

Epistemic modesty in ethics

Nicholas Laskowski¹

Published online: 22 May 2017
© Springer Science+Business Media Dordrecht 2017

Abstract Many prominent ethicists, including Shelly Kagan, John Rawls, and Thomas Scanlon, accept a kind of epistemic modesty thesis concerning our capacity to carry out the project of ethical theorizing. But it is a thesis that has received surprisingly little explicit and focused attention, despite its widespread acceptance. After explaining why the thesis is true, I argue that it has several implications in metaethics, including, especially, implications that should lead us to rethink our understanding of Reductive Realism. In particular, the thesis of epistemic modesty in ethics implies a kind of epistemic modesty about the metaphysical nature of ethics, if Reductive Realism about the metaphysical nature of ethics is true, and it implies that normative concepts are indispensable to practical deliberation in a way that answers an influential objection to Reductive Realism from Jonathan Dancy, David Enoch, William FitzPatrick, and Derek Parfit.

Keywords Normative ethics · Metaethics · Metaphilosophy · Normative Concepts · Reductive Realism · Robust Realism

“Working out the terms of moral justification is an unending task”.
Thomas Scanlon (1998: 361)

“...the process of evaluation and justification [of theories in ethics]
can perhaps never be completely finished”.
Shelly Kagan (1998: 16)

“Taking [ethical theorizing] to the limit, one seeks the conception, or plurality of
conceptions, that would survive the rational consideration of all feasible conceptions
and all rational arguments for them. We cannot,
of course, actually do this...”.
John Rawls (1974: 289)

✉ Nicholas Laskowski
nicholaslaskowski@gmail.com

¹ Institut für Philosophie, Universität Duisburg-Essen, Universitätsstr. 12,
Essen, Nordrhein-Westfalen 45117, Germany

1 Introduction

There are many concerns about the epistemology of ethical theorizing that ethicists tend to bracket. It is not common to find ethicists beginning articles or books with responses to epistemological objections from disagreement or evolution before engaging in ethical theorizing, for example. Ethicists appear to proceed on the reasonable assumption that such objections can be successfully met another day. Yet the epigraphs indicate that there is one concern about the epistemology of ethical theorizing that many prominent ethicists concede—a concern about our capacity to carry out the project of ethics to its limit. But while it is often conceded, it is a concern that has received little explicit and focused discussion. It is perhaps unsurprising, then, that its implications have not been sufficiently appreciated. The goal of this paper is to explore several underappreciated implications of a kind of widely accepted thesis of epistemic modesty in ethics.

In Sects. 2–2.4, I will unpack the thesis of epistemic modesty in ethics that ethicists like Kagan, Rawls, and Scanlon concede, and explain why they are correct in conceding it. In Sects. 3–3.3, I will argue that the thesis of epistemic modesty in ethics has at least three implications that have not been sufficiently appreciated, if they've been appreciated at all. The first implication tells against a suggestion from Michael Smith (2004), Frank Jackson (1998), and Derek Parfit (2011) that the attitudes of suitably idealized agents will converge so long as there isn't any disagreement among them. The second and third implications are more significant; they force us to rethink the commitments of *reductive* theories about the nature of ethics. In particular, the second implication is that we should also be epistemically modest about the metaphysical nature of ethics, if a reductive theory of the nature of ethics is true. This implication turns out to be of surprising benefit to reductivists, as the third implication is that it provides them with a novel line of response to an influential objection that I call the Dispensability Objection, which runs through the work of many high-profile opponents of reductivism, including Jonathan Dancy (2006), David Enoch (2011), William FitzPatrick (2008), and Derek Parfit (2011).

2 The epistemology of ethics

One of the central questions of “normative ethics,” “first-order ethics,” or simply “ethics,” as I will often refer to it, is the question of how to live. Mark Timmons (2012: 16) writes, for example, that normative ethics “attempts to answer very general moral questions about what to do and how to be”. Kagan (1998: 2, his emphasis) says, too, that it “involves substantive proposals about how to act, how to live, and what kind of person to be”. The most general and hence comprehensive proposal for answering the question of how to live would take the form of a substantive theory on which for all x , x is N just in case x is F , where N is a normative feature such as *morally obligatory* and F is a non-trivial condition such as *optimific*. The focus of this paper is the epistemology of such comprehensive theories. To fix ideas, we can start with a strong thesis of epistemic modesty in ethics—one that we will refine as we go along.

Too Strong Necessarily, for any subject S and any comprehensive ethical theory P, S does not know P

Throughout much of the history of ethics it has been thought that moral knowledge comes from reflection on cases and principles, which suggests that whatever moral knowledge it is possible for S to have of a comprehensive ethical theory might come from such reflection.¹ Call a sequence of reflection that begins with judgments about cases the *method of cases*. A close look at the method of cases will reveal why it is the wrong kind of thing to give us knowledge of a comprehensive ethical theory. This doesn't quite vindicate Too Strong, but as I will explain in Sects. 2.2–2.4, it gets us something close to it.

2.1 The method of cases, illustration

Let's begin with a simple case from Fred Feldman (1986) where a subject *s* has three actions available at time t_1 that will lead them to experience various levels of pleasure.

- (a1) work in garden +12
- (a2) go to dump +08
- (a3) start painting house –15

Intuitively, *s* ought to work in the garden at t_1 because doing so will lead them to experience more pleasure than going to the dump and starting to paint the house. A theory that this intuition might lead one to embrace is a version of Act Utilitarianism in the tradition of Bentham, which says that for any action *x*, agent *s*, and any time *t*, *s* is morally obligated to *x* at *t* just in case *s*'s performing *x* maximizes pleasure at *t*. This view seems to capture the intuition that *s* ought to work in the garden at t_1 by predicting that *s* ought to work in the garden at t_1 , because doing so will lead *s* to experience more pleasure than the alternative actions available at the time.

But to get a handle on exactly how Benthamite Act Utilitarianism works and a feel for its extensional adequacy it seems necessary to test it against more cases. Let's continue to follow Feldman in imagining, then, that *s* wouldn't enjoy any of the preparation that they would have to do at t_0 to work in the garden at t_1 such as gathering tools from the shed. Imagine, too, that *s* would enjoy some of the other actions that they could perform at t_0 , such as loading the dump truck. All told, *s* has the following options available at t_0 .

- (a4) gather tools –01
- (a5) load truck +01
- (a6) mix paint –02

Even though *s* would find gathering tools mildly unpleasant, it intuitively seems like they ought to do so, since *s* would then be in a position at t_1 to perform the

¹ As McMahan (2016) notes, in the context of a discussion of the role of cases in ethics, "Hypothetical examples, even when used as a means of understanding the most serious of moral issues, have been deployed by philosophers at least since Plato appealed to the ring of Gyges...".

thoroughly enjoyable act of gardening. However, on Benthamite Act Utilitarianism, *s* has an obligation to load the truck at t_0 , since doing so maximizes pleasure at t_0 . Moreover, *s* also has an obligation to garden at t_1 on Benthamite Act Utilitarianism. But if *s* loads the truck at t_0 then they can't fulfill their obligation to work in the garden at t_1 . Intuitively, *s* can't have an obligation to perform an action at t_0 that puts *s* out of a position to fulfill another moral obligation at t_1 , so it seems like the method of cases reveals that Benthamite Act Utilitarianism is extensionally inadequate, since it fails to accord with intuitions about this case.

