

Moral Uncertainty and Fetishistic Motivation

1. Introduction

Sometimes it's not certain¹ which of several mutually exclusive moral views is correct. There's some chance that utilitarianism is true, and some chance that the Doctrine of Double Effect is true. There's some chance that a person's culpable bad conduct itself provides a reason to punish her, and some chance that it does not. There is some chance that we have some reason to promote sadistic pleasure, and some chance that we have no such reason. Like most everyone, I think that what we should do is a matter of which, if any, of these views are correct. But I also think that what we should do depends on the chances of views like these being correct, regardless of which are, in fact, correct. I do not regard these claims as inconsistent, for in making the first claim, I am using "should" in what's often called its "objective" sense; in making the second, I am using "should" in one of its "subjective" senses.²

There is a burgeoning literature on the question of *how*, exactly, these probabilities bear on what one subjectively should do. Some writers argue that one should do the action with the highest expected objective value (EOV); others claim that this is undesirable, or even impossible, on the grounds that intertheoretic comparisons of value are meaningless. (See Hudson 1989, Gracely 1996, Gustafsson and Torpman 2014, Hedden 2015.) This literature was anticipated by several centuries by debates among

¹ The useful locution "It's not certain" (as opposed to "S is not certain") expresses a claim about epistemic probabilities — specifically, the claim that no spoken-of proposition has an epistemic probability of 1. Statements about epistemic probabilities express, rather than report, credences, and as such epistemic probabilities may well diverge from credences. Metasemantic views about epistemic probability statements are a big deal in the philosophy of language (Yalcin 2007, Egan and Weatherson 2011, Moss 2013), but this work is just starting to penetrate the ethical literature on action under uncertainty. In my opinion, someone sympathetic to the view I defend in this paper avoids several pitfalls by focusing on epistemic probabilities of moral propositions, rather than subjective or objective probabilities thereof. I argue for this in my 2012 and 2014a. At any rate, the reader of this paper should understand "probabilities" and "chances" as referring to epistemic probabilities; by implication, the reader should not make the mistake of understanding such talk as referring to the agent's own credences.

² I also think that that what we should do depends on the probabilities of various non-moral propositions being true. In so saying, I'm using another, probably more familiar, "subjective" sense of "should". To be clear: I am not simply defining the various senses of "should" here, but rather making stronger claims about how, as a matter of substantive moral theory, these senses apply.

Catholic moral theologians on the question of how an action's "formal sinfulness" was a function of the probability and gravity of its "material sinfulness". Answers included "rigorism" — basically, the view that an action is a formal sin if there is a chance it is a material sin; "laxism" — the view that an action is *not* a formal sin if there is a chance it is *not* a material sin, and others less extreme. (Jonsen and Toulmin 1988)

Until recently, though, no one has found it necessary to coin a label for the position assumed by all the participants in these debates — namely, that there is at least *some* important sense in which what's right depends *in some way or other* on the probabilities of objective moral norms. I suspect I speak for many other participants when I say I simply found this position obvious. My thought had always been something like this: it is valuable to have norms by which to guide our actions under uncertainty, including moral uncertainty, and only norms that are sensitive to moral uncertainty are suitable to play this action-guiding role; this struck me as a strong reason to believe that at least some such norms are true.³

But now some critics have started to question this view, arguing that there's no important sense in which the probabilities of objective moral views bear on what one should do. So let me borrow a term from one of these critics, Elizabeth Harman (2015),⁴ for the position that I and others have taken for granted: call it "moral uncertainty", or just "uncertainty" for short. Correspondingly, let us call a theory, norm, or view "uncertainty" if, according to that theory, norm, or view, what I subjectively should do in some important sense depends upon the probabilities of objective moral views. Finally, let us call an agent an "uncertainty agent" insofar as she acts on an uncertainty norm.

In this paper, I want to address a type of anti-uncertainty argument offered by Brian Weatherson, and seconded by Brian Hedden, the crucial premises of which are: (1) that acting on uncertainty norms necessarily involves motivation by reasons or rightness as such, and (2) that such motivation is bad in an important way. I will argue that (1) and (2) are false, and that at any rate, the quality of an agent's motivation is not pertinent to the truth or falsity of uncertainty norms in the way that Weatherson's and Hedden's arguments require. Finally, I will present what I regard as the most powerful intuition in favor of uncertainty, and argue that Weatherson's attempts to capture it are unsatisfactory.

³ It is only fair to mention that this insistence on "guidance" in the face of uncertainty introduces the threat of a regress. This is noted in passing in my 2012 and in Weatherson 2014. I sketch a solution to this regress problem in my 2014b and (ms)

⁴ Harman argues that moral ignorance (of which moral uncertainty is one variety) does not exculpate, and so there is no *important* sense in which what I ought to depends on the probabilities of moral propositions. My reply to Harman may be found in my (ms).

2. Weatherson's Argument

Hedden's version of this argument forms a very small part of a longer anti-uncertainty paper, and thus leaves many premises unstated. For this reason, I will focus on Weatherson's version.

