Canadian Journal of Philosophy (2026), 56: e18, 1–20
doi:10.1017/can.2026.10067

ARTICLE

Trying to Be Moral, Morally

Chelsea Rosenthal

Simon Fraser University, Burnaby, BC, Canada
Email: chelsea_rosenthal@sfu.ca

Abstract
Often, we're unsure what morality requires, and there's a debate about whether our moral uncertainty and moral beliefs matter to what we should do. I argue that (1) our moral-epistemic circumstances (our beliefs, credences, and evidence about morality) matter to what we morally ought to do; (2) many objections to this are only problems if we assume all moral norms are sensitive to moral-epistemic circumstances; and (3) we should avoid this with an account incorporating procedural norms that are sensitive to our moral-epistemic circumstances (analogous to legal due process norms) and substantive norms that aren't (analogous to substantive justice norms).

Keywords: moral uncertainty; moral risk; uncertaintism; substantive oughts; procedural oughts; normative internalism; normative externalism; intertheoretic value comparison; metanormativism; moral ignorance

1) Introduction

Most of us are trying to do the right thing—we want to act morally. But figuring out what morality calls for is challenging. Our moral judgments are fallible, and we are often uncertain. This can happen because we do not know the relevant descriptive facts—for example, whether an item of clothing was made in a sweatshop or whether a grieving friend would find a phone call comforting as opposed to intrusive. But we can also face difficulties figuring out relevant moral norms, even when the descriptive facts are clear. For example, someone may be unsure whether it's okay to loyally conceal a friend's wrongdoing to protect them from serious penalties or, in a case I discuss below, a person may consider whether to break a promise of confidentiality in order get help for a friend struggling with addiction.

This paper argues that being reckless when navigating our moral uncertainty and the fallibility of our moral judgments is, itself, morally wrong. (When I discuss moral uncertainty, here, I'm referring to uncertainty about what is moral to do, which cannot be resolved by settling descriptive questions.) Avoiding this recklessness could include simple steps like steering clear of a morally questionable activity just in case it's wrong or taking actions that reflect careful thought, not drunken impulses.

This much might seem intuitive. But it has become controversial in recent work. A debate has emerged about whether what we should do is sensitive to what I'll call our moral-epistemic circumstances: our beliefs, credences, and evidence regarding moral norms. This includes disagreement about whether what we should do depends on factors like our moral uncertainty. One camp thinks that what we should do is sensitive to our moral-epistemic circumstances (Guerrero, 2007; Lockhart, 2000; MacAskill, et al. 2020; Moller, 2011; Sepielli 2009). The other camp rejects this

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Canadian Journal of Philosophy, Inc. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (http://creativecommons.org/licenses/by-nc-nd/4.0), which permits non-commercial re-use, distribution, and reproduction in any medium, provided that no alterations are made and the original article is properly cited. The written permission of Cambridge University Press or the rights holder(s) must be obtained prior to any commercial use and/or adaptation of the article.

and says that what we should do is independent of our moral-epistemic circumstances—that what is right for us to do can be settled without peering into our heads (Harman, 2015; Weatherson, 2019). Following Elizabeth Harman, I'll call the first camp "uncertaintists" and discuss the second camp as critics of uncertaintism.[1] But a variety of terms have been used for similar ideas. Prominently, Brian Weatherson refers to the first group as a type of normative internalist and the second group as a type of normative externalist (2019). I'll treat "uncertaintist" as largely synonymous with Weatherson's "internalist" (I adopt the first term mainly to avoid confusion with debates about motivational internalism in ethics). While uncertaintism gives moral uncertainty the prominence in its name, moral uncertainty is only one of a variety of important moral-epistemic circumstances that could matter to what we morally ought to do—and I'll take the uncertaintist thesis to be broad enough to encompass others (e.g., being confident in a moral view). What's key is that the uncertaintist believes that some combination of our beliefs, credences, and evidence regarding morality matters to what we morally ought to do.

The literature on this issue is in a peculiar place. On the one hand, some of the most prominent criticisms of uncertaintism have targeted a version of uncertaintism that says that our moral-epistemic circumstances impact what we *morally* ought to do (Harman, 2015; Weatherson, 2014, 2019). On the other hand, there has been surprisingly little direct defense of this type of uncertaintism. Most arguments that our epistemic circumstances matter to what we *morally* ought to do have focused on credences or evidence about morally relevant *descriptive* facts, rather than credences or evidence about *moral* norms.[2] At the same time, there is extensive literature about the normative implications of uncertainty about *moral* norms, but in most cases, these authors either explicitly defend theses about what we *rationally* ought to do in the face of moral uncertainty (Lockhart, 2000; Ross, 2006; Sepielli, 2009; Tarsney, 2018) or leave open whether the relevant norms are rational or moral (MacAskill, 2014) (the exceptions tend to focus on narrower questions (Guerrero, 2007; Moller, 2011) or do not take the moral character of these norms as a central claim to be argued for).[3] This paper fills that gap, defending uncertaintism as an explicitly moral claim: our moral-epistemic circumstances matter to what we morally ought to do—and failing to take them into account appropriately is, itself, morally bad (not only a failure of rationality). I go on to develop a type of uncertaintism that can avoid many of the leading criticisms.

The paper begins, in Section 2, by offering two common sense arguments for the claim that our moral-epistemic circumstances matter to what we morally ought to do. First, acting with respect toward those who could be victims of our wrongdoing requires that we are careful, rather than reckless, as we navigate our moral uncertainty and the possibility that our moral judgments could be mistaken. Second, if we accept commonplace requirements to prepare well for morally important

[1]Harman describes uncertaintism as the view that our moral uncertainty and specific moral credences are important to what we morally ought to do and contrasts it with her favored alternative, actualism, according to which "a person's moral beliefs and moral credences are usually irrelevant to how she (subjectively) should act" (2015, pp. 53–54, 58). In the discussion that follows, I'll understand uncertaintist views to include both this form of uncertaintism and also views on which an agent's moral *evidence* or *justified* moral beliefs are relevant to how they should act (a possibility that Harman, herself, also briefly addresses) (pp. 76–78). But my arguments will apply to Harman's belief and credence-based version of uncertaintism as well.

[2]See e.g., Graham (2010) and Smith (2010). See also related arguments by Jackson (1991) and Railton (1984) and discussion of the notion of a "decision ought" in Jackson and Smith (2006, p. 270).

[3]Field (2019)'s focus is on arguing that normative uncertainty should be handled in the same way as factual uncertainty (10, note 7). Zimmerman (2014) provides an account of moral obligation in the face of moral uncertainty (about both relevant descriptive facts and moral principles). But while he favors discussing our obligations in these cases as moral obligations, he regards proposals to reframe the discussion in terms of what's rational or reasonable as "merely a verbal quibble" (45). So, although he briefly indicates why he finds the framing in terms of morality to be natural, this is not the focus of his argument. In contrast to the view I develop below, Zimmerman's account takes there to be only one type of moral norm that applies to us in these cases (what he terms "the prospective ought"). I'll offer some worries about related single-ought accounts in Section 4. Krister Bykvist also briefly raises the possibility that we morally ought to navigate evaluative uncertainty in a rational way, though he does not argue for it in full (Bykvist, 2020, pp. 324–327).

decisions (with, say, careful forethought rather than drunken impulses), then the moral credences we form through those preparations must also matter to what we should eventually do. Without this, it's unclear why we should bother with good preparations.

Section 3 of the paper addresses some key criticisms of uncertaintism—remember, critics have especially targeted this variety of uncertaintism that says that our moral-epistemic circumstances impact what we *morally* ought to do. I argue that a lot of worries about this view are really problems for uncertaintist theories with particular features—features that we can reject without rejecting the core idea of uncertaintism. In particular, influential objections to uncertaintism have tended to suppose that uncertaintists are committed to a series of controversial claims: that moral ignorance is exculpatory, that we should take our moral-epistemic circumstances into account in specific ways involving expectation maximization, and that *all* moral norms are sensitive to our moral-epistemic circumstances—not just some of them. I outline how uncertaintists can avoid each of these three commitments and why there is often independent reason to think that the best versions of uncertaintism will do so.

In Section 4, I offer a positive account: I defend a framework for our obligations in the face of moral uncertainty and fallible moral judgments that avoids the critics' challenges, developing in particular this third idea—that moral-epistemic circumstances can matter for some moral norms without mattering for all of them. I argue that we are subject to both procedural moral norms that are sensitive to our moral-epistemic circumstances (analogous to the due process norms of a legal system) and substantive norms that do not need to be sensitive to them (analogous to norms of substantive justice). Procedural norms, on this account, are moral norms about how we should go about trying to satisfy other moral norms given that we face challenges in discovering their content —for example, our uncertainty and fallible judgment. For example, we would be endorsing procedural norms if we suggested that people should "act based on careful moral contemplation" or "avoid actions whose moral rightness is unclear, if they can be avoided at low cost." In contrast, substantive oughts are ordinary, first-order moral norms that are not focused on the moral strategies to use when trying to comply with other moral norms (e.g., "don't lie" and "maximize utility" could be substantive oughts).

