A and B to marry must depend on A's and B's races, since objective rightness is relative to all the facts." The uncertaintist is making a mistake of just this kind when she reasons: "Whether it's subjectively right relative to the agent's entire belief-set to do A must depend on the agent's credences in moral propositions, since by definition this sort of rightness is relative to *all* the agent's beliefs." If an 'ought' *relative* to all my beliefs *depends* upon all of my beliefs, then that will have to be established by argument, not by conceptual fiat.

Harman bolsters her view through another epistemic analogy:

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?¹⁹

Harman tells us that she should not:

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.²⁰

Similarly, then, the mere fact that someone believes or has credence that P is a practical reason does not make it the case that P is a practical

¹⁹ Harman (2015), p. 71.

²⁰ Ibid.

reason, even a belief-relative reason. To put it perhaps over-simply: *merely thinking it can't make it so.*

My paper, “What to Do When You Don’t Know What to Do When You Don’t Know What to Do . . .” (2014a) is driven in part by a desire to hang on to something like the intuition that animates Harman—namely, that merely thinking so can’t make bad reasoning good—without giving up on uncertainty. Much of what I shall say here draws on that paper.

While I want to resist the idea that “thinking makes it so” when it comes to good theoretical or practical reasoning, I also want to resist the idea that it wouldn’t *somehow* be better to reason in a way that looks good by the lights of one’s views about what constitutes good reasoning, or, better, by one’s views about what follows from what or what supports what.²¹ Following my 2014a paper, here’s what I’d want to say about Harman’s Mary:

- (1) It is not rational to conclude that P on the basis of her starting beliefs, relative to those starting beliefs.
- (2) It is rational to conclude that P on the basis of her starting beliefs, relative to her belief that P follows from or is evidentially supported by the contents of her starting beliefs.
- (3) Relative to her starting beliefs *and* her beliefs regarding whether P follows from or is evidentially supported by their contents, there is nothing it is overall most rational to do. I say it is not “first-order rational” to believe P, but that it is “second-order rational” to believe it, and that first- and second-order rationality are incommensurable.

I am permitted, within this framework, to classify Mary’s way of reasoning as bad. For were it good reasoning, believing P would be rational on both “orders.” I would want to say something similar about the practical case:

- (1) First-order subjective rightness depends on the probabilities of various ways the world might be non-normatively, as well as the reasons

²¹ For further support of this intuition, see Schechter (2013).

- for action, given each of these ways, according to the moral theory that is *actually true*.
- (2) Second-order subjective rightness depends upon the probabilities of moral theories.
 - (3) There's nothing it's overall most subjectively right to do, because the orders of subjective rightness are incommensurable.

Why say that the different orders of rationality and subjective rightness are incommensurable? My inspiration here is Elizabeth Anderson's "pragmatist" argument in *Value in Ethics and Economics* that we oughtn't to consider two sorts of value commensurable unless there is a practical need to commensurate them, to put them together. I don't think the various orders of subjective normativity need to be put together into some overall measure of subjective normativity, because I don't think the agent will ever face what, from her perspective, is a practical conflict among the orders. This is not to say that the orders will always agree on what to do; they often will often disagree: it may be first-order subjectively right to do A, and second-order subjectively right to do B instead. It's just that they'll never seem to come apart from the agent's perspective. We might say, borrowing liberally from Richard Moran, that each order of subjective normativity is "transparent" to the ones below it.²²

After all, if I am *certain* about some moral theory, then I can simply do what it says, given the non-moral probabilities. In other words, I will do what I take to be first-order subjectively right to do. There need be no ascent to the second order of subjective rightness, and if there is for whatever reason, the answer of what to do at that next order will be the same. So no conflict there. But now suppose that I am *uncertain* among moral theories. In that case, I cannot, as I argued earlier, guide myself by first-order subjective rightness, depending as it does on which moral theory is *actually correct*; only by ascending to the second order can I hope to engage in norm-guided action. So there can be no conflict between orders here, either.

