

Indeterminate Oughts*

J. Robert G. Williams

Sometimes it is indeterminate what an agent morally ought do. This generates a Decision Ought Challenge—to give moral guidance to agents in such a scenario. This article is a field guide to the options for a theory of the decision ought for cases of indeterminacy. Three categories of view are evaluated, and the best representative for each is identified.

I. INTRODUCTION: FROM UNCERTAINTY TO INDETERMINACY

Sometimes there is something that agents ought to do, but epistemic limitations prevent them from knowing what that is. Such *Uncertain Ought* scenarios have been the focus of recent literature. An example:

You have promised to vote for Mary in an election for party president. You then learn that Fred's policies are much better. What to do? . . . Suppose that you are pretty sure but far from certain that you promised to vote for Mary. It was some time ago; it was quite a party; you know you indicated that you would vote for her, but did you actually promise to do so? What should absolutists about promise keeping say about how the element of doubt affects the question of what you ought to do in such a case?¹

There is a reading on which the correct promise-absolutist answer to the question *who ought you vote for?* is that you ought to vote for Mary if and only if (iff) you in fact promised so to do. But such "objective oughts" do

* Special thanks to Thomas Brouwer, Susanna Rinard, Seth Lazar, Daniel Elstein, and the audience of an ancestor of this article at a conference at the Australian National University. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013)/ERC grant agreement 312938.

1. Michael Smith and Frank Jackson, "Absolutist Moral Theories and Uncertainty," *Journal of Philosophy* 103 (2006): 267–83, 267.

Ethics 127 (April 2017): 645–673

© 2017 by The University of Chicago. All rights reserved. 0014-1704/2017/12703-0007\$10.00

not give applicable moral advice for agents in Uncertain Ought scenarios, given their ignorance of the relevant (nonmoral) facts. Even friends of the objective ought should see a role for an account of moral guidance for agents with limited factual information. Following Smith and Jackson, we call the target of such a theory the ‘decision ought’ (aka ‘the subjective ought’). The *Decision Ought Challenge* is to construct a theory of this decision ought.

My concern in this article is with a variant of the Decision Ought Challenge, prompted not by Uncertain Ought scenarios but by *Indeterminate Ought* scenarios. Consider the variant of the above case in which you recollect exactly what you said to Mary at the party, but you know that what happened was a borderline case of promising: your words were ambiguous and your level of intoxication was neither clearly too high to prevent you making contracts with another nor clearly low enough to allow them to be made. In short, you are aware of all the relevant underlying facts, but they don’t settle the matter of whether a promise was made. Again we raise the question whether you ought to vote for Fred or Mary. Again it may be correct to say that you ought to vote for Mary iff you promised that you would. But just as in the original setup, this gives you no moral guidance. Your uncertainty in this case about whether you promised Mary your vote may or may not be the same phenomenon as everyday uncertainty about factual or normative matters (henceforth: *quotidian uncertainty*)—that is up for debate. But whatever the character of the uncertainty prompted by borderline or indeterminate cases, they set up a Decision Ought Challenge to develop a theory of applicable moral advice.

There is no consensus view about how we ought to factor indeterminacy into moral deliberation. There is indeed no consensus view concerning how indeterminacy interacts with rational decision making more generally. But there has been progress on the latter question recently, and this puts us in a position to address the Decision Ought Challenge for Indeterminate Oughts. This article will chart the territory, canvassing first the possibility of reducing the problem to that of (quotidian) Uncertain Oughts and second the possibility of building on distinctive theories of indeterminacy and rational decision making.

This article is a scouting report on the contours of an underexplored philosophical landscape. We are not yet in a position to see which of the approaches surveyed will be vindicated, though some readers will no doubt find themselves with commitments in this area in virtue of their extant commitments elsewhere. But this will be an opinionated report. For each of the three broad categories of approach surveyed below, I will identify the account within that category that has the best chance of success.

The plan of campaign is as follows. Section II takes stock of varieties of moral indeterminacy that may lead to Indeterminate Ought scenarios.

Section III evaluates the proposal that moral decision making under indeterminacy is just a special case of moral decision making under quotidian uncertainty—and so no easier or harder than the latter. The key question is what commitments we have to take on to sustain this reduction. Though the leading account here is epistemicism (as per Williamson), in the light of special problems that epistemicism has in accounting for the stability of moral concepts, I will be recommending a different representative for this class of theories.² In Section IV, I describe a family of nonclassical theories of the decision ought under indeterminacy, building on my previous work on rational nonclassical decision making under indeterminacy.³ Nonclassical theories allow certain kinds of moral indeterminacy to be bracketed—in the sense that they will not lead to Indeterminate Ought scenarios. They promise to allow us to evade apparent instances of the Decision Ought Challenge. However, not every case of moral indeterminacy can be evaded in this way. The second task of this section is to apply the nonclassical theories directly to the residual Indeterminate Ought cases. The accounts that emerge often have some unwelcome features. One particular nonclassical theory looks best placed to avoid these, and in the light of this I recommend it as the best representative of this category of views. In Section V, I offer an alternative account of the decision ought under indeterminacy, building on my previous “classical supervaluational” *Mind Making* model of classical decision making under indeterminacy.⁴ By generalizing an observation I’ve made elsewhere I argue that this is not subject to the problem identified in Section IV for analogous approaches within the nonclassical paradigm.⁵ The conclusion identifies issues to be tackled in future work in the area.

II. VARIETIES OF MORAL INDETERMINACY

Below we take stock of varieties of moral indeterminacy. As well as clarifying our subject matter, this initial survey will allow us to identify the relevance of global or local *skepticism* about moral indeterminacy to subsequent discussion.

A. Indeterminacy in (the Application of) the Moral Code

Suppose one thinks that there is a basic moral principle that (absolutely) prohibits promise breaking. Then the notion *promising* appears in the

2. Timothy Williamson, *Vagueness* (London: Routledge, 1994).

3. See J. Robert G. Williams, “Nonclassical Minds and Indeterminate Survival,” *Philosophical Review* 123 (2014): 379–428, and “Indeterminacy, Angst and Conflicting Values,” *Ratio* 29 (2016): 412–33.

4. See J. Robert G. Williams, “Decision Making under Indeterminacy,” *Philosopher’s Imprint* 14 (2014): 1–34.

5. Williams, “Indeterminacy, Angst and Conflicting Values.”

very content of our basic moral code. It is because that notion admits borderline cases that we get moral indeterminacy, and this moral indeterminacy can be leveraged, as above, to construct a scenario in which it is indeterminate what an agent ought do.

This kind of moral indeterminacy illustrates our first broad class of cases of moral indeterminacy.⁶ Other (absolute or overridable) prohibitions give the same kind of result: a prohibition on killing an innocent person will plausibly generate Indeterminate Ought scenarios if we trade on the vagueness of killing, innocence, or personhood.⁷ The moral code need not even be deontological to deliver cases of this kind. Moral axiology—concerned with assigning moral value to outcomes—is *prima facie* not crystalline. Perhaps moral demerits are assigned to outcomes if a promise breaking or killing of the innocent occurs therein; indeterminacy in these features will then generate indeterminacy in the value of outcomes.

There is an alternative indirect route to moral indeterminacy through promising. It can be indeterminate whether a promise would be broken by a certain outcome (and therefore, whether a prohibition is violated or disvalue accrues) when it's determinate that it was promised that *p*, but subsequent actions leave it indeterminate whether *p*. Thus, vagueness in what counts as “faithfulness” can lead to indeterminacy in whether a vow of fidelity was kept. A commitment to one’s family not to work on the weekend is vulnerable to the indeterminacy of “work” and “week-end,” and a contract to install a new main water pipe safely is vulnerable to the vagueness of “safe.” The moral code need not use the vague vocabulary of *faithful*, *weekend*, or *safe* in order to generate these moral indeterminacies: the role it gives to linguistically framed social institutions of promising allows individuals or society to turn almost any indetermi-

6. The moral code need not be absolutist to get this result. An overridable moral prohibition on breaking a promise may still generate an Indeterminate Ought scenario, in a case in which the putative prohibition would not be overridden.