The arguments I develop have three large implications. First, they suggest that many of the debates about moral uncertaintism have been distorted by a focus on extreme options: either moral-epistemic circumstances always matter to what we morally ought to do or they never matter to what we morally ought to do. But if I'm right, there's a fruitful middle ground that has been largely ignored, in which moral norms are *sometimes* sensitive to our moral-epistemic circumstances. And many worries about uncertaintism can be seen, instead, as reasons to reject an extreme uncertaintism in favor of a more moderate form of uncertaintism, rather than counseling for giving up uncertaintism entirely.

A second set of implications follows from my argument that it's *morally* bad to be reckless in the face of moral uncertainty—not just a problem of rationality.[4] I'll leave open here what distinguishes moral norms from other norms (e.g., whether it is something about their gravity or their structure, etc.). But it's far from a minor, terminological difference. Whether a norm is moral has potential implications for how we think about blame, for instance (we may be to blame for procedural recklessness, itself, even if this does not alter blame for the substantive rightness and wrongness of our actions). And if we *morally* ought to navigate moral uncertainty in particular ways or using particular strategies, then these responsibilities present a wide swath of underexplored moral norms. Taking a closer look at them may shed light on other moral questions. For example, moral

[4]This distinction between moral and rational norms is not meant to rule out the possibility that morality derives from rationality (e.g., in the manner suggested by Kant). The moral norms that Kant describes are more robust than the means-ends rationality discussed in the decision-theoretic analyses of moral uncertainty that I'm rejecting (decision theory is not, for example, a study of the goodness of the will).

norms calling for interpersonal toleration might be best explained as good strategies for being moral when facing disagreement and uncertainty (Rosenthal, 2019). But explanations like this require that using good strategies for being moral is something that we morally ought to do.

Finally, crucially, understanding these strategic responsibilities as *moral* norms has a large impact on how we determine their content. When navigating moral uncertainty, we can—and should—ask, for example, which risks of wrongdoing are compatible with respecting others, and which personal gains we should be willing to sacrifice in order to manifest that respect. It's not only that we have strategic responsibilities that happen to be moral—it's that thinking about these responsibilities' content and demands will involve engaging directly with moral ideas and values.

2) Common Sense Reasons for Uncertaintism

To see why our moral-epistemic circumstances matter to what we morally ought to do, let's consider what it takes to treat others with respect when we find ourselves unsure what's morally called for, or simply aware that our moral judgments might be flawed. This can come up when people are deciding whether to turn in a friend for cheating on an exam, for example, or when they must make end-of-life decisions on behalf of an aging parent. I'll discuss in more detail a case of someone who considers breaking a promise of confidentiality in order to get help for a friend. Geraldine's friend, Ellen, wants to confide in her about recent difficult experiences, but first, Ellen asks Geraldine to promise to keep their conversation a secret. Geraldine agrees, and Ellen reveals that she has been using drugs again after a period of sobriety. Ellen has a history of struggles with addiction, and Geraldine wonders whether it would be moral for her to break her promise of confidentiality in order to involve people who might be able to help her friend.[5]

I will not try to provide an answer for Geraldine here. Instead, I'll start by suggesting something, perhaps obvious, about the way she should approach the decision: she should not avoid thinking about the uncomfortable moral questions involved by, instead, deferring to her pet parrot's advice. We can picture Geraldine asking her parrot, Polly, "should I break my promise to Ellen?" and, having had a long day, figuring that Polly's answer is good enough. But I'll assume that ordinary adult humans, with all our flaws, are still better equipped to think through humans' moral problem than are parrots—who likely do not understand the complexities at stake. So, by trusting Polly rather than engaging in her own careful consideration of the relevant issues, Geraldine seems to be behaving recklessly. She is facing a difficult, weighty moral decision. Making the decision based on a quick chat with a parrot would be morally bad—even if the parrot's answer accidentally led to a good course of action.[6] Taking action with so little contemplation and attention would be treating her friend, Ellen, in a morally objectionable way.

This simple guidance—do not rely on a parrot's moral advice—helps illustrate two common sense reasons that our moral-epistemic circumstances sometimes matter to what we morally ought to do. First, and most straightforwardly, behaving with respect toward others includes recognizing that they can make moral claims on us, and that mistreating them is something that we should care enough to try to avoid (even knowing there are no options that guarantee this). So, taking morally important actions based on haphazard or reckless processes involves acting disrespectfully toward the people who will be impacted—if Geraldine's decision was based on a parrot's advice, her friend Ellen could reasonably complain that she wasn't given morally important forms of respect (even if Geraldine's decision was, as a matter of luck, a good one).

[5]This case is adapted from Rosenthal (2021).

[6]For discussion of moral responsibilities to adequately investigate moral questions before acting, see Guerrero (2007), pp. 63–65; Rosen (2004), p. 301; for discussion of these preparatory duties with respect to uncertainty about morally relevant descriptive facts, see Smith (2014). For recent discussion of the instrumental value of normative information, see Russell (2025).

But we can only say that it's morally bad for Geraldine to rely on Polly's answer, if moral-epistemic circumstances matter to what we morally ought to do. Criticizing Geraldine for following her parrot's guidance involves criticizing her based on the moral-epistemic circumstances that led to her actions—the credences she formed following her chat with the parrot. If Geraldine's moral-epistemic circumstances were irrelevant, then it would only matter whether she kept her promise of confidentiality or not—keeping the promise would mean she has done what she morally ought to do (if, let's suppose, keeping her promise was right), even if she simply followed her parrot's advice and was lucky it coincided with a morally good choice.[7]

It's worth bearing in mind, as well, that when Geraldine acts disrespectfully by deferring to Polly we cannot simply say that this is just an indication of her character (though it may be that too). Geraldine is wronging a particular person—Ellen could reasonably complain that *she, herself*, had been mistreated by Geraldine's recklessness. And this complaint would remain even if making an important decision based on a parrot's advice was, in fact, out-of-character for Geraldine. She still would have done something wrong when she made an important decision in that way and ran a high risk of mistreating Ellen—even if the parrot's advice led to a substantively good decision. For this reason, apologizing to Geraldine for handling the decision so recklessly would be appropriate, for example.

Moral obligations like these, that require us to navigate moral uncertainty responsibly, also aren't a matter of fetishistic focus on our own moral purity (as Weatherson has worried); instead, they involve trying to treat the people in our lives well, even when we are unsure what that involves (2014).[8] What this calls for may vary with the context. It might include choosing a "morally safer" action instead of one of uncertain permissibility—respecting a potential romantic partner might involve erring in the direction of being more forthcoming with them, for example. In other cases, it might call for us to seek out good advice or sensible sounding boards—and at a minimum, to avoid making moral decisions based on drunken impulses or advice from pets, when the stakes are weighty, as in Geraldine's case. Notice that the precautions I'm mentioning are not, for the most part, envelope-pushing suggestions—just reminders of ordinary aspects of common sense morality. In short, we morally ought to treat others with respect, and so (at least sometimes) we morally ought to take steps to make it less likely that we are wronging them. This requires taking some of our moral-epistemic circumstances into account: our uncertainty about the permissibility of some actions, for example, or our improved assessment of a situation following useful advice.

I'll turn now to a second, but closely related, reason to think it's morally bad to use reckless strategies for navigating our moral uncertainty and fallible moral judgment—even if we happen to end up acting well. I'll suggest that if we accept (less controversially) that we have moral responsibilities to prepare thoughtfully for important moral decisions, then we also must accept (more controversially) that our moral-epistemic circumstances matter to which actions we should ultimately take.

To see this, we can notice that strategies to reduce our chances of wrongdoing come in two varieties. On the one hand, there are precautions involving what we do *before* we act, in preparation for acting. For example, suppose Lisa is considering accepting a political job in which she is likely to face moral dilemmas that make it difficult to avoid harming some people and "getting her hands dirty" (Walzer, 1973), but in which she could also be positioned to bring about important social good. When deciding whether to take the job, she might seek out multiple perspectives from people whose moral insight she respects or who have thought about the relevant values when their communities were impacted by political decisions the past. In another case, Justin might position

[7]Of course, a critic of uncertaintism might still say that consulting the parrot is, itself, a morally bad act, even if whether Geraldine should keep her promise is independent of her beliefs. But, as I suggest below, it's unclear why it would be bad to be reckless in forming our moral beliefs (say, by consulting a parrot) if we do not have a responsibility to act in ways that reflect those beliefs.

[8]For discussion of difficulties facing Weatherson's view, see Aboodi (2017), Hicks (2019), and Johnson King (2020).

himself to be less likely to make a morally bad decisions when deciding whether to reveal that a friend cheated on his partner, simply by being sure to consider the issue carefully rather than deciding what to do while rushed or drunk. Of course, there are no guarantees that taking these steps will lead to good actions (perhaps Lisa and Justin's first thoughtless impulses were best, after all), but because these precautions can make it more likely that we'll act morally, they can be morally valuable.