Weatherson begins by introducing us to three characters who must make decisions under uncertainty:

Carla, who is baking a cake, and thinks that the powder she's about to add is probably sweetener, but may be arsenic;

Martha, who is deciding what to eat for dinner, and thinks that eating meat is probably morally permissible, but may be seriously morally wrong; and

Agnes, who is contemplating terminating her pregnancy, and thinks that doing so is probably morally permissible, but may be seriously morally wrong. (Weatherson 2014, p. 143)

Weatherson accepts that Carla would in some sense be doing wrong by putting the powder in the cake, even if it turns out that the powder is sweetener. She's taking an unjustified risk. He denies, however, that Martha would be doing wrong by eating the meat, unless of course it turns out that eating meat is wrong in the objective sense. And he denies that Agnes would be doing wrong by terminating her pregnancy, unless it turns out that abortion is wrong in the objective sense. Simply put, Weatherson does not think that "running risks morally" can be wrong in itself, notwithstanding that running risks *non-morally* sometimes is.

He frames his conclusion in terms of what he calls the "'Might' Argument":

1. In the circumstances that Agnes/Martha are in, having an abortion/eating a steak might be morally wrong.
2. In the circumstances that Agnes/Martha are in, continuing the pregnancy/eating vegetables is definitely morally permissible.⁵
3. Missing Premise
4. So, Agnes should not have the abortion, and Martha should not eat the steak. (p. 145)

⁵ Terminating a pregnancy is probably not the best example of something that might be wrong to do but is certainly permissible not to do; after all, there are plausible arguments that carrying a pregnancy to term is the riskier option. See, e.g., Harman 1999.

Weatherson wants to defend the claim there is no way to fill in the missing premise, (3), such that the argument comes out as sound, but the conclusion does not simply follow from (3) alone. That is, he does not mean to rule out the possibility of any sound argument against abortion or meat-eating — only an uncertaintist one.

For what it's worth, I don't think there's a satisfactory way for the uncertaintist to fill in the "'Might' Argument", either. That's because I don't think you can get conclusions about what you subjectively ought to do from probabilities of propositions that include only binary deontic concepts — 'wrong', 'permissible', and so on. To get conclusions about an action's subjective rightness, we need information about *how* objectively wrong it might be, how strong the objective *reasons against it* might be, and also what the reasons *for* it might be. So one might accept uncertaintism while rejecting the "Might" Argument.

But of course I get the point at which Weatherson is gesturing — that there's *no* sound argument that takes premises about probabilities of objective norms and gives us subjective norms as conclusions.

Weatherson's argument against uncertaintism focuses on the *motivations* that someone acting on an uncertaintist norm allegedly must exhibit. He writes of Martha:

"Why should she turn down the steak? Not because she values the interests of the cow over her dining. She does not. And not because she should have that value. By hypothesis, she need not do so....Rather, she has to care about morality as such. And that seems wrong.

The argument I'm making here owes a lot to a similar argument offered for a somewhat different conclusion by Michael Smith (1994). He compared the person who desires to do what is actually right, as he put it, desires the right *de re*, with the person who desires to do what is right whatever that turns out to be, as he put it, desires the right *de dicto*.

"Good people care non-derivatively about honesty, the weal and woe of their children and friends, the well-being of their fellows, people getting what they deserve, justice, equality, and the like, not just one thing: doing what they believe to be right, where this is read *de dicto* and not *de re*. Indeed, commonsense tells us that being so motivated is a fetish or moral vice, not the one and only moral virtue. (Smith 1994, 75)" (p. 152)

Now, one might expect Weatherson to employ these claims about motivation as premises in something like the following argument:

- 1) Anyone acting on any uncertaintist norm must in so doing be motivated by morality as such.
 - 2) Motivation by morality as such is a bad motivation.
 - 3) If anyone who acts on a norm must in so doing exhibit a bad motivation, then the norm is false.
- C) So all uncertaintist norms are false.

But this is not Weatherson's approach; rather, he argues by analogy. He claims that Agnes's and Martha's cases are less like Carla's, and more like the cases of *Bob* and *Bruce*, each of whom is deciding whether to stop at an art gallery on the way home, and each of whom is quite confident that doing so will be no better for him than the alternatives. Bob, however, maintains a sliver of credence in an "objective list" theory on which art-appreciation is a component of well-being, and Bruce does not rule out a theory that privileges "higher" over "lower" pleasures and counts the pleasure of art-appreciation among the former. Just as Agnes and Martha are morally uncertain, Bob and Bruce are, we might say, "prudentially uncertain". (pp. 148-150)

At this point, a "prudential uncertaintist" might step in and say that, notwithstanding which view of well-being is actually correct, it would be "subjectively prudent"⁶ for Bob, say, to go to the art gallery, on the grounds that the "objective list" theory he considers might be correct. In other words, Bob should hedge his prudential bets.

But Weatherson disagrees. The mere probability that a theory of well-being is true provides no prudential reason to pursue what that theory recommends. (p. 150)

Martha's case is relevantly similar to Bob's, Weatherson tells us, because the motivations that Martha would have to exhibit in eschewing the steak would be akin to the motivations Bob would have to exhibit in attending the gallery. Just as Bob would be motivated by well-being as such (as opposed to those things that, as it turns out, are components of well-being), Martha would be motivated by morality as such (as opposed to those things that, in fact, make actions right or wrong). Bob's would-be motivations are unnatural and unhealthy, Weatherson tells us, and Martha's are wrong. (pp. 151-152).

So just as the prudential uncertaintist is wrong about Bob's case, the moral uncertaintist is wrong about Martha's.

