Ethics and the Question of What to Do

Olle Risberg, Uppsala University

Abstract: Several recent debates in ethics and metaethics highlight what has been called the "central deliberative question." For instance, in cases involving normative uncertainty, it is natural to ask questions like "I don't know what I ought to do—*now* what ought I to do?" But it is not clear how this question should be understood, since what I ought to do is precisely what I do not know. Similar things can be said about questions raised by normative conflicts, so-called "alternative normative concepts," and other similar problems. This paper defends a form of non-cognitivism about these questions that is combined with cognitivism about normative questions proper. A central claim is that we should distinguish the question of what we ought to do from the question of what *to* do, and that this distinction in turn has important consequences for our understanding of normative guidance, decision-making and deliberation. Two challenges to the non-cognitivist view defended are met, and its relationship to realism and "quasi-realism" about normativity is considered.

1. Introduction

The aim of this paper is to present and defend an account of a distinctive form of "practical" or "deliberative" question that is central in several debates in ethics, metaethics, and metanormativity more generally. Most writers assume that this question concerns some special

1

normative issue, such as what we ought to do "all things considered."[1] I will argue against this assumption and instead endorse an alternative view, which combines elements of both metaethical cognitivism and non-cognitivism. A notable consequence of this view is that even if there are truths about how we (all things considered) ought to act—truths which may even be objective, irreducible, and so on—the "central deliberative question," as it is has sometimes been called, does not concern those truths.[2] Instead, that question does not have a true answer.

One debate that highlights the relevant kind of question is the one about normative uncertainty.[3] Since we are not epistemically flawless beings, it seems that we are often (or at least sometimes) not in a position to know what we ought to do. As many have noted, such situations make it natural to ask questions like: "I don't know what I ought to do—*now* what ought I to do?" For obvious reasons, however, it is unclear how this question should be understood. After all, what the agent ought to do is precisely what she does not know!

Another example concerns choices in the face of conflicting normative requirements.[4] If we must choose between promoting the common good and promoting our own good, for instance, the requirements of morality might clash with those of prudence in such a way that we cannot satisfy both. Such situations invite other questions that are difficult to understand, such as: "Which ought—the moral or the prudential one—ought I *really* to satisfy?" Here too it is unclear how to understand the question that is raised, since after all, it is *really the case* that we

---

[1] Similarly to Mark Schroeder and others, I will generally use the term "normative" to mean, roughly, "having to do with value, oughts, reasons, duties, and the like" (Schroeder, "Realism and Reduction: The Quest for Robustness", 3), though see section 10 for a discussion of other things that can be meant by "normative."

[2] For this expression, see, e.g., Lord, "What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)," 1110, and McPherson, "Explaining Practical Normativity," 621.

[3] See, e.g., MacAskill, Bykvist and Ord, *Moral Uncertainty*, Sepielli, "What to Do when You Don't Know What to Do," and Weatherson, *Normative Externalism*.

[4] See, e.g., Chang, "'All Things Considered'", and Baker, "Skepticism about Ought Simpliciter."

ought morally to satisfy the requirements of morality, *and* that we ought prudentially to satisfy the requirements of prudence.

I will argue that the salient question in these and other choice situations does not strictly speaking concern what we ought to do, in any sense of "ought." Nor does it concern any other normative question. The reason, as I will argue, is that this question may well remain unanswered even in choice situations where all the truths, *including all the normative truths*, are known. The best explanation of this fact, I suggest, is that while uncertainty about normative questions amounts to uncertainty about the truth of some normative proposition—concerning, e.g., what one ought to do—the "central deliberative question" is instead the question of what *to* do. I further suggest that we understand the question of what to do along the lines suggested by Allan Gibbard.[5] On this view, roughly, one does not answer this question by forming a belief about what the world is like (not even in normative respects), but by forming an intention to act in a certain way. We should thus adopt cognitivism about normative questions but non-cognitivism about the question that sometimes seems to remain even when all normative questions are answered. A similar non-cognitivist view has recently been defended by Justin Clarke-Doane in response to one of the problems that I will discuss, concerning what Matti Eklund calls *alternative normative concepts*, and one contribution of the paper is to argue that this form of non-cognitivism is also plausible with respect to several other problems in ethics, metaethics and metanormativity.[6]

After briefly introducing the problem of alternative normative concepts and the non-cognitivist view about the question that it raises (sect. 2), I will show how similar questions are also raised by an argument against objective consequentialism due to Frank Jackson (sect. 3–4),

---

[5] Gibbard, *Thinking How to Live*.

[6] Clarke-Doane, *Morality and Mathematics*; Eklund, *Choosing Normative Concepts*; see also Balaguer, "Moral Folkism and the Deflation of (Lots of) Normative and Metaethics."

in the normative uncertainty debate (sect. 5–6), and by normative conflicts and what Christine Korsgaard calls the "normative question" (sect. 7).[7] Along the way, I will consider a number of alternative accounts of how this question should be understood, and argue that they all face important challenges. My final argument to that effect focuses on the possibility that the normative truths are dramatically different from what we take them to be (sect. 8). I then return to the question of how the relevant form of non-cognitivism is best understood, explain how my preferred version of it differs from "quasi-realism" about normative judgments (sect. 9) and consider two challenges to it (sect. 10). Section 11 concludes.

2. Alternative normative concepts

The problem of alternative normative concepts can be introduced by noting that for the most part, historical contingencies are least partially responsible for what concepts we happen to employ—and, more generally, for how we happen to think. If evolutionary processes had shaped our cognition differently, for example, we might well have represented the world using concepts that we do not in fact have. This raises the question: could the same be true of our normative concepts, such as OUGHT, GOOD, and REASON?[8] That is, are there *alternative* normative concepts, which could play the same role in our lives as our actual normative concepts do, but which are true of different actions, attitudes, and so on? If so, is there any suitably neutral way to ask which set of normative concepts we *ought* to use?

Eklund makes the problem vivid by imagining a community of speakers, "Alternative," who use the concept OUGHT* in much the same way that we use the concept OUGHT. That is, while we perform actions that we judge that we ought to perform, they perform actions that they judge

---

[7] Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection"; Korsgaard, *The Sources of Normativity*.
[8] I use small caps to denote concepts.

that they ought* to perform; whereas we criticize and resent people who do things that we believe ought not to be done, they criticize and resent people who do things that they believe ought* not to be done; and so on. But in the imagined case, OUGHT and OUGHT* are not coextensive—there are some actions that ought but ought* not to be done (or vice versa). If this case is possible, then, as Eklund notes,

> [a] first thought one might have is that […] there is some sort of live issue as to whether we or the alternative community get things right. They do what they do based on considerations about what is "good" and "right" in their sense; we do what we do based on considerations about what is "good" and "right" in our sense. Since our normative terms and their normative counterparts aren't coextensive, we then act differently… [But] what set of normative terms ought to be used when we ask ourselves what to do?[9]

In other words, if we learn that there are alternatives to the normative concepts that we actually have, we might want to ask questions like: What ought we to do? Should we "go with" what we ought to do or what we ought* to do? Which set of normative concepts and/or terms ought we use? However, as Eklund goes on to note, it is not plausible that the salient further question literally concerns what normative concepts we ought to use (or any other issue that can straightforwardly be put in terms of our actual normative concepts). The reason is that this question might have an answer that is too easy: perhaps we simply ought to use OUGHT—and perhaps we equally ought* to use OUGHT*! Similarly, perhaps we should go with what we should to do, but should* go with what we should* do—and so on.

---

[9] Eklund, *Choosing Normative Concepts*, 22.

Several other views about the salient further question are possible. One is that the case is impossible as described (and so the question does not even arise), because all concepts that have the same "normative role" with respect to guiding behavior are also coextensive.[10] Another view is that the question is in some sense "ineffable"—it is genuine but cannot be perspicuously expressed in our language, and perhaps not in any possible language either.[11] A third view is that although the case is possible, there is no genuine further question at all—there is only what we ought to do and what we ought* to do and that is that. I mention these views only to set them aside. Instead, as already mentioned, the view that I will ultimately go on to endorse is a kind of non-cognitivism about this question. Drawing on Gibbard, Clarke-Doane proposes that the salient further question is best understood as a question of *what to do*; e.g., whether to do what we ought or what we ought* to do, or whether to use OUGHT or OUGHT* in deliberation, etc.[12] This question is meant to be non-cognitive in the sense that one does not "answer" or "settle" it by forming a belief about some matter of fact (or by forming some other kind of doxastic attitude). Instead, one answers it by forming a non-cognitive attitude of some kind. On Gibbard's view, it is a kind of intention.

What is the relation between the question of what to do and the question of what we ought to do? According to Gibbard's non-cognitivism (as it is usually understood), they are simply identical, given that his analysis is supposed to be true of the normative concepts that we in fact have.[13] But other views may also be had. In particular, the view that I will go on to endorse is that

---

[10] This might follow from certain forms of "conceptual role semantics"; see, e.g., Wedgwood, "Conceptual Role Semantics for Moral Terms." See also Fitzpatrick, "Commentary on Matti Eklund, *Choosing Normative Concepts*" for a suggestion along these lines. The expression "normative role" is from Eklund (e.g. Eklund, *Choosing Normative Concepts*, 10).

[11] For discussion of this view, see Eklund, *Choosing Normative Concepts*, ch. 2.2, and Clarke-Doane, *Morality and Mathematics*, 172.

[12] See Gibbard, *Thinking How to Live*, and Clarke-Doane, *Morality and Mathematics*, ch. 6.

[13] There are some interpretative complications, however; for instance, Gibbard at one point suggests that any analysis is likely to "strain" the concept that is analyzed, and proposes only that his view strains our actual normative concepts less than competing views (Gibbard, *Wise Choices, Apt Feelings*, 32).

the questions are different, in that the question of what we ought to do is a "question of fact" while the question of what to do is not. Thus, the question that I earlier called the "central deliberative question," and which I take Eklund's imagined scenario to highlight, is on this view understood as a non-factual question, rather than a factual, normative one. This view, I will argue, best explains why similarly puzzling kinds of questions—or, as I will often put it, similarly puzzling kinds of *uncertainty*—are raised by several other problems in ethics, metaethics, and metanormativity. One promising explanation of this commonality is that all these questions (and the corresponding forms of uncertainty) are of the same kind, and that some kind of non-cognitivism is thus true in all these cases. Reflection on these other cases accordingly provides independent support for the relevant form of non-cognitivism.

3. Objective Act-Consequentialism and Choices Under Empirical Uncertainty

A quite different topic in ethics that brings the central deliberative question to the fore concerns a prominent worry about objective act-consequentialism, which is the view that we always ought to perform the action that would in fact have the best consequences.[14] The worry is that in most or all real-life situations, it is impossible for us to know which action the view prescribes. This concern is also what motivates Jackson's objection to the view, which departs from the following case:

> Jill is a physician who has to decide on the correct treatment for her patient, John, who has a minor but not trivial skin complaint. She has three drugs to choose from: drug A, drug B, and drug C. Careful consideration of the literature has led her to the following opinions. Drug A is very likely to relieve the condition but will not completely cure it. One of drugs

---

[14] When context does not indicate otherwise, I use "consequentialism" and "objective consequentialism" to refer to objective act-consequentialism.