On the other hand, we might try to improve our chances of acting morally using precautions that involve the actions we ultimately take, rather than the way we prepare. For example, Amy might avoid a maybe-permissible action in favor of a clearly permissible alternative. Perhaps she is unsure if she has a moral obligation to discuss some aspects of her history with a new romantic partner, but she is confident that it's morally permissible to be forthcoming about her past. So, she errs in a morally safe direction and has a candid conversation with her partner.

Moral obligations to take the first type of precaution—involving preparing to act—are less controversial and have not been the main targets of opponents of uncertaintism. Claiming that there's nothing morally good about making our important moral decisions soberly and carefully would be a puzzling rejection of common sense morality.[9]

But, accepting that we have obligations to prepare well for important moral obligations has important wider implications for the cases that do worry uncertaintism's critics. Preparatory obligations would serve no purpose if we were not also obligated to ultimately act in ways that reflect the morally understanding we have formed during that preparation. So, what actions we should ultimately take will depend upon the newfound beliefs that we developed while seeking advice or engaging in contemplation, for example. But the new beliefs, credences, or evidence we gain by preparing for our moral decisions are part of our moral-epistemic circumstances. So, if *what we figure out about morality* while preparing to act matters to what we should eventually do, then our moral-epistemic circumstances matter to what we should do.

This dynamic makes it very difficult to reject the core claim of uncertaintism—the claim that our beliefs, credences, or evidence about morality can matter to what we morally ought to do. Rejecting this would require saying either (i) that we never have moral obligations to try to figure out what is moral to do (even when making very weighty decisions) *or* (ii) that we may have obligations to go through the intellectual exercise of puzzling over morality but no corresponding moral obligations to act in ways that are informed by it. And either of these claims would be very surprising. So far, then, we have two arguments that our moral-epistemic circumstances sometimes matter to what we morally ought to do.

But, in spite of the support for uncertaintism in commonsense morality, objections to uncertaintism sometimes seem to be on to something. There are reasonable concerns that we should try to make sense of. In the next section, I'll suggest that these challenges are better understood as problems for particular, popular varieties of uncertaintism, rather than problems for uncertaintism as a whole. There are a wide variety of ways to develop uncertaintist positions. So, we can, and should, relocate a lot of these debates: let's talk about how and in what ways our moral-epistemic circumstances matter to what we morally should do—rather than whether they matter to it at all. In the next section, I'll identify three features that some criticisms tend to attribute to uncertaintism and spell out why the criticisms depend on these assumptions. But I suggest that these features aren't essential to uncertaintism, and I close the paper by offering an alternative form of uncertaintism that avoids these problems.

[9]The content of these preparatory obligations may, arguably, already depend on our moral-epistemic circumstances (if which preparations we should take depends on whether a decision seems high stakes to us or on whether we are uncertain, for example). For an argument that intuitively required advance contemplation requires de dicto moral motivation, see Aboodi (2017).

In what follows, I often use language suggestive of moral realism, for ease of discussion. But I take my argument to be compatible with a variety of metaethical views. I only need it to be the case that being moral is something we can be better or worse at (for reasons besides *akrasia*). If this holds, I'll suggest we morally ought to try to be "better at" morality, regardless of the underlying metaethical structure. Indeed, insofar as a metaethical view does not leave room for concerns about possible moral mistakes and efforts to avoid those mistakes, that seems, itself, to be a problem for that metaethical theory—an indication that it cannot account for a significant aspect of our moral lives.[10]

3) Criticisms of Uncertaintism as Criticisms of Only Specific Types of Uncertaintism

Some criticisms of uncertaintism focus on its implications for people who are confidently wrong. Brian Weatherson worries that uncertaintism licenses confident zealots to follow their convictions —after all, if what we should do depends on our moral beliefs, it might seem that the confidently mistaken should act on their wrong moral beliefs (having no uncertainty to hold them back) (Weatherson, 2019, p. 6). And Elizabeth Harman suggests that uncertaintism involves exculpating people who have done awful things while sincerely believing in their own righteousness (Harman, 2015).

But these objections to uncertaintism depend on attributing particular commitments to uncertaintism. Below I highlight three features that one or both of these criticisms attribute to uncertaintism: that it is committed to a decision-theoretic analysis of our moral responsibilities (possibly one involving expectation maximization, in particular); that it must involve moral-epistemic circumstances mattering to *all* moral norms—not only some of them; and that it must treat false moral beliefs as exculpating (in part because of these first two commitments). I suggest that none of these features is actually required by the core claims of uncertaintism. So, we can take these concerns on board while viewing them as reasons to adopt alternative versions of uncertaintism, rather than reasons to reject uncertaintism entirely.

A. Decision-theoretic morality

Criticisms of uncertaintism often presuppose that uncertaintism brings with it a decision-theoretic approach to morality—that is, an approach that uses formal models of rational choice under uncertainty (and, that, in its most traditional form, takes rational decisions to be those that maximize expected value, as with those authors in the paragraph below). I share critics' worry that deriving moral norms from decision-theoretic models can lead to implausible conclusions, but while decision-theoretic approaches have unfortunately dominated this literature, uncertaintism does not need to involve this.

Weatherson suggests that his arguments against uncertaintism (or, in his terminology, internalism) are centrally motivated by concerns about uncertaintist views that advocate a particular decision-theoretic approach: expectation maximization (Weatherson, 2019, p. 17). This expectation-maximizing approach has been influential in work on moral uncertainty, but also subject to criticism. On this type of view, when deciding how to act, we should maximize "expected moral value" (Lockhart, 2000), "expected objective value" (Sepielli, 2009, p. 11), or "expected choice-worthiness" (MacAskill, 2014, p. 20). Roughly, the expected moral value of an act can be calculated by taking the moral value of the act under each possible moral view, multiplying it by the probability (from the agent's point of view) that this moral view is correct, and then adding together these figures for each possible moral view.

[10]For discussion of metaethical non-cognitivism's capacity to handle normative uncertainty, see Makins (2022); Robinson and Steele (2023); Sepielli (2012); Staffel (2019).

Expectation maximization does yield troubling guidance in certain cases. When an action would be exceptionally bad or good according to an extremely unlikely moral theory, that unlikely possibility can come to dictate what we should do in counterintuitive ways. And, as Weatherson discusses, expectation maximization seems to support confident zealots who act on misguided convictions. One of Weatherson's key examples is Maximilien Robespierre, the French revolutionary, whose Reign of Terror seemed to stem from sincere conviction. He seemed to genuinely believe his actions very likely had extremely high moral value—and were very unlikely to be morally bad. So, his violence might appear expectation maximizing, and the uncertaintist who advocates expectation maximization might seem to encourage terrible acts in Robespierre's case.

One way for uncertaintists to avoid this is to adopt a form of uncertaintism in which what we should do depends partly on our moral evidence, rather our actual beliefs or credences—after all, Robespierre may not have been justified in his convictions. Harman and Weatherson each offer some discussion of this possible move by uncertaintists, and though I'm sympathetic to it (I'm partial to that evidence-centered variety of uncertaintism), my main focus here will be on the power of another reply.

A good response to Robespierre's case is to abandon expectation maximization, rather than abandoning uncertaintism. Expectation maximization is not the only way for our moral-epistemic circumstances to matter to what we should do. Concerns have been raised about the feasibility of expectation maximization in many cases of moral uncertainty, and a variety of substitutes have been discussed (e.g., MacAskill, 2016). And, in earlier work, I've argued that we should reject expectation maximization because it does not adequately account for factors that are important to how we navigate moral uncertainty (Rosenthal, 2021).[11] Indeed, I've made the case that there is a broader problem: decision-theoretic models on their own cannot settle what we morally ought to do or how we should navigate moral uncertainty because selecting or constructing a decision-theoretic model requires that we *already* have the answer to moral questions like "how much should we be willing to sacrifice to reduce our chances of wronging others?." The most that decision-theoretic models could hope to do is formalize the conclusions of these moral inquiries post hoc, rather than settling the relevant moral questions (I leave open whether this post-hoc formalization would be possible). So, even apart from the worries that Weatherson raises about cases like Robespierre's, there are some initial reasons to doubt whether expectation maximization—or decision-theoretic approaches to moral uncertainty in general—will be fruitful.

The good news is that uncertaintists do not have to be saddled with these approaches. Our moral-epistemic circumstances can matter to what we morally ought to do without any commitment to expectation maximization or other decision-theoretic approaches to morality. It is a thorny, open question whether the moral considerations that matter to what we should do can be flattened into a single formal model without losing much of the richness and complexity of our moral lives. I cannot hope to resolve that here. But I want to at least suggest that this question is orthogonal to the question of whether our moral-epistemic circumstances are some of the factors that matter to what we ought to do.