I should say that, for the same reasons, there can be no conflict from the agent's point of view between *any* of the orders of subjective normativity and objective normativity—i.e., the kind that doesn't

²² See Moran (2001).

depend on probabilities of any sort.²³ I mention this not because it's at the nub of my debate with Harman, but because it's responsive to another set of critics. If you follow this literature on "subjective" and "objective" 'oughts' and so forth, you may have picked up on a schism that runs through it. To put it very simply, some people are what I call elsewhere "dividers": they think there are all of these different sorts of rightness—objective, subjective of various sorts, maybe "prospective"—and no interesting question of what it's right to do in some neutral, or overall, or simply *other* sense. I'm a divider. Others are "debaters"—they think that there is an interesting question about what's right in some neutral or overarching sense.²⁴

Debaters used to talk in terms of rightness or 'ought' *simpliciter*, but lately it's become popular to focus instead on something they call "deliberative" rightness or the deliberative 'ought', which is the sort of rightness to which we must appeal to resolve "deontic dilemmas" that arise between the deliverances of subjective and objective normativity.²⁵ But if what I said above is correct, no question in terms of "deliberative" rightness arises, because we as agents see no need to arbitrate between subjective and objective perspectives. So I don't see what warrant there is for the view that there is an interesting question of what's right in some kind of *ur-sense*.

So why, then, favour a "multiple-orders" view like mine over Harman's actualism?

First, Harman offers no way of capturing the sense in which actions that are prospectively unintelligible from the agent's point of view are irrational or subjectively mistaken. Someone who has settled on A as in any sense the right thing, but then fails to do A, is in a certain way not functioning well as an agent. The ordinary akratic agent exemplifies this failure, but so, it seems to me, do Harman's "Mary," and those morally uncertain agents who don't do what they regard as right in light of the probabilities of moral theories. Second, Harman's views about subjective rightness relative to all the beliefs/probabilities sit uneasily with her

²³ The obvious objection concerns so-called "mineshaft" or "Jackson" (1991) cases, in which the agent can be certain that the subjectively right action is not objectively right. I address this at length in Sepielli (2014a and forthcoming).

²⁴ I discuss this divider/debater distinction in Sepielli (forthcoming).

²⁵ See Kiesewetter (2011) and Lord (2015).

forementioned views about rightness relative to the moral ones specifically—i.e., the kind of rightness she called “uninteresting.” For if moral beliefs are utterly irrelevant to what it’s right to do relative to all of one’s beliefs, then why should they have any right-making power taken in and of themselves? Shouldn’t Harman say instead that *nothing* is right (or normatively valenced in any way) relative to the agent’s normative beliefs?—that saying one ‘ought’ to do A relative to one’s normative beliefs alone is like saying it’s illegal to do A relative to, say, nothing more than some random Internet commenter’s declaring A-ing unconstitutional?

On the other hand, it seems to me that the main reason Harman gives for preferring her view is that you want to avoid saying that “thinking makes it so” when it comes to the quality of practical or theoretical reasoning. But again, my apparatus allows me to say that, too. I can capture the significance of our beliefs about good reasoning or reasons for action, as well as the significance of good reasoning and reasons for action, and I have an account of why acknowledging the significance of both does not ensnare us in anything like a practical dilemma.

Now, I should note that while Harman accepts actualism, she also grants that it is possible to reject uncertainty, and indeed, reject it on the basis of the “Mary” counterexample, without being an actualist. One can think that moral beliefs/probabilities are relevant to subjective rightness without thinking that they are in any important sense fully determinative of it. Indeed, Harman provides another example that militates in favour of this view, and cites my 2014a paper and a forthcoming paper by Miriam Schoenfield as presenting ways of developing it.²⁶ So it’s wrong, I think, to read Harman’s stance as “actualism or nothing”; rather it’s “anything but uncertainty.”

Let’s take stock of the dispute between me and Harman: I believe, and Harman will at least *countenance*, a position on which rightness relative to all the probabilities is partly but not entirely determined by the probabilities of moral claims. And we know from the subsection above that Harman agrees with me that in one sense of ‘ought’, what one ought to do is, in fact, determined by only the probabilities of normative claims. So what fundamentally separates us? It seems to be the “interestingness” question. She thinks this latter ‘ought’ is unimportant because it’s relative

²⁶ Harman (2015), pp. 73–4.

to only a subset of beliefs/probabilities; I think it is important because of its connection to the agent's best try at doing what's objectively valuable.²⁷

6. Trying, guidance, and regress

Earlier we saw that there is a connection between the agent's best try and the notion of action-guidance. A norm can't correctly specify an agent's best try at doing something unless that norm is one she can use to guide her actions, a norm she can act on. The presumption of the moral uncertainty project is that one cannot guide one's actions by a norm about which one is consciously uncertain. So in order for us to guide our actions, and in order for there to be so much as *be* an action that counts as my best try at doing what my objective reasons support, I must be certain of a norm that specifies which action this "best try" is.