B and C will completely cure the skin condition; the other though will kill the patient, and there is no way that she can tell which of the two is the perfect cure and which the killer drug.[15]

The problem stems from the fact that according to objective consequentialism, Jill ought to give John the perfect cure, even though she does not know which drug that is. What Jill knows is only that it is *either* objectively best to give John drug B or to give him drug C. She can thus deduce that it is objectively suboptimal to give him drug A. But since Jill does not know whether drug B or drug C is the perfect cure, she does not know how to realize the best outcome. In view of this fact about her epistemic situation, Jackson writes that

> [the] problem arises from the fact that we are dealing with an *ethical* theory when we deal with consequentialism, a theory about *action*, about *what to do*. […] Now, the fact that an action has in fact the best consequences may be a matter which is obscure to an agent. (Similarly, it may be obscure to the agent what the objective chances are.) In the drugs example, Jill has some idea but not enough of an idea about which course of action would have the best results. […] Hence, the fact that a course of action would have the best results is not in itself a guide to action […][16]

When Jill is uncertain about what to do, the argument goes, learning that she ought to perform the objectively best action is useless since she does not know which action that is. Jackson thus

[15] Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 462–463. See also Regan, *Utilitarianism and Co-Operation*, and Kolodny and MacFarlane, "Ifs and Oughts."
[16] Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 466–467.

concludes that consequentialism—which is a theory about what we *ought* to do—fails to answer the question of what *to* do for agents who do not know how to realize the best outcome.[17]

It is extremely common that we do not know which of our alternative actions that are objectively best, however, and the point of the Jill and John case is not merely to emphasize that fact. Our ignorance about this is more easily illustrated by the fact that many of our actions have "massive causal ramifications."[18] In particular, seemingly mundane actions (like buying coffee) may affect what germ cells that will ever figure in conception, and thus, what people that will ever exist. As a result, their impact on the total amount of future well-being can be both dramatic and unknowable for us. Arguably, however, these considerations do not pose the same problem for objective consequentialism, since in such cases the view at least *suggests* an answer to the question of what to do: namely, to perform the action that is *most likely* to maximize objective value, or to *try one's best* to do so, or something along those lines.[19] By contrast, in the case of Jill and John, such courses of action seem clearly objectionable. Since Jill knows that drug A will not be best, the action that is most likely for her to maximize objective value is perhaps to flip a coin and give John either drug B or drug C, depending on the outcome. Yet it is surely a terrible idea to make the decision in this way. Instead, intuitively, Jill ought to give drug A to John, even though she knows that this does not maximize objective value.

Jackson's argument accordingly supports the view that we sometimes ought to perform actions that we know to be objectively suboptimal. If that view is true, then objective

---

[17] Note that while it is not clear what Jackson takes the expression "what to do" to mean, he probably does not accept the non-cognitivist interpretation of it that is associated with Gibbard, since he defends a form of cognitivism about normative concepts elsewhere (see Jackson, *From Metaphysics to Ethics*).

[18] This expression is from Lenman, "Consequentialism and Cluelessness," 344.

[19] Jackson attributes a view along those lines to Peter Railton, though I note that it is unclear to me whether Railton in fact meant to endorse this view. See Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 466, and Railton, "Alienation, Consequentialism and the Demands of Morality."

consequentialism is false, since objective consequentialism entails that all suboptimal actions are impermissible. Most of the literature on Jackson's argument has thus focused on the question of whether giving John drug A really is what Jill ought to do.[20] But in the current context, there is another aspect of Jackson's argumentation that is more important. What ultimately underlies the argument is a widespread and natural view about the role of normative thinking in practical deliberation.[21] It is illustrated by Jackson's claim that unlike other areas of inquiry, ethics is centrally concerned with the "passage to action":

> It is fine for a theory in physics to tell us about its central notions in a way which leaves it obscure how to move from those notions to action, for that passage can be left to something which is not physics; but the passage to action is the very business of ethics.[22]

In recent discussions about normativity, similar remarks have been frequent. For example, Jacob Ross writes that

> in genuine deliberation, we are guided, at least implicitly, by the question "what sould [*sic*] I do?" or "what ought I to do?" And we ask this question not simply in order to satisfy our curiosity, but in order to make up our minds about what to do, that is, in order to form an intention. Thus, the role of the *ought* of practical deliberation is to guide our intentions, and thereby to guide our actions.[23]

---

[20] See, e.g., Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance*.

[21] Of course, that is not to say that the view is universally accepted. For opposition, see, e.g., Parfit, *On What Matters, Volume Two*; and Zimmerman, *Ignorance and Moral Obligation*. See also Weatherson, *Normative Externalism*, who defends a view on which answers to normative questions need not be guiding. This view is congenial with my conclusion that the central deliberative question does not concern what we ought to do.

[22] Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 467.

[23] Ross, "Rationality, Normativity, and Commitment," 164.

In the same vein, Errol Lord claims that it is "commonly assumed that the answer to the central deliberative question is the thing that you ought to do, full stop," David Faraci writes that "substantive normative claims answer (or at least entail that there is an answer to) the question of *what to do*," and Jonathan Way and Daniel Whiting suggest that "[i]n deliberation, we ask ourselves a single question, 'What ought I to do?'"[24] While the idea that all these claims suggest is perhaps somewhat imprecise, it is also highly intuitive. It is plausible that we normally do not engage in normative thinking with the sole aim of learning more about the world. We also do so in order to reach choices in our lives. And the worry that Jackson highlights is that objective consequentialism suggests that this aim is misguided. For while our lives are unavoidably full of uncertainty and ignorance about empirical facts, consequentialism entails that unless we have knowledge of those facts, we cannot figure out what we ought to do. Perhaps that result would not be so bad if we could instead settle for the action that is most likely to be best. But what Jackson's argument suggests is that we sometimes ought not even to do that—rather, in some situations, we ought to perform actions which we know to be objectively suboptimal. And the worry is that in view of all this, it is hard to see how consequentialism can be reconciled with the role of normative thinking in practical deliberation that Jackson and Ross suggest. Even if objective consequentialism were true and we knew that this was so, our uncertainty about the

---

[24] See Lord, "What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)", 1110, Faraci, "On Leaving Room for Doubt," 248, and Way and Whiting, "Perspectivism and the Argument from Guidance," 362. More generally along the same lines, Mark Timmons holds that normative ethics has both a "practical" and a "theoretical" aim, and Michael Smith claims that a metaethical theory must be able to accommodate both the "objectivity" and the "practicality" of moral judgments (see Timmons, *Moral Theory: An Introduction*, ch. 1, and M. Smith, *The Moral Problem*, ch. 1). For a recent discussion about using morality as a decision-guide under empirical uncertainty, see Holly Smith, *Making Morality Work*. Holly Smith does not focus on the issue of fundamental moral uncertainty (or of fundamental normative uncertainty more generally), however, which will be central in what follows. For an argument that Holly Smith's idea that moral theories should be practically "usable" leads to a non-cognitivist view like the one that I endorse in this paper, see Clarke-Doane, "From Non-Usability to Non-Factualism."

question of what to do would remain unresolved—and yet this is the very question, the thought goes, that consequentialism and other normative theories seek to answer.

In what follows, I will summarize the above claims by saying that objective consequentialism fails to *address* the central deliberative question for agents who, like Jill, lack the relevant empirical knowledge. If Jackson's and Ross's idea is correct, the fact that consequentialism fails to do so is a serious problem for the view.

4. Decision-Theoretic Consequentialism

In view of the problems for objective consequentialism, Jackson instead endorses "decision-theoretic consequentialism," whose main motivation is its purported ability to avoid those problems. According to decision-theoretic consequentialism, every agent ought to maximize "expected moral utility," where an action's expected moral utility is determined (roughly) by summing the probability-weighed values of its possible outcomes.[25] Notably, while the values in question are meant to be objective, Jackson takes the relevant probabilities to be the agent's subjective ones. Thus, in one respect, the view resembles classical decision theory, as the agent's own mental states partially determine what she ought to do, and in this respect it also differs from objective consequentialism.[26] In another respect, however, Jackson's view resembles objective consequentialism and differs from classical decision theory, as the agent's preferences are not taken to determine the relevant ordering of an action's possible outcomes—instead, that ordering is determined by the objective value facts, whatever they turn out to be.

---

[25] See Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 464.

[26] Strictly speaking, classical decision theory states the conditions for representing an individual's preferences with a particular representation function. Decision-theoretic consequentialism adopts the constraints of decision theory on preference orderings for use in a normative consequentialist theory. Thanks to Andrew Reisner for discussion.

Decision-theoretic consequentialism thus involves a combination of objective and subjective elements that is striking and seemingly unstable. Indeed, as a result, this view is susceptible to the very same problem that Jackson takes objective consequentialism to face: that it fails to address the central deliberative question for an agent to whom "the fact that an action has in fact the best consequences [is] obscure" (*cf.* sect. 2). For decision-theoretic consequentialism also centrally appeals to facts that are often "obscure" to us—namely, the objective value-facts about the possible outcomes of our actions.[27] We are often uncertain about what is objectively good, and even when we are not, our views are often mistaken. In particular, the widespread disagreement about value illustrates this point: so many people have conflicting axiological views that at best, only a few of us can be correct.[28] Hence, while Jackson is right that we often do not know what action maximizes objective value, he overlooks the fact that we often do not know what action maximizes expected moral utility either.

One might think that the problem is less serious for decision-theoretic consequentialism if the relevant value facts can be known *a priori*, at least in principle. The problem with this suggestion is that the mere *possibility* of axiological knowledge makes no difference to an agent when she does not in fact have it. The case of Jill and John illustrates this point: while Jill does not know whether drug B or drug C is the perfect cure, it is perfectly possible for her to acquire such knowledge—all she has to do is to give John one of the risky drugs and observe the results. Clearly, however, the principled possibility of such knowledge is useless to her when she does not in fact have it. And the point is that the *way* in which she could acquire such knowledge is in

---

[27] For similar worries, see M. Smith, "Moore on the Right, the Good, and Uncertainty," and Bykvist, "How to Do Wrong Knowingly and Get Away with It."

[28] Moreover, on many plausible views about the epistemology of disagreement, a subject's true axiological beliefs often or always fail to amount to *knowledge* when they are disputed (at least when the opponents are the subject's epistemic "peers"); see further, e.g., McGrath, "Moral Disagreement and Moral Expertise," and Risberg and Tersman, "A New Route from Moral Disagreement to Moral Skepticism," "Disagreement, Defeat, and Higher-Order Evidence," and "Moral Realism and the Argument from Skepticism."

this regard irrelevant. Merely possible knowledge, whether *a priori* or otherwise, cannot help us in our decision-making.

