

What Decision Theory Can't Tell Us About Moral Uncertainty¹
Chelsea Rosenthal
Philosophical Studies 178: 3085–3105 (2021)

Abstract:

We're often unsure what morality requires, but we need to act anyway. There is a growing philosophical literature on how to navigate moral uncertainty. But much of it asks how to rationally pursue the goal of acting morally, using decision-theoretic models to address that question. I argue that using these popular approaches leaves some central and pressing questions about moral uncertainty unaddressed. To help us make sense of experiences of moral uncertainty, we should shift away from focusing on what it's rational to do when facing moral uncertainty, and instead look directly at what it's moral to do about moral uncertainty—for example, how risk averse we morally ought to be, or which personal sacrifices we're morally obligated to make in order to reduce our risk of moral wrongdoing. And orthodox, expectation-maximizing, decision-theoretic models aren't well-suited to this task—in part because they presuppose the answers to some important moral questions. For example, if approaching moral uncertainty in a moral way requires us to “maximize expected moral rightness,” that's, itself, a contentious claim about the demands of morality—one that requires significant moral argument, and that I ultimately suggest is mistaken. Of course, it's possible to opt, instead, for a variety of alternative decision-theoretic models. But, in order to choose between proposed decision-theoretic models, and select one that is well-suited to handling these cases, we first would need to settle more foundational, moral questions—about, for example, what we should be willing to give up in order to reduce the risk that we're acting wrongly. Decision theory may be able to formalize the conclusions of these deliberations, but it is not a substitute for them, and it won't be able to settle the right answers in advance. For now, when we discuss moral uncertainty, we need to wade directly into moral debate, without the aid of decision theory's formalism.

I. Introduction

In our daily lives, we're often unsure what being moral requires, and our moral decision-making is far from perfect. But we need to act anyway. We need to decide, for

¹ For helpful feedback on the ideas presented in this paper (in some cases in earlier forms), I am grateful to James Fritz, B.R. George, Alex Guerrero, Elizabeth Harman, Chris Howard, Zoë Johnson-King, Colin David Jones, Matthew Liao, Alex London, Robert Long, Will MacAskill, Jordan MacKenzie, Ishani Maitra, Liam Murphy, Jake Nebel, Japa Pallikkathayil, Samuel Scheffler, Andrew Sepielli, David Storrs-Fox, Sharon Street, Christian Tarsney, Brian Weatherston, Alex Worsnip, Jake Zuehl, and an anonymous referee, as well as the NYU Graduate Student Extreme Value Theory Group, the Philosophy Thesis Prep Seminar at NYU, the 2018 Foundations of Normativity Workshop at the University of Edinburgh, the 2019 Chapel Hill Normativity Workshop, and the audience at my dissertation defense. During work on this paper, I have been fortunate to receive support from an Andrew W. Mellon Dissertation Fellowship and a Henry M. MacCracken Fellowship.

example, whether to eat meat or whether to lie to a friend for the friend's own good. There is a growing philosophical literature on how to navigate this type of moral uncertainty.² But I'll suggest that the approach taken in most of that work leaves some central and pressing questions about moral uncertainty unaddressed.

Imagine, for example, that I have promised to attend a friend's wedding, but an important professional opportunity has arisen that conflicts with the ceremony. I'm unsure whether it's morally acceptable to break the promise, given the change of circumstances, and I need to determine how to act, given that uncertainty. One thing I might do is forgo the professional opportunity and attend the wedding, *just in case* breaking the promise would be morally wrong (figuring that attending the wedding would at least be morally permissible, even if it's not required). But depending upon the opportunity and the nature of the friendship,

² See e.g., Alexander A. Guerrero, "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution," *Philosophical Studies* 136, no. 1 (2007): 59-97; Claire Field, "Recklessness and Uncertainty: Jackson Cases and Merely Apparent Asymmetry," *Journal of Moral Philosophy* 16, no. 4 (2019): 391-413; Johan E. Gustafsson and Olle Torpman, "In Defense of My Favourite Theory," *Pacific Philosophical Quarterly* 95, no. 2 (2014): 159-174; Elizabeth Harman, "The Irrelevance of Moral Uncertainty," in *Oxford Studies in Metaethics*, Vol. 10, ed. Russ Shafer-Landau (New York: Oxford University Press, 2015): 53-79; Brian Hedden, "Does MITE Make Right?" in *Oxford Studies in Metaethics*, Vol. 11, ed. Russ Shafer-Landau (New York: Oxford University Press, 2016): 102-128; Amelia Hicks, "Moral Uncertainty and Value Comparison," in *Oxford Studies in Metaethics*, Vol. 13, ed. Russ Shafer-Landau, (New York: Oxford University Press, 2018): 161-183; Ted Lockhart, *Moral Uncertainty and its Consequences* (New York: Oxford University Press, 2000); William MacAskill, "The Infectiousness of Nihilism," *Ethics* 123, no. 3 (2013): 508-520; William MacAskill and Toby Ord, "Why Maximize Expected Choice-Worthiness?" *Nous* 54, no. 2 (2020); Dan Moller, "Abortion and Moral Risk," *Philosophy* 86, no. 3 (2011): 425-443; Jacob Ross, "Rejecting Ethical Deflationism," *Ethics* 116, no. 4 (2006): 742-768; Andrew Sepielli, "What to Do When You Don't Know What to Do," in *Oxford Studies in Metaethics* Vol. 4, ed. Russ Shafer-Landau (New York: Oxford University Press, 2009): 5-28; Holly Smith, "The Subjective Moral Duty to Inform Oneself Before Acting," *Ethics* 125, no. 1 (2014): 11-38.; Christian Tarsney, "Intertheoretic Value Comparison: A Modest Proposal," *Journal of Moral Philosophy* 15, no. 3 (2018): 324-344; Brian Weatherson, *Normative Externalism* (New York: Oxford University Press, 2019); Brian Weatherson, "Running Risks Morally," *Philosophical Studies* 167, no. 1 (2014): 141-163; Michael J. Zimmerman, *Ignorance and Moral Obligation* (New York: Oxford University Press, 2014).

I might wonder how much I should be willing to sacrifice in order to reduce my risk of moral wrongdoing. Questions like these are at the heart of our lived experiences navigating moral uncertainty, but approaches taken by philosophical work on the topic often presuppose the answers to questions like these, rather than engaging with them directly.

I ultimately raise doubts about two features that are common in the literature. First, much of the work on moral uncertainty concerns how to *rationally* behave, given moral uncertainty and the aim of acting morally. This is, in practice, a discussion about instrumental rationality and how to effectively achieve a goal: acting morally. Second, most recent discussion of this question answer it using decision-theoretic approaches.

Insofar as the question of what it's *rational* to do under moral uncertainty comes apart from the question of what it's *moral* to do under moral uncertainty, I'll suggest that it's the second one that we should shift to focusing on – and I argue that the standard, decision-theoretic approaches aren't well-suited to addressing those questions. In the case of the skipped wedding, for example, what we really need to know is what sacrifices I *morally* ought to make to reduce my risk of wrongdoing – discussing instrumental rationality won't be enough. When we're unsure what it's right to do, we need to know how we morally ought to navigate that predicament. And I'll argue the decision-theoretic models that have been used to address questions about rationality under moral uncertainty won't be able to settle moral questions like these – even though these will be the key questions to answer.