But a regress problem looms. Let us suppose that I am uncertain among some ordinary moral theories, $T_1 \dots T_n$. Unable to guide my action directly by any of these, I ask what to do given the probability distribution over $T_1 \dots T_n$. But I am uncertain as to the answer, assigning some probability to each of $U_1 \dots U_n$. This prompts me to ask what I ought to do given *this* probability distribution; I am again uncertain—this time among theories $V_1 \dots V_n$. We can imagine this process iterating indefinitely, such that I become normatively uncertain, as it were, "all the way up." Then it appears that, at least in this case, norm-guided action is impossible, and there is no way for me to try my best, or give my best attempt, at doing what I have objective reason to do. The possibility of normative uncertainty all the way up makes the uncertainty project look pointless.

²⁷ Indeed, I think the 'ought' relative to all of the probabilities is not that interesting. Following Kolodny (2005) on "global rationality," I don't actually think it plays much of a theoretical role. I don't think it's helpful in specifying the agent's best try at anything or in guiding one's actions, and I haven't yet seen that it plays a role otherwise. What I really wanted to do in the 2014a paper was defend a view about what we ought to do in the sense that Harman thinks is uninteresting but *I* of course think is interesting. I wanted to defend it against the objection that the grounds for accepting any view of this sort would also support a view of what one ought to do relative to all the agent's beliefs or all the probabilities that would run afoul of at least one very powerful intuition of about rationality. My defence was to give a theory of the latter, more global 'ought' on which it would not—a theory that would capture all of these intuitions in what seemed like a principled way.

This conclusion strikes me as overly pessimistic, though, for two reasons.

First, it may be that, as my mind travels up through norms at the various levels—level T, level U, level V, etc., we might call them—I will eventually become certain about which *action* to do even if I never become certain, at any level, about which *theory* is correct. Here is a slightly simplified description of how I, myself, think about moral uncertainty: Suppose in some situation, I am uncertain between two relevant moral theories, T1 and T2; T1 recommends doing A over doing B, and T2 recommends doing B over doing A. I am slightly more confident in T1 than I am in T2. I am also uncertain between two theories, U1 and U2, which concern what to do under uncertainty among T1 and T2. Since some specificity is important to what I want to say about this case, let's say that U1 tells me simply to act on the theory of T1 or T2 that is most probable, and that U2 tells me to engage in what we call "moral hedging"—to act on the less probable theory if the moral stakes according to that theory are greater. So according to U1, I ought simply to act on T1, and so ought to do A. According to U2, let's suppose, I ought to do B, since the "gap" in moral value between A and B is bigger on T2 than it is on T1. Suppose that I am far more confident in U2 than I am in U1 (as, indeed, I really am).

Now suppose, finally, that I am uncertain among two theories at the next level—V1 and V2. V1 tells me to act on the theory of U1 and U2 which is the most probable, while V2 tells me to engage in hedging with respect to *U1 and U2*.

Here, I find it plausible to suppose that I will end up with some convergence around B being my best try, notwithstanding my uncertainty among the theories. Obviously, V1 counsels acting on U2, and as such, doing B. What about V2? Well, since V2 counsels *hedging*, it depends on how the value-assignments of U1 and U2 compare, which will require us to compare subjective moral value intertheoretically. Anyone who is familiar with the literature on moral/normative uncertainty knows that it is a difficult question how intertheoretic value comparisons are possible.²⁸ But let's suppose that we're able to make these sorts of comparisons. More specifically, let's say—and I think there

²⁸ See, e.g., Ross (2006b) and Sepielli (2009 and n.d.).

are good reasons to say this, actually—that the amount of subjective moral value at stake according to U1 is roughly the same as the amount of subjective moral value at stake according to U2. Tossing around some numbers, let's say that the former assigns A a value of 10, and B a value of -10, and that the latter assigns B a value of 10 and A a value of -10. Since U2 is so much more probable, V2 will, like V1, favour acting on U2, and as such, doing B. We have convergence.

Now, to be sure, convergence of this sort will be harder to come by if there are more available actions and more available theories—if, e.g., we carve up “doing A” into “doing A in way 1,” “doing A in way 2,” and so on, and divide “hedging” theories into those that counsel expected value maximization, those that counsel slight risk aversion according to function 1, those that counsel it according to function 2, and so on. But as a matter of psychological reality, we tend not to individuate actions and theories that finely, at least not in the cases of conscious deliberation which are our concern here. I think about moral uncertainty as much as anyone—probably to an unhealthy degree!—and yet I tend to divide the action-space in moral decisions rather coarsely—into, e.g., “eat meat” and “don't eat meat”; I also divide the theory-space coarsely—pretty much just into uncertainty theories that say something like “do what's most probably best” and those that counsel some sort of hedging.