At this point, some readers might be thinking that they don't share these judgments about Feldman's cases and might reasonably be feeling a bit restless about the discussion as a result. But allow me to urge such readers not to fret. It's true that my discussion in this section is proceeding on the assumption that my judgments about cases are probative. However, it is a assumption that not only makes the exposition easier to follow, but actually makes the job of explaining why we have excellent reason to accept a thesis of epistemic modesty in ethics even harder, since diversity among judgments about cases itself provides strong reason to be epistemically modest in ethics.

Taking a second look at Benthamite Act Utilitarianism, it isn't clear that we have a counterexample to the view that we had in mind. Instead, the second version of Feldman's case seems to serve as a counterexample to a version of Benthamite Act Utilitarianism that is *non-sequentialist*, as I'll call it. A sequentialist version of Benthamite Act Utilitarianism says that for any action *x*, agent *s*, and time *t*, *s* is morally obligated to *x* at time t_n just in case *x* is part of *sequence* of actions available to *s* to initiate at t_n that maximizes more pleasure than any other sequence available to *s* to initiate at t_n . A sequentialist version of Benthamite Act Utilitarianism does not seem to predict that *s* has an obligation at t_0 that puts *s* out of a position to fulfill their obligation at t_1 , because such a version says that *s* is obligated to perform whichever *sequence* of actions produces more pleasure than any other available sequence of actions, which in this case is *a4* and *a1*. So, if anything, it looks like we have a counterexample to Benthamite Act Non-Sequentialist Utilitarianism but not Benthamite Act Sequentialist Utilitarianism.

To ensure that we have a firm understanding of the view under consideration in Benthamite Act Sequentialist Utilitarianism, it seems necessary to apply it to more cases. Consider, then, a case from Michael Zimmerman (2008: 120), who imagines Brenda inviting her ex-fiancé Alf to her wedding. As Zimmerman describes the case, the best thing for Alf to do would be to accept the invitation, show up, and behave, the worst thing would be to show up and misbehave, and the second best thing would be for Alf to decline and not show up at all. However, Alf is far from mature, and he would not behave if he were to accept his invitation to the wedding.

Intuitively, it seems like Alf ought to decline the invitation to avoid misbehaving. But because the sequence of actions available to Alf to initiate at t_0 that maximizes pleasure is the one that involves accepting the invitation, attending the wedding, and behaving, Benthamite Sequentialist Act Utilitarianism seems to predict that Alf ought to initiate it. Thus, Benthamite Sequentialist Act Utilitarianism seems extensionally inadequate, because it yields unintuitive predictions about this case.

Again, however, upon closer inspection, it isn't clear that we have a counterexample to the theory that we had in mind in Benthamite Sequentialist Act Utilitarianism. For it seems that it is a version of Benthamite Sequentialist Act Utilitarianism that is *Possibilist* in form that is extensionally inadequate, on which for any action x , agent s , and time t , agent s is morally obligated to x at t_n just in case x is part of a sequence of actions available to s at t_n that s can initiate and the amount of pleasure that *could* result if s did initiate it is greater than the amount of pleasure that *could* result if s did not initiate it. What this counterexample does not seem to reveal is that a version of Benthamite Sequentialist Act Utilitarianism that is *Actualist* is extensionally inadequate, on which for any action x , agent s , and time t , s is morally obligated to x at t_n just in case x is part of a sequence of actions available to s at t_n that s can initiate and the amount of pleasure that *would* result if s did initiate it is greater than the amount of pleasure that *would* result if s did not initiate it.²

But to ensure that we have a handle on Benthamite Actualist Sequentialist Act Utilitarianism, it again seems necessary to consult our intuitions about more cases. To this end we might follow Allan Gibbard (1965) in supposing that you and I have a pair of buttons in front of us—A and B. If we both push B, then that would produce more pleasure than any other available sequence of actions. If we both push A, then that would produce the second most amount of pleasure out of the sequences of actions available to us. But if I push A and you push B, or vice versa, then that would produce vastly more pain than other sequences of actions.³

Intuitively, it seems like we have an obligation to push B and that we violate this obligation by doing anything else. But suppose that we both push A. According to Benthamite Actualist Sequentialist Act Utilitarianism, we do not fail to carry out any obligations, since at t_0 I am initiating a sequence of actions that would produce more pleasure than the amount of pleasure that would be produced if I did not initiate it, and at t_0 you are initiating a sequence of actions that would produce more pleasure than the amount of pleasure that would be produced if you did not initiate it. So, it seems that we have a counterexample to Benthamite Actualist Sequentialist Act Utilitarianism.

It isn't likely to come as much of a surprise at this point that things are again not quite as they appear. On a second pass of the case, we can see that we don't have a counterexample to a view that is Benthamite Actualist Sequentialist Act Utilitarian in kind, but rather to a Benthamite Actualist Sequentialist Act Utilitarian view that is what I will call *Individualist*, in the sense that moral obligation is not a matter of what we together ought to do, over and above what agents ought to do individually. So, it looks like we have a counterexample to Benthamite Actualist Sequentialist Individualist Act Utilitarianism, but not a version of Benthamite Actualist Sequentialist Act Utilitarianism that I will call *Collectivist*, a view that predicts

² For some defenses of actualism, see Goldman (1976), Sobel (1976), Jackson and Pargetter (1986), and Goble (1993). For some defenses of possibilism, see Goldman (1976), Greenspan (1978), Humberstone (1983), Feldman (1986), and Zimmerman (2008). For some defenses of other views on the topic, see Portmore (2011), Ross (2013), and Timmerman (2015).

³ See also Regan (1980).

that while you and I are each fulfilling our individual obligations, we are failing to fulfill further obligations as a group.⁴

I'll spare the reader from engaging in another round of the method of cases in an assessment of Benthamite Collectivist Actualist Sequentialist Act Utilitarianism. At this point, we have enough in front of us to begin explaining why we have excellent reason to believe a modest thesis about epistemic modesty in ethics.