7. Vague personal identity makes for a particularly rich source of indeterminacy in the content of the moral code. Vagueness in who is the same person as whom can lead to indeterminacy in who has obligations or legitimate demands as a result of earlier happenings or how to value outcomes. For just one example: distributional justice requires that suffering not be concentrated over time in a single individual, and this contributes to the (dis) value of an outcome. But in a situation in which it is indeterminate whether two people in the past or in the future count as the same person (because, e.g., of psychological disruption in the intervening years) it will be indeterminate whether the suffering is so concentrated and so whether an outcome is or is not distributionally unjust. For raw materials for indeterminate personal identity, see Bernard Williams, “The Self and the Future,” *Philosophical Review* 79 (1970): 161–80 (repr., *Problems of the Self* [Cambridge: Cambridge University Press, 1973], 46–63); Peter van Inwagen, *Material Beings* (Ithaca, NY: Cornell University Press, 1990); and Theodore Sider, “Criteria of Personal Identity and the Limits of Conceptual Analysis,” in “Philosophical Perspectives,” suppl., *Noûs* 15 (2001): 189–209. My “Indeterminacy, Angst and Conflicting Values” contains ethical decision problems based on such cases.

nacy in any sentence into a moral indeterminacy, if the appropriate promise is undertaken (or law written or contract made). So the moral code itself might be written in a perfectly precisified vocabulary, and moral indeterminacy could still arise by this indirect route.

B. Structural Indeterminacy

The cases discussed so far trace Indeterminate Oughts back to vagueness in specific prohibitions or axiological rules. But another important category of moral indeterminacies generating Indeterminate Oughts arises from the *structural* features of morality: how value assignments or constraints come together to determine which acts are permissible. A first example involves value aggregation. You have a bag of goods to deliver to a charity shop and have to turn left (to the homeless charity) or right (to the arts charity). It may be quite clear how much marginal value the money earned will realize in relieving suffering on the one hand and how much marginal value they would have in fostering arts on the other. The difficulty is how to aggregate these two different kinds of value—to determine the rate of exchange between them in order to arrive at an overall value to the act of turning right or turning left.⁸ I follow Broome in diagnosing the difficulty in such ‘superhard’ or ‘incommensurable’ decision problems as one of indeterminacy: there is just no fact of the matter concerning what quantity of extra artistic endeavor among poor kids in the community is worth one extra night of someone sleeping rough.⁹ Giving to the arts charity and giving to the homeless charity will both realize some pro tanto goods, but it will be indeterminate whether the overall value of giving to the homeless charity is greater or lower than the overall value of giving to the arts charity, since it is indeterminate how to compare the disparate goods.¹⁰

Aggregation indeterminacy is not the only kind of structural moral indeterminacy. It’s plausible that agents have a range of permissible options about the extent to which they factor their own self-interests into

8. I take this case from Miriam Schoenfield, “Decision Making in the Face of Parity,” in “Philosophical Perspectives,” suppl., *Noûs* 28 (2014): 263–77.

9. See John Broome, “Is Incommensurability Vagueness?” in *Ethics out of Economics* (Cambridge: Cambridge University Press, 1998), 123–44 (originally published in *Incommensurability, Incomparability and Practical Reason*, ed. Ruth Chang [Cambridge, MA: Harvard University Press 1997], 67–89).

10. Another example of aggregation indeterminacy involves trade-offs between overridable prohibitions and the value of consequences—for a deontologist but nonabsolutist about promising, one’s promise to vote for Mary will mean one objectively ought to vote for her when the consequences of voting for Fred are only marginally better; when the consequences are massively better, the promise is overridden. In the middle we can expect borderline cases. A third possible example of aggregation indeterminacy would be when two overridable promises conflict.

evaluating outcomes.¹¹ This is critical to understanding permissible (but not obligatory) self-sacrifice. If I put myself in danger (leap off the bridge onto the tracks) to save a child's life, that is heroism and permissible. If I push you off the bridge, with the same result, that is assault and impermissible. I have the liberty of factoring in prospective benefits to myself to a greater or lesser extent in evaluating a prospect resulting from an action open to me. Thus, even if the content of prohibitions or axiological rules doesn't mention personal identity, the structural role of agents' options means that indeterminacy in whether they are the same person as one of those affected can produce indeterminacy in what it is permissible to do.¹²

C. *Evasive and Eliminative Maneuvers*

In the cases above, it's worth noting the gap between claiming that there is indeterminacy in some morally relevant fact and the claim that this indeterminacy in turn produces indeterminacy in what ought to be done. In many cases, there will be plausible principles bridging this gap. But for all we've said, one's overall theory might seek to insulate evaluation of action from this underlying indeterminacy. For example, an overridable prohibition on breaking promises may say that it is automatically overridden if it is a borderline case. Structural rules might tell us how to weigh borderline prohibitions against good outcomes. An evasion strategy is one which avoids the Decision Ought Challenge for indeterminacy, not by denying moral indeterminacy in the broadest sense but by denying that moral indeterminacy gives rise to indeterminacy in what one objectively ought do. We'll be considering a (limited) evasion strategy in the discussion of truth-centric indeterminate axiological rules and prohibitions in Section IV below.¹³

11. I am fond of (a slight variant of) the treatment of this phenomenon put forward in Ted Sider, "Asymmetry and Self-Sacrifice," *Philosophical Studies* 70 (1993): 117–32.

12. Consider cases in which I know now that certain benefits which I have the option of taking now will damage the interests of a future person x , where it is indeterminate whether I am x . If the benefits are saving the child and the damage is x 's death, we have a recipe for a case in which it is indeterminate whether the act is heroism or assault, because it is indeterminate whether x is someone whose interests moral theory allows me to discount. See also n. 7 above.

13. One evasion strategy in cases of ethical uncertainty is to appeal to *moral safety*—the idea that when it would be wrong to bring about a certain outcome, it's also wrong to risk bringing about that outcome. One might try to extend this from riskiness to indeterminacy, by arguing that if it's wrong to bring about a situation such that p , it is wrong to bring about a situation in which it is not determinately not p . I do not consider such derivative constraints here explicitly, but some of the "dilemma-style" accounts of the decision ought in Sec. IV might be seen as articulating this thought. For discussion in the context of vagueness and the ethics of abortion, see Roy Sorensen, "Vagueness and the Desiderata for Definition," in *Definitions and Definability*, ed. J. H. Fetzer, D. Shatz, and G. Schlesinger (Dordrecht: Springer, 1991), 71–109.

The many kinds of moral indeterminacy illustrate how radical one would have to be to avoid the Decision Ought Challenge for indeterminacy by denying the very existence of moral indeterminacy. The option might be pursued: seeking to eliminate vagueness in the content of the moral code, one might argue that our familiar formulations of moral rules are mere approximations to patterns that arise in virtue of the fundamental moral code, to be written in some crystalline, utterly precise ‘moralese’.¹⁴ But more local skepticism is more plausible: one could question whether one specific kind of indeterminacy put forward above is really rightly classified as indeterminacy. Ruth Chang argues that the distinctive puzzles of moral aggregation do not reflect moral indeterminacy.¹⁵ A consequentialist might reject vague moral prohibitions because she does not admit a theoretically significant category, *moral prohibition*, in the first place.

One who feels tempted by global evasion or global skepticism about Indeterminate Oughts will have to explain away all apparent instances of moral sorites sequences. These consist of an ordered finite sequence of possible actions, where the first case is clearly morally permissible, where the last case is clearly morally impermissible, and where, due to their extreme similarity in morally relevant respects, for each N we seem compelled to deny that case N is morally permissible and case $N + 1$ is morally impermissible. Examples are plentiful in the literature. Constantinescu gives examples of the quantity of lies compatible with a moral imperative to be honest, with the intensity of pain inflicted compatible with a prohibition on being cruel, and (borrowing from Sorensen) with the quantity of blood one can extract from a person without acting wrongly.¹⁶ Prima

14. The notion of moralese (a supposed language in which the most fundamental moral truths are articulated) is designed to echo *ontologese* (a supposed language in which the most fundamental metaphysical truths are articulated). In each case, ordinary moral or ontological truths will be grounded in the truths that can be formulated in moralese or ontologese. On a certain kind of hard-core moral realism, moralese will be a proper part of ontologese (Tristram McPherson, “What Is at Stake in Debates among Normative Realists?” *Noûs* 49 [2015]: 123–46), but I do not build that into the notion. The claim that ontologese is precise is one way to articulate a prominent thought that there is no metaphysical vagueness (Elizabeth Barnes, “Fundamental Indeterminacy,” *Analytic Philosophy* 55 [2014]: 339–62). Despite its currency, it’s hard to identify good reasons for this assumption. The status of moralese is similar. It’s difficult for me to envisage what crystalline moralese would look like. Perhaps impersonal hedonistic utilitarianism would be a candidate but only if it could never be indeterminate how many hedons of pleasure a given outcome realizes—and why believe that?

15. See Ruth Chang, “The Possibility of Parity,” *Ethics* 112 (2002): 659–88, and also her introduction to *Incommensurability, Incomparability and Practical Reason*, 1–34. For discussion and a response to Chang, see my “Indeterminacy, Angst and Conflicting Values.”