This, then seems to be Weatherson's argument:

⁶ Weatherson uses "rational" rather than "subjectively prudent". Since "rational" is used by different philosophers in different ways, I prefer to stick with the neologism. Nothing of significance rides on this choice, though.

- 1) The prudential uncertaintist's proposed norm of "subjective prudence" is false.
 - 2) Bob's motivation in acting on such a norm would have to concern well-being as such.
 - 3) Having a motivation that concerns well-being as such is bad.
 - 4) Martha's motivation in acting on any moral uncertaintist norm would have to concern morality as such.
 - 5) Having a motivation that concerns morality as such is bad.
 - 6) If (2-5), then Martha's case is *analogous in the right way* to Bob's.
 - 7) If Martha's case is analogous in the right way to Bob's, then if the prudential uncertaintist's norm is false, then so are all moral uncertaintist norms.
- C) So all moral uncertaintist norms are false.

Notice that while the first argument includes a premise linking the quality of an agent's motivation in acting on a norm with that norm's truth or falsity — namely, Premise (3) — Weatherson's actual argument does not. Still, it seems to me that Weatherson is committed to accepting such a link. For his analogical argument to work, the similarity between Bob's would-be motivations and Martha's must be the *right kind* of similarity. It is the right kind of similarity only if (2) and (3) imply (1) — that is, only if the fetishism Bob would exhibit in acting on a prudential uncertaintist norm implies that the norm is incorrect; and only if (4) and (5) imply (C) — that is, only if the fetishism Martha would exhibit in acting on a moral uncertaintist norm implies that *that* norm is incorrect. Otherwise, (2-5) represent just some of the many ways an agent who would act on a false norm may resemble some other agent; they would not allow us to infer anything about the moral uncertaintist norm from anything about the prudential uncertaintist one.

Again, these "links" are not premises in Weatherson's argument. But they are claims about the normative significance of motivation that must be true for the analogy to be a good one as opposed to a shoddy one. As such, Weatherson is rationally committed to them.⁷

Against Weatherson, I will argue, first, that the uncertaintist agent needn't be motivated by rightness or reasons as such; second, that such a motivation is not obviously bad in the relevant way; and third, that the quality of an agent's motivation does not have the normative significance Weatherson (and Hedden) need it to have.

3. Must the Uncertaintist Agent Be Motivated by Morality as Such?

⁷ Similarly, someone who executes a material inference does not employ as premises the claims that would transform her inference into a formally valid one. But she is rationally committed to them all the same.

Weatherson and Hedden claim that the uncertainist agent must be motivated by rightness or reasons as such. But I claim that there are two other plausible interpretations of this agent's motivations in acting on an uncertainist norm, each of which entirely obviates the need to impute motivation by rightness or reasons.

First, the uncertainist agent might be ultimately motivated by the features that the *objective* norms amongst which she is uncertain designate as morally significant. If her credence is .5 that we have moral reasons not to cause animal suffering, she might thereby be ultimately motivated, to at least some degree, not to cause animal suffering. Of course, she will probably not be as motivated as she would have been had she been certain of the same. And surely, the strength of her motivation will be determined not only by her credence in this theory, but by the uncertainist norm that she accepts. For example, we would expect a risk-neutral uncertainist and a risk-averse uncertainist to be differently motivated, even if their credences in objective theories are the same. On this picture, the rational agent's desires — insofar as desires are involved at all — are governed not only "from above" by thoughts about objective norms, but also from *further above* by thoughts about uncertainist norms.

Weatherson addresses something like this idea when he writes that Martha cannot be motivated by the interests of the cow, since she does not "value" them. But insofar as this claim strikes us as plausible, I suspect it is largely because it is couched exclusively in terms of the binary notion of *valuing*.⁸ For even if Martha does not value the cow's interests in the sense of having a full belief in their moral significance, she may nonetheless have some credence that they matter. The real question, then, is whether we may be motivated by features that we are not certain are morally significant.⁹

I can imagine two main grounds for saying "no", but neither strikes me as persuasive. First, one might reject the possibility of such motivation on the basis of a plau-

⁸ One might similarly ask, "Would Martha drive all the way to the grocery store if she doesn't believe it's open?", which may well elicit a negative answer. But if we fill in the story a bit more, and say that she nonetheless has some credence that the store is open, her driving there is perfectly intelligible.

⁹ In any case, the claim that one cannot be motivated by F unless one values F *simpliciter* does not entail that no one uncertain about the value of F can be motivated by it. For one can value F *simpliciter* while being uncertain about just how valuable it is relative to other features. And indeed, many cases of moral uncertainty are of just this type -- e.g. a medical review board is uncertain of how considerations of patient autonomy trade off against considerations of patient well-being. It is not clear from Weatherson's paper why an uncertainist agent in such a case could not be correctly described as motivated by autonomy and well-being, rather than by the prospect of rightness. Thanks to a referee for this suggestion.

sible general theory of moral motivation. It would have to be a kind of "hybrid" theory that accounts for moral motivation in morally certain agents without positing a *de dicto* desire for rightness, but explains it in the uncertain agent only by positing such a desire. I would leave it up to my opponents to explain how such a theory might work and why it would be preferable to its more unified competitors, but for now let me simply note that Michael Smith's own view of moral motivation, for which he originally enlisted the "fetishism" argument, is not such a theory. In his (2002), Smith tells us that for a rational agent, the motivational force of a moral judgment is a function of two features — "importance" (i.e. the moral significance the judgment assigns to an action), and "certitude" (i.e. the degree of confidence the judgment manifests). (pp. 309-310) Nor does he regard motivation arising from uncertain judgments as any more fetishistic than motivation arising from certain ones. (p. 312) Smith would agree, then, that an agent may be ultimately motivated by things she thinks might be morally significant, even if, as in Martha's and Agnes's cases, she is not sure that they are.