2.2 The method of cases, discussion

This illustration of the method of cases brings out something important about us: We do not have an antecedent grasp of all the ethically significant features of the world that the true ethical theory captures if there is one (I'll assume there is one from here on for ease of exposition). This was first made apparent in our assessment of Benthamite Act Utilitarianism, when we saw that it didn't account for the ethical significance of *sequences* of actions. It was made more apparent in our assessment of Benthamite Sequentialist Act Utilitarianism, when we saw that it didn't account for the ethical significance of the *modal profiles* of such sequences. And it was made even more apparent in our assessment of Benthamite Actualist Sequentialist Act Utilitarianism, when we saw that it didn't account for the ethical significance of *collective* obligations.⁵

But if we do not have an antecedent grasp of all the ethically significant features of the world that the true ethical theory captures, then it seems that, until we use the method of cases to reflect on such features, our concepts are such that we do not *believe* the true ethical theory. Think back to the moment before we made the Individualist/Collectivist distinction, when we were only employing the Benthamite Actualist Sequentialist Act Utilitarian idea that moral obligation is analyzable in terms of the amount of pleasure in the consequences of sequences of actions that would result from our initiating them. Because we did not yet have a grasp of the Individualist/Collectivist distinction, it wasn't *until* we used the method of cases to reveal it that our concept CONSEQUENCE was fine-grained enough for us to incorporate the Collectivist component in formulating the true ethical theory. But then it seems that it wasn't until *after* we used the method of cases in this way that we were able to believe this ethical theory.

That we do not have a grasp of all the world's ethically significant features that the true ethical theory captures is hardly a revelation in light of the fact, however, that many ethicists still have jobs! Of course, few would likely claim that we

⁴ See Jackson (1987) and especially Dietz (2016).

⁵ Ross (2006) appears to highlight a related phenomenon. "...in general, the more we reflect on questions of ethical theory, the greater is the number of ethical theories among which our credence is divided. What initially appears to be a single ethical theory often turns out to be specifiable in a number of ways, each of which has some plausibility. And when a problem arises for an initial formulation of a theory, it is often possible to solve this problem by modifying the theory in any of several ways, revealing once more a multiplicity of theories, each having some degree of plausibility".

currently believe the true ethical theory. But that we do not have such a grasp raises the important question of whether we will *eventually* have it.

Yet the answer to this question also seems to be negative. Our use of the method of cases over the course of the history of ethical theorizing, a microcosm of which is on display in the previous section, looks like it constitutes excellent inductive evidence for the claim that there will *always* be ethically significant features of which we are unaware.⁶ But if there will always be ethically significant features of the world of which we are unaware, then it seems to be the case that we will never believe the true, fully-specified ethical theory.⁷ And if we will never believe the true ethical theory then it is plausible to think that we will never know it.

In other words, the problem with the method of cases is not that it yields unreliable beliefs, unjustified beliefs, or that we cannot use it to satisfy whichever condition it is that distinguishes having a true belief about a comprehensive ethical theory and having knowledge of such a theory. Indeed, Sect. 2.1 looks like a detailed illustration of how our intuitions about cases might license believing one theory in ethics over another.⁸ The *epistemic* problem with the method of cases is ultimately *alethic*: We will never know the *true* ethical theory because we will never *believe* it.⁹

Note that this explanation doesn't imply that it is impossible to have knowledge of the ethical significance of *any* feature of the world, and hence it doesn't imply that it is impossible to know anything at all about the true ethical theory. For on the simplifying assumption that my intuitive judgments are probative, the application of the method of cases above clearly illustrates its effectiveness. After all, it seems like it is possible to know that the true ethical theory is not a version of Benthamite Act Utilitarianism that is non-sequentialist, possibilist, or individualist. This is a point that comes from Scanlon, Kagan, and Rawls, as the epigraphs of this paper indicate. It is also a point that is hard to overstate. That we will never know the true, *fully-specified* ethical theory is compatible with knowing quite a bit about it.

It might be tempting to think that this line of explanation is merely a new spin on a kind of *under determination of ethical theory by evidence*. After all, I have been focusing on the idea that ethical theories purport to tell us how to live, by telling us which actions are morally obligatory for any action. But on an alternative conception of ethical theorizing,¹⁰ ethical theories also purport to *explain why* such

⁶ In a similar spirit, Tanmsjo (2015: xi) writes, "It is not possible to show that a moral principle is true in the abstract. Moral principles always surprise us in concrete applications".

⁷ See Carlson (1995) for a striking illustration of refinements it is possible to make to existing ethical theories, particularly ethical theories in the tradition of consequentialism.

⁸ While intuitions about cases would seem to play a major role in licensing beliefs about comprehensive ethical theories on the method of cases, it seems open to friends of the method of cases to also admit that consistency, generality, internal and external support, and other epistemically relevant features of theories might factor into such licensing. Thanks to anonymous referee for inviting me to speak to this issue.

⁹ Thanks to an anonymous referee for pressing me to clarify this point, and thanks to Caleb Perl for help with doing so.

¹⁰ See Schroeder (2006).

actions are morally obligatory. And on such a conception, it might be said that our intuitive judgments about which actions we morally ought to perform in particular cases is consistent with many explanations from *different* ethical theories. For example, the intuitive judgment that s morally ought to work in the garden seems consistent with the Benthamite Act Utilitarian explanation that they morally ought to do so because it would maximize pleasure and the Divine Command Theorist explanation that they ought to do so because God forbids them to do otherwise.

But such appearances are misleading. Even if we understand ethical theorizing under this alternative conception, the illustration of the method of cases in Sect. 2.1 suggests that the full set of possible intuitive judgments *maximally discriminates* among the full set of possible ethical theories. In other words, the evidence ultimately supports the one true ethical theory, which means that the idea of underdetermination of (ethical) theory by evidence does not figure in the explanation above of why we will never know the true ethical theory.

It is also worth noting that the explanation does not rely on any sort of *pessimism about ethical induction* on which we are not in a position to have knowledge of the ethical features in case_{n+1} on the basis of our knowledge of the ethical features in case_n. For the supervenience of the normative on the non-normative suggests that we are in a position to know that case_{n+1} has the very same ethical features as case_n insofar as case_{n+1} has the same non-normative features as case_n. Yet the supervenience of the normative on the non-normative is also compatible with expecting different ethical features in case_{n+1} if the non-normative features of case_{n+1} are different from the non-normative features of case_n. The problem isn't that there is no way to guarantee the reliability of induction in ethics, but rather that cases involving non-normative features that we have yet to reflect upon tend to have ethical features that outstrip our best available theories.

Since the explanation of why we will never know the true ethical theory is not obviously an explanation that relies upon familiar ideas concerning the underdetermination of theory by evidence, nor familiar ideas concerning pessimism about induction, it is not an explanation we can obviously resist by appealing to familiar replies to these issues. Nor, too, does it help to point out that ethical theories often have *ceteris paribus* clauses built into them, such that many of them are immune to counterexamples from one-off cases. This is because the true ethical theory specifies *all* of the ways in which *cetera* fail to be *paria*. Yet the very fact that we include *ceteris paribus* clauses into ethical theories is a tacit acknowledgment that specifying how all of the ways in which *cetera* can fail to be *paria* is out of our reach, since if it wasn't so then there wouldn't be any reason not to simply specify all the exceptions.