We see something like this with the ordinary, non-epistemic factors that matter to what we morally ought to do: it might be possible to account for them all in one calculation, in the way some utilitarians imagine—or it might not be possible, as many value pluralists suggest. And the same range of options is open for moral-epistemic factors that impact what we morally ought to do—accepting that these moral-epistemic factors matter to what we should do does not tell us whether decision-theoretic models are able to determine *how* they matter. Perhaps, for example, what's valuable about any particular way of navigating moral uncertainty might be irreducibly pluralist, and its good and bad features might be incommensurable, unable to be neatly weighed against each other—making decision-theoretic approaches ill-suited to guiding us through these circumstances. If this is true, Weatherson's Robespierre concern about uncertaintism will not seem especially

[11]For an opposing argument, though in the context of morally relevant, empirical uncertainty, see Zhao (2023).

relevant because ultimately its target is the counterintuitive moral recommendations of decision-theoretic models that have dominated uncertaintist theories, rather than objections to uncertaintism, itself. We can say that moral-epistemic circumstances matter to what we should do without signing onto models like these.

Once we recognize this, we can begin to identify a range of possible moderate and less worrying uncertaintist principles that would give a role to moral-epistemic circumstances, without generating the counter-intuitive implications that lead to many objections to uncertaintism. For example, Alexander Guerrero has argued in favor of a specific principle that he suggests supports vegetarianism ("Do not Know, Do not Kill") according to which, roughly, it's blameworthy to kill organisms in cases when we do not believe there's something morally important requiring it and we are unsure whether they have moral status (Guerrero, 2007, p. 79).

Like expectation-maximizing uncertaintist principles, Guerrero's "Do not Know, Do not Kill" says that what we should do depends partly on our moral-epistemic circumstances—it's squarely an uncertaintist principle. But Guerrero's principle gives much more cautious advice, telling us to handle uncertainty by steering clear of some actions, just in case they are wrong—without offering any direct support for acting from zealous convictions. Uncertaintists can give a range of sensible, moderate advice like this: for example, "reduce moral mistakes by making important moral decisions under conditions that foster clear thinking (and avoid making these decisions while drunk, rushed, or sleep-deprived)—then act in ways that reflect this clearer thinking"; or "when unsure what's morally right, seek input from impacted parties with opposing views and work to understand the reasons for their positions—then act in ways that take that understanding into account." With each of these principles, we are told to act in ways that depend on our moral-epistemic circumstances—ways that take into account the views we form when seeking input or thinking clearly. And in examples like these, the uncertaintist can actively counsel against rashness. In short, our moral-epistemic circumstances can matter to what we should do in various ways, without them mattering in the specific, expansive ways that support confident zealots—uncertaintists do not have to be committed to anything that troubling. And, once we recognize this, much of the discomfort with uncertaintism's recommendations loses its force. As an added benefit, these moderate uncertaintist recommendations may be less vulnerable to the worries about intertheoretic value comparisons that have plagued many uncertaintist proposals, because implementing the recommendations does not need to require the calculations involved in expectation maximization (though exploring this possibility fully is outside the scope of this paper).[12]

B. No other moral norms

Let's turn now to a second feature that critics tend to attribute to uncertaintism. I'll ultimately suggest that, in this case also, we can retain the core commitments of uncertaintism while avoiding the feature that's causing trouble. Weatherson's criticisms of uncertaintism (in his terms, "internalism") seem to be targeting versions of uncertaintism in which our moral-epistemic circumstances matter to *every* moral norm and in which there are no moral norms that are independent of our moral-epistemic circumstances.

If we attribute this feature to uncertaintism, it closes off one promising way for uncertaintists to avoid the problems with cases like Robespierre. People often believe they are doing the right thing, while, in fact, engaging in terrible acts. And we may want to say they have still done something morally wrong—perhaps something horribly wrong—despite their beliefs that they are acting rightly. As we have seen, Weatherson worries that uncertaintists cannot say this. On his view, uncertaintists are committed to saying that these well-intentioned actors have not done anything

[12]For challenges to uncertaintism based on intertheoretic value comparison, see Gracely (1996), Gustafsson and Torpman (2014), and Hedden (2016).

wrong because they have followed the guidance their moral-epistemic circumstances provided (Weatherson, 2019, p. 84).

But one way for uncertaintists to avoid this is to adopt a moderate version of uncertaintism, in which our moral-epistemic circumstances only matter to *some* moral norms. The space of moral norms may simply be more crowded than Weatherson's picture suggests. It's possible to say, on the one hand, that the well-intentioned person who causes harm did some things morally right: depending on the details, perhaps they really did navigate their moral-epistemic circumstances well, given what they were working with, and really minimized their risk of wrongdoing as best as they could—and those things genuinely are morally important. But, on the other hand, they also got something very important morally wrong—despite their best intentions, they did the substantively wrong thing. The first set of norms can be sensitive to their moral-epistemic circumstances, without the second, substantive set of norms also being sensitive to their moral-epistemic circumstances. This is much like saying that a trial met due process standards, but reached the wrong verdict (e.g., convicting an innocent person). Once we recognize that there can be multiple moral norms that are relevant to our actions, it's open to uncertaintists to say, both, that someone navigated their moral-epistemic circumstances well and that they did something substantively wrong. (I'll sketch a view along these lines in Section 4, below.) So, here, again, uncertaintism does not have to come with the features that would lead to uncritical endorsement of confident zealots.

C. Exculpating the morally ignorant

Finally, Elizabeth Harman has argued that uncertaintists are committed to exculpating people who act on false moral beliefs, not because of particular features of decision-theoretic approaches to moral uncertainty but simply because of the way that even epistemically responsible moral credences can lead us astray. Harman argues uncertaintists are committed to the view that moral ignorance exculpates when people act badly while believing in their own righteousness (or, in an evidence-focused variation, when they act badly while *epistemically justified* in believing they are acting rightly).[13] According to Harman, this commitment to exculpate follows from the combination of two claims: uncertaintists' claim that what we should do depends on our credences, and the claim that we cannot be blameworthy if we have done what we should. If we should do the things we *believe* are good, for example, then it appears that we cannot be blamed for doing exactly that, even if our beliefs are mistaken.

But uncertaintists do not need to be committed to the view that moral ignorance exculpates[14] if we allow that multiple norms might apply to our actions—only some of which depend on our credences or evidence. Our ignorance or misleading moral evidence might mean that we aren't to blame for violating some of the norms—after all, we may have complied with the norms concerning how we navigate moral uncertainty. But there may be other important norms—e.g., concerning the substantive rightness of our actions—that we have still violated. And, we might be blameworthy for those violations. (I'll leave open whether complying with some norms should reduce our blameworthiness for violating other norms—uncertaintists could take a variety of positions on this.)[15]

Harman considers one related possibility—that uncertaintists are only talking about a specific, narrow reading of ought claims, and people might be blameworthy under other senses of ought (Harman, 2015, pp. 70–73). According to Harman, this avoids leaving misguided wrongdoers

[13]There are a wide variety of arguments made against the idea that moral ignorance is exculpatory (including by Weatherson in Chapter 5 of *Normative Externalism*), but my focus here is specifically on Harman's argument because she directly makes the case that uncertaintists are saddled with this commitment. It's, of course, an open question whether this commitment is, itself, a problem (see e.g., Rosen (2003) for an argument that moral ignorance is, indeed, exculpatory).

[14]For additional replies to Harman, see Sepielli (2017).

[15]I am, as it happens, sympathetic to the possibility that this dynamic might lead to somewhat reduced blameworthiness in certain types of cases, but uncertaintists who are keen to avoid this are able to do so without any inconsistency.

blameless, but it does this at the cost of being an uninteresting position that just carves out idiosyncratic readings of our ought claims (Harman, 2015, pp. 70–73). We could ask what we ought to do, taking into account *only* our own moral views and nothing else, but we could also ask what we ought to do taking into account only advice gleaned from the soundtracks of Wes Anderson films. And neither question seems that important.

But my suggestion here can avoid Harman's charge of uninterestingness. I'm not offering multiple interpretations of ought claims or asking what we should do under idiosyncratic limitations. Instead, I'm suggesting that we often have a variety of moral responsibilities that apply to the same decision, in much the same way that legal decision-making can involve procedural responsibilities alongside responsibilities of substantive justice. (I defend this possibility in section 4, below.) If this is true, then the norms the uncertaintist is discussing aren't quirky, alternative readings of our moral responsibilities but additional, important responsibilities that we have, much like due process requirements. They're straightforwardly important because they are some (though not all) of our moral responsibilities—and considering them may also help to explain a variety of specific moral obligations (e.g., if we ought to be tolerant because it is a good way to navigate moral uncertainty and disagreement, as I suggested in Rosenthal, 2019, pp. 60–87). Allowing for multiple types of moral norms, with moral-epistemic circumstances mattering for some but not others, seems to provide a way to avoid Harman's concerns about exculpatory moral ignorance.