16. See Cristian Constantinescu, “Moral Vagueness: A Dilemma for Non-naturalism,” in *Oxford Studies in Metaethics*, ed. Russ Shafer-Landau (Oxford: Oxford University Press,

facie, sorites susceptibility shows that a term is vague and has borderline cases. And the cases above seem to show that thick and thin ethical terms, and in particular permissibility, have borderline cases. The latter case means that not just moral indeterminacy in general but Indeterminate Oughts in particular will be hard to avoid. Moral sorites are not the best starting point for positive theorizing, since they are nonconstructive—unlike the examples in our list above, they do not tell us what aspect of moral theory is producing the indeterminacy. But they are an effective rebuttal to global skeptics and evaders.

My project below is to explore how decision oughts work under whatever moral indeterminacy there may be. For that project, I need to cast the net as wide as possible, to provide a toolkit of interest to as many moral theorists as I can, and for this dialectical reason I will set aside local as well as global skepticism about moral indeterminacy.

III. REDUCTION TO QUOTIDIAN UNCERTAINTY

If our uncertainty over p , when we know that p is indeterminate, were quotidian, then this article would be a short one. The Decision Ought Challenge for indeterminacy would just be a special case of the Decision Ought Challenge for uncertainty. Consider the case in which we know it is borderline whether what we did counted as a promise to vote for Mary, and hence indeterminate whether we objectively ought to vote for Mary. The suggestion is that this produces a situation in which we are uncertain whether we promised and so uncertain whether we objectively ought to vote for her. Indeterminacy may be a distinctive source of quotidian uncertainty, but it's unclear why the source of the uncertainty should matter any more than whether the uncertainty results from sources such as unreliable testimony, inexact observation, or lack of confidence in recall. Whatever we say about the general case should just be applied to the special case of indeterminacy-prompted uncertainty. There would be no independent Decision Ought Challenge for moral indeterminacy.

The reduction of the puzzle of moral indeterminacy to the puzzle of moral uncertainty would be attractively economical, but its starting assumption is prima facie peculiar. When one is (in the ordinary sense) uncertain whether p , one can wonder whether p , gather evidence for or against p , hope or fear that p will turn out to be the case. But in the case of indeterminacy, one could simply know that there's no fact of the matter whether p , and hoping, wondering, evidence gathering, and the other accompaniments of ordinary uncertainty seem prima facie out of place

2014), 9:152–85; and Roy Sorensen, *Vagueness and Contradiction* (Oxford: Oxford University Press, 2001).

when one knows already that there is no fact out there to be hoped for or wondered about.

The hypothesis to be explored in this section is that the above betrays a mistaken prototheory of indeterminacy. To substantiate such a charge, we need a conception of indeterminacy which predicts and explains quotidian uncertainty in the relevant cases. There is a natural candidate here: Williamson's version of the *epistemicist* theory of vagueness.¹⁷ As I will understand this proposal, it asserts there is a fact of the matter concerning which pebble addition makes a nonheap into a heap or which color shade marks the boundary between blue and nonblue—or indeed, as to who is the same person as whom, how values aggregate into overall moral value, under what conditions one has promised, and so forth. According to Williamson we do not and practically cannot know where these sharp boundaries lie, which generates an irresolvable but quotidian uncertainty over whether a borderline case of promising is genuinely a promise, and so forth. Williamson offers an explanation of why knowledge cannot be obtained in such cases: he argues that any polar judgments concerning borderline cases could easily be false, consistent with all our evidence, and for that reason will never count as knowledge (they're "unsafe").

What entitles Williamson to the claim that judgments in the borderline area are unsafe? He offers a distinctive explanation: the boundaries of vague terms according to Williamson are sharp but unstable—"heap," "blue," or "permissible" will pick out distinct properties, with different sharp cutoffs, in nearby possible situations, depending primarily on the usage of the term across the whole community. Someone who inspected a sorites sequence and formed a judgment about the location of the cutoff for blue—believing that patch 545 was blue and 546 was not blue—might in fact believe truly. But that belief, formed on that basis, is very easily false, since the cutoff for the term would shift with just minor variations of remote usage. So the Williamson picture is this: our attitude to x is red when x is a borderline red/orange should be quotidian uncertainty, where this is uncertainty with a peculiar, distinctive source in *metasemantic plasticity*—the fact that they have an unstable extension that sensitively depends on the details of usage.

For our purposes, the significance of epistemicism is just to underwrite the central reductive claim that attitudes to indeterminacy are just special cases of quotidian uncertainty. And the significance of that is to pass the buck to whatever independent theories of quotidian Uncertain Oughts we have on the shelf. This is not the place to tackle that general debate, but I will make two remarks about choice points in this reductive account of decision oughts under indeterminacy.

17. See Williamson, *Vagueness*.

A. *Factual versus Normative Uncertainty*

The first choice point is whether the uncertainty induced by indeterminacy should be represented as factual uncertainty relevant to the application of known moral principles or normative uncertainty about what those moral principles are. This is vital when the two cases are handled differently. For example, if I am factually uncertain about the consequences of pressing a button—whether it will give us all nasty headaches or make good food appear on the table, I may still be certain about the value of each possible outcome. Plausibly what I ought to do, in the decision-ought sense, is that which maximizes the subjective expected value of the terms: the average value across possible outcomes of the act, weighted by my confidence that the action would bring about the respective outcome. On the other hand, suppose I know the outcome of a possible act but am uncertain about the value of that outcome. I may know what happens to Alpha if she steps into a virtual reality simulation of her perfect life but be uncertain whether the correct axiology would say that this is a good outcome (for Alpha or in general) or a poor outcome, because I do not know how much it normatively matters that Alpha's pleasure is derived from social relations with real-world people rather than figments of the program. Rival coherent axiologies would rate the scenario differently, and my uncertainty is normative uncertainty if it's over which one of these is correct.

My own view is that normative uncertainty should be treated very differently from factual uncertainty; indeed, I think normative uncertainty should simply be set aside in determining what an agent decision-ought to do.¹⁸ A simple model for my view would be the view that you decision-ought to do something if that thing maximizes the subjective expectation of objective value. Factual uncertainty is factored into this quantity—it determines what is expected. But an agent's uncertainty over objective value itself is irrelevant to this calculation.

If that is right, determining which category of uncertainty moral indeterminacy produces is important.¹⁹ Schematically, suppose that an ax-

18. See Gideon Rosen, "Culpability and Ignorance," *Proceedings of the Aristotelian Society* 103 (2003): 61–84, and "Scepticism about Moral Responsibility," *Philosophical Perspectives* 18 (2004): 295–313; and Elizabeth Harman, "Does Moral Ignorance Exculpate?" *Ratio* 24 (2011): 443–68, for recent contributions to the debate on whether moral ignorance is exculpatory. I am taking it here that it is not, though all I really need to make the point is that moral ignorance/uncertainty and factual ignorance/uncertainty receive different treatments. There is a further topic—the focus of discussion of 'normative uncertainty'—about what an agent (who desires to be moral) is rationally required to do, given uncertainty about what morality demands. That raises some interesting questions but is not the question that we are dealing with here. For this view, see Andrew Sepielli, "What to Do When You Don't Know What to Do," *Noûs* 48 (2013): 521–44.

19. The exception is indeterminacy outside the moral code. Where one has promised that *p*, and there is a moral rule prohibiting promise breaking and an outcome where (it is

iological rule tells us that any outcome gets value 1 iff that outcome is such that D holds, and value 0 otherwise, and that w is an outcome in which it is borderline/indeterminate whether D—which according to the current analysis means we’re uncertain (under the supposition that w is actual) whether D holds, that is, whether being in w is being in w^* (like w but where D obtains) or being in w^{**} (like w but where D fails to obtain). Now, if the vagueness of D is factual uncertainty, and if I divide my credence evenly across w^* and w^{**} , conditionally on w obtaining, then an act known to produce w will have expected value 0.5. But if it is normative uncertainty, then the expected value will be either +1 or 0 (and my normative uncertainty will prevent me knowing which). If I have to choose between bringing about w and an action with expected value 0.8, then the former but not the latter model will imply that I’m in a position to know what I decision-ought to do.

I am inclined to think of indeterminacy in the moral code as issuing (given epistemicism) in factual not normative uncertainty, but I am not aware of any compelling arguments either way. Some considerations both ways are given in a footnote.²⁰

B. Models of Quotidian Uncertainty

The second choice point concerns the character of quotidian uncertainty produced by indeterminacy. One thought is that when we know that something is paradigmatically borderline blue, or a borderline case of a person or of a permissible act, then we should have an intermediate

known that) it will be indeterminate whether p , then it seems to me that the quotidian uncertainty over p that the epistemicist predicts will be factual rather than normative uncertainty.