While I'll defer more detailed discussion of the distinction between rational and moral norms until Part IV, it's worth emphasizing that I'm not primarily concerned with the terminology. My worry, instead, is that framing the issue in terms of rationality has allowed us to presuppose the answers to many of the important questions about moral uncertainty – such

as those about how much we ought to sacrifice – and a successful treatment of this topic needs to grapple more directly with those questions, regardless of the terminology.

In Part II, I'll briefly suggest that there can be moral norms (not only rational norms) concerning what we should do when we encounter moral uncertainty – being reckless in the face of moral uncertainty can, itself, be morally bad. I argue, in Part III, that orthodox decision theory can't settle the content of these norms, and that there are more general problems with decision-theoretic accounts of moral uncertainty that would also face alternative decision-theoretic models. With these challenges on the table, I'll be able to make the case in Part IV that it's these moral questions that should be our focus when we think about moral uncertainty.

To lay the foundation for this discussion, I want to begin by describing the role that decision-theoretic analysis tends to play in theories of moral uncertainty. Andrew Sepielli, for example, asks “what it's rational to do ... under normative uncertainty.”³ And the approach he favors involves choosing the action with the highest level of what he calls “Expected Objective Value (EOV).”⁴ On Sepielli's view, “we get the expected value of an action by multiplying the subjective probability that some [claim about which actions are favored by the balance of reasons] is true by the objective value of that action if it is true, doing the same for all of the other [such claims], and adding up the results.”⁵ This is not a first-order, consequentialist argument. When Sepielli talks about one action having more value than another, he means roughly, that the balance of reasons for acting favors doing it, or that it is a better action.⁶ So he's describing a way to optimize for doing better actions or actions more favored by the balance of reasons – whatever the right criteria for reasons for action turns out to be, and

³ Andrew Sepielli, “What to Do When You Don't Know What to Do,” 9.

⁴ Sepielli, 11.

⁵ Sepielli, 11.

⁶ Sepielli, 7.

leaving open that the right criteria might be anything from utilitarianism to avoiding violations of the categorical imperative.⁷ Along similar lines, Ted Lockhart argues that it's all-things-considered, rational for us to "maximize expected moral rightness" under moral uncertainty,⁸ and Jacob Ross makes the case that it's rational to reject views such as ethical nihilism, by drawing on the idea that decision theory can show us what it's rational to do when faced with moral uncertainty.⁹

Although these arguments focus on rationality, it might seem that a similar approach could be applied to the analogous question of what it's *moral* to do under moral uncertainty. Indeed, there is sometimes a sense that the right approach to normative uncertainty won't depend upon the type of norms in question. William MacAskill suggests that it's *appropriate* for us to maximize expected choice-worthiness under normative uncertainty, using "appropriateness" explicitly "to remain neutral on the issue of whether [the relevant norms – what he calls metanormative norms] are rational norms, or some other sort of norms" (with moral norms understood to be among the options).¹⁰

But I'll argue here that decision theory can't tell us what it's *moral* to do about moral uncertainty. Orthodox decision-theoretic models, like those used by most of these authors,

⁷ Sepielli, 7.

⁸ Lockhart, *Moral Uncertainty and Its Consequences*.

⁹ Jacob Ross, "Rejecting Ethical Deflationism." See also Jacob Ross, "Acceptance and Practical Reason," (PhD diss., Rutgers University, 2006).

¹⁰ William MacAskill, "Normative Uncertainty," (PhD diss., Oxford University, 2014) 16, 16n9, 20. In later work with Toby Ord, MacAskill offers a different definition of appropriateness: "an appropriate action is what would be selected by a rational and morally conscientious agent who had the same set of options and beliefs" (MacAskill and Ord, "Why Maximize Expected Choice-Worthiness?" 329). But this formulation also has implicit in it that higher-order rational norms and higher-order moral norms will have the same content. It's worth noting that, while MacAskill endorses an expectation-maximizing approach where one is possible, he offers an alternative for when this isn't an option (e.g., due to difficulties making value comparisons across different proposed moral theories). For these cases, he proposes that we treat normative uncertainty as a voting problem. See William MacAskill, "Normative Uncertainty as a Voting Problem," *Mind* 125, no. 500 (2016): 967-1004.

don't take into account certain factors that do matter morally (as I suggest in Part III below).¹¹ But even if the specific challenges of orthodox decision-theoretic frameworks are avoided by choosing a different decision-theoretic model, using decision theory to address moral uncertainty encounters a more general problem. In order to select a decision-theoretic model that is well-suited to handling these cases, we first would need to settle more foundational, moral questions – about, for example, what we should be willing to give up (from abortions to business opportunities) in order to reduce the risk that we're doing wrong. Decision theory may be able to formalize the conclusions of these deliberations, but it won't be able to settle the right answers in advance.

II. Moral Norms for Navigating Moral Uncertainty

First, though, it must be the case that there are moral norms about how to navigate moral uncertainty – that this is the sort of thing that morality can speak to. To begin, consider Geraldine's case:

Geraldine's friend, Ellen says that she wants to discuss something personal, but she asks in advance that Geraldine promise not to repeat what Ellen will tell her. Geraldine makes this promise. But once their conversation begins, Geraldine discovers that Ellen, who has a history of addiction, is using drugs again, she considers breaking her promise in order to involve other people who may be able to help.

¹¹ The worries I raise are, by no means, the only challenges facing expected value maximization accounts of moral uncertainty. There is an extensive literature on the question of whether it's possible to make the type of intertheoretic value comparisons that such accounts require. (See e.g., Gustafsson and Torpman, "In Defense of My Favourite Theory;" Hedden, "Does MITE Make Right?;" Hicks, "Moral Uncertainty and Value Comparison;" Ted Lockhart, *Moral Uncertainty and its Consequences*; William MacAskill, "Normative Uncertainty as a Voting Problem;" Andrew Sepielli, "What to Do When You Don't Know What to Do;" Christian Tarsney, "Intertheoretic Value Comparison: A Modest Proposal.") And Brian Hedden has argued that some moral theories, including those involving supererogation, cannot be adequately incorporated into the framework used by expected value maximization accounts (Hedden, "Does MITE Make Right?"). But even if these obstacles can be overcome, there are more general problems, as I suggest here.

What Geraldine should do will depend partly on non-normative facts that I haven't filled in, but it will also depend, in part, on normative, moral questions (e.g., concerning the relationship between well-being considerations and obligations to keep promises). This is a difficult case, but here's one thing that Geraldine clearly should not do: she should not be flippant about the decision she needs to make. Given the gravity of the issues at stake, Geraldine morally ought to take the question seriously, focusing carefully on the relevant particulars of the case and considering the competing values. If she failed to do this -- e.g., bypassing this work, and instead flipping a coin to decide whether to break her promise -- that would, itself, be a moral transgression. She would have shown insufficient respect, both for her friend's well-being and for the commitment that she made to her. There seems, in short, to be a moral reason for us to use good procedures for trying to be moral (in addition, to our ordinary, first-order, moral reasons to treat others well, for example). And, as I'll suggest below, this will ultimately give rise to obligations concerning the actions we take, not only our preparation for them.