Ultimately, then, it looks like the force of many challenges to uncertaintism depends on the type of uncertaintism we adopt. Uncertaintism does not need to embrace expectation maximization or other decision-theoretic models of our moral responsibilities; it does not need to involve moral-epistemic circumstances shaping all moral norms; and (because of this) it does not need to treat moral ignorance as exculpatory. As a result, Harman and Weatherson's challenges may give us a reason to favor some uncertaintist views over others, rather than a reason to suppose that our moral-epistemic circumstances do not matter to what we should do.

4) An Alternative, Moderate Uncertaintism: Substantive and Procedural Oughts

I want to turn now to offering a positive account that builds on this recognition that uncertaintism can take many forms and may be able to avoid key challenges as a result. I develop a structure for uncertaintist accounts in which our moral-epistemic circumstances matter to some norms, but not others—and provide independent motivation for this type of uncertaintism. I'll argue that our moral responsibilities in the face of our moral uncertainty and fallible moral judgments include both substantive requirements and procedural requirements, mirroring the structure of norms of substantive justice and legal due process norms. Rather than competing with each other, these norms address different moral questions—answering procedural moral questions will require information about our moral-epistemic circumstances and answering substantive moral questions may not.

In Section 2, I defended one aspect of this: navigating moral uncertainty recklessly is morally bad and this reveals that our moral-epistemic circumstances sometimes matter to what we morally ought to do. But, why think that there are both moral norms that *do not* need to be sensitive to our moral-epistemic circumstances and these moral norms that *are* sensitive to them?

It's difficult to understand the purpose of uncertaintist norms against moral recklessness without this structure, and an account with both types of norms makes the best sense of ordinary processes of moral deliberation. To see this, let's look at the role each type of moral norm is playing. Sometimes we are unsure of the content of one set of norms (Geraldine, for example, is unsure whether she should break her promise of confidentiality), and our moral judgments are always fallible. At the same time, we usually behave as though there *are* ways we can do a better or worse job navigating that kind of uncertainty and fallibility (both through what we do in advance and by ultimately making decisions that are shaped by that preparation). Depending on the case, this could include considering the issue carefully, seeking input from people who may have wise advice or

insight into relevant perspectives, avoiding morally questionable actions when more clearly permissible alternatives exist, or avoiding making high-stakes moral decisions while drunk or distracted. My aim here is not to settle the best strategies. Instead, I'm suggesting that if we deny there *are* better and worse approaches—if we deny there *are* norms for how we should navigate our uncertainty given our moral-epistemic circumstances—that would be a radical departure from how we ordinarily think about moral decision-making.[16] And we treat these norms as making moral claims on us: if Geraldine decided whether to reveal her friend's drug use by deferring to a parrot, it would be reasonable for her to feel ashamed or for her friend to resent her—even if the parrot led her to a good decision. (Keep in mind that norms about how we make decisions need to come with requirements that we implement those decisions when acting.)

This, together with my arguments in Section 2, already suggests there are some norms that make moral claims on us that are sensitive to our moral-epistemic circumstances—norms concerning what strategies we should use for trying to be moral. But our ordinary practices also recognize the norms that we aim to satisfy with those strategies—which might *not* be sensitive to our moral-epistemic circumstances. When Geraldine worries about making the right decision about whether to break her promise, she *is not* thinking that as long as she does what seems best from her limited moral-epistemic circumstances, she will have behaved morally rightly. The reason Geraldine bothers to consider carefully, in the first place, is because she thinks there's some further set of norms that are important to comply with, but whose content she's uncertain of. Indeed, it's difficult to understand the point of deliberating carefully otherwise.

So, it looks like our ordinary ways of thinking about decision-making under moral uncertainty reflect two sets of norms: some first-order moral norms (that might *not* be sensitive to our moral-epistemic circumstances) and some higher-order norms (sensitive to our moral-epistemic circumstances) that tell us what strategies we should use for trying to comply with those first-order norms. In Section 2, we saw arguments that there are some moral norms concerning the strategies we should use for trying to be moral, that are sensitive to our moral-epistemic circumstances (e.g., norms against recklessness in the face of uncertainty). That these norms make moral claims on us was supported by (i) considerations of respect for those we might be wronging and (ii) the relationship between moral obligations to prepare to act and moral obligations to act in ways that actually reflect that preparation. Now, we have seen that those norms regarding good strategies for trying to be moral cannot exist on their own. Just as obligations to prepare to act only make sense alongside obligations to act in ways that reflect the preparation, these higher-order strategic norms only make sense if they exist alongside first-order norms that they aim to satisfy. Without this, the good strategies would have no purpose and our ordinary ways of thinking about moral decision-making would be difficult to account for. Recognizing both sets of norms does not just avoid particular criticisms of uncertaintism—it's the only way for uncertaintism to adequately motivate the efforts to do the right thing that are central to the view.

I'll refer to these two types of norms as procedural oughts and substantive oughts:

> **Procedural oughts** are higher-order moral norms that tell us that we ought to act in particular ways, on the basis that doing so is the moral way to go about trying to satisfy *another set of moral norms*, given the challenges we face in discovering their content (e.g., our moral uncertainty and our fallible moral judgments). Roughly, these are "moral norms about how to pursue other moral norms" or "strategic norms." Examples could include "act based on careful moral contemplation" or "avoid actions whose moral rightness is unclear, if they can be avoided at low cost."

[16]For a discussion of efforts to manage moral risks in the face of moral luck, see Bagchi (2011).

Substantive oughts are all nonprocedural, moral norms (that is, those moral norms that do not concern how to morally pursue another set of moral norms). Examples could include "don't lie" or "maximize utility."[17]

(For ease of discussion, I'm using "oughts" and "norms" interchangeably.)

Norms of both types make moral claims on us. The contrast between them can be thought of as paralleling the contrast, in legal contexts, between norms of substantive justice (e.g., a trial should convict only the guilty) and norms of due process (e.g., the trial should give the defendant an opportunity to present arguments in their defense). I'll draw on this parallel to highlight two key features of the account.

First, multiple norms (substantive and procedural or substantive justice and legal due process) can apply to the same act, showing different responsibilities that we have when acting. This is already commonplace in legal contexts. To see this, notice that (like strategies for being moral) legal due process norms can have implications for *both* the processes that are used *and* the ultimate decisions that are made. For example, if due process requires that hearsay evidence is not presented to the jury, then due process *also* requires that we only reach and enforce verdicts that aren't based on hearsay evidence. Following a process involves adhering to the results of that process.

Complying with these due process norms might coincide with substantive justice. After all, part of the reason we (typically) do not allow juries to convict based on hearsay evidence is that prohibiting (potentially unreliable) hearsay evidence is a good strategy for avoiding wrongful convictions (that is, we hope this strategy makes us more likely to achieve substantive justice). Convicting based on hearsay evidence might run a reckless risk of wrongful conviction in a way that's unfair to the defendant.

But avoiding recklessness and using good strategies will not always be successful. Sometimes we can follow good strategies and still end up doing something substantively unjust (e.g., convicting the innocent). Likewise, sometimes we can bypass good strategies and be lucky enough to end up achieving substantive justice anyway. But, when this occurs, the procedural violations would still, themselves, be wrong—it would still be unfair to have subjected the defendant to an increased risk of wrongful conviction, even if the conviction was ultimately avoided.

Examples like this help to show that due process norms and norms of substantive justice can make claims about the same decisions while bringing out different responsibilities that we have when acting: we have a responsibility to adhere to due process norms and we also have a responsibility to do what's substantively just. This is not an exotic metaethical claim—it's just the ordinary way that we already think about our responsibilities in legal settings.

Procedural and substantive oughts about morality are similar to legal norms in this respect: one decision can be subject to multiple norms including obligations to do the substantively right thing and also obligations to use (and act on) responsible strategies for trying to. So, in the example above, Geraldine might worry about whether revealing her friend's drug use is substantively right. And she has to decide what strategies to use (or not use) to navigate that uncertainty—those strategies are what she procedurally ought to do. Procedural oughts can apply both to how we should position ourselves to act well (e.g., how much trouble we should go through to seek more moral knowledge) and to questions about what course of action to pick in light of the (perhaps still incomplete)

[17]Some strategies for satisfying moral norms might run afoul of first-order, substantive moral norms as well. For example, torturing someone to extract their moral insights would be an immoral strategy for trying to be moral, but the *basis* of norms against torture is e.g., the pain caused by torture or the degrading nature of torture itself. The impact of torturing on our ability to fulfill some other set of moral norms (not related to torture) will not be what generates reasons not to torture (even if those reasons incidentally apply to what we do to try to find out about morality). For this reason, torturing someone to extract their moral insights would be a violation of substantive norms (or norms of both types, if e.g., it's also a flawed strategy). Ultimately, though, the location of the line between substantive and procedural oughts is less important for my argument than the fact that there are these two sets of norms, one that is sensitive to our moral-epistemic circumstances and one that is not.

knowledge we have acquired (e.g., should we reduce our chances of wrongdoing by hedging our bets and avoiding a potentially wrong act, just in case). On this view, violating procedural oughts is, itself, morally bad, even if a substantively good decision gets made, and following procedural oughts will not always ensure substantively good actions—just as legal due process violations are wrong, even if a substantively just verdict is reached, and complying with due process norms does not guarantee substantive justice. (See footnote 18 for discussion of the relationship between this framework and work on descriptive uncertainty that distinguishes actions satisfying the criteria of right action from actions resulting from following the best decision procedures.)[18]

This interplay between different types of norms also helps show a second, key feature of the account that I'll turn to now. Substantive and procedural oughts are distinct norms that cannot, by themselves, settle each other's content—they are addressing different moral questions, and moral-epistemic circumstances will matter to the answers in the case of procedural norms, even though they do not have to impact substantive norms. To bring this out, let's continue the due process analogy. The exact extent to which due process norms are shaped by the goal of achieving substantive justice is a subject of disagreement—and the right answer may depend upon the particular due process norm. But, for many legal due process norms, it's both the case that the procedural requirements are designed to generate substantively better outcomes, and that this goal will not, on its own, settle what due process requires.