5. Choices Under Normative Uncertainty

Decision-making under axiological uncertainty is a special case of decision-making under normative uncertainty more generally. We sometimes face hard choices, not because we are uncertain about the relevant empirical or axiological facts, but because we are uncertain about the fundamental normative facts. For example, many people must at some point decide whether to have children. When we try to figure out what we ought to do in such situations, we face difficult problems about our obligations toward future people, present people, merely possible people, and so on.[29] Perhaps we sometimes solve those problems. Very often, however, we fail to do so. For instance, maybe it is just too hard for us to determine whether we are obliged to create a person with a good life rather than a person whose life would be worse, but nonetheless worth living, or no person at all. To answer that question, we must take a stance on the many controversial issues in population ethics. Due to their difficulty, some degree of uncertainty about their answers, or perhaps even suspension of judgment, seems to be warranted. Yet even somebody who is uncertain about those questions might one day have to decide whether to become a parent. She cannot wait until the true moral theory has been discovered, since she has to act now. Thus, she will have to deliberate about what to do, even though she has failed to determine what she ought to do.

---

[29] For two classic discussions of these problems, see Parfit, *Reasons and Persons*, part IV; and Arrhenius, *Future Generations: A Challenge for Moral Theory*.

The recent debate about choices under normative uncertainty has primarily been motivated by the aim of providing some sort of guidance in such situations.[30] The hope is that even when we are uncertain about fundamental normative matters, there is a form of normative theorizing that can help us reach actions or decisions. While many different theories about these issues have been proposed, their details need not concern us here.[31] I will instead focus on the question that theories about choices under normative uncertainty are supposed to answer.

Participants to this debate often introduce their topic by noting, as I did above (*cf.* section 1), that situations of normative uncertainty make it natural to ask questions like: "I can't figure out what I ought to do; *now* what ought I to do?" However, they then usually note (as I also did above) that it is not clear how this question should be understood. After all, most traditional moral theories, like utilitarianism and Kantianism, entail that we ought to maximize happiness or to treat humanity as an end in itself (etc.) whether or not we believe this to be the case. Thus, on those views, what we ought to do is simply independent of our beliefs about the matter. This has led some to think that the question just posed has a trivial answer: we simply ought to do what the true normative theory entails that we ought to do, whether or not we know what that

---

[30] Michael Zimmerman's work on this topic is an exception. See, e.g., Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance*.

[31] The currently most popular view is that normatively uncertain agents ought to maximize "expected choiceworthiness." Unlike Jackson's concept of expected moral utility, the concept of expected choiceworthiness is supposed to be sensitive both to the agent's normative probabilities and her non-normative probabilities. The viability of this strategy is the subject of an ongoing debate. One major concern is that it requires that "inter-theoretical" comparisons of choiceworthiness are meaningful. In other words, the degree to which an action is right according to utilitarianism must be comparable to the degree to which it is wrong according to Kantianism, for example, as its expected choiceworthiness is supposed to be the probability-weighed sum of those values. It is still unclear when, if ever, such comparisons are meaningful; in particular, as William MacAskill notes, the matter is especially complicated for agents who have some (justified) degree of belief in nihilism, on which the moral value of every action is not zero but undefined (see MacAskill, "The Infectiousness of Nihilism"). See for an overview of the debate and a discussion of how the strategy of maximizing expected choiceworthiness can be expanded to handle cases that involve incomparability, see MacAskill, Bykvist, and Ord, *Moral Uncertainty*. See also Harman, "The Irrelevance of Moral Uncertainty", for an argument that moral uncertainty is not normatively important.

is.[32] On the one hand, this claim seems close to platitudinous and thus hard to deny. On the other hand, however, there is an obvious sense in which this answer fails to address the agent's uncertainty in the situation just considered, just like objective consequentialism fails to address Jill's uncertainty in the case of Jill and John.

To avoid this result, the most popular strategy has been to hold that in the question "What ought I to do when I don't know what I ought to do?", the different occurrences of "ought" have different meanings. Following Andrew Sepielli, let us call this the *dividers'* strategy.[33] Dividers usually take the first occurrence of "ought" in this question to stand for the so-called "objective" ought. Traditional first-order theories like utilitarianism and Kantianism concern what we ought to do in this sense. The second occurrence of "ought," by contrast, is supposed to stand for something that is not the concern of such theories. It is less clear what that ought is like, however, because dividers disagree about how many oughts there are. Some dividers stop at two—on that view, there is just an objective ought and a "subjective" ought and that is that.[34] This has been a minority view, however, and many dividers instead posit a much larger number of oughts. In part, this is due to the fact that an important argument for the dividers' view relies on the idea that seemingly incompatible "ought"-sentences can be jointly true relative to different states of information.[35] Since there are clearly many different states of information, dividers are pushed towards positing many different oughts as well. For example, Andrew Sepielli writes that:

---

[32] See, e.g., Weatherson, "Review: Moral Uncertainty and Its Consequences," and *Normative Externalism*.

[33] Sepielli, "Subjective and Objective Reasons." Sepielli adopts this terminology to distinguish between "dividers" and "debaters" about the question of how we ought to act under uncertainty. For present purposes, we need not consider what semantics for "ought" that dividers should adopt. While it has sometimes been said that "ought" is genuinely ambiguous, like "bat" or "bank", a more plausible view is that the lexical entry for "ought" has an informational parameter that is supplied by context.

[34] Harman, "The Irrelevance of Moral Uncertainty," and Parfit, "What We Together Do," both seem to endorse this view (though they also seem to endorse different theories about what we subjectively ought to do).

[35] For discussion of this idea, see Kolodny and MacFarlane, "Ifs and Oughts."

[We] may speak of the belief-relative sense of "ought", the reasonable-belief-relative sense, the degree-of-belief-relative (or credence-relative, or subjective-probability-relative) sense, the evidence-relative sense, and the objective-probability-relative sense, each of which depends for its proper application on the feature mentioned in its label. We could ramify even further. There are, for example, different "interpretations" of objective probability— the long-run frequency interpretation, the propensity interpretation, the logical interpretation, etc.—and there could be an OUGHT corresponding to each interpretation. Finally, there is a subjective OUGHT that I call the minimal-probability-relative OUGHT.[36]

While the topic of normative uncertainty does not figure in Jackson's discussion, he nonetheless anticipates the dividers' strategy by positing what he considers "an annoying profusion of 'oughts'":

I think that we have no alternative but to recognize a whole range of oughts—what [Jill] ought to do by the light of her beliefs at the time of action, what she ought to do by the lights of what she later establishes (a retrospective ought, as it is sometimes put), what she ought to do by the lights of one or another onlooker who has different information on the subject, and, what is more, what she ought to do by God's lights […][37]

The idea is that by God's lights, Jill ought to give John the perfect cure, but by her own lights, she ought to give him the safe cure. And dividers seek to make sense of the normative uncertainty debate in a similar way: Objectively, they think, we ought to satisfy the true first-order normative

---

[36] Sepielli, "What to Do When You Don't Know What to Do," 48. (Sepielli uses capital letters to denote concepts, but he also notes that his idea does not strictly require that there are many distinct OUGHT-concepts.)

[37] Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 471–472.

theory, but when we cannot determine what we objectively ought to do, we can at least try to determine what we ought to do in some other sense of "ought." It is the latter, non-objective kind of ought that the normative uncertainty debate is taken to concern.

Clearly, regress threatens. While it is true that the traditional, "objective" questions of normative ethics are sometimes hard, the current controversies in the normative uncertainty debate suggest that those questions are not easier.[38] If we cannot figure out what we ought to do in the sense of "ought" that is central to that debate, are we then supposed to try to figure out what we ought to do in yet a new sense of "ought"? But why should we expect that to be easier? Does this ever stop?[39]

However, while the regress problem is important, in what follows I will focus on another problem which (for reasons that will emerge) I take to be more fundamental. The problem concerns the apparent stalemate that arises between all the oughts that dividers posit. Recall that theories about normative uncertainty are supposed to provide some sort of guidance to agents like the potential parent, who must decide whether to have children. The idea is to posit many different oughts to make sense of the question that such agents may naturally ask. In the relevant cases, however, these oughts will often prescribe different action—otherwise figuring out what we non-objectively ought to do would be just as hard as figuring out what we objectively ought to do (since those questions would simply have the same answer). And the existence of such conflicts seems only to give rise to the central deliberative question once again, for we may now also be uncertain about which of all these oughts to satisfy. What should we do when they

---

[38] See also Feldman, "Actual Utility, The Objection from Impracticality, and the Move to Expected Utility," for a convincing argument that maximizing expected value is normally not significantly easier than maximizing objective value.

[39] For further discussion of the regress problem, see Sepielli, "What to Do When You Don't Know What to Do When You Don't Know What to Do…", who seeks to solve it by distinguishing between "perspectival" and "systematic" notions of rationality and between different "orders" of rationality. The discussion in section 6 will indicate why I find this solution unconvincing.

diverge? Is there any genuine sense in which one of them can be said to be privileged, or more important than the others?

6. The Tie-Breaking Problem

Michael Zimmerman presents the relevant worry when commenting on a variation of the case of Jill and John:

> [Jill] seeks your advice, telling you that she believes that drug B would be best for John but that she isn't sure of this. "So," she says, "what ought I to do?" You are very well informed. You know that A would be best for John, that Jill believes that B would be best for him, and that the evidence available to Jill (evidence of which she is apparently not fully availing herself, since her belief does not comport with it) indicates that C would be best for him. You therefore reply, "Well, Jill, objectively you ought to give John drug A, subjectively you ought to give him B, and prospectively you ought to give him C." This is of no help to Jill. It is not the sort of answer she's looking for. She replies, "You're prevaricating. Which of the 'oughts' that you've mentioned is the one that *really* counts? Which 'ought' *ought* I to act on? I want to know which drug I am morally obligated to give John, *period*. Is it A, B, or C? It can only be one of them. It can't be all three."[40]

Of course, Jill's questions here are imprecise. If there are many different oughts, she cannot make progress by asking which ought she really ought to act on. For it is *really the case* that she objectively ought to act on the objective ought. The problem is that it is equally the case that she subjectively ought to act on the subjective ought. Imprecision aside, however, there is surely *some*

---

[40] Zimmerman, *Living with Uncertainty: The Moral Significance of Ignorance*, 7.

important, non-trivial form of uncertainty that Jill is trying to express here. It is very similar to

what the prospective parent tried to ask about the choice of whether to have children. Jill must

decide which drug to give to John, but what she is told in the dialogue above does not take her

closer to action. (Compare: Being told that we ought to do what is objectively best or objectively

right similarly fails to take us closer to action when we do not know what is objectively best or

objectively right.) In relation to the aim of providing guidance to uncertain agents, this result is a

disaster.

Jackson anticipates this problem too. In an attempt to avoid it, he stipulates that by

"ought," he means "the ought most immediately relevant to action, the ought which I urged it to

be the primary business of an ethical theory to deliver."[41] However, whether this stipulation

solves the problem depends entirely on what the word "relevant" is supposed to mean. Since

Jackson does not say, let us consider some possibilities.

Jackson probably did not intend his ought to be relevant in some merely descriptive sense

of "relevant." The reason is that descriptive facts do nothing to address Jill's uncertainty in the

dialogue above. For example, perhaps one of the oughts is the one that we *in fact* tend to focus on

in deliberation. In Jill's situation, however, this is surely beside the point—for whatever it is that

her uncertainty concerns, it is not which ought we *do* tend to satisfy. Rather, as her questions

suggest, it is something closer to the question of what ought she *ought* to satisfy.