In other work, I've argued in more detail that we have this type of procedural moral responsibility.¹² Here I'll only defend this idea in brief, before turning to look at what decision theory can, and cannot, tell us about the content of these procedural norms.

First, a note about terminology. I'll be using "procedural oughts" or "procedural norms" to refer to the higher-order norms I've been discussing so far, and I'll refer to ordinary, first-order moral responsibilities as "substantive oughts" or "substantive norms."¹³ Put more precisely,

Procedural oughts are 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,

¹² Chelsea Rosenthal, "Trying to Be Moral, Morally," in *Ethics for Fallible People* (PhD diss., New York University, 2019). See also Chelsea Rosenthal, "Ethics for Fallible People," [article manuscript].

¹³ For ease of discussion, I'll be using the terms "oughts" and "norms" interchangeably.

given the challenges we face in discovering their content (roughly, “norms about how to pursue other norms” or “strategic norms”).¹⁴

Substantive oughts are all non-procedural, moral norms (that is, those moral norms that don’t concern how to morally pursue another set of moral norms).

On this picture, procedural oughts aren’t reducible to substantive oughts: not only can you violate one without violating the other, determining the content of procedural norms will also require addressing issues outside the scope of substantive, first-order morality – for example, how risk averse to be when facing moral uncertainty.

Geraldine’s case began to show that we have moral responsibilities concerning how we prepare for decision-making, when we’re unsure what morality calls for. Depending on the particular case, this could include avoiding thoughtless or drunken choices, devoting additional time to research or contemplation, or seeking additional perspectives from wise friends. And in some low-stakes cases, it might require very little. Not all cases of moral uncertainty demand our unwavering attention.

But, in discussing procedural oughts, I have in mind more than just obligations to prepare for decision-making under moral uncertainty. Trying to satisfy first-order, moral norms (“substantive oughts”) involves using good preparatory approaches, but it also involves making decisions that don’t, themselves, carry an inappropriate risk of wrongdoing. This could mean, for example, “hedging your bets,” and avoiding activities that seem morally questionable,

¹⁴ Gideon Rosen has discussed the related phenomenon of what he terms “procedural epistemic obligations”: “as you move through the world you are required to take certain steps to inform yourself about matters that might bear on the permissibility of your conduct,” and these steps are your procedural epistemic obligations. Rosen’s focus is different than mine – he is specifically looking at the way that taking, or failing to take, appropriate steps in advance of acting can impact our blameworthiness. But his “procedural epistemic obligations” might be seen as a subset of the “procedural oughts” I discuss here. Gideon Rosen, “Skepticism about Moral Responsibility,” *Philosophical Perspectives*, 18, no. 1 (2004): 301. See also Elizabeth Harman’s discussion of procedural moral obligations in “Ethics is Hard! What Follows?” [manuscript].

just in case they're wrong. We commonly have at least some uncertainty about what's morally right, and may need to make decisions quickly, so versions of this "hedging our bets" strategy are often good ways of trying to satisfy substantive moral norms. And my claim, here, is that failing to use good strategies for trying to be moral (whether in preparation or in the acts we ultimately choose) is, itself, morally bad.

This is easiest to see in the case of preparatory responsibilities. It seems clear that not thinking carefully or not preparing well for a morally weighty action will sometimes, itself, be a failure to do what we morally ought to do. And, often, this bad preparation will manifest a disrespect for those we're at risk of wronging. This was true with Geraldine's case. But if it's morally important to act rightly, it's also morally important to choose acts that are likely to be right.

And if we are to make sense of the type of preparatory obligations that Geraldine has, there must also be procedural norms that apply to the ultimate actions we choose, not only to our preparatory activities. Once we recognize the preparatory obligations, there are two possibilities. Either these preparatory, procedural norms have accompanying moral norms calling for us to ultimately act in ways that reflect this preparation, or they don't. On the one hand, if there aren't such ultimate norms (e.g., morally obligating us to act in ways that are shaped by our contemplation or research), then it's puzzling why we'd be compelled to engage in this preparation at all. On the other hand, if there are moral norms calling for these preparations to guide our actions, that means that we're not only morally compelled to prepare well for our decisions under moral uncertainty – we also morally ought to choose eventual actions that are good strategies for being moral.

This might seem surprising because it means that substantive and procedural norms can, in some sense, place conflicting demands on us.¹⁵ Sometimes using a good strategy for being moral will still lead us to an action that turns out to be substantively immoral. But conflicts between first-order and higher-order norms are less worrying than they might seem, in part because, as Andrew Sepielli has suggested, we don't simultaneously experience first-order and higher-order norms as practically guiding us.¹⁶ When we know what first-order morality calls for, then higher-order norms just tell us to do that – there is no conflict.

But even if conflicts between first-order and higher-order norms don't pose a threat to action guidance, they do suggest that sometimes morality will place a combination of demands on us that it's impossible to fulfill. By satisfying procedural norms, we may violate substantive norms, or vice versa. Here too, however, it's our incomplete moral knowledge that generates the conflict (if we knew what substantive norms called for, procedural norms would call for the same thing). So, we won't be able to comply with all of morality's demands due to incomplete moral knowledge, but it's commonplace to leave moral obligations unfulfilled because we don't realize we have them.

One way to understand these moral requirements is by analogy with the norms of justice that apply in a courtroom. We are already accustomed to accepting that the same act in a court, for example, might further substantive justice while compromising the requirements of procedural legitimacy (or vice versa). (Indeed, these courtroom cases present somewhat starker conflicts, because full knowledge of what is substantively just may not remove conflicts with

¹⁵ Discussion of this point below draws heavily on Chelsea Rosenthal, "Trying to Be Moral, Morally," in "Ethics for Fallible People" (PhD diss.); see also Chelsea Rosenthal, "Ethics for Fallible People," [article manuscript].

¹⁶ Andrew Sepielli, "What to Do When You Don't Know What to Do When You Don't Know What to Do ..." *Nous* 48, no. 3 (2014): 536.

procedural justice – for example when dispositive evidence is excluded from a trial for procedural reasons.)

If there are moral norms of the type I’ve been describing – norms about how we try to act morally when we’re gripped by moral uncertainty – they represent a large, under-examined class of moral norms. And exploring these norms may be powerful for addressing a variety of other moral questions. (I’ve written elsewhere, for example, about whether norms surrounding threshold deontology or interpersonal toleration could be explained in terms of morally obligatory, good strategies for trying to be moral.)¹⁷ So we’ll need to get clear on the content of these norms – and the existing literature on moral uncertainty might seem like a natural place to start. While most of that work looks at what it’s rational to do when we face moral uncertainty, we might think that similar approaches could be used to address questions about what it’s moral to do about moral uncertainty. The trouble is that, while these accounts can offer us some insights into the issue, they’re not well-suited to resolve it on their own.