For example, substantive justice requires convicting only the guilty, but processes that reduce the risk of mistaken convictions can be costly in other ways (e.g., because of financial expense or increased jury service requirements). Our actions toward a defendant may ultimately be determined by whether we accept only unanimous jury verdicts in criminal trials (our current practice) or convict defendants based on, for example, a two-third supermajority of jurors

[18]Frank Jackson (1991) has argued that when trying to act morally under conditions of *descriptive* uncertainty, we should do what would "maximiz[e] … expected moral utility" (464) given our epistemic constraints, and he briefly acknowledges that this account may require multiple moral oughts, including the decision-theoretic norms he is discussing and norms regarding what would be best from an omniscient perspective (471). Because Jackson is discussing descriptive uncertainty (and presupposing that consequentialism is the correct moral theory), a defense of my procedural/substantive oughts account requires different considerations (the opponents of uncertaintism I've been engaging with generally already accept that uncertainty about descriptive facts can matter to what we morally ought to do, but argue that uncertainty about moral norms should be treated differently). My aim here is not to try to address the substance of Jackson's account, but only to clarify where there are and aren't analogies available between his treatment of descriptive uncertainty and my treatment of moral uncertainty, for readers who may find this helpful in situating my proposal. I share Jackson's view that the best strategies for trying to be moral might not satisfy objective criteria for rightness. While Jackson's decision-theoretic consequentialist oughts are, like procedural oughts, concerned with the strategies we should use for navigating epistemic limits, his approach to their content differs significantly (beyond his assumption of consequentialism). In part, this is because I reject decision-theoretic frameworks for determining what we procedurally ought to do (as I've argued in 3A). But, a deeper difference in our approaches to uncertainty is this: because, for Jackson, the best strategies for being moral are determined by decision-theoretic models, there does not seem to be room in his account for the best strategies to depend on further—distinctively procedural—moral questions: for example, questions like how can we morally distribute the risk of being a *victim* of our wrongdoing, and how large of a personal sacrifice that's certain should we be willing to make in order to avoid potential wrongdoing that's uncertain. By contrast, I take the content of procedural oughts to centrally depend on how we resolve moral questions like these (not just on importing a decision-theoretic model).

For more recent discussion of a different type of two-tier system of moral norms, see Smith (2018). Holly Smith, like Jackson, is concerned with uncertainty about moral decision-making in the context of uncertainty about *descriptive facts* rather than uncertainty about moral norms and develops an in-depth, two-tier framework for such cases. Ultimately, the approach I defend has a different structure from what Smith develops though: I defend an account in which procedural oughts aren't just ways of indirectly using substantive norms—instead, complying with procedural oughts is, in itself, an important moral responsibility (analogous to the responsibility to adhere to due process norms in a legal setting) and upholding that responsibility reflects respect for other people.

For a discussion of the possibility of a two-tier system of norms of rationality, rather than morality, in which there are norms regarding the rationality of attitudes (like belief) and separate norms regarding the rationality of processes of forming those attitudes, see Thorstad (2024, pp. 398–399).

(which might avoid some retrials, and so modestly reduce the court resources and number of jurors needed). Knowing that substantive justice requires convicting only the guilty will not settle how to trade-off increased risks of wrongly convicting a defendant against costs in time or money that are more certain. There are distinct normative questions about the trade-offs we should make and the strategies we should use for pursuing substantive justice—questions like "How much time and resources must we be willing to give up in order to achieve a small reduction in the risk of wrongful conviction? What about a large reduction in that risk?" These are procedural questions that require answers, and that aren't addressed by the substantive requirement to convict only the guilty.

The moral context has a similar structure. Procedural oughts can address questions about what strategies we morally ought to use when trying to satisfy substantive oughts—questions like, "if I'm morally uncertain whether taking an exciting career opportunity would make me unacceptably complicit in unethical behavior, should I decline the job just in case?"[19] or "if my moral uncertainty leaves me unsure whether to take a particular action, should I spend time seeking advice and additional perspectives (and then act in ways that take that input into account)—or can I shrug and 'wing it'?"[20] Substantive oughts are silent on these questions. They address what we substantively ought to do, not what hard trade-offs to make when we're not sure what we substantively ought to do—moral questions like these are the domain of procedural oughts.

I will not aim to exhaustively settle the factors that determine procedural norms here, but some examples will be illustrative. The likelihood that a particular action is substantively wrong will matter to procedural norms, yes—but it's not the only consideration that's relevant. What I procedurally ought to do—that is, how I should go about trying to do the substantively right thing (given my uncertainty and fallibility)—might depend, for instance, on how I should weigh possible substantive norms when they have different levels of plausibility and severity, or how much I should be willing to sacrifice in order to modestly reduce my risk of substantive wrongdoing. And they may require knowing which options were met with skepticism when I sought advice from wise friends, for example, or which actions seem, to me, like they might stray into a moral grey area—information about my moral-epistemic circumstances (that is not relevant to what I *substantively* ought to do).

Our moral-epistemic circumstances matter to the answer to procedural questions, like those about whether to seek additional perspectives and advice. The questions, themselves, may not arise under idealized conditions of infallible moral judgment and complete moral certainty. And determining, say, whether we should spend time getting additional input may depend on what we believe so far, for example. But this only means we are taking into account the information that's most relevant for answering our initial question—and that happens to include some information about our moral-epistemic circumstances. Because procedural and substantive oughts depend on different information, it's possible for one, but not the other, to be sensitive to our moral-epistemic circumstances.

Understanding that substantive and procedural oughts are answering different moral questions (and relying on different information to do so) can also help show why procedural oughts aren't the same as "subjective oughts," at least not in the usual sense. I'm not primarily concerned with policing how far "subjective oughts" terminology from subjective/objective oughts discussions can

[19]Avoiding the job opportunity just in case it's morally wrong would be a strategy for trying to comply with substantive moral norms, while uncertain about their content. So, procedural oughts would tell us whether that is a strategy that we morally ought to use.

[20]Whether we procedurally ought to seek this type of advice may depend partly on whether it's appropriate to revise our credences in a moral view based on moral testimony and peer disagreement. This is, itself, the subject of disagreement among moral epistemologists. (For overviews of this debate, see Cosker-Rowland (2017) and Hills (2013).) Resolving the debate is outside the scope of this paper, but for readers who believe we should sometimes revise our credences in light of others' moral testimony, this can serve as a helpful example of something we procedurally ought to do.

be made to stretch, and I will leave open whether to characterize procedural oughts as special cases of subjective oughts or a fully distinct phenomenon. But there are important reasons not to run procedural oughts together with typical subjective oughts discussions because the features that distinguish procedural oughts from typical subjective oughts are also features that make the content of procedural oughts "interesting" in Harman's sense and important to attend to.

First, though less relevantly here, procedural oughts can be specified in more and less subjective ways, depending upon whether the procedural oughts take into account our actual beliefs and credences or only our evidence. And Amelia Hicks has even argued that belief-sensitive norms about how we should navigate moral uncertainty are all, in fact, objective (Hicks, 2022).

Second, more crucially, objective/subjective oughts debates are typically concerned with different interpretations (objective and subjective interpretations) of the same norm, addressing the same question—say, should a doctor give a patient the medicine in bottle A—but answering that question with different limitations on the information used (roughly, subjective oughts draw only on the credences of the person acting, and objective oughts draw on all relevant information). In contrast, substantive and procedural oughts aren't different interpretations of the same obligations or answers to the same moral question with different restrictions on our information. Instead, substantive and procedural oughts identify different moral obligations and address different questions, using the information that is relevant to those questions. We ought to do the substantively right thing, and *also* we ought to be moral in the way we go about trying to do the substantively right thing. So, showing why procedural oughts are interesting or important (in spite of Harman's concerns) is straightforward: we should care about procedural oughts simply because they are additional obligations that we have, and we genuinely wrong people if we violate them.