Neither does the explanation above trade on any idiosyncrasies of using the method of cases to assess theories in the *consequentialist* tradition. We need not look further than debates in ethics about the nature of principles of reasonable rejection and whether aggregative principles are consistent with them to see that *non-consequentialist* theories are also subject to indefinite levels of refinement in accounting for all of the world's ethically significant features.¹¹ The true ethical

¹¹ See Taurek (1977), Kamm (1998), and Walden (2014) for a sampling of one such line of debate.

theory makes correct predictions about *all* possible cases, whether or not such a theory ultimately takes a consequentialist or non-consequentialist shape.

It might also be tempting to invoke Rawls (1972) to cast doubt on the explanation. In particular, it might be tempting to invoke the idea that is sometimes attributed to him that we need not heed all of our judgments about cases. Why, one might wonder, for example, couldn't a proponent of Benthamite Non-Sequential Act Utilitarianism stick to their view by dismissing the judgment in the second case from Feldman that the protagonist intuitively could not have an obligation at t_0 (to load the truck) that takes them out of a position to fulfill a different obligation at t_1 (to garden)?

But sticking to one's guns in this way is to fail to appreciate the excellent inductive evidence from the history of ethical theorizing that there will always be ethically significant features of which we are unaware.¹² Indeed, it seems like an especially egregious failing in light of the explosion of ethical theorizing in the second half of the twentieth century. Take the ethics of war, for example.¹³ Some ethicists¹⁴ have thought that we should divide up the subject matter of the ethics of war into the categories of *jus ad bellum* (resort to war) and *jus in bello* (conduct in war). But others¹⁵ more recently suggest that we should also add the category of *jus post bellum* (actions after war), while still others¹⁶ even more recently suggest that we should also add *jus ex bello* (exiting war). In each of these categories, too, there are a dizzying array of subtle issues for which many fine distinctions are currently being made to accommodate them. No less than six principles are traditionally said to be necessary and sufficient for *jus ad bellum*. Yet further scrutiny has yielded even finer distinctions. And the ethics of war is just one among very many examples. At this point in time, W.D. Ross's (2002) idea that there were less than a dozen ethically significant features of the world seems quaint.¹⁷

¹² Moreover, as the second epigraph of this paper indicates, it's not clear that Rawls (1974: 289) doesn't accept epistemic modesty in ethics. He writes, "Taking this process to the limit, one seeks the conception, or plurality of conceptions, that would survive the rational consideration of all feasible conceptions and all rational arguments for them. We cannot, of course, actually do this..." It's true that he goes on to write, "...but we can do what seems like the next best thing, namely, to characterize the structures of the predominant conceptions familiar to us from the philosophical tradition, and to work out the further refinements of these that strike us at most promising." But even if one were to think that Consequentialism, Deontology, Virtue Ethics, and other "predominant conceptions familiar to us from the [Western] philosophical tradition" were exhaustive of the space of conceptions available to us for rational consideration, the illustration of the method of cases in Sect. 2.1 still suggests that there is no limit to possible refinements to these views that would strike as "most promising".

¹³ War is a familiar recent topic of intense interest that seems to expand our sense of the space of possible views in ethics. But new and fascinating topics seem to crop up in ethics regularly, at least on our best days. See Horton (forthcoming) for a discussion of an underappreciated topic that seems likely to powerfully further illustrate this same phenomenon.

¹⁴ See Walzer (1992).

¹⁵ Bass (2004).

¹⁶ See Moellendorf (2008).

¹⁷ MacAskill (2016: 1000) echoes a similar sentiment: "Despite thousands of years of thought, we are little closer to knowing what constitutes a good life than when we started. Indeed, progress in moral philosophy seems to have found more problems than it has solved. This is true, for example, of progress

So far I have been explaining why there is excellent inductive evidence for the claim that we will never use the method of cases to know the true ethical theory. But I have been understanding the method of cases as a sequence of reflections on cases and principles that *begins with cases*. This leaves it open that it is possible to know the true ethical theory by initiating a sequence of reflections on cases and principles that *begins with principles*. In the next section, I will complete the explanation of why we should be epistemically modest in ethics by explaining why we shouldn't expect such sequences of reflections to provide any relief.

2.3 The role of principles

Instead of reasoning from particular cases to principles and to more judgments about particular cases and back again, as we do when we use the method of cases, it might be said that an alternative approach starts with principles. To fix ideas, I will focus on an exemplar of such an approach, Henry Sidgwick (1907), who attempts to extract the true ethical theory from “axioms” that he took to be self-evident. In particular, Sidgwick (380–382) famously claims that Utilitarianism derives, in some sense, from the following two principles:

Principle of Benevolence: “...each one is morally bound to regard the good of any other individual as much as his own, except in so far as he judges it to be less, when impartially viewed, or less certainly knowable or attainable by him”.

Principle of Justice: “It cannot be right for A to treat B in a manner in which it would be wrong for B to treat A, merely on the ground that they are two different individuals, and without there being any difference between the natures or circumstances of the two which can be stated as a reasonable ground for difference of treatment”.

For our purposes, Sidgwick seems to be suggesting that it is possible to know the true ethical theory in virtue of it being possible to know self-evident principles, and it being possible to derive the true ethical theory from them.

One problem with this idea is that it is a matter of tremendous controversy whether it is possible to derive, in any sense, the true ethical theory from such principles.¹⁸ But even if we grant that it is possible to do so, it seems as though the picture of the epistemology of ethical theorizing on display from Sidgwick suffers

Footnote 17 continued

in population ethics and animal ethics. It may even be that, given the difficulty of the subject matter, we should never be certain of one particular normative view—our normative evidence and experiences will always be limited, always open to many reasonable interpretations, and there will always be judgment calls involved in weighing different epistemic virtues”.

¹⁸ See Schneewind (1977), Nakano-Okuno (2011), Parfit (2011), Phillips (2011), de Lazari-Radek and Singer (2014), Shaver (2014), Hurka (2014), and Crisp (2015).

from the same sort of problem as the method of cases. As we saw above, the fundamental issue is that there is excellent inductive evidence for the claim that there will always be ethically significant features of which we are unaware, which suggests that it is not possible to believe and hence know the true ethical theory. Similarly, it seems that there is also excellent inductive evidence for the claim that there will always be axioms or refinements to such axioms for capturing ethically significant features of the world of which we are unaware.

This is no more apparent than in Sidgwick's very own refinement of the principles of benevolence and justice over the course of *The Method of Ethics*. Indeed, Sidgwick devotes much of his seven editions of *The Method of Ethics* at nearly 600 pages a piece to tweaking his formulations of the principles of Benevolence and Justice, which powerfully suggests that even if it is possible to derive the true ethical theory from self-evident principles or axioms, that there will always be self-evident principles or axioms or refinements to them of which we are unaware. But if so then it is plausible to think that we will also never know the true ethical theory by initiating a sequence of reflections starting with principles.

That neither initiating a sequence of reflections starting with particular cases nor initiating a sequence of reflections starting with principles does the trick, still leaves it open that we will know the true ethical theory in virtue of initiating a sequence of reflections starting with *both* particular cases and principles. But since it is hard to see why such a sequence wouldn't suffer from the same limitations as the other two, we are left without any other sequences of reflections that could underwrite the possibility of knowing the true ethical theory.