III. Limitations Facing Decision Theory

To start, let’s consider what orthodox decision theory would tell us about what it’s *moral* to do under moral uncertainty, if we applied it to this question without modifications. I’ll take an orthodox decision-theoretic analysis of this question to be one that suggests we *morally* ought to maximize expected moral rightness (or to maximize any similar value, such as MacAskill’s expected choice-worthiness or Sepielli’s expected objective value – I’m using the language of expected moral rightness because I see it as the most intuitive phrase; I take my

¹⁷ Chelsea Rosenthal, “Why Desperate Times (But Only Desperate Times) Call for Consequentialism,” in *Oxford Studies in Normative Ethics*, Vol. 8, ed. Mark Timmons (New York: Oxford University Press, 2018); Chelsea Rosenthal, “Tolerating Each Other,” in “Ethics for Fallible People” (PhD diss., New York University, 2019).

arguments to apply to this set of concepts as a whole). In this context, expected moral rightness can be understood as a function of an act's possible levels of moral rightness and each one's probabilities (put more precisely: the sum of the possible degrees of moral rightness of the act, each multiplied by their respective probabilities).¹⁸ This approach will be (higher-order) risk neutral, in the sense that it is indifferent between more and less risky options, provided that they have the same expected moral rightness. For example, on this picture, if we're unsure whether a high-stakes act is extremely morally good or extremely morally bad, this need not be treated differently than cases where the uncertainty concerns whether a low-stakes act is a little bit good or a little bit bad.

Distribution of the Risk of Victimization

I want to argue that there are a number of features of our choices under moral uncertainty that should matter morally, but which these orthodox decision-theoretic approaches aren't well-equipped to account for. First, when we take on a higher-order risk of acting wrongly, we typically subject others to a risk of *being wronged* by us. The cases I have in mind are not, primarily, cases in which there's empirical uncertainty about the results of my actions, but rather cases of uncertainty about the moral status of those actions – where I'm unsure whether an action is wrong, and there is someone who I would be wronging, if it is wrong. But it may be morally important who is at risk of being a victim of my wrongdoing, and how that risk of being a victim is distributed.

Consider the following two cases.

First, Erica is an involved citizen deciding what policies to advocate for, and she is unsure whether it would be morally good to support or oppose Policy A. Erica realizes that Policy A would have very little impact on most people, beyond, for example, the type of minor variations in taxes that always happen from year to year. But it's clear to

¹⁸ See Lockhart, *Moral Uncertainty and Its Consequences*.

Erica that, if she is wrong about whether it is good to support Policy A, acting on that mistake would involve behaving very unjustly toward members of a particular, marginalized group. Erica hasn't bothered to think very carefully about Policy A before taking a side and becoming involved in local political organizing surrounding it, and when she acts in this context, there's a sizeable chance that she's acting wrongly. (On some accounts, reparations for slavery could be like Policy A, though my argument doesn't depend upon particular features of reparations.¹⁹)

Second, Mark is also a local citizen who likes to get involved in political issues. He is deciding whether to oppose or advocate for Policy B, and he's unsure which would be morally good. Mark realizes that Policy B is relevant to most members of his society, and getting it right would mean treating them more justly. But it's already clear that getting Policy B wrong would wrong these people in only a modest way – it doesn't seem plausible that it would be a grave harm. Mark hasn't bothered to think very carefully about Policy B before taking a side and becoming involved in local political organizing surrounding it, and when he acts in this context, there's a sizeable chance that he's acting wrongly. (Some policies designed to curb mild government corruption might provide examples like Policy B.)

In both cases, there's a sizeable chance that the agent is acting wrongly, although they're facing moral uncertainty. If Mark is mistaken about what morality calls for, his action risks wronging a very large number of people in a modest way, while Erica's action risks wronging a smaller number of people more severely, if she is mistaken about what morality calls for. Orthodox decision-theoretic accounts of moral uncertainty focus on the agent's risk of acting as they morally shouldn't. And depending upon how we fill in the particulars, these accounts may give similar evaluations of Mark's and Erica's actions, because the likelihood that the agent is doing wrong is similar (sizeable in both cases), and the total wrong they're at risk of doing may be similar as well (whether because of aggregation of potential wrongdoing across a large number of people in Mark's case or aggregation of more serious wrongs across a smaller number in Erica's case).

¹⁹ I'm grateful to Daniel Wodak for suggesting that I incorporate an example involving reparations.

But I want to suggest that it's worse for Erica to take the risks that she takes – for reasons that orthodox decision theory isn't well-suited to account for. Erica's risk of doing wrong may be similar to Mark's. But some people shoulder a much greater risk than others of being a victim of her wrongdoing. And it seems morally worse to be reckless in ways that impose a larger risk of victimhood specifically on a marginalized community, rather than distributing a low risk of victimhood out across the population. And this can be true even when the expected moral rightness of the two actions is the same.

A similar problem arises when we assess someone's pattern of recklessness or caution over time, in cases where they encounter moral uncertainty. For example, suppose both Jeff and Carly are local political officials, who are sometimes unsure whether it would be morally good to back particular policy proposals. Both of them are sometimes cautious in the face of this moral uncertainty, while at other times, they make hasty decisions that are a bit reckless just to get their work wrapped up. For Carly, this recklessness and caution tend to be distributed a bit randomly (or to depend only on whether she has weekend travel plans). But Jeff tends to be reckless in the face of uncertainty about whether it would be morally good to support policies impacting low-income neighborhoods, while being much more careful about his stances on policies impacting the wealthy neighborhoods where his friends live. Over time, Carly wronging some of her constituents is as likely as Jeff wronging some of his, and the severity of the wrongs at stake are similar. The expected moral rightness of their actions over time could be identical. But Jeff's risk-taking seems morally worse, because of the way that the risks are distributed. A good account of procedural moral norms should be able to say that this type of distribution matters when determining how we should navigate moral uncertainty, and that Jeff and Erica's behavior was worse than Carly and Mark's.

Orthodox decision theory doesn't have a natural way of distinguishing these cases or accommodating these concerns about the distribution of the risk of being wronged. Using these models stipulates that there is no important difference between these cases, but doesn't really provide an argument for it. One approach may be for decision theorists to keep orthodox decision theory – and continue to say that we should maximize expected moral rightness (or e.g., expected choice-worthiness) – but to build anti-discrimination norms into the first-order account of which sorts of actions are more likely to be right. For example, they might say that satisfying certain fairness norms is one of the qualities of a right act (or might otherwise build these priorities into our assessment of the likelihood that an action is desirable) – saying that actions that impose disproportionate risks on some groups are simply more likely to be wrong.²⁰ In this way, expected moral rightness calculations would try to distinguish between Carly and Jeff's actions, characterizing each one as being arrived at through particular processes, with some processes being preferable to others.

The difficulty is that the fairness I'm concerned about isn't part of first-order morality – I'm concerned about bias infecting the way we navigate uncertainty *about the content of first-order morality*. This can't be addressed by building the criteria for fairness into first-order morality. If it matters whether a higher-order standard distributes risk fairly, that's a criterion for picking a good higher-order theory – not a criterion for picking a good, first-order theory. It would be circular to say that one of my first-order criteria for a good act is that it was done in accordance with higher-order norms that pursue the satisfaction of first-order criteria in a good way. And even under this circular account, we would need to change the content of the higher-order norms in order to address discrimination in how we navigate uncertainty. We couldn't simply stick with orthodox decision theory.