For this reason, it is natural to think that Jackson rather intended his ought to be relevant

in some *normative* sense of "relevant." While Jackson does not elaborate on this point, he would

perhaps agree with Mark Schroeder that there is an "important *deliberative* sense of 'ought', which

is the central subject of moral inquiry about what we ought to do and why."[42] Schroeder mentions

---

[41] Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," 472.
[42] Schroeder, "Ought, Agents, and Actions," 2. Similarly, Lord writes: "[an agent] doesn't seem to learn what she ought to do, full stop, by learning what she subjectively and objectively ought to do. There is

several features that, in his view, distinguish this ought from others. The currently most important feature is that the deliberative ought, according to Schroeder, is "the right kind of thing to *close deliberation.*"[43] This seems congenial to what Jackson has in mind.

Importantly, however, to insist that the deliberative ought is the right kind of thing to close deliberation serves in this context only to relocate the problem.[44] For, if there are very many senses of "ought," why should there not also be very many senses of "right"? In particular, even if it is subjectively right to close deliberation using the subjective ought, it presumably objectively right to close deliberation using the objective ought. Similarly, and more directly in Jackson's terms, even if his ought is subjectively normatively relevant, it seems undeniable that the question of what Jill objectively ought to do is objectively normatively relevant—normatively relevant, that is, by God's lights.

The upshot is that the stalemate that arose among the different kinds of oughts, and which the appeal to normative relevance was supposed to get rid of, now arises among the different kinds of normative relevance instead. This is an instance of what I will call "the tie-breaking problem." The problem is that if we believe that the deliberative uncertainty that I have highlighted concerns a normative question, it is problematic to think that the normative realm is divided into distinct "domains" or "spheres." For to answer the question that this form of uncertainty concerns, we must somehow single out *one* action as the one to be performed. And since distinct oughts normally prescribe different actions, we face the question of which of these oughts to satisfy. For somebody who is uncertain about this question, it is useless to learn

---

another question that hasn't been answered yet: what ought [she] to do? When we theorize about what answers this question, we theorize about the deliberative ought." (Lord, "What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)", 1138).

[43] Schroeder, "Ought, Agents, and Actions," 9.

[44] Of course, Schroeder's characterization of the deliberative ought may still be apt for his own purposes, which is to distinguish the deliberative ought from the "evaluative" ought that is involved when we claim, e.g., that there ought to be world peace.

trivialities such as that she objectively ought to satisfy the objective ought, or that she subjectively ought to satisfy the subjective ought.

We might attempt to avoid this stalemate by appealing to some tie-breaking notion that distinguishes one ought from all the others. Zimmerman's Jill tries to do that by asking which ought she *ought* to satisfy. Jackson instead suggests that one of the oughts is "most relevant to action." But all such proposals face a dilemma. On the one hand, if the notion that plays the tie-breaking role is descriptive, then it is beside the point. It is simply plain that Jill's uncertainty does not merely concern whether the objective ought or the subjective ought has a certain descriptive feature. On the other hand, if the tie-breaking notion is normative, then we should expect it to be just as divided as the other normative notions. Thus, rather than breaking the tie, this move only reinstates the stalemate that we faced among the diverging oughts.[45]

This also puts us in a position to see why the tie-breaking problem is more fundamental than the regress problem (*cf.* sect. 4). Recall that a regress of oughts threatens when we are uncertain about what we ought to do in the sense of "ought" that is supposed to be central to the normative uncertainty debate. It is normally assumed that it would be problematic to simply bite the bullet and accept that such a regress does indeed arise. But why? Regresses are not *intrinsically* problematic; for example, we can all agree that if it is true that *p*, then it is also true that it is true

---

[45] Graham, "In Defense of Objectivism about Moral Obligation," argues that moral obligations are objective (in the sense, roughly, that they are independent of our evidence) on the ground that it is the objective moral sense of "ought" that concerns a morally conscientious person. However, Graham also holds that in Jackson's case of Jill and John, Jill is morally conscientious only if she does something (i.e., giving John drug A) that she knows that she ought objectively not to do. This makes it hard to avoid the conclusion that there is in fact another kind of ought that tracks what a morally conscientious person *does* (where this may depart from the ought she is *concerned with*), which is (in some sense) the one that we *really* ought to satisfy. At any rate, since Graham assumes that "[a] morally conscientious person is solely concerned with her moral obligations" ("In Defense of Objectivism about Moral Obligation," 98), his suggestion sheds no light on cases in which moral requirements conflict with other kinds of normative requirements, which I discuss in section 7. For further critical discussion of Graham's view, see Mason, "Objectivism and Prospectivism About Rightness," sect. 4.c.

that *p*, and true that it is true that is true that *p*, and so on. So why would it be so bad to accept a regress in this particular case?

I suspect that the regress of oughts seems problematic only given the further assumption that of all the oughts that dividers posit, one of them is supposed to be the ought that addresses the central deliberative question in cases of normative uncertainty. And it is this assumption that the tie-breaking problem calls into question. *If* that assumption is accepted, then the regress of oughts is a problem because it suggests that we might be ignorant of what we ought to do at each point of the regress. We could try to figure out what we ought to do, fail to do so, move on to figure out what we ought to do in some other sense of "ought," fail again, and so on. Far from being guiding, this process would never result in action. But the tie-breaking problem calls the crucial assumption into question at an earlier stage, before worries about regress even arise.

In other words, while the regress problem is an important epistemological worry for dividers, the tie-breaking problem is a conceptual worry which is prior to it. That problem is to make sense of the question of which ought we *really* ought to act on, as Zimmerman's Jill puts it, rather than the epistemic problem of whether we can know what we ought to do, for some given sense of "ought."

7. The Normative Question

So far, I have discussed a number of problems concerning whether normative theories can guide us in choice situations that involve different forms of uncertainty. However, a possible reaction to the discussion so far is to hold that if a normative theory fails to address the deliberative uncertainty of an agent who lacks relevant information, it is (so to speak) the agent and not the

theory that is to blame.[46] The idea is that we can acknowledge that normative theories cannot address the central deliberative question for every agent, no matter their epistemic situation, but insist that they should at least address that question for agents who know all the relevant truths.

A problem for this proposal, however, is that considerations that relate to different states of information are not the only possible reason to posit many oughts. Another possible reason is that different "sources" of normativity, such as morality, prudence, and so on, may generate distinct normative requirements.[47] If that is so, the tie-breaking problem arises again when those requirements cannot be jointly satisfied. To illustrate, suppose that prudence requires you to maximize your own well-being while morality requires you to sacrifice yourself for the sake of others. When you ask for advice, you are told only that prudentially, you ought to be selfish, but morally, you ought to be altruistic. As I suggested above (sect. 1), questions like those from Zimmerman's Jill are natural here too: "Which of the "oughts" that you've mentioned is the one that *really* counts? Which "ought" *ought* I to act on?"[48] There is a further salient question about which both you and Jill are uncertain. But once again, it is difficult to argue that your uncertainty

---

[46] For this suggestion, see, e.g., Bykvist, "Violations of Normative Invariance: Some Thoughts on Shifty Oughts," 113. In this vein, both Krister Bykvist and Erik Carlson hold that moral theories must be practically useful for "ideal" agents only (see Bykvist, "Violations of Normative Invariance: Some Thoughts on Shifty Oughts," and Carlson, "Deliberation, Foreknowledge, and Morality as a Guide to Action"). However, I think we should generally be suspicious about appealing to idealized agents in normative theorizing; see further Risberg, "Weighting Surprise Parties: Some Problems for Schroeder" and "The Entanglement Problem and Idealization in Moral Philosophy."

[47] There are many possible views about the structure of normative conflicts, however, and not everyone agrees that there are genuinely distinct sources of normativity (for discussion, see Reisner, "Normative Conflicts and the Structure of Normativity"). Philosophers who do not think so often hold that there is ultimately only one kind of normative question, such as what we *all things considered* ought to do (for suggestions along these lines, see, e.g., Crisp, *Reasons and the Good*; and Tännsjö, *From Reasons to Norms: On the Basic Question in Ethics*). I will return to this suggestion shortly.

[48] Interestingly, Zimmerman elsewhere suggests that there is no comprehensible question concerning what one "really" ought to do when the moral ought conflicts with a non-moral ought (see Zimmerman, *The Concept of Moral Obligation*, 1–2). He thus seems to take conflicting oughts that are due to different sources of normativity to be less problematic than conflicting oughts that are relative to different states of information. In light of the obvious similarities between the two problems, however, this strikes me as an unattractive view.

literally concerns whether you ought to act morally or prudentially. For, again, it is *really the case* that morally, you ought to act morally. The problem is that it is equally the case that prudentially, you ought to act prudentially.

The challenge of understanding the salient further question arises particularly clearly in the debate about the so-called "normative question," which is associated with Christine Korsgaard. She formulates this question as follows:

> When we seek a philosophical foundation for morality […] we are asking what *justifies* the claims that morality makes on us. This is what I am calling 'the normative question'. [49]

Korsgaard continues to write that an answer to this question

> must actually succeed in *addressing* someone in [the "first-person" position from which the normative question is asked]. It must not merely specify what we might say, in the third person, *about* an agent who challenges or ignores the existence of moral claims. Every moral theory defines its concepts in a way that allows us to say something negative about people who do that—say, that they are amoral or bad. But an agent who doubts whether he must really do what morality says also doubts whether it's so bad to be morally bad, so the bare possibility of this sort of criticism settles nothing. [50]

While Korsgaard's reasoning here is supposed to present a problem for moral realism, and for normative realism more generally, there has been a lot of confusion about what the problem is

---

[49] Korsgaard, *The Sources of Normativity*, 9–10.
[50] Korsgaard, *The Sources of Normativity*, 16.

supposed to be.[51] For what question, more exactly, is it that the relevant sort of criticism fails to settle? Surely it is not literally whether moral claims are morally justified or whether it is morally bad to be morally bad. In this vein, Derek Parfit writes:

> According to what Korsgaard calls normative realism, when we know the relevant facts, we are rational if we want, and do, what we have decisive reasons to want, and do. So Korsgaard seems here to suggest that, if realism were true, we might need a reason to want, and do, what we knew that we had decisive reasons to want, and do. That is clearly false. If you should do something, it is not an open question whether you should do it.[52]

While Parfit's claims are undeniable as far as they go, it would be surprising if they were to settle the doubts of the agent that Korsgaard has in mind. For although the question of whether it is so bad to be morally bad is imprecise, there does seem to be an important question which that agent is trying to express. Insofar as that question concerns a non-trivial issue, as it seems to do, it cannot be answered by the trivial facts that Parfit notes.

According to a popular view, Korsgaard's question about morality is best understood in terms of a normative concept that is not "indexed" to any particular source of normativity. This concept has variously been suggested to concern either reasons, rationality, correctness, the "favoring-relation," or a special kind of ought (which among other things has been called the "all things considered ought," the "ought full stop," the "ought period," and the "ought *simpliciter*").