20. One might think of the axiological rule in the main text as producing uncertainty about value. After all, an agent can have full knowledge of a canonical description of an outcome w and still be ignorant (in a quotidian sense) of whether or not D. The reason for that is seemingly that there are two possible sharp value functions that the vague specification may, for all the agent knows, pick out (one which maps w to 1, the other that maps w to 0). Ignorance of the One True value function is normative uncertainty. That line of thought suggests that under the epistemicist reduction, axiological indeterminacy is a source of normative uncertainty, not factual uncertainty. But there are aspects of the case that sit oddly with this diagnosis. For example, ex hypothesi we know a canonical formulation of the moral principles in question—to be sure, there’s uncertainty about how to apply it, but our knowledge of the moral rule itself is impeccable. That suggests the epistemicist ignorance is not moral but factual ignorance. We can model this by taking ‘outcomes’ to be doxastic possibilities for the agent. There are two relevant outcomes in the vicinity: one in which the underlying facts are as specified by w , and D is the case, and another with the same underlying facts, where not-D is the case. The vague axiological rule applied to outcomes individuated in this way yields a determinate function from outcomes to values—e.g., we know that the first has value +1 and the latter value 0. But it will be an open question for us which of these outcomes an action that yields D will produce. Framed in this way, our uncertainty is uncertainty about a factual matter.

credence that it is blue/a person/permissible. Another, for skeptics of the degrees-of-belief model of uncertainty, may think of uncertainty as residing in a positive state of suspending judgment.²¹ And in the recent literature, a third intermediate model has been proposed. According to the *imprecise belief* model an agent's cognitive attitudes are graded but not modeled by real-valued measures of confidence (the typical formal model of this involves *confidence intervals* or more generally a set of real-valued credences that the agent is open to).²² Each of these rival models of quotidian uncertainty requires a different treatment of the decision ought under uncertainty, and there's more controversy in some of the cases than in others concerning what this should be.

Some resolutions of this choice point bring new arguments against the reduction of moral indeterminacy to quotidian uncertainty. For example, Chang, in her discussion of aggregating values, identifies a number of characteristics of decision making in such 'superhard cases'.²³ If you are torn by the conflicting values realized by giving to the arts charity or to the homeless, then, for Chang, you are permitted to give to either. On the Bayesian model of uncertainty over which is better, this dual permissibility must be represented as a case in which the expected utilities of each, for the agent, are tied. But on that model an extra penny donated to one or the other option (a moral "sweetener") would break the symmetry. But this is not how the quandary generated by aggregation of values works—even after the sweetener is added, the quandary remains. The Bayesian model fails to capture this, and this might be an argument that something other than quotidian uncertainty is going on. So if quotidian uncertainty is Bayesian, there is reason to reject the epistemicist reduction of value aggregation indeterminacy, at least, to that of quotidian uncertainty.²⁴

C. *Evaluating Moral Epistemicism*

The reduction of Indeterminate Oughts to quotidian Uncertain Oughts needs to be underwritten by a tenable theory of indeterminacy. We've

21. See Jane Friedman, "Rational Agnosticism and Degrees of Belief," in *Oxford Studies in Epistemology*, ed. Tamar Gendler and John Hawthorne (Oxford: Oxford University Press, 2013), 4:57–81.

22. A very helpful survey of formal models of (quotidian) uncertainty can be found in Joseph Y. Halpern, *Reasoning about Uncertainty*, rev. ed. (2003; repr., Cambridge, MA: MIT Press, 2005).

23. See Ruth Chang, "Parity, Interval Value and Choice," *Ethics* 115 (2005): 315–50.

24. Schoenfield, "Decision Making in the Face of Parity," develops an interesting argument against using imprecise beliefs (at least as standardly modeled with sets of probability functions) as a model of uncertainty characteristic of value aggregation. Her point, if correct, would also refute the indeterminacy-based model defended in my "Indeterminacy, Angst and Conflicting Values" and discussed in Sec. V below.

been taking the underwriting theory to be epistemicism. But rather than replay general animadversions against epistemicism, I want to highlight a single but central worry, local to the moral case.

The basic problem can be expressed straightforwardly. The reference of basic moral terms like “permissibility” remains constant under variations of usage—they are *metasemantically stable*. On the other hand, Williamson’s epistemicism makes the denial of metasemantic stability (what we earlier called metasemantic plasticity) the heart of what vagueness is.²⁵ So this account is inapplicable to the vague term “permissibility.”

That “permissible” is not only stable but massively stable is the lesson of Horgan and Timmons’s *moral twin earth*.²⁶ Recall the setup: Earth and Twin Earth are similar in many respects. Visitors from one to the other would assume they were speaking the same language, in the same natural environment. But there’s a crucial difference: the property that causally regulates the deployments of “permissible” in Twin Earth is *according with the categorical imperative*, whereas the property that plays that role on Earth is, we’re supposing, some suitably nuanced consequentialist property. Other than that, “permissibility” plays analogous roles in both communities—a basis for praise and blame, a guide for action, and so forth. Visitors from Earth to Twin Earth may note the difference in opinion about what counts as permissible by the lights of the new society. But—the claim is—they would regard this as a first-order disagreement about what is permissible, not as a discovery that the two societies are talking about different things when they use the counterpart words. The moral of moral twin earth is that even drastic differences in communal dispositions to deploy “permissible” do not shift its boundaries—it is metasemantically stable.

Macro-stability of the referent of “permissibility” is in principle compatible with *microplasticity*. Like a stone loosely stuck in a hollow, a position may be unstable on the scale of millimeters but not be shiftable by meters or miles. However, the Twin Earth case seems to be just as strong if we imagined the two communities had, for example, systematically different dispositions for applying “permissible” to borderline cases. The ur-phenomenon here is that a term possessing the kind of role in practical reasoning of “permissibility” will continue to pick out *permissibility* under both macro- and microvariations of usage.²⁷

25. Compare the discussion of moral twin earth and moral vagueness in Miriam Schoenfield, “Moral Vagueness Is Ontic Vagueness,” *Ethics* 126 (2015): 257–82.

26. See Terence Horgan and Mark Timmons, “Troubles on Moral Twin Earth: Moral Queerness Revived,” *Synthese* 92 (1992): 221–60.

27. Another argument in the literature, found in Constantinescu’s “Moral Vagueness” and prefigured in Tom Dougherty’s “Vague Value” (*Philosophy and Phenomenological Research* 89 [2013]: 352–72), is the charge that epistemicism about moral matters, with its implica-

D. Summary

This part has evaluated whether a theory of vagueness can (i) underwrite the reduction of attitudes to the indeterminate with ordinary uncertainty and (ii) allow vague terms to be metasemantically stable. As Schoenfield notes, many other accounts that cohere with (i) have the same problem with (ii)—among them Kearns and Magidor’s doctrine of arbitrary reference and Sorensen’s truth-maker gap epistemicism.²⁸

I do know of one theory that promises both (i) and (ii). This is the account of metaphysical indeterminacy outlined by Barnes and Williams as applied by Barnes and Cameron to the case of indeterminate future contingents (e.g., tomorrow, the trains will be halted by flooding in the valley).²⁹ It is important to Barnes and Cameron that their account does not demand that we revise our quotidian attitudes to future contingents; those attitudes being, in paradigm cases, uncertain. The theory of indeterminacy that this is built on takes indeterminacy to be a primitive feature of the world, and there’s nothing comparable to the plasticity required by the theories of Williamson and others.³⁰ Despite some concerns I have about the Barnes and Cameron development of the view,³¹

tion that there exist unknowable moral truths, is in tension with epistemic constraints on morality, i.e., *Epistemic constraint on reasons*.—If R is a normative reason for X to Φ , then R can feature in a rational justification of the claim that X ought to Φ , a justification which X knows or could come to know if X’s reasoning were maximally improved. Crispin Wright, “Vagueness: A Fifth Column Approach,” in *Liars and Heaps*, ed. J. C. Beall (Oxford: Oxford University Press, 2003), 84–105, provides an elegant argument against epistemicism about vagueness in any domain subject to an epistemic constraint. Interestingly, his argument would generalize to any ‘verdict excluding’ theory of vagueness, including, e.g., the truth value gapper theory below.

28. Schoenfield, “Moral Vagueness Is Ontic Vagueness.” See Stephen Kearns and Ofra Magidor, “Semantic Sovereignty,” *Philosophy and Phenomenological Research* 85 (2012): 322–50; and Sorensen, *Vagueness and Contradiction*.

29. See Elizabeth Barnes and J. Robert G. Williams, “A Theory of Metaphysical Indeterminacy,” *Oxford Studies in Metaphysics*, ed. Karen Bennett and Dean W. Zimmerman (Oxford: Oxford University Press, 2011), 6:103–48; and see Elizabeth Barnes and Ross P. Cameron, “The Open Future: Bivalence, Determinism and Ontology,” *Philosophical Studies* 146 (2009): 291–309, and “Back to the Open Future,” *Philosophical Perspectives* 25 (2011): 1–26.