²⁰ I'm grateful to Jake Nebel for raising this point.

Stakes

Similarly, the *morally* best way to handle moral uncertainty may depend partly on the stakes of the different options (e.g., how bad an act could be at worst, or, relatedly, how polarized the possible levels of wrongdoing are). Orthodox decision-theoretic approaches have difficulty explaining some aspects of this, however. They can say that not acting in ways that are probably severe wrongs is more important than not acting in ways that are probably trivial wrongs. But they're less able to account for some other, key features of the stakes of our actions. Consider what orthodox decision-theory can say about a choice between the following two acts, when both acts' moral status is uncertain. The first act might be extremely morally choice-worthy or it might be morally horrific – it's unclear which, and each possibility is equally plausible. Acting to trigger a violent uprising might have this status, for example, if there is uncertainty about the goodness and justifiability of that violence. A second act is much lower stakes, with an equal possibility that it's mildly good or mildly bad, but little chance of anything more significant (perhaps minor, flattering lies are like this). Depending upon how we specify the particulars, an orthodox decision-theoretic approach to moral uncertainty may end up treating these two cases as equivalent or near equivalent, because the first act's extreme options should be able to balance each other and yield a fairly moderate expected moral rightness. But the stakes seem to make a moral difference. It will sometimes be *morally* reprehensible to risk moral catastrophe – in some cases this will be too reckless – while low-stakes risks may be acceptable, despite having the same expected moral rightness.

If we follow Sepielli and Lockhart, and use an orthodox, risk neutral approach to higher-order moral risks,²¹ we won't be able to distinguish two acts with the same expected moral rightness.²² Nevertheless, much of the current work on moral uncertainty appears to assume this type of risk neutrality. Stipulating some level of risk aversion does better,²³ and allows us to classify some of these highly-polarized choices as too reckless. The difficulty here is that whether we should be risk neutral, and what the appropriate level of risk aversion is, will, themselves, be significant moral questions. Because we are asking how we morally ought to go about trying to be moral, these questions amount to asking how much, and what sort of, risk of moral wrongdoing it's *morally* appropriate to take when navigating moral uncertainty. So the reasons available to address this will have to be moral reasons. Any particular, decision-theoretic model (risk neutral or risk averse) builds in a presupposed answer to that debate, rather than providing us with the moral reasons needed to resolve it.

Sacrifice

²¹ Lockhart, *Moral Uncertainty and Its Consequences*; Sepielli, "What to Do When You Don't Know What to Do."

²² To some extent, MacAskill and Ord leave open the possibility of using a risk averse account instead of this type of risk neutral approach. Whether this is called for, according to them, will depend upon whether risk aversion is also rational in the case of empirical uncertainty, as Lara Buchak has argued that it is (*Risk and Rationality* (New York: Oxford University Press, 2013)). Their primary commitment is to treating normative uncertainty like we treat empirical uncertainty. (See MacAskill and Ord, "Why Maximize Expected Choice-Worthiness?" 338 and MacAskill, "Normative Uncertainty," 34.) This gives them a potential avenue for avoiding the "stakes" worry, by rejecting risk neutrality. But doing so would require accepting much more general claims about decision-making under *empirical* uncertainty – like Buchak's – that are quite controversial (although I'm sympathetic to Buchak's account). And MacAskill and Ord don't seem inclined to take this approach – although they don't weigh in on the issue, they treat risk-neutral, expected value maximization as the default position and adopt it for purposes of discussion (MacAskill and Ord, 338; MacAskill, "Normative Uncertainty," 51).

²³ For arguments in favor of alternative decision-theoretic analyses that can treat risk aversion as rational, see Buchak, *Risk and Rationality*.

A further problem for orthodox decision theory concerns the personal sacrifices associated with mitigating moral risk.²⁴ What steps we should take to reduce our risk of moral wrongdoing will depend, partly, on the personal sacrifices involved in our various options.

As we've seen, we sometimes choose our eventual actions in order to "hedge our bets" and reduce our risk of moral wrongdoing. But this can involve significant sacrifices. When we're engaged in a contract negotiation or a legal battle, for example, we may face questions about what we morally ought to disclose, and while some disclosures can seem not to be morally required, we may not be fully certain. One way to handle this uncertainty is to disclose, just in case. And where disclosure carries few costs, this will often be what we procedurally ought to do, even if the risk of wrongdoing is very small. But when disclosure becomes costly – perhaps certain aspects of our medical history would be used against us, if disclosed, and would dramatically change settlements – a higher moral risk may be required before we procedurally ought to hedge our bets and disclose just in case.

The difficulty is that orthodox decision theory is framed in terms of maximization. If we morally ought to maximize expected morally rightness, there isn't a clear way of taking into account the personal sacrifices that are sometimes involved in mitigating our risk of wrongdoing – that is, the sacrifices involved in following procedural oughts (as distinct from the sacrifices involved in following substantive, first-order morality). Again, this seems to be an area where employing decision theory presupposes the answers to difficult moral questions –

²⁴ For a similar point, see Dan Moller, "Abortion and Moral Risk." Moller suggests that the costs to the agent will have to be among the factors relevant to determining how we morally ought to handle moral uncertainty (440). Moller, however, is not mainly focused on "working out the details of a complete theory of moral risk," instead focusing primarily on showing that considerations of moral risk give us a significant reason not to get abortions (whether or not that reason is ultimately outweighed). Worries about demandingness are also implicit in Brian Weatherson's review of Lockhart's *Moral Uncertainty and Its Consequences* (*Mind* 111, no. 443 (2002): 693-696).

in this case about what morality requires us to give up in order to mitigate the risk of wrongdoing. If approaching moral uncertainty in a moral way requires us to do everything we can to maximize expected moral rightness, that's, itself, a contentious moral claim, and one that requires a separate defense.

This brings out one of the difficulties with anti-abortion arguments that have been made using decision-theoretic accounts of moral uncertainty (see e.g., Lockhart, MacAskill).²⁵ The proposal in these anti-abortion arguments is, roughly, that we shouldn't get abortions because there is a risk that they could be morally horrible acts (close to, or equivalent to murder), while carrying a pregnancy to term doesn't seem to carry a comparably dramatic risk of wrongdoing. On these views, that is, getting an abortion seems to be morally riskier than not getting one. And the expected moral rightness of that choice may be lower than for alternatives. But forgoing an abortion can involve extraordinary sacrifices – potentially months of nausea or vaginal tearing, for example – and carrying a pregnancy to term can dramatically change someone's life for the worse. Sacrifices like these cannot be required just to achieve a mild reduction in our risk of wrongdoing. Perhaps they're required when it is all-but-certain that the alternative is severe wrongdoing. But, this is only true of abortion decisions if we presuppose that abortion is likely to be very wrong – an assumption that I don't take to be very plausible. So, a moral uncertainty argument against abortion appears to need either a further argument about how much we procedurally ought to sacrifice or recourse to substantive, first-order debates about the morality of abortion. The uncertainty, on its own,

²⁵ Lockhart, *Moral Uncertainty and Its Consequences*, 50-73; William MacAskill, "Moral Recklessness and Moral Caution" (manuscript); William MacAskill, "Practical Ethics Given Moral Uncertainty," *80,000 Hours Blog*, January 31, 2012, <https://80000hours.org/2012/01/practical-ethics-given-moral-uncertainty/>. For a claim along these lines that is developed outside of a decision-theoretic framework, see Dan Moller's argument that considerations of moral uncertainty give us reasons not to have abortions (though those reasons may be outweighed) (Moller, "Abortion and Moral Risk").

won't be enough, and as I discuss below, there are reasons to doubt that we must be willing to make sacrifices of this magnitude anytime those sacrifices will maximize expected moral rightness.