Second, one might reject my account of the uncertainist's motivations on phenomenological grounds — i.e. that it does not feel, in acting under moral uncertainty, as though we are motivated by "weal and woe" and the like. But it is unclear to me how my opponents could specify the phenomenology that is allegedly absent here.

This is partly because the phenomenology of desire is generally heterogeneous. As Hume noted, some of our passions are "calm" and others are "violent". The feeling associated with my motivation to take out the trash is not the feeling associated with my motivation to escape from a burning building. The feeling associated with picking chocolate ice cream over a close competitor — say, strawberry — is not that associated with picking chocolate over a clearly inferior flavour — say, brine. In light of this, how are we supposed to informatively describe the phenomenology that is "missing" from the uncertainist agent on account of her supposedly fetishistic motivation?

Even if we wish to say that the phenomenology associated with the uncertainist agent's motivations is different than that associated with the motivations of more confident agents, we may resist the further claim that the former phenomenology indicates fetishism. Why not instead say that it indicates motivation by the objects of uncertain judgments rather than the objects of confident ones? It is not clear what the anti-uncertainist could say here, since it is difficult, even in the case of agents who are morally certain, to distinguish on phenomenological grounds between *de re* and *de dicto* motivation to do what's right. (On this point, see, e.g. Schroeder 2015). As evidence for this, note that Smith himself does not provide a phenomenological argument against his main externalist opponent, putting all of his argumentative eggs in the "fetishism" basket instead.

If appeals to theory and phenomenology seem unpromising, why, then, might one resist my proposal regarding the uncertainist agent's motivation? I suspect that resistance rests on a tendency to read too much into the *explanations* we provide to

ourselves and others for our actions. Suppose you asked Martha why she's not eating meat. It is unlikely, perhaps, that she'd cite the interests of the cow straightaway. Rather, it is more probable that she'd say something like, "Well, if it turns out that the interests of the cow are morally important, I'd be doing something very wrong in eating it!"

But we should not conclude that she's motivated by the prospect of wrongness rather than by the interests of the cow; there is a better explanation of why she responds as she does: When we are asked, in the context of a morally-freighted decision, why we did A or would do A, and we say "P", the implication is that we have a full belief, or all-in judgment, that P is at least a reason to do A. Or at any rate, that is how it would be received in ordinary conversation. But such a belief is obviously a much stronger judgment than a credence that P is a reason to do A; correspondingly, it is a more difficult citadel to defend. So there is some pressure for those who only have a credence, especially a low credence, that P is a reason to do A, to say so -- to others, but also, as a result, to oneself. There is no pressure of the same sort for someone who has a full belief that P is a reason to say so; her simply saying "P" in response to "Why would you do A?" is enough. But this does not indicate that one is motivated by P in the case of certainty and by the probability that P is a reason in the case of uncertainty; rather, it indicates that saying "P" best communicates what is salient and informative in the former case, while saying both "P" and "there's some probability that P is a reason" does so in the latter.

An uncertainty agent might also be motivated not by rightness or reasons, but by "thicker" moral considerations. To his credit, Weatherson considers this objection, and indeed expresses agreement with the following claim:

"Even if it is fetishistic to be motivated by the good as such, this doesn't extend to thick moral properties....So [the] objections to arguments from moral uncertainty don't extend to arguments from what we might call virtue uncertainty." (Weatherson 2014, p. 158)

However, he thinks that this opens the door for only a very weak version of the "Might' Argument":

"The fact that some action might manifest a vice can hardly be a decisive consideration against doing it. If the vice in question is relatively small, or the chance of manifesting it is relatively small, it is easy to see how this kind of consideration could be overridden." (p. 159)

Notice the slide, though, from “thick moral properties” to “virtues” and “vices”. This is understandable, since much of the discussion in the “thick concepts” literature revolves around virtue/vice concepts. But it results in Weatherson giving short shrift to this objection.

For there are thick moral features other than virtues and vices, and many of these features are difficult to override. We might think of Martha as motivated not to *murder*, or not to do something *tantamount to murder*. We might think of Agnes as motivated not to *severely wrong*¹⁰ the cow. For what it’s worth, these were the sorts of motivations that I self-ascribed when I first became persuaded by uncertainty. But surely the reasons against murder and severe wrongings are not as easy to override as the reasons against exhibiting a vice such as (to use Weatherson’s example) cowardice.

And in fact, some thick features are *impossible* to override. Judith Jarvis Thomson (1990) famously distinguishes between “infringing” and “violating” a right. To do the former is to do what someone has a right that you not do; to do the latter is to do the former and act wrongly in so doing. But of course, nothing can justify or provide sufficient reason for violating a right, since violating a right is by definition wrong.

Perhaps Weatherson could argue that being motivated not to violate a right is fetishistic, since violation implies wrongness and motivation by wrongness is fetishistic. But I don’t see how he could do this while maintaining his (entirely correct) position that motivation by other thick features is non-fetishistic. For many such features, including most of the virtues and vices I can think of, are reason-implicating, and motivation by reasons is something Weatherson would regard as fetishistic. So it cannot be that, generally, motivation by F is fetishistic on the grounds that motivation by G is fetishistic and something’s being or having F implies its being or having G.