At this point, it would be instructive to take a step back. In Sects. 2.2 and 2.3, I explained why we will know the true ethical theory via (1) sequences of reflections beginning with cases, (2) sequences of reflections beginning with principles, or (3) sequences of reflections beginning with both cases and principles. It is easy to come away with the impression that the explanation vindicates the thesis I introduced in the beginning of Sect. 2.

Too Strong Necessarily, for any subject S and any comprehensive ethical theory P, S does not know P

But as I'll clarify in the next section, the explanation in Sects. 2.2 and 2.3 vindicates a weaker and much more plausible thesis about the epistemology of ethical theorizing.

2.4 Clarifying the idea of epistemic modesty in ethics

As we've seen, our use of standard epistemological approaches to ethical theorizing constitutes strong inductive evidence for the claim that there will always be ethically significant features of the world of which we are unaware, which suggests that we'll never have the conceptual resources to believe and hence know the true ethical theory. Recall, however, that standard epistemological approaches to ethical

theorizing are not epistemologically bankrupt. Indeed, Sect. 2.1 highlights that the method of cases is an effective tool for cataloguing the world's ethically significant features. It seems, then, that the reason why it is plausible to think that there will always be ethically significant features of the world of which we are unaware, and hence the reason why it is plausible to think that we will never believe and hence know the true ethical theory, doesn't solely trace to some kind of deficiency with standard approaches to the epistemology of ethical theorizing. Rather, it seems that it traces to some feature of *our* use of such approaches.

To see this, we might imagine God using the method of cases to become aware of all the world's ethically significant features. For if we can use the method of cases to make ourselves aware of some of the world's ethically significant features, then surely God can use the method of cases to become acquainted with many more such features. Indeed, it seems as though the nature of God is such that God could use the method of cases indefinitely. And if so, it seems as though all the world's ethically significant features would be within God's reach. This suggests that the explanation from Sect. 2.2 and 2.3 *doesn't* imply that it is *metaphysically impossible* to know the true ethical theory. Instead of Too Strong, the explanation on offer supports a more modest thesis of epistemic modesty in ethics.

Modest Necessarily, for any subject *S like us in a world like ours* and any comprehensive ethical theory *P*, *S* does not know *P*

This thesis says that our limitations as finite creatures are such that we cannot use standard approaches to the epistemology of ethical theorizing to know the true ethical theory. It does not, however, rule out that non-finite entities like God could use such approaches to know the true ethical theory. Yet if God could indefinitely reflect in such a way as to know the true ethical theory, isn't it natural to think that we could also know it in virtue of God *telling* it to us?

Perhaps. But a number of conditions would have to be met for this to be the case. In addition to it being the case that God exists, it has to be the case that God's nature is such that God interacts with the world. Moreover, and perhaps more interestingly, it has to be the case that God could tell someone the true ethical theory in a way where the person could *recognize* it as such. But it is plausible to think that the class of people who could recognize that God is communicating the true ethical theory to them in a way where they could recognize it is exceptionally limited. Such a person would plausibly not only have to be a theist, but a theist who believes that God's nature is such that God interacts with the world. Such a theist would plausibly also have to understand what ethical theories even look like, and moreover, that she has some antecedent sympathy on her part for Consequentialism, Deontology, or whichever theory it is that God is telling them is true. This *isn't* to say that a scenario in which God communicates the true ethical theory to a Theistic Interactionist Consequentialist (or Deontologist or whatever) Ethicist in a way where they believe it is metaphysically impossible. But it is to say that such a scenario doesn't seem to lend very strong support for the claim that it is possible for us to know the true ethical theory.

Of course, many people believe that a scenario like this in fact took place. After all, it is often said that God told Moses the true ethical theory upon communicating the Ten Commandments. Yet in addition to such a scenario depending on all the same controversial assumptions above, there is another problem with it that is more in the spirit of the explanation of why we should be epistemically modest in ethics. It seems like it would take an awful lot of work to come to know the Ten Commandments as a an ethical theory that tells us which actions are morally obligatory for *any* action. For example, it's hard to tell which actions the Ten Commandments prescribes to perform in Feldman's garden case, and which actions you and I morally ought to perform in Gibbard's case. To figure this out, it seems like s, and us, would need to reflect on such cases, and likewise for any case we have yet to reflect upon. This looks like the same problem all over again.

There is another reason to think that the explanation in Sects. 2.2 and 2.3 supports Modest over Too Strong. Imagine an intelligent being from outside our solar system who has been observing the location of molecules in our solar system over the last 150 thousand years, such that it knows the correct theory of reference determination and all the facts about our world, or at least all the facts about our world that are relevant to determining the referents of 'obligatory', 'wrong', and other paradigmatic normative words in English. Assuming that the correct theory of reference determination is naturalistic, such a being would plausibly know whether any particular action instantiates the property that 'wrong' picks out.¹⁹ But if the intelligent being is a speaker of our language, and it is appropriate for speakers to use sentences just in case they know that such sentences are true, and speakers know that the sentences that they are using are true just in case they know that the entities picked out by them are instantiated, then it is plausible to think that the intelligent being knows all the word's ethically significant features. After all, such a creature would know all the conditions under which uses of 'wrong' are true. Thus, via semantic descent, such a being would know how to formulate a theory that accounts for all of the world's ethically significant features.

It could be, then, that we will know the true ethical theory in virtue of such a being telling it to us. But such a scenario seems to lend as little support to the idea that we will know the true ethical theory as the divine testimony scenario above. Not only does it depend on the controversial idea of a highly intelligent being knowing all the facts that are relevant to reference determination, it also depends, perhaps more interestingly, on the controversial assumption that if an entity knows all the facts about how words in English are used, then it knows all the ethically significant features that figure in the true ethical theory.²⁰

To be clear, nothing about the explanation of Modest on offer rules out the metaphysical possibility of knowing the true ethical theory, nor does anything about it rule out that we will know the true ethical theory in virtue of divine or alien testimony. All the explanation says is that there is excellent inductive evidence for the claim that there is a limit on the capacity of creatures like us in a world like ours

¹⁹ See Gibbard (2014) for doubts about such an assumption.

²⁰ See McKeown-Green et al. (2015).

to use the epistemological methods for ethical theorizing that are most plausibly available to us, such as initiating sequences of reflection starting from particular cases or principles. In the following sections, however, I will show that Modest does not fail to be interesting because of its modesty.

3 Epistemic modesty in ethics, implications

That we will never know the true ethical theory has at least three implications, which I will discuss in (roughly) ascending order of significance. First, it implies that disagreement is not the only or perhaps even the primary obstacle to the familiar idea that sufficiently informed and rational agents converge in their attitudes toward ethics. Second, it implies that we have excellent inductive evidence to believe that we will never know the true *reductive* theory of the nature of ethics if there is one (I'll assume there is for ease of exposition from here on). Third, it implies that it is possible to rebut a recently prominent objection to reductive views about the nature of ethics that I call the Dispensability Objection.