---

[51] Dreier, "Can Reasons Fundamentalism Answer the Normative Question?", helpfully identifies some misunderstandings in the debate. However, for reasons that I will present in sections 8–9, I do not share Dreier's view that the problem is that realists cannot explain why it is irrational to act contrary to one's normative judgments.

[52] Parfit, *On What Matters: Volume Two*, 418. Note that Parfit assumes the controversial view that it is always rational for us to do and want what we have most reason to do and want. In particular, if there are "state-given" reasons for attitudes, we might sometimes have most reason to be irrational; for further discussion, see Reisner, "Is There Reason to Be Theoretically Rational?".

For present purposes it does not matter which of those concepts we invoke, so I will focus on the concept ALL THINGS CONSIDERED OUGHT (though I will sometimes omit the "all things considered"-qualifier in what follows). The idea is that when an agent faces conflicting normative requirements, like moral and prudential ones, she may acknowledge both that she morally ought to perform a certain action and that she prudentially ought to perform some other action. The salient further question is then what she all things considered ought to do. The all things considered ought is supposed to be the tie-breaker that resolves her uncertainty.

While it has sometimes been doubted whether the concept ALL THINGS CONSIDERED OUGHT is comprehensible, I will here set this worry aside.[53] Instead, in the next section, I will argue that even if there is such a special ought, the deliberative question that I have highlighted does not concern it. The reason is that even facts about what we all things considered ought to do can in principle be subjected to a certain form of practical questioning. That such questioning is comprehensible even when all normative questions are settled shows, I believe, that the central deliberative question is not a special normative question.

8. Outrageous Normative Truths

Many seemingly trivial questions have figured in the discussion so far: Whether it is morally bad to be morally bad, for example, and which ought we ought to satisfy. Another trivial question is whether the normative truth will turn out to be normatively outrageous. Of course it will not! However, a non-trivial question in the same neighborhood is whether the normative truth will

---

[53] For such doubts, see, e.g., Copp, "The Ring of Gyges: Overridingness and the Unity of Reason," Tiffany, "Deflationary Normative Pluralism," and Baker, "Skepticism about Ought Simpliciter." Perhaps Sidgwick's "dualism of practical reason" should also be understood as a version of skepticism about the all things considered ought (see Sidgwick, *Methods of Ethics*). However, another understanding of Sidgwick's view is that while the concept ALL THINGS CONSIDERED OUGHT is itself comprehensible, it is simply not satisfied by any action when morality and prudence conflict.

turn out to outrage *us*, in the purely descriptive sense of striking us as outrageous. That this is at least conceptually possible follows from the commonly accepted view that the thinnest normative concepts, like ALL THINGS CONSIDERED OUGHT, do not have enough descriptive content to adjudicate between various competing views about first-order normative questions.[54] This view is supported by versions of G. E. Moore's open question argument, Hume's law, the "is/ought-gap" and several related ideas. For example, even if consequentialism is true—and even true by metaphysical necessity—it is at least conceptually possible that a staunch absolutist theory is true, on which the consequences of our actions are irrelevant to their normative status. On such a theory, what is normatively important is not whether an action has a good outcome. All that matters is that it does not violate a certain set of rules. On this view, it is always forbidden to lie, for instance, no matter the consequences of telling the truth. Indeed, this view is often attributed to Kant.

While the staunch theory about lying is probably false, it is nonetheless possible for us to reason under the hypothesis that it is true. For instance, it is clear that given the truth of the staunch theory, the consequences of our actions are normatively irrelevant, and thus, most of us are seriously mistaken about ethics. We can confidently accept such conditionals while rejecting their antecedents. Similarly, it is clearly not the case that given that the staunch theory is true, the staunch theory is false, so what our intuitions are tracking here is not just the trivial fact that a material implication is true if its antecedent is false. Rather, even when we know that *p* is false,

---

[54] What I mean by this, roughly, is that competence with such concepts is not sufficient for knowing which first-order normative theory is true. Note also that what I say here is compatible with the idea of "moral fixed points" that Cuneo and Shafer-Landau, "The Moral Fixed Points: New Directions for Moral Nonnaturalism" endorse. The reason is that this idea pertains specifically to *moral* concepts, rather than to *normative* concepts in the more inclusive sense, and one of the consequences of this idea is precisely that moral concepts are much "thicker" than what is ordinarily supposed (*cf.* Cuneo and Shafer-Landau "The Moral Fixed Points: New Directions for Moral Nonnaturalism", 406). Indeed, as Cuneo and Shafer-Landau note, the idea of moral fixed points is not supposed to help with the question that arises when morality conflicts with some other source of normativity, such as prudence, or perhaps "shmorality" ("The Moral Fixed Points: New Directions for Moral Nonnaturalism", 406–407).

we may non-trivially evaluate claims of the form "given p, then q." (Or, in the jargon: even when we know that p is false, we can still "conditionalize" on p.) We may also be uncertain about whether to accept such claims, in the sense that we may be uncertain about whether to accept the consequent given the truth of the antecedent. I will now argue that in a similar way, we may remain uncertain about the deliberative question even given that all the normative questions are settled.

The argument relies on the following thought experiment. You face a choice situation where you can prevent great suffering by telling a lie. By telling the truth, on the other hand, you will cause even more suffering. Suppose now that it is true that, consequences notwithstanding, you are forbidden to tell the lie. In other words, morally, all things considered, and so on, you ought to tell the truth. Contrary to what you used to think, it has turned out to be normatively irrelevant that you could prevent great suffering by acting otherwise. (To be clear, what I want to imagine is not merely that someone *tells you* that you ought to tell the truth, or that receive some other type of evidence for that claim; I want to imagine that *it is the case that* you ought to tell the truth. Again, this assumption is surely coherent, even if it false of metaphysical necessity.)

In this case, at least three reactions are possible. The first is to "go with" the normative truth even though it is outrageous. "If that is what I ought to do," you could say, "then it is also what I shall do," hence proceeding to tell the truth. The second possibility is simply to give up on the commitment to doing what you ought to do. "If *that* is what I ought to do," you might say, "then I shall instead do what I ought not to do," thus going on to lie. The point is not that you may conclude that the ethical truth has turned out to be unethical (or that the normative truth has turned out to be "unnormative")—that remains an incoherent view. The point is rather that you may turn your back on the ethical truth, so to speak, because the trivial fact that the ethical truth is ethical might strike you as no more impressive than the fact that immoral actions are

legally required in countries whose laws are also immoral. Finally, the third possibility is to remain deliberatively uncertain. If you learn that you ought to cause great suffering, you might try to question ethics itself—"I ought to do something that strikes me as outrageous; *now* what ought I to do?" But it is now clear that this question does not literally express what your uncertainty concerns. You *know* what you ought to do—morally, all things considered, and so on. This is stipulated. Even so, you might remain uncertain about the deliberative question.

It does not matter what reaction we are in fact disposed to have. What matters is just the first reaction is not the only comprehensible one. The possibility of the other reactions shows that the deliberative question is not settled even by the assumption that all the truths, including all the normative truths, are known. On the view that I will go on to suggest in the next section, this is also the type of uncertainty that is made salient, in different ways, by the ethical and metaethical debates I have considered above. However, an underlying assumption in those debates is that this type of uncertainty must concern a special, puzzling normative question that it is difficult to express: What we all things considered ought to do, for instance, or what we ought to do in a sense of "ought" that is relevant when we do not know what we objectively ought to do. In view of the argument just presented, I believe that we should reject this assumption.

The argument just presented can be helpfully contrasted with two related ones from the literature. First, Clarke-Doane supports his non-cognitivist view of the "further question" by appeal to an argument that involves conditionalizing on what he calls "evaluative pluralism," which is roughly the view that there are alternative normative concepts in the sense characterized earlier (*cf*. sect. 2).[55] In short, the idea is that under the assumption that we ought to perform some action, *A*, but also ought* not to perform *A*, it seems that we can remain deliberatively uncertain about whether to perform *A*. While I am sympathetic to Clarke-Doane's argument, an important

---

[55] Clarke-Doane, *Morality and Mathematics*, 167–168.

difference is that mine does not involve conditionalizing on pluralism but on the first-order normative claim that we always ought not to lie. This is an advantage since not everyone finds the relevant form of pluralism even intelligible; for instance, William Fitzpatrick writes that "there are no intelligible alternative notions of 'value*' or 'shmalue', or 'good*' or 'appropriate*' […] We shouldn't rush to think we have the foggiest idea what such things would even mean."[56] If Fitzpatrick is right, it is not clear that we can even coherently conditionalize on pluralism. By contrast, as I have emphasized, the view that we should never lie is perfectly comprehensible (albeit implausible).[57]

Another interesting argument has been offered by Matthew Bedke in a critique of metaethical non-naturalists—roughly, those who think that normative facts are mind-independent and different in kind from those that are studied by the sciences.[58] Simplifying somewhat, Bedke's central claim is that non-naturalists are committed to revising their moral beliefs in immoral ways. He asks us to imagine being told by a reliable oracle that there is no non-natural property that human pain and non-human pain have in common. If we are non-naturalists and trust the oracle, we are forced to conclude that human pain and non-human pain are not both intrinsically bad (at least insofar as we do not abandon our non-naturalism), since non-naturalism implies that intrinsic badness is a non-natural property. And according to Bedke, being disposed to revise one's moral views on the basis of such "non-natural information" is

---

[56] Fitzpatrick, "Commentary on Matti Eklund, *Choosing Normative Concepts*," 6.

[57] Clarke-Doane claims that it is "hard to see" how pluralism, understood as a metaphysical thesis about normative properties, "could be false" (*Morality and Mathematics*, 166), and even that it is "almost trivial" (*Morality and Mathematics*, 163, 175). For criticism of these claims, see Eklund, "The Normative Pluriverse," sect. 3. In particular, as Eklund emphasizes, it is highly non-trivial that the plurality of normative properties are all *instantiable*—especially given a non-naturalist view of their nature. In more recent work, Clarke-Doane writes that "since properties' identity conditions entail instantiation conditions, there is no doubt about [non-natural normative] properties being instantiated if they exist" ("From Non-Usability to Non-Factualism," fn. 12). However, this is also too quick, since it is still a non-trivial question whether the relevant instantiation conditions are satisfiable.

[58] Bedke, "A dilemma for non-naturalists: irrationality or immorality?"

morally objectionable. The merits of Bedke's argument need not concern us here, but three differences between his argument and mine are worth noting.[59] First, Bedke focuses on a case in which the information we receive is formulated in non-normative terms. He does not discuss a case in which we learn that human pain and non-human pain have no *normative* property in common, which would be closer to the case discussed here. With respect to such a case, the argument I presented above can plausibly be run again, since in such a case we may well be deliberatively uncertain about, e.g., whether to care more about human pain than about non-human pain—but that is not Bedke's point. Second, while Bedke's argument involves imagining that we *receive evidence* that a certain (non-natural) claim is true, my argument does not have that epistemic aspect—as emphasized above, it focuses on imagining that a certain (normative) claim is *in fact* true.