30. This suggestion would cohere with the lesson Schoenfield draws from her discussion in “Moral Vagueness Is Ontic Vagueness”: that (for a moral realist) moral vagueness is ontic vagueness. Constantinescu argues in “Moral Vagueness” against the combination of nonnaturalist moral realism and ontic vagueness, though I believe the treatment of primitive metaphysical indeterminacy/ontic vagueness developed in Barnes and Williams, “Theory of Metaphysical Indeterminacy,” is not vulnerable to his charges.

31. My concern relates to a challenge directed at supervaluationism in chap. 5 of Williamson’s *Vagueness*—whether it can distinguish itself from epistemicism. Field suggests that we appeal to the distinctive cognitive role of indeterminacy, i.e., the supposed fact that attitudes to indeterminacy are not quotidian uncertainty, to respond to Williamson’s concerns. See Hartry H. Field, “Indeterminacy, Degree of Belief, and Excluded Middle,” *Noûs*

it is a more promising theoretical option for trying for a reductive approach to the Decision Ought Challenge for indeterminacy than the more familiar epistemicist strategy.

IV. NONCLASSICISM

Epistemicism and the other theories we have been looking at so far are logically and semantically conservative. If one of those theories holds, standard classical logic and classical bivalent semantics apply to vague language just as much as the precise language of the exact sciences. Non-classical theories of indeterminacy, on the other hand, think that indeterminacy requires some revision of that canon. Here is Parfit, for example, floating a nonclassical approach to moral indeterminacy:

Some questions may be indeterminate, in the sense that they have no answer. That is sometimes true, for example, of the question “Is he bald?” If some man has no hair, he is bald. If some man has a full head of hair, he is *not* bald. But we cannot plausibly assume that, in all cases between those two extremes, any man must either *be*, or *not* be bald. In many cases, though it is not true that some man is bald, it is also not true that this man is *not* bald. Similar claims might apply to normative questions.³²

The classical law of excluded middle says that for any proposition p , p or not p , and Parfit is saying that in borderline cases we cannot assume this. The classical principle of bivalence is that for any proposition p , either p is true or p is false. Parfit asserts counterinstances, or *truth value gaps*: propositions which are neither true nor false.

If we have a gappy account of normative indeterminacy, as Parfit does, then I see no plausible way of treating Indeterminate Oughts as Uncertain Oughts. If a proposition p is untrue, then you shouldn't have any confidence in it, and certainly not middling or imprecise confidence of quotidian uncertainty. The link between revisionary theories of truth values and positive accounts of the cognitive role of indeterminacy is at the heart of a family of 'nonclassical theories of mind' I've been developing in recent work.

Nonclassical theories of vagueness are many and varied. For present purposes I concentrate on three broad categories. There are gap theories such as Parfit's, on which indeterminate cases are neither true nor false. There are glut theories, associated with philosophers like Dominic Hyde and Graham Priest, on which indeterminate cases are both true

34 (2000): 1–30 (repr. *Truth and the Absence of Fact* [Oxford: Oxford University Press, 2001], 278–311). But of course this is exactly what Barnes and Cameron must deny.

32. Derek Parfit, *On What Matters* (Oxford: Oxford University Press, 2011), 2:559–60.

and false.³³ And there are scalers, for whom truth comes in degrees, with indeterminate cases intermediately true.³⁴

A. *The Nonclassical Axiological Approaches: Gluts, Gaps, Scales*

One of the most striking features of a nonclassical approach is that it has the capacity to silo certain kinds of moral indeterminacy, allowing indeterminacy in the content of the moral code to persist without generating Indeterminate Oughts. An *axiological rule* is something that specifies the value of an outcome. There may be several such rules that apply to a single outcome—if so, they will need to be aggregated in some way (this in turn can lead to structural moral indeterminacy, as we saw earlier). But for the purposes of a simple model, let us suppose that a single axiological rule determines the moral value of an outcome.

We want a sample axiological rule that's simple to state and indeterminacy involving. I will use an artificial example, so one can see the mechanisms in play without the distraction of worrying about the plausibility of the moral axiology. Once the mechanisms are appreciated, the transfer to the case of moral axiology will be straightforward.

Consider the bower bird. The bower bird loves blue things. The bower value of an outcome is equal to the number of blue objects that exist in that scenario. Bower value can and probably will be indeterminate, for outcomes very often have borderline cases of blue things in them.

At first it appears that this axiological indeterminacy generates ought indeterminacy. Suppose that the bower bird has a choice between pecking to the right to create one additional blue object or pecking to the left to create two objects that are borderline blue or refraining from

33. See Dominic Hyde, "From Heaps and Gaps to Heaps of Gluts," *Mind* 106 (1997): 641–60; and Graham Priest, *In Contradiction: A Study of the Transconsistent*, 2nd expanded ed. (Dordrecht: Nijhoff, 1986; Oxford: Oxford University Press, 2006).

34. See, e.g., Kenton F. Machina, "Truth, Belief and Vagueness," *Journal of Philosophical Logic* 5 (1976): 47–78, repr. *Vagueness: A Reader*, ed. Rosanna Keefe and Peter Smith (Cambridge, MA: MIT Press, 1997), 174–204; Nicholas J. J. Smith, *Vagueness and Degrees of Truth* (Oxford: Oxford University Press, 2008); Williams, "Nonclassical Minds." My personal preference is for supervaluational-style variants of all three approaches (supervaluationism, subvaluationism, and degree supervaluationism, respectively). Such approaches always endorse excluded middle and never endorse explicit contradictions. There are more radical gappers, glutters, and scalers—and Parfit, denying excluded middle, is among them—however, the differences make no difference for present purposes, and they are covered by my taxonomy. Not covered are those who posit non-real-valued (e.g., partially ordered) truth values—though in Dummett's terminology, it is *assertoric* rather than *ingredient* truth value that matters here, e.g., the structure of what plays the role of a designated value (see Michael Dummett, *The Logical Basis of Metaphysics* [Cambridge, MA: Harvard University Press, 1991], 47–49). Even intuitionism can be represented within the taxonomy—the status of being forced in a Kripke model acting as truth value 1 and otherwise 0. Field's skepticism over the very existence of semantically relevant truth values is the most interesting and well-developed exception to my approach that I'm aware of.

pecking and sustaining the status quo. Let's suppose that the status quo features N determinately blue things, and no borderline cases, for a value of N . Pecking right then has bower value $N + 1$, refraining from pecking has bower value N , and pecking left is indeterminate in bower value. If the new objects are blue, then it has bower value $N + 2$ and is optimal. If the new objects are not blue, then it has bower value N and is suboptimal. But it is indeterminate whether the new objects are blue and accordingly indeterminate whether the bower bird should peck to the right or left.

Everything turns on the bower value that our rule assigns to the outcome produced by pecking left. Notice that the following two rules specify the same function, given standard classical assumptions:

(Rule 1) $BV(w) = N$ iff in w there are exactly N blue things.

(Rule 2) $BV(w) = N$ iff in w there are exactly N things such that it is true that they are blue.

What is crucial to bear in mind is that these are inequivalent in many nonclassical settings. This is because their classical equivalence derives from the principle:

(Truth Equivalence) p is true iff p .

Anyone who wants to (a) maintain modus tollens, (b) reject contradictions, and (c) assert truth value gaps must give up this equivalence schema. Suppose for reductio we have Truth Equivalence. The existence of a truth value gap tells us that it is not true that p , nor true that not p . Tollens on the first conjunct to get not- p . Tollens on the second to get not-not- p . Putting them together results in a contradiction. Formally:

1. $\text{True}(p) \leftrightarrow p$	assumption for reductio
2. $\neg\text{True}(p) \wedge \neg\text{true}(\neg p)$	truth value gap
3. $\neg\text{True}(p)$	from 2
4. $\neg p$	modus tollens on 1 and 3
5. $\neg\text{True}(\neg p)$	from 2
6. $\neg\neg p$	modus tollens on 1 and 5
7. $\neg p \wedge \neg\neg p$	from 4 and 6

Something must go, and the truth value gapper pins the blame on the equivalence schema.³⁵ In doing so, they make rule 1 and rule 2 inequivalent.

35. See Crispin Wright, *Truth and Objectivity* (Cambridge, MA: Harvard University Press, 1992). Wright, holding fixed Truth Equivalence, uses this line of reasoning against truth value gaps.