4. Is Motivation by Morality Always Bad in the Relevant Way?

But let us assume for argument’s sake that the uncertainty agent must be motivated by thin moral considerations as such. I do not think that Weatherson has adequately made the case that such motivation is bad in the important sense.

Weatherson says that his “fetishism” argument “owes a lot” to Michael Smith’s. However, one can easily accept Smith’s position while rejecting Weatherson’s. What Smith criticizes in the passage above is having rightness as one’s *only* ultimate motivation — i.e. not *also* having “weal and woe” and so forth among one’s ultimate motivations. But even if an agent must be motivated solely by morality when she acts on uncertainty norms, this does not mean that she cannot be motivated by other things elsewhere. She can be motivated directly by weal and woe when she’s not morally uncertain.

¹⁰ The contrast at work here is between the second-personal notion of “wronging” someone and the third-personal notion of “doing wrong”. See, e.g., Darwall (2007).

Indeed, this strikes me as a morally attractive pattern of motivation: when you're confident enough about which considerations would render your actions objectively right, you act directly on those considerations; but then you're not so confident, you do (what we're now assuming is) the next best thing and aim for objective value itself. (Of course, I also think the patterns of motivation bruted above are attractive.)

Weatherson considers something like this response. Here is the objection as he presents it, and his answer:

“Objection: It may be wrong to be only concerned with right and wrong, but it isn't wrong to have this be one of your considerations.

Reply: I don't think you get the 'Might' Argument to work unless concern with right and wrong, whatever they turn out to be, are the only considerations. Assume that they are only one consideration among many. Then even if they point in one direction, they may be overridden by the other considerations. And if the 'Might' Argument doesn't work, then normative internalism, in its strongest forms, is false. So I really only need to appeal to the plausible view that right and wrong as such shouldn't be our only motivations to get the conclusions I want.” (Weatherson 2014, p. 160)

But notice that this possibility is not quite what I had in mind. My suggestion was that an agent who is motivated exclusively by morality when she's morally uncertain may nonetheless be moved by other considerations when she's more confident. Weatherson is describing an agent who is motivated both by morality and also by these other considerations on the very *same* occasions of moral uncertainty. The allegedly problematic motivational conflict that confronts his agent is absent for mine.

Of course this prompts the question of why this conflict is supposed to be problematic in the first place. Weatherson tells us that the “'Might' argument' doesn't work” unless the agent is exclusively motivated by right and wrong in cases of moral uncertainty. But no premise in the “'Might' argument” appeals to our motivations, nor is the conclusion a claim about our motivations. It's a conclusion about what's subjectively right. However, perhaps Weatherson can argue that if the uncertaintist agent's direction of motivation is not fully determined by her judgment of what she (subjectively) ought to do, then she is, by the uncertaintist's lights at least, practically irrational. He can then present the uncertaintist with a dilemma: agents who act on uncertaintist norms are *either* objectionably fetishistic (if they are motivated solely by those norms), *or else* practically irrational (if they are motivated by other things, too).¹¹

¹¹ Thanks to a referee for help here.

I am also tempted to ask Weatherson and Hedden something to the effect of: “*What do you want from me?! — qua* uncertaintist agent, that is. I’m uncertain about certain moral issues, and let’s suppose reasonably so (i.e. my credences line up with the epistemic probabilities.) You say that in acting on an uncertaintist norm, I thereby exhibit bad motivations. But is there anything I can do instead, such that I don’t exhibit bad motivations? If the answer is ‘no’, then there must be something wrong with your argument somewhere.”

Let us try to spell out this charge more carefully. When I am consciously uncertain about morality, it seems that I must either act on an uncertaintist norm, or else do an action that is not fully norm-guided — what I call in my (2014b) a “leap of faith”.¹² But, just intuitively, the motivations of an agent in taking a leap of faith can hardly be called virtuous; such an agent seems to be acting recklessly by her own lights. And indeed we can go beyond intuition. I argue in (2014b) that to take a leap of faith in the face of conscious uncertainty among norms is to implicitly accept and act upon a norm about what to do in the face of such uncertainty. So we must act on uncertaintist norms either way, it would seem; the only question is whether to do so explicitly, and perhaps with some conscious thought and deliberation, or implicitly, without any such thought. It seems, then, that the former is the better strategy; at the very least, it is not clear to me how it could be the worse.

If that’s right, then it does not seem that the motivations of the uncertaintist agent are bad in the *comparative* sense — i.e. worse than the other motivations she could have exhibited. But normally when we call a motivation “bad”, we do think it is bad in this comparative sense. I act on the norm “Do whatever’s best for me” and thereby exhibit a worse motivation that I’d have exhibited had I acted on “Help the afflicted” instead. More controversially: I act on “promote my childrens’ happiness” and thereby exhibit a worse motivation that I’d have exhibited had I acted on “Promote my childrens’ ends” instead. And so on.