Third and finally, whereas Bedke's argument targets non-naturalists specifically, my argument succeeds also given other views about the nature of normative truths. For example, consider the version of constructivism on which normative beliefs are not "fully representational," in the sense that they seek to represent robust, mind-independent facts, but are instead true just in case (and because) they accord with the rules or procedures that are "constitutive of agency."[60] Surely, the assumption that those rules always require us to tell the truth is at least intelligible—indeed, on some interpretations, it is an assumption that Kant in fact accepted. Thus, we can imagine facing a choice situation where we can prevent great suffering

---

[59] For critical discussion of Bedke's argument, see instead Enoch, "Thanks, We're Good: Why Moral Realism is not Morally Objectionable."

[60] This is one way to understand the constructivist view of Korsgaard, *The Sources of Normativity*. As Gibbard notes, Korsgaard can also be read as a non-cognitivist (see Gibbard, "Morality as Consistency in Living: Korsgaard's Kantian Lectures"). On that interpretation, my argument does not clearly work against her view (and is not meant to do so), since it is not clear what it means to conditionalize on a normative claim if such claims express non-cognitive attitudes. However, I take it that most constructivists want to distance themselves from non-cognitivism; see, e.g., Skorupski, *The Domain of Reason*, 4.

by telling a lie, but in which the rules that are constitutive of agency require us to tell the truth. In such a case, we might still be uncertain about the central deliberative question (or, indeed, even disposed to answer it in a way that the rules of agency forbid). This is so independently of whether normative truths are construed as non-natural or otherwise mind-independent.[61]

9. A Non-Cognitivist View of Deliberative Uncertainty

If not even the all things considered ought is what "settles" the central deliberative question, then what does this question concern? In this section I will present the account that I favor.

Consider the version of metaethical non-cognitivism according to which normative judgments are intentions, plans, decisions, or some similar kind of mental state.[62] According to such a view (most straightforwardly understood at least), to judge that an action ought to be performed is not to have a belief about it—instead, to make such a judgment is, roughly, to intend or decide to perform the action in question if the opportunity should arise.[63] Hence, on this view, uncertainty about what one ought to do does not amount to uncertainty about the truth of any proposition. The reason is simply that unlike beliefs, the relevant mental states—decisions, intentions, and the like—cannot be true or false.[64] Instead, uncertainty about what one ought to

---

[61] Further, as Enoch, "Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action" has emphasized, the status of the rules that are constitutive of agency can also be challenged directly, since we can wonder whether to be an agent rather than a "shmagent" (where a shmagent is an agent-like creature that is governed by different constitutive rules). For a discussion of the shmagency worry in the context of evaluative pluralism, see Clarke-Doane, *Morality and Mathematics*, 168.

[62] Like many others, I assume that it is true of at least some kind of non-cognitive attitude that if I have that attitude towards performing a given action right now, then I will perform that action right now if I can. (For instance, Paul Grice ("Intention and Uncertainty," 263–264) endorses this for intention, and Gibbard (*Thinking How to Live*, 152–153) endorses it for planning.) If there are several types of non-cognitive attitudes that are related to action in this way, then the differences between them will not matter in what follows.

[63] While this view is associated with Gibbard, *Thinking How to Live*, note that whereas Gibbard endorses a quasi-realist version of this view about normative judgments, I do not favor it either as a version of quasi-realism or as a view of normative judgments, for reasons that I will get to in a moment.

[64] I am assuming the falsity of extreme forms of cognitivism about intentions, according to which my intention to perform a certain action is simply identical to my belief that I will perform that action. Such views face well-known problems; see, e.g., Bratman, "Intention and Means-End Reasoning."

do amounts to a kind of non-cognitive uncertainty about what decision to make. It is better characterized as a state of practical indecision (i.e., as a state of not having decided which action to perform in the relevant situation, or not having formed an intention about it, etc.) than as a state of uncertainty about what the world is like.

The non-cognitivist account of normative uncertainty just sketched has serious problems.[65] Hence, throughout the paper, I have assumed that it is false. Instead, I have taken for granted the cognitivist views that normative judgments are beliefs that can be true or false, that normative uncertainty amounts to uncertainty about the truth of normative propositions, and so on. However, even assuming the truth of those views, what I wish to propose is that the non-cognitivist view just outlined is true of something else: namely, of the deliberative uncertainty, or the "central deliberative question," that we have all along struggled to express literally. On this "divided" view of normative and deliberative uncertainty, uncertainty of the latter sort does not concern a puzzling normative question that is special in some seemingly inexpressible sense. Rather, I suggest, such uncertainty simply concerns what we may call the question of *what to do*. As Jamie Dreier puts it, this type of question is the one that "you answer when and only when you have decided what to do. It is answered with an intention, perhaps, or a plan."[66] It is not answered by a belief or some similar kind of mental state. Thus, *a fortiori*, it is not answered by a belief whose content is a proposition about what you ought to do, or any other normative proposition.[67]

---

[65] See, e.g., Bykvist and Olson, "Expressivism and Moral Certitude," and MacAskill, Bykvist, and Ord, *Moral Uncertainty*, chapter 7. Briefly, the problem is that while normative judgments can vary in at least two independent dimensions—how good we judge that something is, for example, and how confident we are that it is good to that degree—paradigm non-cognitive states, like desires, vary only in one dimension; i.e., with respect to their strength.

[66] Dreier, "Can Reasons Fundamentalism Answer the Normative Question?", 172.

[67] While I find it natural to talk about "questions" and "answers" in this way, it is worth noting that these expressions are ambiguous in ways that can cause confusion. In one sense, the question of whether *p* has two answers (at least disregarding indeterminacy and the like); these are its "possible" or "candidate" answers. One possible answer to this question is that *p*, and the other is that *not-p*. In another sense,

Dreier suggests that the relevant type of question can be expressed by the interrogative sentence, "What shall I do?" Dreier's claim may or may not be correct, but this need not concern us here.[68] What matters is that we can informatively characterize the relevant mental state—it is a separate question whether we can perspicuously express it in English or some other natural language. What is plausible, however, is that many of the not-easily-understood interrogative sentences that I have discussed in this paper are naturally understood as *attempts at* communicating this type of state, even if they do not do so precisely. These include the following:

(i) I don't know what I ought to do; *now* what ought I to do? (sect. 4);

(ii) Which "ought" *ought* I to act on? (sect. 5);

(iii) Are the claims that morality make on us really justified? (sect. 6);

(iv) What I ought to do strikes me as outrageous; *now* what ought I to do? (sect. 7)

What I propose is that, on a natural interpretation, these questions are all strictly-speaking-unsuccessful attempts at expressing uncertainty about the question of what to do in the relevant situation; i.e., of what to do when one does not know what one ought to do, of what to do when different "oughts" are in conflict, of whether to act in accordance with the claims that morality make on us, and of what to do if the normative truths turn out to be outrageous, respectively. Unlike alternative views, this view explains why this type of uncertainty may remain even given

---

however, the question of whether *p* has only one answer—this is its *true* or *correct* answer. If *p* is true, then *p* is the true answer to the question of whether *p*, whereas if *p* is false, then the true answer to that question is *not-p*. To *answer* (verb) the question of whether *p*, moreover, can also mean different things: in one sense, to answer a question is to perform the speech act of asserting a candidate answer to it (in a suitable context), whereas if one ponders the question for oneself, then one answers it by accepting one of its candidate answers; in this case by forming either a belief that *p* or a belief that *not-p*. It is this latter, "first-personal" sense of "[to] answer" that I have in mind in the main text.

[68] Perhaps, as suggested to me by Michael Zimmerman (in personal communication), the sentence "What *am* I to do?" better captures the relevant question.

that we have knowledge of all the truths, including all the normative ones. Whether we have knowledge of those truths ultimately does not matter, because the relevant question never directly concerned them in the first place. Instead, because the deliberative question does not even concern what is true, it does not have a true answer.

Another advantage is that this view avoids the tie-breaking problem by steering between the horns of the dilemma (*cf.* sect. 5). The problem, recall, is that when different oughts diverge (e.g., the subjective and the objective ought, or the moral and the prudential ought, or even OUGHT and OUGHT*), it is hard to make sense of the salient further question in the neighborhood of which ought we *really* ought to act on. On the one hand, descriptive truths seem obviously beside the point. On the other hand, appealing to further normative truths seems only to relocate the problem. This dilemma is avoided if the salient question is neither about some descriptive truth or some normative truth, but is instead non-cognitive and thus not about any truth at all.

It is also worth noting that unlike most contemporary forms of non-cognitivism about normative judgments, the non-cognitivist view I have developed here is not wedded to the research program of "quasi-realism."[69] While the quasi-realist's position has always been difficult to state precisely, her aim is to explain most or all realist-seeming notions—i.e., normative truths, beliefs, knowledge, reasoning, argumentation, uncertainty, and so on—in non-cognitivist friendly terms, and thus avoid the "heavyweight," supposedly problematic commitments that genuine realists incur. Whether this program is successful is at best extremely controversial.[70] And since the form of non-cognitivism I have presented does not concern normative judgments

---

[69] The program was first endorsed by Simon Blackburn (see, e.g., Blackburn, *Essays in Quasi-Realism*) and many others have since followed suit; for example, Toppinen, "Non-Naturalism Gone Quasi: Explaining the Necessary Connections Between the Natural and the Normative," even suggests that quasi-realists should endorse non-naturalism about normative truths.

[70] For two influential critical discussions of quasi-realism, see Dreier, "Meta-Ethics and the Problem of Creeping Minimalism," and Schroeder, *Being For*.

proper, it does not require for its truth that the quasi-realist program succeeds. Accordingly, my account allows (though it does not entail) that normative truths, beliefs and the like are best understood in realist terms, so that, e.g., normative truth requires correspondence with reality, normative beliefs are "fully representational," and so on.

While the non-cognitivist view of deliberative uncertainty thus diverges from non-cognitivism about normative judgments, I think that it nonetheless captures an important intuition that has often been invoked in support of the latter view. This is the intuition that certain practical questions seem not to answer to matters of fact. No matter how much we learn about the world, those questions may in principle remain open.[71] The divided view vindicates this intuition, since it entails that the question of what to do is not a question of fact. In contrast, whether quasi-realist versions of metaethical non-cognitivism ultimately vindicate this intuition as well is far from clear. For what the quasi-realist assumes is that the relevant practical questions are questions about what we ought to do. Thus, when she goes on to try to accommodate the possibility of normative truth, knowledge, and so on, she no longer has the resources to explain why the relevant practical questions could remain open even given that we have knowledge of all the truths, including the normative ones.

10. Two Challenges

Before concluding, I will consider two possible challenges for my view.[72] The first challenge is that it might fail to capture the "normativity" of ought-truths and/or ought-judgments (or

---

[71] For example, this intuition arguably underlies Nowell-Smith's remark that "learning about 'values' or 'duties' might well be as exciting as learning about spiral nebulae or waterspouts. But what if I am not interested? Why should I *do* anything about these newly-revealed objects?" (Nowell-Smith, *Ethics*, 41). It is also illustrated by non-cognitivists' frequent reliance upon Moore's open question argument (*cf.* Darwall, Gibbard, and Railton, "Toward *Fin de siècle* Ethics: Some Trends") and the assumption that moral disagreements could remain even in "ideal" epistemic conditions (see, e.g., Tersman, *Moral Disagreement*).
[72] Thanks to two anonymous referees for presenting the two challenges considered in this section.