3.1 First implication: lack of disagreement does not guarantee convergence

It is a familiar idea in ethics that the normative content of the attitudes of sufficiently rational agents who are sufficiently informed about the non-normative facts *converge*. It is also an idea that plays an important role in much ethical theorizing. For example, Smith (2004: 338) maintains that the failure of such convergence entails error theory, Jackson (1998: 137) suggests that convergence is necessary to ensure that we're all talking about a common subject matter when we use normative words, and Parfit (2011) even implies that his life would be meaningless without such convergence.

Another familiar idea in ethics is that *disagreement* is one of the primary obstacles to convergence. For example, Smith (2004: 110) writes, "...it may well be the case that if normative ethics progresses but without making any significant impact on the deep-seated disagreement that exist in the community on evaluative matters, then we will, for good reasons, come to lose our conviction that we would all converge on a single set of desires...". Jackson writes, (1998: 137), "I take it that it is part of current folk morality that convergence will or would occur. We have some kind of commitment to the idea that moral disagreements can be resolved by sufficient critical reflection...Indeed some hold that we know enough as of now about moral disagreement to know that convergence will (would) not occur". Indeed, Parfit dedicates the bulk of *On What Matters* to showing that no one really disagrees with him.

It is easy to read these statements and come away with the impression that convergence would result if there were no disagreement. But this impression might be mistaken. On the explanation of Modest on offer, we will never know the true ethical theory because we will never believe it. And we will never believe it because we lack the capacity to catalogue all of the world's ethically significant features. Insofar as sufficiently rational agents who are sufficiently informed about the non-

normative facts are like us in also lacking such a capacity, it is the case that such agents will never believe the true ethical theory, too. But then two such agents could fail to converge on the true ethical theory, in the sense that they could both fail to believe it, even if there isn't any disagreement between the two of them.

3.2 Second implication: epistemic modesty about the nature of ethics

It is common to distinguish “normative ethics” or “ethics” as I have been calling it from “metaethics”. As we’ve seen, ethics is said to address the question of how to live. In contrast, metaethics is commonly said to address questions *about* the question of how to live, like the *linguistic* question of how we manage to talk about how to live and the *metaphysical* question of the nature of how to live. Indeed, it is sometimes said that answers to metaethical questions do not carry any commitments about how to live, and vice versa.

Yet some philosophers take issue with this claim. In his discussion of the issue, Tristram McPherson (2008: 3) remarks that the linguistic claim that ‘right’ means ‘conducive to happiness’ “transparently has implications for the content of the correct normative theory”. Mark Schroeder ([forthcoming](#), emphasis mine) makes a similar point in discussing claims about the nature of ethics. In particular, in commenting on *reductive* analyses of wrongness, he writes, “once we know *what it is* to be wrong, we will know a condition that is necessary and sufficient for something to be wrong...we will have attained the holy grail of the most ambitious and general kind of explanatory normative ethical theory”.

McPherson’s and Schroeder’s idea that the distinction between ethics and metaethics is not as clean as tradition would have it seems plausible. In particular, it seems plausible to hold that theses about something’s nature entail the conditions under which generalizations about it are true. Consider the reductive thesis that *what it is* to be pure grain alcohol is to be an ionic compound of carbon, hydrogen, and oxygen. This thesis looks like it entails the generalization that for all *x*, *x* is pure grain alcohol just in case *x* is an ionic compound of carbon, hydrogen, and oxygen. Similarly, the Reductive Benthamite Act Utilitarian view that what it is to be morally obligatory is to maximize pleasure looks like it entails the substantive normative generalization that for all *x*, *x* is morally obligatory just in case *x* maximizes pleasure.

In light of the points made by McPherson and Schroeder, it is natural to think that there are other connections between ethics and metaethics. For example, it might be that if we are or will eventually be in a position to know the true, fully-specified reductive theory about the nature of ethics, then we are or will eventually be in a position to know the true, fully-specified ethical theory that follows from it. But as we’ve seen, it seems like we will never know the true, fully-specified ethical theory. If not, however, then it follows that we will never be in a position to know the true, fully-specified reductive theory about the nature of ethics. Epistemic modesty in ethics entails epistemic modesty about the nature of ethics.

3.3 Third implication: the indispensability of normative concepts

Reductive Realists (Reductivists, from here on) about the nature of ethics claim that while normative properties figure in a metaphysical explanation of everything, no such properties figure in a metaphysical explanation of everything at the most fundamental level. On such views, normative properties are not among the basic building blocks of the world; instead, normative properties are fully constituted by, fully analyzable in terms of, fully grounded in, or “fully reducible to,” as I will say for ease of exposition, other properties that are among the basic building blocks of the world.²¹ For example, on the most prominent version of reductive realism—Reductive Naturalism²²—the basic building blocks of the world are properties that are discoverable via the empirical sciences, and normative properties are fully analyzable in terms of these so-called natural properties.

There is a prominent high-level objection to reductive naturalism and reductivism more generally that purports to show that no reductive view even could be true. Dancy, FitzPatrick, Enoch, and Parfit each suggest that if reductivism is true, then we do not *need or have reason* to use normative concepts to think about the world in deciding on how to act in it. Their motivating thought seems to be that if all there is, so to speak, to the normative property of being morally obligatory is (say) the natural property of maximizing pleasure, then settling whether a subject has an obligation to ϕ or decisive reason to ϕ is just a matter of settling whether ϕ -ing maximizes pleasure.

But the idea that we do not need or have reason to use normative concepts strikes these opponents of reductivism as profoundly mistaken. As Dancy (2006: 142) writes, “If we do not use normative concepts, we cannot address the question what is practically relevant and what is not”. Similarly, FitzPatrick (2008: 176, emphasis mine) claims that “we are not aiming at discovering [any non-normative property in deliberation], but at discovering...the truth about what is good (or right, or what there is reason to do, and so on), *as such*”. As Enoch (2011: 108) also puts it: “When I ask myself what I should do, it seems that just answering ‘Oh, pressing the blue button will maximize happiness’ is a complete non-starter, it completely fails to address the question [of whether to push it]”. Parfit (2011: 367), too, claims that “We should expect that, on [reductivism], we don’t need to make irreducibly normative claims”.

According to these opponents of reductivism, we need or have reason to use normative concepts in deliberation. But then it follows from this claim and the claim that if reductivism is true then we do not need or have reason to use normative concepts that reductivism is false. I call this high-level objection to reductivism *The*

²¹ Compare Jackson (1998) and Shafer-Landau (2003), who understand reductive realism in the metaphysical ideology of identity, to Chang (2013) and Schroeder (2007), who understand it in the ideology of metaphysical constitution, ground, and analysis, and who are broadly in line with contemporary metaphysicians like Bennett (2011) and Schaffer (2009). See also Schroeder (2006) and Dunaway (Manuscript) for discussions of the relationship between reductivism and realism.