At this point, Weatherson and Hedden might grant that the uncertaintist agent’s motivations are not bad in the comparative sense, but suggest that they are bad, nonetheless, in some more “absolute” sense. That is, they are bad, but not worse than other motivations the agent might have exhibited in the same situation. I will consider this sort of suggestion in the next section, where I shall reply that the the absolute

¹² (Weatherson’s focus on certain class of uncertaintist norms (i.e. those enjoining moral recklessness) may lead some readers to think that other uncertaintist norms (e.g. those prescribing action in cases where there is no morally “safe” option) are beyond the reach of his argument. But it is a mistake to think this. His argument applies to all uncertaintist norms that we can use to guide action. As such, if his motivational arguemnt were correct, the only alternative to fetishism would be what I’ve called a “leap of faith”.

badness of a motivation an agent exhibits in acting on a norm of the sort uncertainists care about gives us little reason to think the norm is false.

For now, though, I want to consider a way for Weatherson and Hedden to rescue the claim that the uncertainist agent's motivations are bad in the comparative sense. They might say: "It's true that once you're in a position of moral uncertainty, you can't act from good motivations. You're doomed to act on unseemly ones by dint of being in such a state in the first place. But while the uncertainist agent's motivations may be no worse than the "leaper's" motivations, both are worse than those of someone who avoids moral uncertainty in the first place, and can act in a way that is neither reckless from his own perspective nor fetishistic. So the uncertainist agent's motivations are bad in the comparative sense after all."

While this is a way of showing that the uncertainist agent must necessarily exhibit bad motivations, it is not a way of showing that she must exhibit them as a result of acting on uncertainist norms. But it's really the latter question we're interested in. When it comes to the truth or falsity of uncertainism, what matters is not whether uncertainist norms accompany bad motivation, even necessarily so. What matters is whether, in acting on uncertainist norms, one thereby makes one's motivational profile worse -- whether the uncertainist norms are responsible for the bad motivation. And so far, we've seen no reason to think that they are.

5. From Bad Motivations to False Norms?

Weatherson and Hedden want to infer the falsity of all uncertainist norms from the supposed badness of the motivation a person exhibits in acting on such norms. But I don't think there exists the sort of link between quality-of-motive and correctness-of-norm that would license such an inference.

As far as I can tell, there are two things Weatherson and Hedden might say about what such a link amounts to. (Neither actually says anything about this.) First, they might say that the badness of the uncertainist agent's motivations can be the sort of thing that determines the rightness or wrongness of actions, and as such makes uncertainist norms false. Second, and less ambitiously, they might say that the badness of these motivations is mere evidence or reason to believe that the norms are false.

My intuition, at least, is that the agent's motivations just don't seem to be sorts of things that matter in the first way. While the quality of my motivation may partly determine what Kant called the "moral worth" of my action, it cannot influence whether the action is, as Kant would put it, in conformity with duty. To say otherwise is to place the agent's motivations alongside features like pleasure, autonomy, and flourishing, as something that plausibly counts in favor of, or against, our doing an action. But I can't see how it belongs in such privileged company. We would not say of an otherwise advisable action that an agent ought not to do it because, were she to do it, she couldn't help but exhibit bad motivations.

And note that assigning such a normative role to motivation leads to absurd conclusions about which motives are good or bad in the first place. For if something is relevant to the rightness of an action, it is acceptable to be motivated by it. But it does not seem acceptable to be motivated by my own would-be motivations. It seems, if anything, a more twisted fetishism than anything Weatherson and Hedden describe to think, "If I did A, or followed rule R, I'd do it with a bad motivation; so I'll do B or follow rule S instead."

But really, all we need is the weaker thesis that the badness of an agent's motivation does not have "trumping power" — that it does not override or defeat any and all (subjective) reasons for an action grounded in the chance of that action being (objectively) right or its complement being (objectively) wrong.¹³ The uncertaintist "wins", after all, if there are *any* actions that we ought subjectively to do on the grounds of the epistemically possible objective reasons in its favor. Unless one has some *independent* argument that such subjective reasons are (a) non-existent, (b) easily outweighed, or (c) easily defeated, it is difficult to see what could be said for assigning quality-of-motive such a powerful normative role.

At this point, Weatherson and Hedden might instead opt for the view that the badness of the uncertaintist agent's motivations is simply *evidence* that, or otherwise *support the claim*, that uncertaintism is false. For of course one ethical claim can support another without the former's truth being any part of what makes the latter true. Suppose we're not sure whether we think A's use of force against B was legitimate. We might try to get at an answer by asking "Would it be permissible for B to forcefully resist?"; we would not thereby be committing ourselves to the implausible claim that what *makes* someone's use of force legitimate is that someone else could not permissibly resist.

But I think that sort of strategy is open to some serious challenges. First, if the idea is not that bad motivations *make* uncertaintist norms false, then one would like to know what *does* make them false. To adopt the stance that motivational considerations are mere evidence of their falsity is to concede, in effect, that the deep, fundamental argument against uncertaintism, the real reason why it is wrong, lies elsewhere.

Second, I wish to make good on a promissory note from the end of the last section. I argued that, in acting on an uncertaintist norm, an agent does not thereby exhibit worse motivations than she might have exhibited otherwise. I left it open, however, that she might thereby exhibit a motivation that is nonetheless bad in an absolute sense. But now I want to claim that this would not constitute evidence against the truth of uncertaintism. Normally, when we condemn a norm on grounds that one who acts on it is badly motivated, we have it in mind that she would not have been (as) badly

¹³ I owe this astute point to an anonymous referee.

motivated had she not acted on the norm. I gave some examples in the previous section. Here, we have a case where the agent is apparently doomed to exhibit a bad motivation just by being morally uncertain; that she acts on some uncertainist norm rather than taking a "leap of faith" determines only the form this bad motivation will take.