"oughts," for short). This challenge can be spelled out in different ways, depending on how the relevant notion of normativity is understood. However, I will argue that each version of the challenge can be met. For many things that can be meant by "normative," my account allows that oughts are normative. For some possible senses of "normative," the account may well rule out that oughts are normative, but there is also no strong independent support for thinking that oughts are normative in those senses in the first place. Either way, then, the challenge fails.

To begin with, one possible idea is that a truth or judgment is normative just in case it is related to some suitable normative notion, such as *ought*, *good* or *reason*, in the right way. For instance, Schroeder suggests that what is "distinctive of the normative" is that it is "all about reasons," and John Broome suggests that the term "normative" means "to do with ought," where the relevant "ought" "is a normative one."[73] This may be called the *trivial* sense of "normative," since it implies that at least one normative notion—i.e., reasons for Schroeder and ought for Broome—counts as normative simply by fiat, by being related to itself in the right way. I myself suspect that it is difficult to provide a more informative characterization of normativity than one along these lines, but my response does not rest on this assumption. What I want to emphasize is just that clearly, nothing in my account excludes that oughts are normative in this sense: oughts may well count as normative because they are analyzable in terms of normative reasons, for instance, or (as Broome's view suggests) simply because they are oughts. What matters is just that even if oughts are normative in this sense, we may still ask what to do with them, as it were.

A version of the idea just presented is to take oughts to be normative in the sense of standing in some relation to the normative notions of *rationality* and/or *coherence*.[74] For instance, perhaps oughts count as normative because they figure in some true "enkratic" principle, such

<hr />

[73] Schroeder, "Realism and Reduction: The Quest for Robustness," 13; Broome, *Rationality through Reasoning*, 10.

[74] A version of this challenge was offered to me by Jonathan Way (in personal communication).

as: if a subject judges that she ought to do *A* but does not intend to do *A*, then she is incoherent or irrational. This idea is also compatible with my account—just as we can ask whether to do what we ought to do, we can also ask whether to be incoherent, whether to be irrational, and so on.

Another popular idea is that we should distinguish between *robust* and *merely formal* normativity.[75] This distinction departs from the intuitive difference between the oughts of (e.g.) morality, epistemology and prudence on the one hand, which are usually taken to be robustly normative, and those of (e.g.) etiquette, chess, and grammar on the other hand, which are usually taken to be merely formally normative. How this intuitive difference should be cashed out in more detail is controversial. One view is that robustly normative requirements differ from merely formally normative ones in that they entail the existence of genuine (or genuinely normative) oughts, reasons, or the like. This view takes us back to the first suggestion considered above—I have already argued that my account allows that some oughts are normative in this sense. Another view is that robustly normative oughts differ from merely formally normative ones in that they are in some suitable sense not "up to us." For instance, maybe the oughts of etiquette, chess and grammar depend on our attitudes and conventions in a way that the oughts of morality, prudence and epistemology do not. (This might, but need not, in turn be because robustly normative oughts are "non-natural.") Nothing in my account rules out that some oughts are robustly normative in this sense either, as again, the supposition that some oughts are not up to us (and perhaps also non-natural) does not prevent us from asking what to do with them.

A somewhat different idea is that ought-judgments are normative in the sense that they are necessarily connected to *motivation*. The most straightforward version of this kind of "motivational internalism" states that if a subject judges that she ought to perform *A*, then she is

---

[75] See further, e.g., Finlay, "Defining Normativity," sect. 3.2.

at least defeasibly motivated to perform *A*. To begin with, I think the arguments in this paper provide at least some reason to deny such a strong, unqualified form of internalism. It seems plausible, for instance, that an ought-judgment might leave a subject motivationally "cold" if she finds it outrageous, or if she cares only about what she ought* to do. That said, the account I have presented is consistent with the idea that ought-judgments necessarily provide even a very high degree of defeasible motivation to perform the relevant action. The reason is that even if this is true, we still face the question of whether to act in accordance with the motivation that the ought-judgment provides.

The only version of motivational internalism that might pose problems for my proposal is the extremely strong view that ought-judgments always provide *overriding* or *indefeasible* motivation, so that if a subject judges that she ought to perform *A*, then she performs *A* (at least if she can). The reason is that this view entails that we always do what we think we ought to do, which in turn could make it hard to see how we might answer the practical question of what *to* do and the normative question of what we ought to do in different ways. However, such extreme forms of internalism are arguably *too* extreme; indeed, as Fredrik Björklund, Gunnar Björnsson, John Eriksson, Ragnar Francén Olinder and Caj Strandberg note, "[i]n contemporary metaethics, it is regularly assumed" that even the view that ought-judgments entail *defeasible* motivation is "too strong," since counter-examples to it seem "possible to conceive."[76] Thus, I am happy to simply assume that this extreme form of motivational internalism, according to which normative judgments always provide overriding motivation, is false.

The final idea that I will consider is that ought-judgments are normative in the sense that they constitute answers to questions about what to do. As it stands, this suggestion does not

---

[76] Björklund, Björnsson, Eriksson, Francén Olinder, and Strandberg, "Recent Work on Motivational Internalism," 126.

amount to much more than the denial of the view that I have offered. I do not deny that this suggestion has often been taken for granted—on the contrary, as I have emphasized, the assumption that something like this view is true arguably underlies several important debates in ethics, metaethics and metanormativity. But unless some independent support for this suggestion is presented, the mere fact that my view contradicts it surely cannot itself be seen as an objection.

The second challenge that I will consider is that if my account is correct, it is not clear why we so often ask ourselves normative questions, especially when we are trying to reach choices. Why do we not simply ask ourselves what to do, rather than what we ought to do, if these questions really are distinct? While this challenge is not identical from the first one, they might still be related, as this challenge can also be understood as an expression of the more general worry that oughts are in some sense more practically significant than my account allows.

The question of why we ask ourselves what we ought to do rather than what to do (when we do so) may well be at least partially an empirical question. Accordingly, what I will have to say about it is bound to be somewhat speculative. But that said, three considerations are worth noting before closing.

First, a kind of error theory might in some cases be plausible. After all, philosophical theorizing often allows us to draw distinctions that we do not usually recognize in everyday life. So it might be that we sometimes ask ourselves what we ought to do rather than what to do simply because we have not realized that these questions are distinct. If we were to realize this, perhaps we would care less about what we ought to do and more about what to do. I am not suggesting that this kind of error theory fully explains why we so frequently ask ourselves what we ought to do (rather than what to do), but it may at least play a role in such an explanation.

Second, another partial explanation might be that we are sometimes simply interested in what the normative truths are. In particular, even Ross is right that we do not ask ourselves

normative questions *only* to satisfy our curiosity (*cf.* sect. 3), that does not entail that curiosity is never even *part* of the reason why we do so. Indeed, although I have focused on normative and metanormative debates that highlight the question of what to do, many other debates in these fields are less closely connected to action. For instance, it is far from clear how questions such as whether the betterness-relation is transitive, whether the good is more fundamental than the right, and so on, could even have a bearing (except perhaps very indirectly) on the question of what to do. With respect to such questions, the reason why we investigate them might be the same as when we try to find out whether the causation-relation is transitive or whether the brain is more fundamental than the mind: We are simply interested in their answers.

Third and finally, as externalists about moral motivation have emphasized, it might be a metaphysically contingent but still quite modally robust fact about us that we often want to do the right thing.[77] This suggestion interacts with the second one made above, since it sheds further light on why we might often be interested in figuring out answers to normative questions: Doing so might not only satisfy our intellectual curiosity but also help us achieve something we want. If this idea is correct, it also helps explain why we often do not consider the question of whether to do what we ought to do—we might simply be happy with figuring out what we ought to do and do our best to act accordingly.

Though there is more to say about these issues, the three considerations I have presented together go at least some way towards explaining why we so often ask ourselves what we ought to do, even if that question is distinct from that of what to do.

---

[77] See, e.g., Copp, "Belief, Reason and Motivation: Michael Smith's *The Moral Problem*," 49–51.

11. Concluding Remarks

In this paper, I have focused on a distinctive deliberative question that many debates in ethics, metaethics and metanormativity highlight. I have argued against the common view that this question concerns a special normative notion. Instead, I have offered a combination of cognitivism about normative questions and non-cognitivism about the question of what to do. An upshot of this divided view is that even if there are truths about what we all things considered ought to do, the central deliberative question does not have a true answer.

As I have noted (*cf*. sect. 9), my view is strictly speaking consistent with the "robust" or "ardent" realist position that there are objective, irreducible, heavyweight truths about how we ought to act.[78] That said, however, I do think that my view threatens to undermine an important argument for normative realism. For if the question of what to do does not even concern the truths that realists posit, then there is one way in which those truths seem much less interesting than we often take them to be. At least in many cases, as Jackson and Ross suggest (sect. 3), we do not ask ourselves what we ought to do with the sole aim of learning more about the world. We also do so to reach choices in our lives. An attractive feature of ardent realism is that it promises to take such questions seriously, by positing objective truths that are supposed to constitute their answers.[79] However, this argument is undercut if there is an open practical question that remains even when we acquire knowledge of those truths: Whether to do what we ought to do.[80]

References

Arrhenius, Gustaf. *Future Generations: A Challenge for Moral Theory* Uppsala: Acta Universitatis Upsaliensis, 2000.

Balaguer, Mark. "Moral Folkism and the Deflation of (Lots of) Normative and Metaethics." In *Abstract Objects: For and Against*, edited by José L. Falguera and Concha Martinez-Vidal, 297–318. Cham: Springer, 2020.

Baker, Derek. "Skepticism about Ought Simpliciter." In *Oxford Studies in Metaethics Vol. 13*, edited by Russ Shafer-Landau, 230–252. Oxford: Oxford University Press, 2018.

Bedke, Matthew. "A dilemma for non-naturalists: irrationality or immorality?" *Philosophical Studies* 177, no. 4 (April 2020): 1027–1042.

Blackburn, Simon. *Essays in Quasi-Realism*. Oxford: Oxford University Press, 1993.

Bratman, Michael. (1981). "Intention and Means-End Reasoning." *The Philosophical Review* 90, no. 2 (April 1981): 252–265.

Broome, John. *Rationality through Reasoning*, Malden: Wiley-Blackwell, 2013.

Bykvist, Krister. "How to Do Wrong Knowingly and Get Away with It." In *Neither/Nor: Philosophical Papers Dedicated to Erik Carlson on the Occasion of His Fiftieth Birthday*, edited by Frans Svensson and Rysiek Sliwinski, 31–47. Uppsala: Department of Philosophy, Uppsala University, 2011.

Bykvist, Krister. "Violations of Normative Invariance: Some Thoughts on Shifty Oughts", *Theoria* 73, no. 2 (April 2007): 98–120.

Bykvist, Krister, and Jonas Olson. "Expressivism and Moral Certitude." *The Philosophical Quarterly* 59, no. 235 (April 2009): 202–215.