Consider what the gapper says about the axiological rule specifications above. Applied to the key scenario, rule 1 as before tells us that the value of pecking left is indeterminate between N and $N + 2$. But since the two new items are indeterminate cases, for the gapper it is not true that they are blue. So they do not contribute to bower value, according to rule 2. According to rule 2, the bower value of that scenario is N , just the same as the status quo. So rule 2 makes pecking right the only option that increases bower value and so determinately the one to choose. The lesson, first and foremost, is that single axiological rules undergo fission when we move to a nonclassical setting—typically into one that produces indeterminate value assignments and others that do not.

We have seen my suggestion for one nonclassical theorist—the truth value gapper. But the point extends to other theorists. For the glutter, the value assigned by rule 1 is again indeterminate, but rule 2 will assign value $N + 2$ to the outcome produced by pecking left. After all, that outcome contains two additional borderline blue things, and by the lights of the glut theory, that is two more things such that it is true that they're blue. Again, indeterminacy in what one ought to do is removed, this time replaced by the advice to choose to peck left.

What of those who, rather than positing gaps or gluts in truth value, posit truth that comes in degrees? Such scalars can appeal to a generalization of rule 2 above:

(Rule 3) $BV(w) = N$ iff in w , $N =$ the sum of the degree to which x is blue, for every x .

If we think of the classical case as positing just two degrees of truth—1 (fully true) and 0 (fully false)—then in the classical case rule 3 is yet another equivalent reformulation of rule 1 and rule 2. But on the scalar's theory, an indeterminate blue thing (blue to degree 0.6, say) will contribute 0.6 to the overall bower value of an outcome. In the case at hand, two 0.6 blue items would give $N + 1.2$ as the bower value of pecking left; two 0.4 blue items would give $N + 0.8$ bower value to pecking left. So which act is recommended to the bowerbird depends on the exact degree of truth involved. Yet again, no indeterminate value is induced.

Individuating rules by semantico-logical equivalence, what we have is two ways of picking out one and the same rule in a classical setting. There is no reason to even choose one over the other, until we introduce considerations of indeterminacy. In the context of a nonclassical treatment of indeterminacy, however, rules 1 and 2 pick out different axiological rules. And one of those rules does not assign indeterminate value. There's no replacement of one rule by another but a refinement of a new ambiguity. Dialectically, this means that someone who wishes to argue that we get indeterminate value (and Indeterminate Oughts) from

indeterminacy in this part of the content of the moral code needs to give positive reasons for the favored disambiguation. The key point here is not that nonclassical settings enforce the truth-centric, indeterminacy-evading formulation of axiology (rule 2), but they show how indeterminacy in the moral code may be siloed and not produce any indeterminacy in value or in what one ought to do.

The point generalizes to deontological prohibitions and moral indeterminacy outside the code, with minor variations. In the nonclassical setting, a moral constraint not to break promises is one thing; a moral constraint that it never be true that one breaks promises is another; and a moral constraint that, on promising that p , you ensure p is true is a third. But all are reformulations of a single rule in the classical setting. Suitably chosen truth-centric formulations will allow us again to keep the moral indeterminacy from infecting the question of what one ought to do.

For the theorist seeking to understand moral indeterminacy, these points are central, as they show in a nonclassical setting that (i) the puzzle of Indeterminate Oughts need not arise simply from indeterminacy in axiology or prohibitions but also that (ii) what values outcomes have, or what actions are prohibited, can depend critically on which theory of indeterminacy (gap, glut, or scale) we have as a backdrop. These are points to take to the bank straight away. The strategy, however, does not generalize very naturally to all cases of moral indeterminacy—for example, aggregation indeterminacies. Even the truth-centric principles may admit borderline cases, if there is higher-order vagueness in the concepts involved. And, recall, moral sorites for permissibility seem to show us that we've good reason to think that ultimately Indeterminate Oughts are unavoidable. For all these reasons, we need to consider what the nonclassicalist can say about Indeterminate Ought scenarios themselves.

B. Nonclassical Permission: Gluts, Gaps, Scaled Permission

Consider again the decision whether to turn left (to the arts charity) or right (to the homeless charity). We will suppose this delivers an Indeterminate Ought scenario—a case in which it is indeterminate whether it is better to give to the arts or to the homeless and so indeterminate whether one (objectively) ought to turn left or instead right.

The propositions *one ought to turn right* and *one ought to turn left* are indeterminate in this scenario, as indeed are *it is permissible to turn right* and *it is permissible to turn left*. The respective nonclassical theories of indeterminacy have distinctive things to say about their alethic status. Gappers will say none of these is true. Glutters will say that all of them are. Scalpers will say that each is true to a degree.

What attitude should the rational agent take to the proposition *that they ought to turn left*, given all this? The earlier reductive account had it that this would be ordinary quotidian indeterminacy. The nonclassical

theory of mind for which I argue elsewhere deploys truth-centric normative principles for belief to answer this question in the present setting.³⁶ The net result is that gappers should have no confidence either in an indeterminate proposition or its negation (after all, they know that neither is true); glutters should have full confidence in both (they know each is true); and scalers should have partial confidence, aligned to the known degree of truth. Simply applying these general principles to the Indeterminate Ought scenario, gappers will have no confidence at all that either option is permissible, glutters will have full confidence that both are, and scalers will have partial confidence in each claim.

Does this help us make progress with the Decision Ought Challenge under indeterminacy? I will concentrate on gappers and glutters first. We find two natural ways of extracting an account of the decision ought in these settings. My view is that (in light of the purpose of having a decision ought in the first place) the gapper should choose one of these accounts, and the glutter the other, with the upshot that they converge on a common account (in extension) of what one decision-ought to do. I will argue that that common account is problematic, however.

Note that if a glutter turns right, then she will be turning right while fully believing that what she does is (objectively) permissible—indeed, obligatory.³⁷ From the agent's point of view, turning right is fully normatively endorsed. Contrast the gapper: if she turns right, she will be doing so while utterly rejecting the claim that what she does is permissible. The gappers' position seems far less happy than the glutters under this description. A first proposal, then, is that options are decision permissible just in case it is true that they are (objectively) permissible. When faced with situations in which it is indeterminate what is objectively permissible, the glutter will see multiple acts that are decision permissible, while the gapper will see nothing that is decision permissible. On this proposal, for the glutter, the situation resembles cases of moral ties; for the gapper, a moral dilemma.

Hard on the heels of this account of the decision ought is its dual. Rather than focusing on attitudes to permissibility, we might turn to attitudes to impermissibility. The glutter thinks all available acts are forbidden, while the gapper rejects each impermissibility claim. The glutter

36. See my "Generalized Probabilism: Dutch Books and Accuracy Domination," *Journal of Philosophical Logic* 41 (2012): 811–40, "Gradational Accuracy and Non-classical Semantics," *Review of Symbolic Logic* 5 (2012): 513–37, "Nonclassical Minds," and "Indeterminacy, Angst and Conflicting Values."

37. It may seem strange that the glutter believes that turning right is obligatory, when she also believes that an alternative—turning left—is permissible. But the glutter equally believes that turning right is not obligatory and that turning left is not permissible. The truth value glutter believes indeterminate cases are those in which both a proposition and its negation are true and so endorses them both. This is certainly strange, but it is a feature of the glutter's approach to indeterminacy in general.

who takes an action now seems to be in the unhappy position—of doing what is by his or her own lights impermissible. So a second candidate theory of the decision ought would say that options are decision impermissible just in case it is true that they are impermissible. Now the situations are reversed. Under this account, it is the glutter who sees a moral dilemma, in the decision-ought sense, and the gapper who sees the case as one of a moral tie.

I suggest that the package deals that result in decision-ought dilemmas should be rejected. The point of a decision ought was to give implementable advice. Faced with two options, “do either” is implementable advice; “do neither” is not. Accordingly the gapper has strong reason to favor the impermissibility-focused account, and the glutter has reason to favor the permissibility-focused account. What account of the decision ought is fit depends on the background theory of indeterminacy in play. The upshot, as promised, is that both glutters and gappers would see options in Indeterminate Ought scenarios as decision permissible.

But this sort of story has problems. Suppose we face a decision between two options in which one is determinately objectively permissible and the other is an indeterminate case of objective permissibility. Let’s suppose for the sake of argument that there is a moral obligation to keep promises one has made. I have promised to bring the bower bird something blue. In front of me, I have a blue plastic spoon, and a greenish-bluish plastic fork, and these are the only two things I could bring. If I bring him the former, it is determinate that I keep my promise. This is determinately objectively permissible. If I bring him the latter, it is indeterminate whether the promise is kept. It is indeterminate whether this is objectively permissible. In such a case, the thing to do, surely, is to give the bower bird the spoon. Our favored package deals, that is, glutters with the permissibility-centric account of the decision ought or gappers with the impermissibility-centric one, have some explaining to do. It is very unclear how they can account for this case.³⁸

I leave this as a puzzle for glutter and gapper theorists of moral indeterminacy to tackle. It is a distinctive puzzle about their account as applied to indeterminacy in oughts, so even if this is fatal, nonclassical theories may play a role in limiting the number of such cases we have to deal with, by the evasion tactics of the previous section.