I struggle to see why this should count against uncertainism, but that's partly because I struggle with how to even think about this sort of case, where no moral claim upon which you could act will afford a good motivation. The closest analogy seems to be something like this: Let us suppose that everyone is indeed fetishistically motivated in all cases, not just those of moral uncertainty. And let us suppose this is a necessary truth. Here someone like Smith or Weatherson would say that whatever norm one acts on, one exhibits a bad motivation. The only way to avoid fetishism would be to avoid acting on norms at all -- to simply pursue one's unmotivated desires. If we are to read off the falsity of a norm from the fact that one who acts on it necessarily exhibits a bad motivation in the absolute sense of "bad", then we would have to say that in this scenario, all practical norms are false. But that is implausible. It makes more sense to say that this situation is lamentable, but that some norms we may act on in this sad condition may nonetheless be true. And that is what I wish to say in the moral uncertainty case.

Indeed, it is even more plausible to say that in the moral uncertainty case. In the scenario where everyone is fetishistically motivated whenever they act on a norm, the alternative to acting on a norm does not seem obviously to involve unseemly motivation. For what is objectionable about acting on an unmotivated desire? So there is at least some case to be made that all practical norms would indeed be false, since they cannot be followed except with bad motivation, and that we should eschew norm-guided action entirely. But in the normative uncertainty case, the alternative, given that one is normatively uncertain, is just to take a leap of faith in the face of one's uncertainty, which I've suggested actually involves exhibiting an even worse motivational profile. It is not just that there would be no way of engaging in norm-guided action without being badly motivated; it's that there would be *no way of going forward at all* without being badly motivated.

But finally, suppose we grant that the absolute badness of motivation is evidence of the falsity of the uncertainist's norms. One might nonetheless maintain that this evidence is easily outweighed. Weatherson writes as though the dialectical situation between him and the uncertainist is one of "duelling analogies" (p. 151): we say Martha's case is more like Carla's; he says it's more like Bruce's. But this misrepresents the sorts of arguments that uncertainists actually give for their position. For instance, as I mentioned in the introduction, I've always presented the main consideration in support of uncertainism as one of action-guidance — that it is only with the aid of uncertainist norms that we can hope to guide our actions under conditions of moral uncertainty, and that gives us reason to believe that some such norm is true. Now, of

course one might discount this consideration. My point here is just that we shouldn't be too worried about Weatherson's and Hedden's motivational arguments until we see how they stack up against all of the kinds of considerations that uncertaintists actually marshal in support of their position.

6. Coda: The Value of Moral Inquiry

I've spent most of this paper on the defensive, trying to persuade you that the "fetishism" argument against uncertainty fails. I'd like to conclude by briefly going on the offensive and pressing what I regard as the single most powerful intuition in support of uncertainty. I do this partly because it is an intuition of the sort that Weatherson claims is consistent with his anti-uncertainty.

The intuition is that, given the probabilities of the various objective norms, sometimes the right action is the *cognitive* one of thinking more about which of these norms is correct. This is not always the right choice, of course. Sometimes the answers to moral questions are so obvious that the gains from further inquiry will be too slight. Sometimes circumstances are so exigent that the losses from omitting to do some ordinary, non-inquiry action will be too great. But in at least some cases, the right response to moral uncertainty is to think more about the issues.¹⁴

Weatherson seems to agree. He writes:

"I agree that [moral deliberation] is good practice. But I think it is consistent with what I've said so far. Start with an observation also by Michael Smith, that moral inquiry has "a certain characteristic coherentist form" (Smith, 1994, 40–41). I think (not originally) that this is because we're not trying to figure out something about this magical thing, the good, but rather because we're trying to systematise and where necessary reconcile our values. When we're doing moral philosophy, we're often doing work that more at the systematising end, trying to figure out whether seemingly disparate values have a common core. When we're trying to figure out what is right in the context of deciding what to do, we're often trying to reconcile, where possible, conflicting values. But as long as we accept that there are genuinely plural values, both in moral and prudential reasoning, we shouldn't think that a desire to determine what is right is driven by a motivation to do the right thing, or to live a good life, as such." (pp. 160-161)

Now, there are a few parts of this passage with which I'm not entirely happy: What work is being done by "magical" here? And is it really so crucial to Weatherson's

¹⁴ I argue for this in my 2010; in doing so, I draw on arguments in Good 1967 and Skyrms 1990, chapter 3.

argument that “we” “accept” that there are genuinely plural values — such that, say, someone with substantial credence in utilitarianism is beyond the argument’s reach? Why, for that matter, should we think there’s any difference between trying to “figure out something about...the good” and trying to “systematise...and reconcile” our values”? (I presume he means “systematise and reconcile well” — that not just any feel-good sytematisation or reconciliation will do.) Finally, even if there is a difference between these two aims, why should we think the aim of systematising and reconciling values is *not* fetishistic if the aim of figuring out about the good *is*? Both are pretty far from “weal and woe”, or so it seems to me.

But my major concern is that I don’t think Weatherson is actually quieting the worry he needs to quiet. He is not showing how the intuition I’ve expressed above is consistent with the conclusion of his original argument; he is, at most, showing how a weaker intuition is consistent with an altogether different motivation-based argument.