A couple of related worries arise at this point. First, doesn't first-order morality already take our sacrifices into account – and if we take sacrifice into account when determining higher-order norms (as I've suggested we should), then are we double-counting any personal prerogatives built into first-order morality?²⁶ The concern is that first-order moral theories already contain claims about what sorts of sacrifices morality does or doesn't demand.

There are multiple ways in which first-order views about sacrifice or demandingness can shape what we procedurally ought to do. First, if more demanding moral theories are less plausible, this will, of course, matter to what we procedurally ought to sacrifice for morality – because the likely content of substantive norms matters to what we procedurally ought to do. But this won't be a form of illicit double-counting any more than avoiding acts that are likely to be wrong double-counts their badness. It only means that appropriate strategies for achieving a goal (acting morally) will take into account what that goal consists of (which acts we morally ought to undertake).

But on the account I've presented, while the content of these substantive, first-order norms – and how they handle demandingness – will matter to what we procedurally ought to do, it can't fully determine it. Peter Singer, for example, thinks that we morally ought to donate most of our income toward the alleviation of global poverty.²⁷ This is a demanding, first-order view. But it doesn't tell us how to handle uncertainty about the view, in particular

²⁶ I'm grateful to Jake Nebel for discussion of this point. See also MacAskill, "Normative Uncertainty," 40-41 and MacAskill and Ord, "Why Maximize Expected Choice-Worthiness?" 342-343.

²⁷ Peter Singer, "Famine, Affluence, and Morality," *Philosophy and Public Affairs* 1, no. 1 (1972): 229-243.

how much time and energy we should sacrifice trying to determine if he's correct, or how plausible his view has to be before we donate "just in case" it's required.²⁸ When faced with questions like these, we ask e.g., which trade-offs are appropriate between personal costs that are nearly certain and moral costs that are less likely. We'll need to know something about sacrifice that isn't accounted for within substantive, first-order morality. We'll need to know, that is, about the sacrifices that we should make to mitigate the risk of transgressing substantive norms, whatever those substantive norms are, and however demanding they happen to be. These determinations will be the second context in which sacrifices matter to procedural norms. They're not double-counting first-order prerogatives because first-order morality never had anything to say about this type of sacrifice.

This also helps to bring out why the content of first-order principles concerning the demandingness of morality will not necessarily be the same as the content of higher-order demandingness principles. What it's reasonable to sacrifice to reduce our risk of wrongdoing can be different from what it's reasonable to sacrifice to avoid actions that are certain to be wrong. Incorporating considerations of sacrifice into second-order moral norms isn't double-counting because it's addressing a different question (a question that orthodox decision-theoretic accounts of moral uncertainty aren't well-suited to address).

One way to think about this is that orthodox decision-theoretic approaches may actually leave sacrifices *under-counted*.²⁹ In many cases, personal sacrifices will be treated as optional, but permissible, by some plausible moral theories, and treated as morally required by the rest. Indeed, this is how many authors see the abortion case above. But if we maximize expected

²⁸ For additional discussion of demandingness, moral uncertainty, and Peter Singer, see MacAskill, "Normative Uncertainty," 39-42 and MacAskill and Ord, "Why Maximize Expected Choice-Worthiness?" 342-343.

²⁹ I'm grateful to an anonymous referee for making this point.

moral rightness (or a similar value), we'll often need to make personal sacrifices *just in case* they're required and be left without a role for the personal prerogatives that many moral theories afforded us at the first-order level. So, by making it difficult for these prerogatives to play a role, it may be that orthodox decision-theoretic approaches undercount sacrifices, rather than alternatives over-counting them.

My account faces another, related, objection, however. I've been arguing that orthodox decision theory won't address moral uncertainty well, in part because, if we maximize the expected moral rightness of our actions, this fails to take personal sacrifices seriously. But decision theorists can object that they're able to take sacrifice into account in another way. Rather than asking what we *morally* ought to do about moral uncertainty and maximizing expected moral rightness, we can ask what we all-things-considered ought to do about moral uncertainty, and instead maximize expected, all-things-considered choiceworthiness.³⁰

Some decision-theoretic accounts of moral uncertainty take precisely this approach (while others focus specifically on how we should pursue a single goal: acting morally). And these all-things-considered accounts are positioned to give greater weight to personal sacrifices. But while these norms give sacrifice a greater role, they don't answer the questions I've been posing here. I'm interested in what it's moral to do about moral uncertainty, and what types of sacrifices we morally ought to make to mitigate our risk of moral wrongdoing. These higher-order, all-things-considered norms don't tell us this, in much the same way that all-things-considered, first-order norms are different from questions about whether first-order morality incorporates personal prerogatives.³¹ The former tell us about the relationship

³⁰ See MacAskill and Ord, "Why Maximize Expected Choice-Worthiness?" 342-343 for some discussion of this idea.

³¹ See, for example, the contrast between Samuel Scheffler's discussion of personal prerogatives within morality in *The Rejection of Consequentialism* (New York: Oxford University Press, 1994),

between moral and non-moral values – the latter are questions entirely internal to morality. Much as it's reasonable to ask what sorts of prerogatives are permitted by substantive morality, it's reasonable to ask what sorts of sacrifices we morally must make when there are opportunities to reduce our risk of wrongdoing. The all-things-considered norms are simply addressing a different question.

The role I'm suggesting for sacrifice might, initially, seem surprising. After all, even if there are all-things-considered norms that take sacrifice into account, I'm asking what *morality* tells us about how to handle cases of moral uncertainty. And it might seem that morality should call on us to do everything we can to end up acting morally. But there are good reasons to think it doesn't.

First, morality's claims are often much more modest in analogous situations of ordinary, first-order risk of harm to others. Although there are some differences between the cases, when we think that a harm is, itself, impermissible, that doesn't mean that any risk of that harm is impermissible. And while we often have a responsibility to reduce the risk of harming, typically we don't think that we're obligated to optimize, even when the values weighing against risk reduction have little moral worth. For instance, it's usually considered permissible to drive to unimportant activities while sober and well-rested, despite a small risk that this will result in devastating harm, up to and including killing. (Indeed, the decision theorists would agree with this much.)