The parallel to legal processes is again illuminating. We've seen that the epistemic circumstances of a trial impact what procedures we should use—for example, we need to consider how likely different procedures are to yield correct verdicts, given that we do not know who has actually committed a crime. But saying that a requirement of juror unanimity gives a subjective interpretation of our duty to convict only the guilty, for example, would be misleading. It obscures the way in which these due process norms are duties of their own ("interesting" in Harman's sense, as a result) and address distinct questions, using the information relevant to those questions.

Similarly, when what we procedurally ought to do depends on our moral-epistemic circumstances, this only means we are taking into account the information that's most relevant for answering our procedural questions—and that happens to include some information about our moral-epistemic circumstances. This is not a case where the most relevant information is excluded in order to focus on our own perspective, as it is with typical subjective oughts. Given this, procedural oughts are worth distinguishing from the norms that subjective oughts terminology usually identifies. Procedural oughts are their own moral norms, addressing different moral questions than substantive oughts, and using the most relevant information to do so.

Before moving forward, it's worth briefly addressing a couple of possible concerns about the structure that I've been proposing. Although substantive and procedural oughts depend on different factors, it might still seem surprising that substantive and procedural oughts can directly contradict. What one forbids, the other may require, or vice versa (e.g., if our best strategy for reducing moral risks turns out to lead us to act in ways that are substantively morally wrong, despite our best efforts). But it's already recognized that conflicting moral norms can sometimes bear on the same physical act, as with an act that is both required by legitimacy and forbidden by norms of substantive justice, for example. And, as Andrew Sepielli has noted, in practice in morality, we typically only receive guidance from oughts at one level at a time (2017, p. 110). Put in the language of my account, we do not normally encounter the demands of substantive morality and then decide how to handle conflicts with procedural norms—if we genuinely knew exactly what substantive morality required, then there would be no conflict with procedural norms because the best procedure would involve using that knowledge of substantive morality. But in real life, we cannot be sure if we have that substantive knowledge, so we are often dealing with procedural oughts.

Along similar lines, these conflicts mean that we may sometimes be unable to satisfy all of morality's demands because we face conflicting requirements. But this is a variation on something that is already commonplace. The conflicting moral requirements only arise because we lack full moral knowledge, and if gaps in our moral knowledge make it difficult for us to comply with all of morality's demands, that's an ordinary and unsurprising predicament.

Recognizing that guidance from substantive and procedural oughts can point in different directions also helps address a problem facing obligations to use good procedures: even the very best procedures sometimes lead us to act badly. If these procedures were our only obligations, that might seem to make a lot of bad actions permissible. Indeed, concern over this possibility motivates many critics of uncertaintism, as we saw in Section 3. But once we recognize that there are also substantive moral norms that may give different guidance, this concern loses much of its bite. If we have done something *morally right* by following good procedures and nevertheless been led astray, there is still something important (and, arguably, more important) that we have done *morally wrong*—we violated substantive moral norms.

So far then, we have a structure in which our moral epistemic circumstances impact what we morally ought to do, without necessarily shaping every moral norm (e.g., they might matter for what we procedurally ought to do without mattering for what we substantively ought to do). This avoids two of the "problem features" we examined above. It directly removes the assumption that moral-epistemic circumstances must matter to all moral norms, if they matter to any. And, indirectly, it addresses the related worries about exculpating people based on their moral ignorance that depended on that assumption. If we adopt my earlier recommendation (in Rosenthal (2021)) to reject decision theory as a guide for navigating moral uncertainty, we can also avoid the remaining "problem feature" (decision-theoretic approaches). With this alternative structure for thinking about moral uncertainty on the table, I want to turn now to briefly addressing a final worry that might arise about this type of uncertaintism.

5) Infinite Regresses

This framework can invite concerns about infinite regresses. I've been focusing on uncertainty and fallible judgment about substantive morality's content. But we could also ask how to handle uncertainty and fallibility regarding what second-order procedural norms require.[21] Questions about procedural norms are hard, and it's by no means obvious what the correct second-order, procedural account calls for. So perhaps we'd need a third-order theory of how to approach these second-order questions, and a fourth-order theory to address uncertainty about the right third-order theory—with the theories multiplying ad infinitum.

Procedural norms, so the worry goes, are supposed to be action guiding—they are supposed to sort out what we should do when we are mired in uncertainty.[22] But this infinite pile of third-, fourth-, or seventeenth-order theories may seem to fail at that. So, you might wonder why we are bothering with anything beyond the first-order, substantive questions at all.

But this is less of a problem than it seems because these layers of norms do not have to prevent reasonable and tractable moral guidance in most cases but also because, if my argument so far is right, procedural norms are important regardless of whether they are action-guiding. We need a theory of procedural oughts simply because we need theories about the ways in which our actions

[21]See also Andrew Sepielli's treatment of related issues (though in the context of norms of rationality) in Sepielli (2014); and MacAskill's approach to addressing this objection (2014, pp. 217–219).

[22]See Hicks (2022) for an uncertaintist view that incorporates strong requirements that all moral norms be able to actually guide action.

might fail to be moral—and failing to follow procedural norms are some of those ways. But I also do not need to give up action guidance so quickly for two reasons.

First, action guidance does not require settling all questions about what to do, with no room for further uncertainty about whether the guidance, itself, is correct. No theory of procedural norms can do this because no theory can do this at all. But we need a working theory of procedural norms to move forward when we face substantive uncertainty, even if that working theory is not the last word. And this type of deliberative role for procedural norms is suggested by the way that actual deliberative processes reflect distinct substantive and procedural norms.

Second, while there may be infinite higher-order, procedural questions, we usually do not need to get mired in questions about seventeenth-order norms just to respect other people and take care not to wrong them. We can often do that without any attention to these very high order norms at all. There are third-order, procedural norms that (at least sometimes) require us to carefully try to comply with second-order norms and explore their content, precisely because this does not seem like wasted effort—we are sometimes able to identify actions that would be bad strategies for being moral and adjust our behaviors. (Indeed, sometimes we are so unsure about first-order, substantive norms that we *must* attend to second-order norms to figure out what to do at all.) But it's not clear that investing additional time in eighth- or ninth-order norms would be morally called for in a similar way. It's very unlikely that the effort would yield important new guidance that gave us reasons to materially change our behavior. So, without evaluating the content of eighth-order norms, for example, we might say that ninth-order norms (about how we approach eighth-order norms) do not require us to cautiously try to determine the content of eighth-order norms and increase our chances of complying with them—the best, ninth-order procedures might be ones that involve ignoring eighth-order norms! And we could generalize this to say that, past a certain level, it is not the case that we morally ought to attend to these higher-order norms (at least barring special circumstances that provide evidence to the contrary).

Even if we do not need to attend to these norms, are there still infinitely many levels of higher-order moral norms that we might be violating? In principle, there can be. But this need not mean that the higher-order moral claims at stake are especially weighty or especially easy to transgress. Indeed, being completely inattentive to these higher-order norms will often be all we need to do.

6) Conclusion

In sum, then, there are common sense reasons to think that what we morally should do sometimes depends on our moral-epistemic circumstances—that is, there are reasons to be uncertaintists. Respecting others includes reducing our risk of wronging them. And, if it's morally bad to make morally weighty decisions without taking steps that position us to act well (e.g., considering the relevant values carefully and not deciding drunk)—then this requirement only makes sense if we also ought to act in ways that reflect this preparation and decision-making. While a number of significant concerns have been raised about uncertaintism, I've argued that these concerns are better seen as problems for specific varieties of uncertaintism—concerns that give us a reason to adopt alternative varieties of uncertaintism, rather than abandon it. The structure of substantive and procedural oughts that I've offered provides the framework for one type of alternative that is also an independently more attractive approach to uncertaintism.

Examining procedural oughts also suggests that we have a large body of moral responsibilities that are under-explored. Looking more closely at procedural norms opens the door for making better sense of a variety of moral and political phenomena, perhaps allowing us to explain the value of practices like interpersonal toleration, or exceptions to deontological rules in high-stakes cases, in part through their capacity to reduce moral risk, for example (Rosenthal, 2018; Rosenthal, 2019). And, when we see reckless navigation of moral uncertainty as, itself, morally bad, this begins to

show how figuring out what we should do under moral uncertainty may require very different conversations than the ones that have dominated this literature. To identify the best procedures for navigating moral uncertainty, we will need to engage directly with questions about the nature of respect and what sacrifices we should make to reduce our risk of wronging others (although resolving these questions is far outside the scope of this paper).[23]

Chelsea Rosenthal is an assistant professor of philosophy at Simon Fraser University. Her research focuses on ethics, political philosophy, and philosophy of law.

References

Aboodi, R. (2017). One thought too few: Where *De dicto* moral motivation is necessary. *Ethical Theory and Moral Practice*, 20(2), 223–237. https://doi.org/10.1007/s10677-016-9742-5

Bagchi, A. (2011). Managing moral risk: The case of contract. *Columbia Law Review*, 111(8), 1878–1931.