²² For canonical defenses of reductive naturalism see Richard Boyd (1988) and David Brink (1989).

Dispensability Objection,²³ since the crux of it is the idea that reductivism implies that normative concepts are *dispensable*, in the sense that creatures like us in a world like ours can correctly, without suffering any kind of loss or committing any kind of error, stop using them and replace using them to settle deliberative questions about how to act with those non-normative concepts that are associated with those properties that are the most fundamental on reductivism.²⁴

Observe, however, that if reductivism implies anything at all about the possibility of us correctly dispensing with normative concepts, it is plausible to think that reductivism implies that it is possible for us to correctly dispense with using normative concepts on the assumption that reductivism implies that we have *knowledge* of reductivism. After all, if we were to replace our normative concepts with non-normative concepts that did not pick out the normatively fundamental properties, then we would end up acting radically immorally!

This is where epistemic modesty in ethics comes into play as implying a possible response to the *Dispensability Objection*.²⁵ We already saw that epistemic modesty in ethics entails epistemic modesty about the nature of ethics. Since we will never know the true, fully-specified ethical theory, we will never know the true reductive theory from which it follows. But if we will never have such knowledge, then we will never satisfy the knowledge condition on correctly dispensing with normative concepts suggested above.²⁶ But if we will never satisfy such a condition, then reductivism doesn't imply that it is possible for us to correctly dispense with using normative concepts in settling deliberative questions about how to act, contra the *Dispensability Objection*.

At this point, some readers will no doubt think that I haven't taken the *Dispensability Objection* seriously enough. Indeed, some readers might believe that it is possible to revive the *Dispensability Objection*. The core premise of such an argument might be not that reductivism implies that creatures like us in a world like ours can correctly dispense with using normative concepts, but instead that it follows from reductivism that, so long as it is metaphysically possible to know it, it is possible to correctly dispense with using normative concepts *in principle*. The

²³ 'The *Dispensability Objection*' is a wink to 'The *Indispensability Objection*' from Enoch, who argues that we ought to conclude that there are primitive normative properties since we can't avoid committing ourselves to the truth of them.

²⁴ Note that proponents of the *Dispensability Objection* do not unpack the sense in which they think reductivism implies that we can dispense with using normative concepts. Since the true reductive theory is merely a metaphysical thesis about the nature of normative properties, it is hard to see why proponents of the *Dispensability Objection* would think that it implies that we *psychologically* can stop using normative concepts, let alone replace using them with non-normative concepts. It's true that reductive theories are often packaged with auxiliary hypotheses about the nature of normative language, thought, and concepts, but strictly speaking, reductive theories do not obviously have anything to say about normative concepts and the psychological possibility of dispensing with them. It is for this reason that I am understanding the sense of 'can' at issue in the argument as 'can correctly', which I am also understanding broadly as 'without suffering any kind of loss or committing any kind of error'.

²⁵ For another, less ecumenical response to the *Dispensability Objection*, see Laskowski (2015).

²⁶ The decision to treat knowledge as the norm on dispensing with normative concepts is incidental. Reductivism doesn't imply anything at all about *any* of our rational attitudes towards it.

thought then might be that since it is not in principle possible to correctly dispense with using normative concepts, reductivism is false.

Moreover, recall that earlier in this paper, we saw that the explanation of Modest on offer is compatible with the metaphysical possibility of knowing the true ethical theory. In light of this, it is natural to think that it is also metaphysically possible to have knowledge of reductivism. After all, if God and super intelligent aliens can know the true ethical theory, then it seems plausible to think that they can also have knowledge of the reductive theory from which it follows. At the end of the day, perhaps one might believe that reductivism is in fact vulnerable to the Dispensability Objection.

Since I think the first version of the Dispensability Objection represents the full force of the objection, I'm inclined to think it is a mistake to believe that it is possible to revive the Dispensability Objection in this way. Since it is hard to even get a grip on the idea of it being possible to correctly dispense with using normative concepts in principle, it is hard to take claims about the alleged implications of reductivism in cases that depend on it seriously. But even if we were to have a grip on such an idea, and even if reductivism were to imply that it is possible to correctly dispense with using normative concepts in principle, it is hard to see why this would be a problem for reductivism, and nothing that proponents of the original Dispensability Objection have said provide us with any indication.²⁷

4 Conclusion

That we have excellent inductive evidence for the claim that we will never know the true ethical theory if there is one is a modest thesis about epistemic modesty in ethics that many ethicists accept. Yet even though many ethicists accept it, it is a thesis that has not received much if any explicit and focused attention. In the first half of this paper, I explained why it is true and in the second half of this paper, I argued that the thesis doesn't fail to be interesting because of its modesty. In particular, I argued that it implies that suitably idealized agents might not converge in the content of their attitudes toward ethics even if there is no disagreement among them. Then I argued that we also have excellent inductive evidence to believe that we will never know the true reductive theory if there is one. This implication is interesting not just because reductivists can use it to answer a prominent objection

²⁷ It might be said that this is where Parfit's (2011: 368) objections to "hard naturalism," or reductivist views that accept the dispensability of normative concepts, kick in. In discussing Brandt, Parfit makes it clear (375) that his problem with such views is that claims involving normative concepts ultimately come out as "trivial." But Parfit either means that such claims would be trivial for creatures like us in a world like ours or not. If he means the former, it's hard to see why reductivists (hard or otherwise) should worry about what Parfit has to say, since such reductivists can simply claim that because analyses can be non-obvious, claims involving normative concepts aren't guaranteed to be trivial. If he means the latter, it's still hard to see why reductivists should worry. After all, the conditions under which claims involving normative concepts would be trivial for God or highly intelligent aliens are far from clear. Moreover, again, it's even less clear why it would be a problem if such claims were trivial for them, which is precisely what we were hoping to find in this appeal to Parfit's objections to hard naturalism.

to their view, as I argued also follows from the thesis of epistemic modesty in ethics, but because it seems implicit in some discussions of reductivism that reductivists can be distinguished by whether they accept that it is possible for us to know the true ethical theory a priori or a posteriori.²⁸ But if the explanation and arguments of this paper are on the right track, it might be that such discussions are based on the false presupposition that it is possible for us to know the true ethical theory, at all.

Acknowledgements Special thanks to David Copp, Terence Cuneo, Alexander Dietz, Stephen Finlay, Joe Horton, Nathan Robert Howard, Tanya Kostochka, Janet Levin, Michael Milona, Mark Schroeder, Caleb Perl, Abelard Podgorski, Ralph Wedgwood, and Daniel Wodak for insightful feedback on various drafts of this paper. Thanks also to audiences at the USC Speculative Society and National Autonomous University of Mexico for helpful comments.