There is another reason not to be too pessimistic about the possibility of convergence as we go up through the levels of normative uncertainty. It may be that we settle at some point on a *partial* ordering of actions in terms of subjective value even if we never settle on a single subjectively *best* thing to do. We might find that at some level of subjective normativity, all theories prefer both A and B over all of C, D, E, and F, even if these theories disagree about how to rank A vs. B and each of C, D, E, and F vs. one another. As I put it in other work, it is a simplification to speak of an action as either “guided” or not. It is better to say that an action is guided *relative to* another action or not. In the case envisaged, one's doing A would not be norm-guided relative to doing B, but would be norm-guided relative to doing any of C, D, E, and F. Doing A rather than doing C, for instance, would be fully intelligible from the agent's point of view.²⁹

²⁹ Thanks to the audience at the Arizona Workshop in Normative Ethics for helpful feedback, and also to Elizabeth Harman, Eden Lin, Sarah Stroud, and an anonymous referee.

References

- Anderson, E. (1995). *Value in Ethics and Economics* (Cambridge, MA: Harvard University Press).
- Christensen, D. (2010). "Higher-Order Evidence", *Philosophy and Phenomenological Research* 81(1): 185–215.
- Gustafsson, J. E. and Torpman, O. (2014). "In Defence of My Favourite Theory", *Pacific Philosophical Quarterly* 95(2): 159–74.
- Harman, E. (2011). "Does Moral Experience Exculpate?", *Ratio* 24(4): 443–68.
- Harman, E. (2015). "The Irrelevance of Moral Uncertainty". In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. X (Oxford: Oxford University Press).
- Hedden, B. (2016). "Does MITE Make Right: On Decision-Making Under Normative Uncertainty". In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. XI (Oxford: Oxford University Press).
- Jackson, F. (1991). "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection", *Ethics* 100(3): 461–82.
- Kiesewetter, B. (2011). "'Ought' and the Perspective of the Agent", *Journal of Ethics and Social Philosophy* 5(3): 1–24.
- Kolodny, N. (2005). "Why Be Rational?", *Mind* 114(455): 509–63.
- Lockhart, T. (2000). *Moral Uncertainty and its Consequences* (Oxford: Oxford University Press).
- Lord, E. (2015). "Acting for the Right Reasons, Abilities, and Obligation". In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. X (Oxford: Oxford University Press).
- MacAskill, W. (2014). "Normative Uncertainty". DPhil Thesis, University of Oxford.
- Mason, E. (2003). "Consequentialism and the 'Ought Implies Can' Principle", *American Philosophical Quarterly* 40(4): 319–31.
- Moran, R. (2001). *Authority and Estrangement* (Cambridge, MA: Harvard University Press).
- Moss, S. (2013). "Epistemology Formalized", *Philosophical Review* 122(1): 1–43.
- Parfit, D. (2011). *On What Matters* (Oxford: Oxford University Press).
- Ross, J. (2006a). "Acceptance and Practical Reason". PhD Dissertation, Rutgers University.
- Ross, J. (2006b). "Rejecting Ethical Deflationism", *Ethics* 116(4): 742–68.
- Schechter, J. (2013). "Rational Self-Doubt and the Failure of Closure", *Philosophical Studies* 163(2): 428–52.
- Sepielli, A. (2009). "What to Do When You Don't Know What to Do". In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. IV (Oxford: Oxford University Press).
- Sepielli, A. (2010). "Along an Imperfectly-Lighted Path: Practical Rationality and Normative Uncertainty". PhD Dissertation, Rutgers University.
- Sepielli, A. (2012). "Subjective Normativity and Action Guidance". In M. Timmons (ed.), *Oxford Studies in Normative Ethics*, vol. II (Oxford: Oxford University Press).
- Sepielli, A. (2014a). "What to Do When You Don't Know What to Do When You Don't Know What to Do . . .", *Noûs* 47(1): 521–44.
- Sepielli, A. (2014b). "Should You Look Before You Leap?", *The Philosophers' Magazine* 66(3): 89–93.
- Sepielli, A. (forthcoming). "Subjective and Objective Reasons". In D. Star (ed.), *The Oxford Handbook of Reasons and Normativity* (Oxford: Oxford University Press).
- Sepielli, A. (n.d.). "Normative Uncertainty and Intertheoretic Comparisons of Value".
- Weatherson, B. (2014). "Running Risks Morally", *Philosophical Studies* 167(1): 141–63.
- Yalcin, S. (2007). "Epistemic Modals", *Mind* 116(464): 983–1026.