38. The other package deals do no better, facing a dual challenge. Change the case so that we have the borderline blue plastic fork and a determinately red plastic knife and have undertaken the same promise to the bower bird. Now, surely, the thing to do is to take the borderline blue fork; bringing the knife should determinately be ruled out. Again, it is not determinate that taking the fork is objectively better, so capturing this intuition is the job of the decision ought. However glutters with an impermissibility-centric account of the decision ought, or gappers with the permissibility-centric one, are committed to a view of the decision ought on which bringing the fork is decision impermissible.

Before turning to other approaches, we should consider what the scaler might say about the decision ought. Here there is a more promising story to be told. In an Indeterminate Ought scenario, the scaler holds that it is true to various intermediate degrees that each option is permissible (and obligatory). Accordingly, a rational scaler will have more or less confidence in the permissibility of each action, matching its known degree of truth. There is a very natural account of decision permissibility for the scaler—an option is decision permissible iff the degree of truth that it is objectively permissible is maximal (and equally common alternatives to maximizing rules, such as satisficing could also be investigated). The maximizing rule handles elegantly the troubling cases for gappers and glutters given above—any determinately objectively permissible option will outrank an option that is an indeterminate case. But equally an indeterminate case of permissibility will outrank a determinately impermissible option. Scaled decision permissibility delivers the right results in these cases. Of all the nonclassical theories, scaling emerges from this survey with the best marks.³⁹

V. INDETERMINATE NORMS

The final approach to Indeterminate Oughts that I will consider builds on the model of decision making under indeterminacy developed in my “Decision Making under Uncertainty.” The logico-semantic approach that lies behind this model I call *classical supervaluational*.⁴⁰ Like the supervaluationist, it posits a range of “sharpenings” of indeterminate terms—and it is indeterminate whether p exactly when “ p ” is true on one sharpening and false on another. It shares this ideology with supervaluational gappers, glutters, and scalars, since a definition of truth as truth on all sharpenings produces truth value gaps, a definition of truth as truth on some sharpening delivers truth value gluts, and a definition of truth to degree d as truth on d percent of sharpenings delivers real-valued degrees of truth. However, the classical supervaluationist re-

39. This answer to the Decision Ought Challenge for indeterminacy predicts that (under full factual information) the decision-permissible option is one which we have most credence is objectively permissible. Famously, the analogous suggestion for quotidian uncertainty, though quite natural, is wrong—as drug or miner’s cases show. See Smith and Jackson’s “Absolutist Moral Theories” for a paradigmatic case against it.

40. It is sometimes also called “nonstandard supervaluationism.” See Vann McGee and Brian McLaughlin, “Distinctions without a Difference,” suppl., *Southern Journal of Philosophy* 33 (1994): 203–51, for early advocacy. A version of this is used in Barnes and Williams, “Theory of Metaphysical Indeterminacy.” My “Decision Making under Indeterminacy” represents my preferred development of the cognitive role of the theory, a rival to the development of the theory in Barnes and Cameron, “Open Future,” and “Back to the Open Future,” discussed in Sec. III.

fuses to endorse any one of these options or their accompanying distinctive accounts of the cognitive role of indeterminacy. Instead, for the classical supervaluationist, truth—that at which belief aims—is truth on the one correct sharpening. But which sharpening is correct is indeterminate. So it is indeterminate whether p is true, when p itself is indeterminate.

Recall that classically equivalent specifications of a single indeterminate axiological principle split into specifications of distinct rules when we adopted one of the nonclassical accounts of indeterminacy, for example,

(Rule 1) $BV(w) = N$ iff in w there are exactly N blue things.

(Rule 2) $BV(w) = N$ iff in w there are exactly N things such that it is true that they are blue.

This is a quite general phenomenon. For example, the following two evaluative schemes for belief are classically equivalent but diverge in the nonclassical setting:

(Accuracy 1) A belief that p is accurate iff p .

(Accuracy 2) A belief that p is accurate iff it is true that p .

From the first we can derive that it is indeterminate whether believing that p is accurate in cases in which p itself is indeterminate. From instances of the second, by contrast, we can derive that a belief that p is inaccurate in the same circumstances, and if the second principle holds determinately, then it would be determinate that a belief that p is inaccurate in those circumstances, explicitly contradicting the first principle. Truth-centric norms on belief such as the second principle are what I use in work on nonclassical theories of mind to underpin the cognitive attitudes to the indeterminate earlier associated with gappers, glutters, and scalers.⁴¹

Dialectically, endorsing classical supervaluationism strips away this phenomenon. We are left just with the original principles and can no longer play with distinctive truth-involving rules. We are left with equiv-

41. The formulation used in Williams, “Generalized Probabilism,” “Gradational Accuracy,” and “Indeterminacy, Angst and Conflicting Values,” is that the accuracy of a credence d in p is measured by the distance of d from p 's truth value. Here the “accuracy measure” is drawn from the literature on foundations for (classical) Bayesianism (e.g., James M. Joyce, “A Non-pragmatic Vindication of Probabilism,” *Philosophy of Science* 65 [1998]: 575–603), and “truth values” are specified for each nonclassical theory in a natural way. For any theory, if p is true, then it has truth value 1; if untrue, then it has truth value 0, and if it has a degree of truth, the truth value is that degree.

alents to the basic material mode principles as the only (nonrevisionary) game in town, for both axiological principles and norms on belief.

A. *Indeterminate Beliefs*

Before turning back to axiological principles and other moral indeterminacies, consider the impact on belief norms. From (the determinate truth of) the above accuracy principle, when p is indeterminate, and hence it is indeterminate whether it is true, then it is indeterminate whether one is accurate if one believes p or inaccurate if one does so. In sum: it is indeterminate what an omniscient agent would believe. It is tempting at this point to think that what this principle is recommending is that one be such that it is indeterminate whether one believes p . Working within a supervaluation setting, we can determinately make it the case that we satisfy the above principle if we believe p on those sharpenings where p and disbelieve p on those sharpenings where not p . So classical supervaluational accuracy norms fit naturally with a psychology of “indeterminate beliefs”—something often posited in passing in the indeterminacy literature. As Susanna Rinard emphasizes, these indeterminate beliefs are an attractive interpretation of the formal models of imprecise belief that are increasingly used for quotidian uncertainty.⁴²

In the literature on indeterminacy, the idea that the appropriate attitudes to indeterminacy are indeterminate beliefs has been defended by several authors, such as Cian Dorr and Michael Caie.⁴³ This goes beyond what we can extract from a truth norm on belief. If we ought to be accurate, that principle tells us that it is indeterminate what we ought to believe—an Indeterminate Ought scenario for belief. By contrast, the principle endorsed by Dorr and Caie is one in which we ought to have indeterminate beliefs, a reversal of the relative scope of indeterminacy and ought.⁴⁴

If one did pair classical supervaluationism with the Dorr-Caie recommendation of indeterminate beliefs, one could try to build an account of decision making under indeterminacy on extant work on deci-

42. Susanna Rinard, “A Decision Theory for Imprecise Credences,” *Philosopher’s Imprint* 15 (2014): 1–16.

43. See Cian Dorr, “Vagueness without Ignorance,” *Philosophical Perspectives* 17 (2003): 83–113; and Michael Caie, “Belief and Indeterminacy,” *Philosophical Review* 121 (2012): 1–54.

44. It seems to me that the Dorr-Caie proposal may be inconsistent with the principle laid down. Take an agent who, determinately, believes that p , and p is indeterminate. It is indeterminate whether this agent is believing accurately, according to the principles I start from. But the agent does not have an indeterminate belief concerning p and so would violate the kind of norm that Dorr, Caie, et al. favor. So even though it isn’t determinately false that this agent’s belief is accurate, they have to argue that it is nevertheless determinately inappropriate.

sion making given indeterminate beliefs. Happily for this approach, there is an extant literature on imprecise belief as a particular kind of quotidian uncertainty, and Rinard argues that such mental states (modeled by sets of probability functions) be interpreted as representing indeterminacy in what beliefs the agent has.⁴⁵ So this holds out the prospect for a novel kind of reduction to quotidian uncertainty, via the claim that some quotidian uncertainty was indeterminacy in belief all along. Less happily for the approach, there's no consensus in the literature on the right model of decision making given imprecise beliefs.⁴⁶ On the upside, Rinard argues persuasively that the indeterminate belief interpretation of imprecise beliefs can resolve this dispute in the literature. On the downside, her suggestion for a decision theory for indeterminate belief, when applied to our kind of decision situation, would simply be the verdict that it is indeterminate how we should act. If Rinard is right about the decision rule appropriate to indeterminate belief, then this is just another way of getting us back to an Indeterminate Ought scenario. Her approach will lead to indeterminacy in rational oughts and moral decision oughts under quotidian uncertainty, so if anything it intensifies the Decision Ought Challenge for indeterminacy.