Bykvist, K. (2020). Consequentialism, ignorance, and uncertainty. In D. Portmore (Ed.), *The Oxford handbook of consequentialism* (pp. 310–330). New York: Oxford University Press. https://doi.org/10.1093/oxfordhb/9780190905323.013.8

Cosker-Rowland, R. (2017). The epistemology of moral disagreement. *Philosophy Compass*, 12(2), e12398. https://doi.org/10.1111/phc3.12398

Field, C. (2019). Recklessness and uncertainty: Jackson cases and merely apparent asymmetry. *Journal of Moral Philosophy*, 16(4), 391–413. https://doi.org/10.1163/17455243-20182687

Gracely, E. J. (1996). On the noncomparability of judgments made by different ethical theories. *Metaphilosophy*, 27(3), 327–332. https://doi.org/10.1111/j.1467-9973.1996.tb00212.x

Graham, P. (2010). In defense of objectivism about moral obligation. *Ethics*, 121(1), 88–115. https://doi.org/10.1086/656328

Guerrero, A. A. (2007). Don't know, Don't kill: Moral ignorance, culpability, and caution. *Philosophical Studies*, 136(1), 59–97. https://doi.org/10.1007/s11098-007-9143-7

Gustafsson, J. E., & Torpman, O. (2014). In Defence of my Favourite theory. *Pacific Philosophical Quarterly*, 95(2), 159–174. https://doi.org/10.1111/papq.12022

Harman, E. (2015). The irrelevance of moral uncertainty. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 10, pp. 53–79). New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198738695.003.0003

Hedden, B. (2016). Does MITE make right? In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 11, pp. 102–128). New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198784647.003.0005

Hicks, A. (2019). Moral hedging and responding to reasons. *Pacific Philosophical Quarterly*, 100(3), 765–789. https://doi.org/10.1111/papq.12274

Hicks, A. (2022). Non-ideal prescriptions for the morally uncertain. *Philosophical Studies*, 179(4), 1039–1064. https://doi.org/10.1007/s11098-021-01686-1

Hills, A. (2013). Moral testimony. *Philosophy Compass*, 8(6), 552–559. https://doi.org/10.1111/phc3.12040

Jackson, F. (1991). Decision-theoretic consequentialism and the nearest and dearest objection. *Ethics*, 101(3), 461–482. https://doi.org/10.1086/293312

Jackson, F., Smith, M., & Journal of Philosophy, Inc. (2006). Absolutist moral theories and uncertainty. *Journal of Philosophy*, 103(6), 267–283. https://doi.org/10.5840/jphil2006103614

Johnson King, Z. (2020). Praiseworthy motivations. *Nous*, 54(2), 408–430. https://doi.org/10.1111/nous.12276

Lockhart, T. (2000). *Moral uncertainty and its consequences*. New York: Oxford University Press. https://doi.org/10.1093/oso/9780195126105.001.0001

MacAskill, W. 2014. *Normative Uncertainty*. PhD diss., Oxford University.

MacAskill, W. (2016). Normative uncertainty as a voting problem. *Mind*, 125(500), 967–1004. https://doi.org/10.1093/mind/fzv169

MacAskill, W., Bykvist, K., & Ord, T. (2020). *Moral uncertainty*. New York: Oxford University Press. https://doi.org/10.1093/oso/9780198722274.001.0001

[23]For feedback and discussion of this project (or its distant ancestors), I am grateful to Amanda Askell, Ben Bagley, Nic Bommarito, Claire Field, Jane Friedman, B.R. George, Alex Guerrero, Elizabeth Harman, Joe Harper, Zoë Johnson-King, Joanna Langille, Dan Layman, Adam Lerner, Matthew Liao, Sandra Lindgren, Alex London, Rob Long, Tienmu Ma, Will MacAskill, Jordan MacKenzie, Ishani Maitra, Kris McDaniel, Emilio Mora, Liam Murphy, Tom Nagel, Jake Nebel, Japa Pallikkathayil, Asya Passinsky, John Richardson, Sam Scheffler, Andrew Sepielli, Jason Stanley, Sharon Street, David Storrs-Fox, Christian Tarsney, Christina Van Dyke, Dan Waxman, Brian Weatherson, Jake Zuehl, and anonymous referees, as well as audiences at Notre Dame University, Simon Fraser University, the Bellingham Summer Philosophy Conference, the NYU Philosophy Lunch Talk Series, NYU's Graduate Student Political Philosophy Working Group, and the Philosophy Thesis Prep Seminar at NYU.

Makins, N. (2022). Attitudinal ambivalence: Moral uncertainty for non-cognitivists. *Australasian Journal of Philosophy*, 100(3), 580–594. 10.1080/00048402.2021.1908380
Moller, D. (2011). Abortion and moral risk. *Philosophy*, 86(3), 425–443. https://doi.org/10.1017/S0031819111000222
Railton, P. (1984). Alienation, consequentialism, and the demands of morality. *Philosophy & Public Affairs*, 13(2), 134–171.
Robinson, P., & Steele, K. (2023). Moral uncertainty, noncognitivism, and the multi-objective story. *Nous*, 57(4), 922–941. https://doi.org/10.1111/nous.12435
Rosen, G. (2003). IV-culpability and ignorance. *Proceedings of the Aristotelian Society*, 103(1), 61–84. https://doi.org/10.1111/j.0066-7372.2003.00064.x
Rosen, G. (2004). Skepticism about moral responsibility. *Philosophical Perspectives*, 18(1), 295–313. https://doi.org/10.1111/j.1520-8583.2004.00030.x
Rosenthal, C. (2018). Why desperate times (but only desperate times) call for consequentialism. In M. Timmons (Ed.), *Oxford studies in normative ethics* (Vol. 8, pp. 211–235). Oxford: Oxford Academic.
Rosenthal, C. 2019 Ethics for Fallible People. Ph.D. dissertation, New York University.
Rosenthal, C. (2021). What decision theory can't tell us about moral uncertainty. *Philosophical Studies*, 178(10), 3085–3105. https://doi.org/10.1007/s11098-020-01571-3
Ross, J. (2006). Rejecting ethical Deflationism. *Ethics*, 116(4), 742–768. 10.1086/505234
Russell, J. S. (2025). The value of normative information. *Australasian Journal of Philosophy*, 103(2), 469–485. https://doi.org/10.1080/00048402.2024.2320290
Sepielli, A. (2009). What to do when you Don't know what to do. In R. Shafer-Landau (Ed.), *Oxford studies in metaethics* (Vol. 4, pp. 5–28). New York: Oxford University Press. https://doi.org/10.1093/oso/9780199566303.003.0002
Sepielli, A. (2012). Normative uncertainty for non-cognitivists. *Philosophical Studies*, 160(2), 191–207. https://doi.org/10.1007/s11098-011-9713-6
Sepielli, A. (2014). What to do when you Don't know what to do when you Don't know what to do …. *Nous*, 48(3), 521–544. https://doi.org/10.1111/nous.12010
Sepielli, A. (2017). How moral Uncertaintism can be both true and interesting. In M. C. Timmons (Ed.), *Oxford studies in normative ethics* (Vol. 7, pp. 98–116). New York: Oxford University Press.
Smith, H. (2010). Subjective rightness. *Social Philosophy and Policy*, 27(2), 64–110. https://doi.org/10.1017/S0265052509990161
Smith, H. (2014). The subjective moral duty to inform oneself before acting. *Ethics*, 125(1), 11–38. https://doi.org/10.1086/677024
Smith, H. (2018). *Making morality work*. New York: Oxford University Press.
Staffel, J. (2019). Expressivism, normative uncertainty, and arguments for probabilism. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford studies in epistemology* (Vol. 6, pp. 161–189). New York: Oxford University Press. https://doi.org/10.1093/oso/9780198833314.003.0007
Tarsney, C. (2018). Intertheoretic value comparison: A modest proposal. *Journal of Moral Philosophy*, 15(3), 324–344. https://doi.org/10.1163/17455243-20170013
Thorstad, D. (2024). Why bounded rationality (in epistemology)? *Philosophy and Phenomenological Research*, 108(2), 396–413. https://doi.org/10.1111/phpr.12978
Walzer, M. (1973). Political action: The problem of dirty hands. *Philosophy and Public Affairs*, 2(2), 160–180.
Weatherson, B. (2014). Running risks morally. *Philosophical Studies*, 167(1), 141–163. https://doi.org/10.1007/s11098-013-0227-2
Weatherson, B. (2019). *Normative externalism*. New York: Oxford University Press. https://doi.org/10.1093/oso/9780199696536.001.0001
Zhao, M. (2023). Ignore risk; maximize expected moral value. *Nous*, 57(1), 144–161. https://doi.org/10.1111/nous.12398
Zimmerman, M. J. (2014). *Ignorance and moral obligation*. New York: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199688852.001.0001

Cite this article: Rosenthal, C. 2026. Trying to Be Moral, Morally. *Canadian Journal of Philosophy* 56: e18, 1–20, doi:10.1017/can.2026.10067