Let’s distinguish between these two intuitions:

STRONG: The act of moral deliberation is sometimes the subjectively right thing to do, practically speaking, in virtue of the fact that it’s uncertain what’s objectively right.

WEAK: The act of moral deliberation is sometimes the right thing to do.

STRONG is the intuition I presented above — that not only is moral deliberation sometimes the right thing to do, but it is sometimes the *subjectively* right thing to do, practically speaking, in virtue of the epistemic probabilities spread out over the various objective moral theories. WEAK is weaker in two notable respects. It is compatible with the claim that moral deliberation is never the *practically* right thing to do, but merely sometimes the *epistemically* right thing to do. But Weatherson makes it clear that this is not his view, and at any rate, it’s implausible. More importantly, WEAK is compatible with the claim that moral deliberation is never the *subjectively* right thing to do, but merely sometimes the *objectively* right thing to do — that is, never the right thing to do because it’s uncertain which moral theory is correct, but rather only because the deliberating agent is sometimes mistaken about what’s objectively right, and deliberation is way of rectifying that mistake such as to prompt right action.

The plausibility of Weatherson’s position hinges on his ability not only to capture the objective rightness of deliberation, but also its subjective rightness. No one will deny that there is some sense in which we sometimes ought to think more about moral issues about which we’re mistaken. But there also seems to me to be a sense in which we ought to think more about certain moral issues — distributive justice, the ethics of war, the act/omission distinction — just because they’re *hard*: there’s some chance that each of several mutually exclusive views is correct.

This is similar to what we'd say in the non-moral case: There is some sense in which the researchers with the wrong views about the health effects of DHA ought to do and read more research, while those with the correct views ought to consider the matter no further; but there is also some sense in which all researchers in this area ought to do and read more research, just because, at this point in history, it has not been established what the health effects of DHA are. Similarly, I would say, it has not been established whether agent-responsibility is sufficient to render one liable to defensive killing, or whether, *ceteris paribus*, harmful acts are worse than harmful omissions. There is some sense, then, in which all who pursue this sort of inquiry have at least some reason to pursue it further.

However, Weatherson's reinterpretation of moral inquiry as "synthesizing and reconciling" at best allows him to capture WEAK, by showing that moral inquiry can be done without having one's sights set upon rightness or goodness as such. To capture STRONG, though, he would have to allow that moral inquiry, however we characterize its aims, is not only right in some sense or other, but is subjectively right in light of the fact that it's often uncertain which moral views are correct. But the conclusion of his original motivational argument is just that *nothing* is right in that sense, and so he is foreclosed from capturing STRONG.

References

- Darwall, S. L. (2006). *The Second-person Standpoint: Morality, Respect, and Accountability*. Harvard University Press.
- Good, I. J. (1967). On the Principle of Total Evidence. *British Journal for the Philosophy of Science*, 17(4), 319–321.
- Gracely, E. J. (1996). On the Noncomparability of Judgments Made by Different Ethical Theories. *Metaphilosophy*, 27(3), 327–332.
- Gustafsson, J. E., & Torpman, O. (2014). In Defence of My Favourite Theory. *Pacific Philosophical Quarterly*, 95(2), 159–174.
- Harman, E. (1999). Creation Ethics: The Moral Status of Early Fetuses and the Ethics of Abortion. *Philosophy and Public Affairs*, 28(4), 310–24.
- Harman, E. (2015). The Irrelevance of Moral Uncertainty. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Vol. X. Oxford University Press.
- Hedden, B. (2015). Does MITE Make Right: On Decision-Making Under Normative Uncertainty. In R. Shafer-Landau (Ed.), *Oxford Studies in Metaethics*, Vol. XI. Oxford University Press.
- Hudson, J. L. (1989). Subjectivization in Ethics. *American Philosophical Quarterly*, 26(3), 221–229.
- Jonsen, A., & Toulmin, S. (1988). *The Abuse of Casuistry*. Berkeley. University of California Press.
- Moss, S. (2013). Epistemology Formalized. *Philosophical Review*, 122(1), 1–43.
- Schroeder, T. (2015) Desire. *Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/entries/desire/>. Accessed September 18, 2015.

Sepielli, A. (2010). 'Along an Imperfectly-Lighted Path': Practical Rationality and Normative Uncertainty. Ph.D. Dissertation, Rutgers University.

Sepielli, A. (2012). Subjective Normativity and Action Guidance. In M. Timmons (Ed.) *Oxford Studies in Normative Ethics*, Vol. II. 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.... *Nous*, 48(3): 521-44.

Sepielli, A. (2014b). Should You Look Before You Leap? *The Philosophers' Magazine* 66: 89-93.

Sepielli, A. (ms). How Moral Uncertainty Can Be Both True and Interesting.

Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Harvard University Press.

Smith, M. (1994). *The Moral Problem*. Blackwell.

Smith, M. (2002). Evaluation, Uncertainty and Motivation. *Ethical Theory and Moral Practice: An International Forum*, 5(3), 305–320.

Thomson, J. J. (1990). *The Realm of Rights*. Cambridge, MA: Harvard University Press.

Weatherson, B. (2014). Running risks morally. *Philosophical Studies*, 167(1), 141–163.

Yalcin, S. (2007). Epistemic Modals. *Mind*, 116(464): 983–1026.