Second, insofar as we have intuitions about the types of procedural norms I'm interested in, sacrifices seem to matter for them. Moral procedural norms simply don't seem as demanding as an expected rightness maximizing account would require. We don't, for

and Susan Wolf's suggestion that we shouldn't aim to be perfectly morally good in "Moral Saints," *Journal of Philosophy* 79, no. 8 (1982).

example, think that all questions about how we morally ought to act require extensive contemplation, even when more contemplation might be the action that maximizes expected moral rightness. And sometimes this can be the case when the alternative to more contemplation isn't any activity of great moral value that would get much weight in strong first-order theories – forgoing additional contemplation to watch reality television can be acceptable.

Finally, requiring that we maximize expected moral rightness carries with it a strange implication. It can make procedural norms more demanding than any first-order norms under consideration. For example, if we are unsure which of two moral theories is correct, and the theories forbid different activities (think, perhaps, of religious views that prohibit different foods), it will sometimes be feasible to “play it safe” and avoid both sets of activities. In many cases, it looks like maximizing expected moral rightness will require precisely this. And while it's plausible that this will sometimes be called for, it sounds less and less reasonable as the sacrifices involved increase.

I've been worried, here, that an account built on orthodox decision theory wouldn't take into account certain important factors, and that it would become too demanding as a result. Work on moral uncertainty has faced other concerns about the demandingness of maximization. So it's worth considering some replies that have been offered. Ted Lockhart addressed demandingness objections by arguing against the notion of supererogation more generally. On his view, we morally must “maximize expected moral rightness” – doing so can't be merely supererogatory because nothing is supererogatory.³² Andrew Sepielli takes a different tack, framing his theory as one that only provides a criteria for ranking actions from more rational to less rational. But there is no cut-off within the ranking, distinguishing acceptable

³² Lockhart, *Moral Uncertainty and Its Consequences*, 98-110.

acts from impermissible acts – just a ranking from better to worse. He gives a higher ranking to actions with higher expected values and a lower ranking to those with lower expected values, with actions that maximize expected objective value (Sepielli’s variation on expected moral rightness) at the top of the ranking.³³ But because there is no claim that a particular segment of those ranked actions is prohibited or required, there is no way for the account to be too demanding, he suggests.

But even if demandingness, as such, were not a problem for Lockhart, Sepielli, and like-minded theorists, overlooking sacrifice would still present a difficulty. If we adapt Sepielli’s reply to demandingness worries, to apply it to the issue of what it’s *moral* to do under moral uncertainty (e.g., by saying that we only have a ranking of morally good and bad acts), then we’ll still need to ensure that the ranking is, itself, appropriate. Similarly, we can suppose that Lockhart is correct in thinking that nothing is superogatory, and he will still have a good reason to want to rank better and worse options well. But Lockhart and Sepielli can’t create good rankings of their options without taking sacrifice into account.

A good ranking must take all important factors into consideration to ensure that the best options receive the highest rankings, and these approaches seem not to do this – ignoring the extent of personal sacrifice can result in an inaccurate ranking. Consider two scenarios. In each scenario, we face a choice between Act X, a low-risk option, and Act Y, a high-risk option that carries a risk of serious moral wrongdoing. The difference in expected moral rightness between Act X and Act Y is the same in both scenarios. But in one scenario, very little personal sacrifice is involved in choosing the low-risk act (Act X) – it is, on a personal level, comparable to choosing the higher-risk alternative, Act Y. In the other scenario, choosing Act

³³ Andrew Sepielli, “Along an Imperfectly-Lighted Path!: Practical Rationality and Normative Uncertainty” (PhD diss., Rutgers University, 2010), 104.

X over Act Y would require making extensive personal sacrifices. It seems that failing to avoid the risk in the first case is morally worse than failing to avoid it in the second. Indeed, depending upon the particulars, declining to make the sacrifice in the second case may be morally permissible. If we take this distinction seriously, it looks like extent of sacrifice will matter to the content of procedural norms in a way that orthodox decision theory isn't well set up to accommodate. And even if we allow morality to be quite demanding (in keeping with Lockhart), we still have reason to want moral theories to appropriately distinguish better and worse shortcomings – especially since, realistically, we all live in the realm of shortcomings.

A More General Problem

We've now seen three, key factors that should matter to the moral navigation of moral uncertainty: distribution of the risk of victimhood, stakes, and sacrifice. These should be enough to raise doubts about accounts of procedural oughts that use risk neutral, basically orthodox, decision-theoretic models. But decision theorists might want to address these worries by switching to alternative, less orthodox, decision-theoretic models (including e.g., those that aren't risk neutral). There is also a larger underlying problem, however, that can't be addressed by switching to a different, decision-theoretic model.

As the examples above have begun to suggest, if a decision-theoretic model endorses a particular way of navigating moral uncertainty, that doesn't settle whether we *morally* ought to navigate moral uncertainty in the way that model recommends. Orthodox decision-theoretic models leave out morally relevant factors, and different alternative decision-theoretic models would each yield different recommendations. Choosing from among them would, itself, require moral argument – and that moral argument has to be resolved before we can rely on the decision-theoretic model. Decision theory is ordinarily a structure for formalizing views about

rationality, not morality. Where it's adopted for first-order, moral theorizing, it's typically in a context that takes many of the relevant moral assumptions to already be argued for (e.g., discussions internal to consequentialism that presuppose that we ought to maximize some good).

I'll leave open whether it's ultimately possible to build a decision-theoretic model of procedural moral norms that takes into account all necessary factors. But what's crucial is that any feature of such a model will need to be justified by moral argument. We saw above, for example, that choosing a level of risk aversion for a model will depend upon moral arguments about which risks are morally appropriate. Similar arguments would be necessary for any features that tell us how to distribute the risk of victimhood or which sacrifices are called for. Ultimately, decision theory can't provide us with answers to hard moral questions about how to morally approach moral uncertainty – those require direct engagement in often messy, and less precise, moral theorizing. At best, decision theory can provide us with suggestions to consider (insofar as our problems share features with problems of rational decision-making, we'll have good reasons to borrow their insights) and a framework for formalizing the results of our inquiries. But unless we find precise answers to the antecedent moral questions, we should approach formalization timidly, because it carries with it the risk of achieving precision through unjustified assumptions.

IV. Focusing on What's Moral

I've suggested that there are procedural norms with genuine moral force, and argued that there's a great deal that decision theory can't tell us about the content of these moral norms. But, before closing, it's worth offering some reasons for focusing on these moral norms, given the possibility that there also exist norms of rationality addressing the same types of

cases, and which are the focus of most of the current work on moral uncertainty. Like procedural, moral norms, these rational norms concern how we go about trying to be moral, in the face of uncertainty. But rather than telling us which approaches it would be moral to use, they tell us which approaches it would be rational to use – that is, how to *rationally* pursue the goal of acting morally. In the discussion below, I'll refer to them as “rational, procedural norms.”

Even if my arguments up to this point are correct, it might seem that I should be devoting more attention to these rational, procedural norms when thinking about moral uncertainty. And, in keeping with recent work, decision-theoretic analyses may seem well-suited for exploring what it's *rational* to do about moral uncertainty, even if they can't settle which approaches are moral.

But those who want to focus on questions of rationality under moral uncertainty face a dilemma. I'll argue that they must give up one of the following:

- (i) Rational, procedural norms are settled independent of the moral questions I've been discussing (e.g., by a direct application of decision theory, rather than depending upon the weight that we're morally permitted to give to personal sacrifices).
- (ii) Rational, procedural norms are providing us with important guidance that's worth focusing on in navigating moral uncertainty.