Björklund, Fredrik, Gunnar Björnsson, John Eriksson, Ragnar Francén Olinder, and Caj Strandberg. "Recent Work on Motivational Internalism." *Analysis* 72, no. 1 (January 2012): 124–137.

Carlson, Erik. "Deliberation, Foreknowledge, and Morality as a Guide to Action." *Erkenntnis* 57 (July 2002): 71–89.

Chang, Ruth. (2004). "'All Things Considered.'" *Philosophical Perspectives* 18, no. 1 (December 2004): 1–22.

Clarke-Doane, Justin. "From Non-Usability to Non-Factualism." *Analysis* (forthcoming). Draft available at: https://philpapers.org/archive/CLAFNT-2.pdf (retrieved 24 July 2021).

Clarke-Doane, Justin. *Morality and Mathematics*. New York: Oxford University Press, 2020.

Copp, David. "Belief, Reason and Motivation: Michael Smith's *The Moral Problem*." *Ethics* 108, no. 1 (October 1997): 33–54.

Copp, David. "The Ring of Gyges: Overridingness and The Unity of Reason." *Social Philosophy and Policy* 14, no. 1 (Winter 1997): 86–106.

Crisp, Roger. *Reasons and the Good*. New York: Oxford University Press, 2006.

Cuneo, Terence, and Russ Shafer-Landau. "The Moral Fixed Points: New Directions for Moral Nonnaturalism." *Philosophical Studies* 171, no. 3 (December 2014): 399–443.

Darwall, Stephen, Allan Gibbard, and Peter Railton. "Toward *Fin de Siècle* Ethics: Some Trends." *The Philosophical Review* 101, no. 1 (January 1992): 115–189.

Dreier, James. "Can Reasons Fundamentalism Answer the Normative Question?", in *Motivational Internalism*, edited by Fredrik Björklund, Gunnar Björnsson, John Eriksson, Ragnar Francén Olinder, and Caj Strandberg, 167–181. New York: Oxford University Press, 2015.

Dreier, James. "Meta-Ethics and the Problem of Creeping Minimalism." *Philosophical Perspectives* 18, no. 1 (December 2004): 23–44.

Eklund, Matti. *Choosing Normative Concepts*. New York: Oxford University Press, 2017.

Eklund, Matti. "The Normative Pluriverse." *Journal of Ethics and Social Philosophy* 18, no. 2 (August 2020): 121–146.

Enoch, David. "Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action." *The Philosophical Review* 115, no. 2 (April 2006): 169–198.

Enoch, David. *Taking Morality Seriously*. New York: Oxford University Press, 2011.

Enoch, David. "Thanks, We're Good: Why Moral Realism is not Morally Objectionable." *Philosophical Studies* 178, no. 5 (May 2021): 1689–1699.

Faraci, David. "On Leaving Room for Doubt." In *Oxford Studies in Metaethics, Vol. 12*, edited by Russ Shafer-Landau, 244–264. Oxford: Oxford University Press, 2017.

Feldman, Fred. "Actual Utility, The Objection from Impracticality, and the Move to Expected Utility." *Philosophical Studies* 129, no. 1 (May 2006): 49–79.

Finlay, Stephen. "Defining Normativity." in *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*, edited by David Plunkett, Kevin Toh, and Scott J. Shapiro, 187–220. New York: Oxford University Press, 2019.

Fitzpatrick, William. "Commentary on Matti Eklund, *Choosing Normative Concepts*." 2018. Available at: http://www.sas.rochester.edu/phl/fitzpatrick/PeaSoup.pdf (retrieved 4 December 2019).

Gibbard, Allan. "Morality as Consistency in Living: Korsgaard's Kantian Lectures." *Ethics* 110,

no. 1 (October 1999): 140–164.

Gibbard, Allan. *Thinking How to Live*. Cambridge: Harvard University Press, 2003.

Gibbard, Allan. *Wise Choices, Apt Feelings*. Oxford: Clarendon Press, 1992.

Graham, Peter. "In Defense of Objectivism about Moral Obligation." *Ethics* 121, no. 1 (October

2010): 88–115.

Grice, Paul. "Intention and Uncertainty." *Proceedings of the British Academy* 57 (1971): 263–279.

Harman, Elizabeth. "The Irrelevance of Moral Uncertainty." In *Oxford Studies in Metaethics, Vol.

10*, edited by Russ Shafer-Landau, 53–79. Oxford: Oxford University Press, 2015.

Jackson, Frank. "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection."

Ethics 101, no. 3 (April 1991): 461–482.

Jackson, Frank. *From Metaphysics to Ethics*. Oxford: Clarendon Press.

Kolodny, Niko, and John MacFarlane. "Ifs and Oughts." *Journal of Philosophy* 107, no. 3 (March

2010): 115–143.

Korsgaard, Christine. *The Sources of Normativity*. Cambridge: Cambridge University Press, 1996.

Lenman, James. "Consequentialism and Cluelessness." *Philosophy and Public Affairs* 29, no. 4

(Autumn 2000): 342–370.

Lord, Errol. "What You're Rationally Required to Do and What You Ought to Do (Are the Same

Thing!)" *Mind* 126, no. 504 (October 2017): 1109–1154.

MacAskill, William. "The Infectiousness of Nihilism." *Ethics* 123, no. 3 (April 2013): 508–520.

MacAskill, William, Krister Bykvist, and Toby Ord. *Moral Uncertainty*. Oxford: Oxford University

Press.

Mason, Elinor. "Objectivism and Prospectivism About Rightness." *Journal of Ethics and Social

Philosophy* 7, no. 2 (March 2013): 1–22.

McGrath, Sarah. (2008). "Moral Disagreement and Moral Expertise." In *Oxford Studies in Metaethics, Vol. 4*, edited by Russ Shafer-Landau, 87–108. Oxford: Oxford University Press.

McPherson, Tristram. "Explaining Practical Normativity." *Topoi* 37, no. 4 (December 2018): 621–630.

Nowell-Smith, Patrick. H. *Ethics*. Harmondsworth: Pelican Books, 1954.

Parfit, Derek. *On What Matters: Volume Two*, Oxford: Oxford University Press, 2011.

Parfit, Derek. *Reasons and Persons*. Oxford: Oxford University Press, 1984.

Parfit, Derek. "What We Together Do." Unpublished manuscript. Available at: https://philarchive.org/archive/PARWWT-3 (Retrieved 27 July 2021)

Railton, Peter. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13, no. 2 (Spring, 1984): 134–171.

Regan, Donald. *Consequentialism and Co-Operation*. Oxford University Press, 1980.

Reisner, Andrew. "Is There Reason to Be Theoretically Rational?" In *Reasons for Belief*, edited by Andrew Reisner and Asbjørn Steglich-Petersen, 34–53. Cambridge: Cambridge University Press, 2011.

Reisner, Andrew. "Normative Conflicts and the Structure of Normativity." In *Weighing and Reasoning*, edited by Iwao Hirose and Andrew Reisner, 189–206. Oxford: Oxford University Press, 2015.

Risberg, Olle. "The Entanglement Problem and Idealization in Moral Philosophy." *Philosophical Quarterly* 68, no. 272 (July 2018): 542–559.

Risberg, Olle. "Weighting Surprise Parties: Some Problems for Schroeder." *Utilitas* 28, no. 1 (March 2016): 101–107.

Risberg, Olle, and Folke Tersman. "A New Route from Moral Disagreement to Moral Skepticism." *Journal of the American Philosophical Association* 5, no. 2 (Summer 2019): 189–207.

Risberg, Olle, and Folke Tersman. "Disagreement, Indirect Defeat, and Higher-Order Evidence." In *Higher-Order Evidence and Moral Epistemology*, edited by Michael Klenk, 97–114. London: Routledge, 2020.

Risberg, Olle, and Folke Tersman. "Moral Realism and the Argument from Skepticism." *International Journal for the Study of Skepticism* 10, no. 3–4 (November 2020): 283–303.

Ross, Jacob. "Rationality, Normativity, and Commitment." In *Oxford Studies in Metaethics Vol. 7*, edited by Russ Shafer-Landau, 138–181. Oxford: Oxford University Press, 2012.

Schroeder, Mark. *Being For*. Oxford: Oxford University Press, 2008.

Schroeder, Mark. "Ought, Agents, and Actions." *The Philosophical Review* 120, no. 1 (January 2011): 1–41.

Schroeder, Mark. "Realism and Reduction: The Quest for Robustness," *Philosophers' Imprint* 5, no. 1 (February 2005): 1–18.

Sepielli, Andrew. "Subjective and Objective Reasons." In *Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Oxford University Press, 2018. DOI: 10.1093/oxfordhb/9780199657889.013.34

Sepielli, Andrew. "What to Do When You Don't Know What to Do." In *Oxford Studies in Metaethics Vol. 4*, edited by Russ Shafer-Landau, 5–28. Oxford: Oxford University Press.

Sepielli, Andrew. "What to Do When You Don't Know What to Do When You Don't Know What to Do…" *Noûs* 48, no. 3 (September 2014): 521–544.

Sidgwick, Henry. *Methods of Ethics*. London: Macmillan, 1874.

Skorupski, John. *The Domain of Reasons*. Oxford: Oxford University Press, 2010.

Smith, Holly. *Making Morality Work*. Oxford: Oxford University Press, 2018.

Smith, Michael. "Moore on the Right, the Good, and Uncertainty." In *Metaethics After Moore*, edited by Terence Horgan and Mark Timmons, 133–148. Oxford: Oxford University Press, 2006.

Smith, Michael. *The Moral Problem*. Oxford: Blackwell, 1994.

Tännsjö, Torbjörn. *From Reasons to Norms: On the Basic Question in Ethics*. Dordrecht: Springer, 2010.

Tersman, Folke. *Moral Disagreement*. Cambridge: Cambridge University Press, 2006.

Tiffany, Evan. "Deflationary Normative Pluralism", *Canadian Journal of Philosophy* 37, supplementary volume 1 (2007): 231–262.

Timmons, Mark. *Moral Theory: An Introduction*. Lanham: Rowman & Littlefield Publishers, 2002.

Toppinen, Teemu. "Non-Naturalism Gone Quasi: Explaining the Necessary Connections between the Natural and the Normative." In *Oxford Studies in Metaethics Vol. 13*, edited by Russ Shafer-Landau, 26–47. Oxford: Oxford University Press, 2018.

Way, Jonathan, and Daniel Whiting. "Perspectivism and the Argument from Guidance." *Ethical Theory and Moral Practice* 20, no. 2 (April 2017): 361–374.

Weatherson, Brian. *Normative Externalism*. Oxford: Oxford University Press, 2019.

Weatherson, Brian. "Review: Moral Uncertainty and Its Consequences." *Mind* 111, no. 443 (July 2002): 693–696.

Wedgwood, Ralph. "Conceptual Role Semantics for Moral Terms." *The Philosophical Review* 110, no. 1 (January 2001): 1–30.

Zimmerman, Michael. *Ignorance and Moral Obligation*. Oxford: Oxford University Press, 2014.

Zimmerman, Michael. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge: Cambridge University Press, 2008.

Zimmerman, Michael. *The Concept of Moral Obligation*. Cambridge: Cambridge University Press,

1996.