References

- Bass, G. (2004). Jus post bellum. *Philosophy and Public Affairs*, 32(4), 384–412.
- Bennett, E. (2011). By our bootstraps. *Philosophical Perspectives*, 25, 27–41.
- Boyd, R. (1988). How to be a Moral Realist. In G. Sayre-McCord (Ed.), *Essays on Moral Realism* (pp. 181–228). Ithaca: Cornell University Press.
- Brink, D. (1989). *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Carlson, E. (1995). *Consequentialism Reconsidered*. Netherlands: Springer.
- Chang, R. (2013). Grounding practical normativity: Going hybrid. *Philosophical Studies*, 164, 163–187.
- Crisp, R. (2015). *The Cosmos of Duty*. Oxford: Oxford University Press.
- Dancy, J. (2006). Nonnaturalism. In D. Copp (Ed.), *The Oxford Handbook of Ethical Theory*. Oxford: Oxford University Press.
- de Lazari-Radek, K., & Singer, P. (2014). *The Point of View of the Universe: Sidgwick and Contemporary Ethics*. Oxford: Oxford University Press.
- Dietz, A. (2016). What we together ought to do. *Ethics*, 126(4), 955–982.
- Dunaway, B. (Manuscript). “The Metaphysical Conception of Realism”. Available at http://www.personal.umich.edu/~dunaway/MetaphysicalConception_MT2014.pdf.
- Enoch, D. (2011). *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Feldman, F. (1986). *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht: D. Reidel Publishing Company.
- Fitzpatrick, W. (2008). Robust ethical realism, non-naturalism, and normativity. *Oxford Studies in Metaethics*, 3, 159–205.
- Gibbard, A. (1965). Rule-utilitarianism: Merely an illusory alternative? *Australasian Journal of Philosophy*, 43, 211–220.
- Gibbard, A. (2014). *Meaning and Normativity*. Oxford: Oxford University Press.
- Goldman, H. S. (1976). Dated rightness and moral imperfection. *The Philosophical Review*, 85, 449–487.
- Greenspan, P. S. (1978). Oughts and determinism: A response to goldman. *The Philosophical Review*, 87, 77–83.
- Heathwood, C. (2013). Reductionism in ethics. In H. LaFollette (Ed.), *The International Encyclopedia of Ethics*. Hoboken: Wiley.
- Horton, J. (forthcoming). The all of nothing problem. *The Journal of Philosophy*.
- Hurka, T. (2014). *British Ethical Theorists from Sidgwick to Ewing*. Oxford: Oxford University Press.
- Jackson, F. (1987). Group morality. In J. J. C. Smart, P. Pettit, R. Sylvan, & J. Norman (Eds.), *Metaphysics and Morality: Essays in Honour of J.J.C. Smart*. Oxford: B. Blackwell.
- Jackson, F. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.
- Jackson, F., & Pargetter, R. (1986). Oughts, options and actualism. *The Philosophical Review*, 95, 233–255.

²⁸ See Heathwood (2013).

- Kagan, S. (1998). *Normative Ethics*. Boulder: Westview Press.
- Kamm, F. (1998). *Morality, Mortality: Volume I: Death and Whom to Save From It*. Oxford: Oxford University Press.
- Laskowski, N. (2015). Non-analytical naturalism and the nature of normative thought: A reply to parfit. *Journal of Ethics and Social Philosophy*. <http://jesp.org/articles/download/nonanalytical-naturalism.pdf>.
- MacAskill, W. (2016). Normative uncertainty as a voting problem. *Mind*, 125(500), 967–1004.
- McKeown-Green, J., Pettigrove, G., & Webster, A. (2015). Conjuring ethics from words. *Noûs*, 49(1), 71–93.
- McMahan, J. (2016). Torture and method in moral philosophy. In S. Anderson & M. Nussbaum (Eds.), *Torture, Law, and War*. Chicago: University of Chicago Press.
- McPherson, T. (2008). Metaethics and the autonomy of morality. *Philosophers' Imprint*, 8(6), 1–16.
- Mollendorf, D. (2008). Jus ex bello. *Journal of Political Philosophy*, 16(2), 123–136.
- Nakano-Okuno, M. (2011). *Sidgwick and Contemporary Utilitarianism*. Basingstoke: Palgrave Macmillan.
- Parfit, D. (2011). *On What Matters*. Oxford: Oxford University Press.
- Phillips, D. (2011). *Sidgwickian Ethics*. Oxford: Oxford University Press.
- Portmore, D. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford: Oxford University Press.
- Rawls, J. (1972). *A Theory of Justice*. Wotton-under-Edge: Clarendon Press.
- Rawls, J. (1974). The independence of moral theory. *Proceedings and Addresses of the American Philosophical Association*, 48, 5–22.
- Regan, D. (1980). *Utilitarianism and Co-operation*. Oxford: Oxford University Press.
- Ross, W. (2002). *The Right and the Good*. Oxford: Oxford University Press.
- Ross, J. (2006). Rejecting ethical deflationism. *Ethics*, 116(4), 742–768.
- Ross, J. (2013). Actualism, possibilism, and beyond. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics*. Oxford: Oxford University Press.
- Scanlon, T. (1998). *What We Owe to Each Other*. Cambridge: Harvard University Press.
- Schaffer, J. (2009). On what grounds what. In D. Manley, D. J. Chalmers, & R. Wasserman (Eds.), *Metametaphysics: New Essays on the Foundations of Ontology* (pp. 347–383). Oxford: Oxford University Press.
- Schneewind, J. B. (1977). *Sidgwick's Ethics and Victorian Moral Philosophy*. Wotton-under-Edge: Clarendon Press.
- Schroeder, M. (2006). Cudworth and normative explanations. *Journal of Ethics and Social Philosophy*, 1(3), 1–27.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, M. (forthcoming). Normative ethics and metaethics. In T. McPherson and D. Plunkett (Eds.), *Routledge Handbook of Metaethics*.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. New York: Oxford University Press.
- Shaver, R. (2014). Sidgwick's axioms and consequentialism. *Philosophical Review*, 123(2), 173–204.
- Sidgwick, H. (1907). *The Methods of Ethics*. Hackett.
- Smith, M. (2004). *Ethics and the A Priori*. Cambridge: Cambridge University Press.
- Sobel, J. H. (1976). Utilitarianism and past and future mistakes. *Nous*, 10(2), 195–219.
- Tannsjo, T. (2015). *Taking Life: Three Theories on the Ethics of Killing*. Oxford: Oxford University Press.
- Taurek, J. (1977). Should the numbers count? *Philosophy and Public Affairs*, 6(4), 293–316.
- Timmerman, T. (2015). Does scrupulous securitism stand-up to scrutiny? Two problems for moral securitism and how we might fix them. *Philosophical Studies*, 172(6), 1509–1528.
- Timmons, M. (2012). *Moral Theory: An Introduction*. Lanham: Rowman & Littlefield.
- Walden, K. (2014). The aid that leaves something to chance. *Ethics*, 124(2), 231–241.
- Walzer, M. (1992). *Just and Unjust Wars: A Moral Argument with Historical Illustrations*. New York: Basic.
- Zimmerman, M. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge: Cambridge University Press.