(i) and (ii) cannot both be true. To bring this out, first let's get a clearer handle on what it would mean to rationally pursue the goal of acting morally. Both moral, procedural norms and rational, procedural norms can be seen as providing an answer to the question, “how would it be best to pursue the aim of acting morally?” But getting a determinate answer to this question involves specifying “best for what” (or, perhaps, “best according to what criteria”). Moral,

procedural norms tell us which way of pursuing the aim would be morally best. But when we talk about how to rationally pursue the aim of being moral, there are two ways we might understand this.

The first option is to stipulate in advance that we are using the term rational in a specific, technical way, and rationally pursuing the goal of acting morally *just is* maximizing expected moral rightness. If this is correct, (i) is easily satisfied – the content of rational procedural norms won't depend upon the moral questions I've been concerned with above. But it's less clear why the norms, understood in this way, should be particularly interesting or important. When pursuing the goal of acting morally, I care about which approaches are best by moral standards – morality is the reason that using a good strategy or approaching this goal well matters. Other standards for what counts as a good approach will mainly be important insofar as they coincide with moral standards. (I, here, grant decision-theorists their framing of issues of moral uncertainty in terms of pursuing the goal of acting morally; if we reject that framing, we depart even further from the approaches they favor.)

Of course, the reasons I care about satisfying a set of first-order norms won't always match the first-order norms themselves.³⁴ For example, I may want to do a good job satisfying aesthetic norms when purchasing artwork, but my reason for caring about this may be prudential (perhaps I want to exercise good aesthetic judgment partly for financial reasons). If so, these prudential concerns will shape our standards for evaluating the different ways that I might pursue the goal of good aesthetic judgment (how risk averse I should be, for example). In the case of moral uncertainty though, the import of being successful at pursuing morality really does seem to come down to the importance of morality, itself. (Indeed, trying to be moral for reasons outside of morality – e.g., the perception of others – seems, itself, to be a

³⁴ I'm grateful to Rob Hopkins for a question about this point.

moral mistake.) Given this, moral standards should be our guide when deciding how to pursue the goal of acting morally. Which approaches to moral uncertainty satisfy other non-moral, standards will not be important in the same way.

But even if this gives us a reason to move away from rational norms of this specific and narrower type, we might understand rational, procedural norms in another, quite general way. They might simply be the norms asking us to pursue being moral in whatever way is most effective, with orthodox decision theory only being relevant insofar as it's a good tool for determining this. The difficulty is that, before we can determine which approaches are most effective, we'll need to specify our priorities a bit more. In particular, we'll have to ask what sorts of risks are acceptable, given that we can't guarantee success. Should we, for example, maximize the chances of achieving our goal fully or prioritize avoiding the very worst scenarios (those in which we engage in moral horrors)? But the most natural way to select appropriate levels of risk aversion will involve treating these as higher-order moral questions – as I've been discussing, decision theory on its own won't determine good levels of risk aversion (they have to be settled in advance), and it seems to be morality, itself, that sets the standards for what it means to pursue morality in the best way. So if we understand rational, procedural norms to simply be norms about how to *effectively* pursue the goal of acting morally, this won't offer a true alternative to moral theorizing about moral uncertainty. It would only be a framework to be filled in and specified using antecedent moral investigations. And insofar as these norms diverged from *moral*, procedural norms, these departures would make our attention to rational, procedural norms less important, if it's pursuing morality in a moral way that we ultimately care about.

We're now in a position to address one final concern about the view I've been defending: it might seem to treat moral and empirical uncertainty differently in impermissible ways.

MacAskill and Ord have argued that we should treat empirical uncertainty and normative uncertainty alike – that the standards for good decisions under empirical uncertainty should be the same as the standards for good decisions under normative uncertainty.³⁵ On this approach, if we want to pursue a goal effectively (like being moral), and we’re unsure which course of action is best, it won’t matter whether we’re unsure due to uncertainty about empirical facts (e.g., facts about fetal neurological development), or due to uncertainty about underlying norms and principles (e.g., the moral status of fetuses with particular neurological features) – the best strategies for navigating uncertainty will stay the same, regardless of the source of that uncertainty. So, according to this view, if an orthodox, value-maximizing approach is the correct approach to empirical uncertainty – MacAskill and Ord assume it is, for purposes of discussion – then it (or some relevant analogue) would also tell us what it’s best to do in the face of normative uncertainty, including moral uncertainty.³⁶ Ordinarily, decision theorists don’t think that the criteria for good decisions under uncertainty depends on the types of propositions we’re uncertain about. It’s not the case, for example, that one set of standards applies to how we navigate uncertainty about a priori propositions, while different standards apply when we’re uncertain about a posteriori propositions. But having distinctive standards for normative and empirical uncertainty seems to require an analogous division. So, they argue, we should default to using the same standards for choice under normative and empirical uncertainty unless a compelling case can be made for departing.

And, on MacAskill and Ord’s account, there doesn’t seem to be such a case. Instead, we have some affirmative reason to think that normative and empirical uncertainty should be

³⁵ MacAskill and Ord, “Why Maximize Expected Choice-Worthiness?” 338-340. See also MacAskill, “Normative Uncertainty,” 34-36.

³⁶ Of course, this leaves open the possibility that expected value maximization is the wrong approach to all types of choice under uncertainty, but defending this would require adopting a much larger body of contentious claims than I address here.

handled in the same way. People may not know whether their uncertainty about what to do derives from normative or empirical uncertainty (as with someone who is unsure whether to become a vegetarian, without knowing if this is due to lingering questions about animal neurology, the moral status of animals, or some combination). And it would seem bizarre if the standards for evaluating their choices depended upon determining this.

In one sense, I agree with MacAskill and Ord. How I should behave in the face of uncertainty doesn't depend upon whether that uncertainty is ultimately due to doubts about empirical or normative propositions. But, while the *source* of our uncertainty may not alter which decision-making criteria apply, other distinctions may warrant differences in decision-making criteria. In particular, different reasons for caring about using a good strategy or pursuing a goal well can make different decision-making standards fitting, as the discussion above suggested. When we try to be moral in the face of uncertainty, this calls for distinctive decision-making standards, not because that uncertainty isn't empirical (sometimes it is), but because we have a distinctive type of reason to care about navigating the uncertainty well: a moral reason. Put differently, in these cases we want a decision standard that tells us the *morally best* way to pursue our aim, not one that tells us which way is rationally best.

V. Conclusion

I've argued that when determining what we procedurally ought to do, there is a great deal that decision theory alone can't tell us. Decision theory can be a fruitful source of relevant ideas and considerations. But applying orthodox decision theory to the question of what it's moral to do about moral uncertainty seems to come up short: maximizing expected rightness or expected choice-worthiness doesn't adequately account for the sacrifices that different options require of us, or their varying stakes and distributions of the risk of victimhood. More

importantly, any decision-theoretic model we use will need to be justified by how well it formalizes the answers to moral questions. The model will not be able to answer key, moral questions on its own, even if it can formalize the right answers after the fact.