The Dorr-Caie normative principle goes beyond what we can extract from classical supervaluationism from independently motivated truth norms and does not give us a new line on the Decision Ought Challenge. Accordingly, I will not be assuming it, and what follows begins from the assumption that when it is indeterminate whether p , it is indeterminate whether we ought to believe p .

B. *Mind Making and Decision Oughts*

The mind-making account of decision making under indeterminacy offers the following picture.⁴⁷ When it is indeterminate what you objectively ought to do, and correspondingly what you ought to believe you ought to do (and you're aware of this), you have no option but to make up your mind one way or another. In our original vote-for-Mary scenario, you may make up your mind to vote for Mary or make up your mind not to do so. The distinctive claim here is that you must "make up your mind" in a stronger sense than one who, faced with two equally good options, simply picks one over the other. In "Decision Making under Indeterminacy" I argue, first, that in the relevant sense we are not choosing how to act in this one instance—to make up your mind to vote for Mary in this situation is a resolution to act on any future occasion in accordance with sharpenings of the indeterminacy on which those kinds of sayings count

45. Rinard, "Decision Theory."

46. See Halpern's *Reasoning under Uncertainty* for a survey of formal models in the area.

47. See Williams, "Decision Making under Indeterminacy."

as a promise; making up your mind in the opposite way is similarly an open-ended resolution to act in accordance with the sharpenings on which they do not so count.⁴⁸ Second, I take it that adopting an open-ended resolution to act in this way is to judge that such sayings count as a promise (*mutatis mutandis* for the latter). So in making up your mind how to act in the relevant sense, you make up your mind in a (more literal and) doxastic sense too—you make a judgment call. Both the practical and the doxastic profile of the posited “mind making” in the face of indeterminacy go well beyond the arbitrary picking we’re familiar with from cases in which our options are tied for best.

I have just described what agents will do in an indeterminate scenario if they accord with my theory. But what of the normative question whether they should so act? The bottom line is just as before: it’s indeterminate whether the act you choose is what you objectively ought to have chosen. And analogously, it’s indeterminate whether the belief so formed is objectively accurate and so indeterminate whether it ought to have been formed.

I’ll be defending this picture of groundless but binding judgment calls or acts of mind making as what we decision-ought to do in Indeterminate Ought scenarios. So—at least when there is no accompanying subjective uncertainty to be factored in—the claim is:

A choice to X is decision permissible iff it is not determinately objectively impermissible to X.⁴⁹

This is very close to the decision ought available to the gapper and glutter earlier—I will make the case that it is better motivated and not vulnerable to the objections that sunk those approaches.

Decision permissibility on my account gains its status through being one in which we are immune to neutral sanction. A neutral agent in this context is someone who doesn’t take a stance on relevant indeterminate questions. The agent in an Indeterminate Ought scenario is forced to be

48. The open-ended commitment arises on pain of becoming a “value pump.” If one resolves an Indeterminate Ought on one way on a first occasion, and then on the opposite way on a second, then the net effect can be to leave one determinately worse off. The cases have essentially the same structure as the value-pump arguments for a similar “binding” in value aggregation problems given in Chang, “Possibility of Parity,” and “Parity, Interval Value and Choice,” and the same phenomenon is at the heart of Elga’s challenge to imprecise beliefs (Adam Elga, “Subjective Probabilities Should Be Sharp,” *Philosophers Imprint* 10 [2010]: 1–11). As Williams (“Decision Making under Indeterminacy”) and Rinard (“Decision Theory”) both observe, applying standard decision theoretic roles to sequences of choices (or plans) delivers the result that determinately the value-pump sequences ought not be undertaken.

49. Compare the relation of “weak” to “strong” permissibility in Williams, “Decision Making under Indeterminacy,” and “Indeterminacy, Angst and Conflicting Values.”

nonneutral by circumstance—our question is whether the nonneutrality is subject to any significant sanction. So, for example, in our original Indeterminate Ought scenario, consider people who voted for Mary in this situation. It is indeterminate whether they did as they objectively ought. But to sanction them in the ways appropriate to someone who did not do as they objectively should, one would be committed to saying that voting for Mary was not what they objectively should have done. But that is the very indeterminate question once again. The third party could impose sanctions but could not do so while remaining neutral on the indeterminate question. If the only source of criticism of your groundless judgment call is from other equally groundless but opposite judgment calls, then we are no longer fighting over the permissibility of making groundless judgment calls in principle (whether they are decision permissibility) but just reopening the first-order question of which option one objectively ought to take.

So the first half of the mind-making story is a motivation for the definition of decision permissibility above, in terms of the special status of immunity from neutral criticism. Notice that this is not a rationale that the glutters or gappers could give for their analogous definitions. The glutters believe that taking either option is objectively impermissible. The gappers reject the claim that taking either option is objectively permissible. These attitudes can form the basis of criticism of actions taken that the agent themselves would agree with. So in the glitter or gapper approach to decision oughts, the discussion is what the proper basis for criticism in these strange cases might be, and before we resolve that discussion the agent can't claim immunity from criticism.

But acts that are decision permissible while failing to be determinately objectively permissible should leave us queasy. In my "Indeterminacy, Angst and Conflicting Values" I offer a diagnosis of this angst. Faced with an Indeterminate Ought of the kind we are considering, agents must make an arbitrary judgment call, for example, about whether what they did counted as a promise. The process that they go through to form that judgment and choose an option, then, could equally well have ended up with them taking the opposed line. So it is determinately the case that they have judged and acted on the basis of a process which could easily have led to them doing the opposite. They themselves now judge that, for example, voting for Mary would be objectively permissible since they now judge that what they did counted as a promise. So by the same token, they judge that not voting for Mary was impermissible. And yet they were just an arbitrary choice away from doing just that. Even though they are committed to thinking that what they did was okay, they are committed to thinking that the process that led to it was an unreliable means to that end. No wonder that it is uncomfortable, especially when the stakes are high, to make such judgment calls.

Unreliability in forming moral judgments or taking moral actions is a vice of the agent, though not a demerit in the act (an unreliable coffee machine can make great coffee on occasion—the unreliability is part of the process, not the product). Moral agents will avoid making decisions of this character where they can, but sometimes they have no choice. This explains the two kinds of cases that posed challenges for the glitter and gapper theory of the decision ought. The first was a choice between something that is determinately objectively permissible and something which is only indeterminately so. The mind-making account of the decision ought says that both are decision permissible but in addition will say that a moral agent will take the first, since this is the only reliable way to be objectively permissible. The second challenge was a choice between something that is determinately objective impermissible and something which is only indeterminately so. In such scenarios, the best one can do is to do the right thing unreliably, and so the moral agent will (with angst) take the latter option. Mind making, in sum, predicts and explains the right results in these test cases.

VI. CONCLUSION

Given the apparent ubiquity of Indeterminate Ought scenarios, we need a response to the relevant Decision Ought Challenge. The current article has been a field guide to some of the leading options. I have not offered a definite recommendation, but each part above has resulted in a leading option to pursue. The best strategy identified here for reducing Indeterminate Oughts to quotidian Uncertain Oughts is the Barnes et al. combination of a primitive metaphysical moral indeterminacy with quotidian uncertainty about indeterminate matters. The leading non-classical strategy I think is offered by the scalers, who can offer an evaluation of the degree of objective permissibility of indeterminate options, which delivers an extremely natural theory of the decision ought in terms of maximally objectively permissible actions. In the classical supervaluational setting, I advocate the mind-making approach, and I've set out how it provides a rationale for decision permissibility (in terms of immunity to neutral sanction) but also a nuanced account of the angst that attaches to Indeterminate Ought scenarios and can guide action.

Several issues for future research in the area arise from the field guide. We left open in Section III the question whether the reduction of Indeterminate Oughts to Uncertain Oughts produces factual or normative uncertainty. That is a critical choice point, given the divergent treatments of the two in the offing. Case studies are needed to evaluate the predictions of the three leading accounts I have picked. For example, I've defended the Mind Making account as applied to the structural indeterminacy of value aggregation trade-offs (in "Indeterminacy, Angst

and Conflicting Values”)—but if we instead prefer to treat indeterminacy in the scaling way, then we will need to revisit whether the recommendations it delivers cohere with the phenomenon being modeled. Finally, uncertainty and indeterminacy can easily both be present in a high-stakes decision situation. So ultimately, we would need not a decision ought of each in isolation but a combined theory that offers guidance that speaks to agents suffering from both limited information and indeterminacy.