

Normative standards and the epistemology of conceptual ethics

Tristram McPherson^a and David Plunkett^b

^aPhilosophy, The Ohio State University, Columbus, OH, USA; ^bPhilosophy, Dartmouth College, Hanover, NH, USA

ABSTRACT


This paper addresses an important but relatively unexplored question about the relationship between conceptual ethics and other philosophical inquiry: how does the epistemology of conceptual ethics relate to the epistemology of other, more “traditional” forms of philosophical inquiry? This paper takes as its foil the optimistic thought that the epistemology of conceptual ethics will be easier and less mysterious than relevant “traditional” philosophical inquiry. We argue against this foil by focusing on the fact that that conceptual ethics is a form of normative inquiry. Because of the epistemic difficulties that face normative inquiry, we should not expect conceptual ethics to constitute an epistemic panacea. Instead, although the epistemological upshots can vary from case to case, there are systematic reasons why this shift may exacerbate, rather than mitigate, the epistemic difficulties we face in pursuing philosophical inquiry.

ARTICLE HISTORY Received 4 April 2022; Accepted 18 September 2022

KEYWORDS Conceptual ethics; conceptual engineering; epistemology; pragmatism; epistemology of the normative

Introduction

In recent years, there has been increasing explicit interest amongst philosophers in the projects of “conceptual ethics” and “conceptual engineering”.¹ Put roughly, conceptual ethics concerns certain normative and evaluative questions about thought and talk, such as questions about which concepts we should use, and why, and what we should mean by our words, and why. In turn, conceptual engineering (again, put roughly) incorporates such

CONTACT Tristram McPherson  dr.tristram@gmail.com  Philosophy, The Ohio State University, Columbus, OH 43210-1132, USA

¹See, for example, the papers collected in (Burgess, Cappelen, and Plunkett 2020).

© 2022 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

normative and evaluative inquiry into projects that also involve introducing or reforming concepts (or other representational or inferential devices) and trying to implement the use of those new or revised concepts.²

A central cluster of questions about conceptual ethics and conceptual engineering concerns how these projects relate to other forms of philosophical inquiry. For example:

- How much of existing philosophical inquiry already involves conceptual ethics and conceptual engineering (perhaps implicitly)?
- Can explicit engagement in conceptual ethics and conceptual engineering help make progress in philosophical inquiry? (And if so, how much?)

This paper addresses an important but relatively unexplored question about the relationship between conceptual ethics and other philosophical inquiry: how does the epistemology of conceptual ethics relate to the epistemology of other, more “traditional” forms of philosophical inquiry? (For our purposes here, we will understand “traditional” philosophical inquiry as inquiry that doesn’t centrally involve conceptual ethics.)³

One way to see the significance of this question is to note that some answers to it could serve to motivate engaging in conceptual ethics. For example, suppose one thinks that the basic idea of conceptual ethics makes sense, and links up in relevant ways to key aspects of “traditional” philosophical inquiry. Now suppose that inquiry into issues in conceptual ethics is easier to carry out successfully and less epistemologically mysterious than “traditional” philosophical inquiry. Put briefly, this epistemological hypothesis could motivate conceptual ethics inquiry insofar as we have more reason to engage in philosophical projects that are more apt for intelligible success than those that are not.

This epistemological hypothesis might also seem quite plausible. Consider, for example, philosophical inquiry concerning free will. The difficulty of such inquiry is suggested by the endemic and seemingly

²These glosses are from McPherson and Plunkett (McPherson and Plunkett 2021d), which in turn draws on (Burgess and Plunkett 2013a), (Burgess and Plunkett 2013b), (Burgess and Plunkett 2020), and (Cappelen and Plunkett 2020).

³We make this assumption solely for the purposes of expository simplicity. In fact, we are sympathetic to the idea that much philosophical inquiry involves work in conceptual ethics, either explicitly or implicitly. For an overview of some of the many places it shows up explicitly, see (Burgess and Plunkett 2013a), (Cappelen and Plunkett 2020), and (Cappelen 2018). For discussion of the idea that significant parts of philosophy might well involve implicit arguments in conceptual ethics, see (Plunkett 2015) and (Thomasson 2016). Once we relax the assumption that “traditional” philosophy inquiry doesn’t centrally involve work in conceptual ethics, we can better formulate our central question in terms of the epistemological relationship between the conceptual ethics and non-conceptual ethics *elements* of philosophical inquiries. We will return to this issue in the conclusion of the paper.

intractable philosophical disagreements about free will. And the epistemological mystery of inquiry about this topic is suggested by the fact that it is not at all clear how we could settle these debates using empirical investigation, and nor is it clear what the alternative epistemology for addressing these debates is supposed to be. It might seem that the project of identifying a good *FREE WILL*-ish concept to use for central purposes in our lives is, comparatively, more epistemically tractable and less mysterious.⁴ This comparative hypothesis will serve as our foil in this paper.

As we explain in the next section, there are good reasons to take this foil seriously. However, against this foil, we argue that shifting to engage in conceptual ethics projects does not *generally* lighten our epistemic burdens. Rather, although the epistemological upshots can vary from case to case, there are systematic reasons why this shift may instead exacerbate the epistemic difficulties we face in pursuing philosophical inquiry. This is true, we argue, both in terms of the question of whether conceptual ethics is “easier” than “traditional” philosophical inquiry and the question of whether it is “less mysterious”.

We proceed as follows. We begin by introducing our question in more detail and motivating our foil (§1). We then set out the core of our argument. This argument begins from the observation that conceptual ethics is a kind of *normative* and *evaluative* inquiry. Because of this, we should expect the epistemology of conceptual ethics to inherit the difficulties of the epistemology of the normative and evaluative more generally (§2). We then consider certain substantive assumptions that would greatly mitigate the epistemological challenges that we press (§3). We argue that, in at least some cases, these same assumptions will entail that “traditional” philosophical inquiry is also more epistemically tractable than it might initially seem. We conclude by drawing some broader lessons for our understanding of the project of conceptual ethics.

1. Preliminaries: the projects and our foil

In this section, we begin by more carefully introducing the projects of conceptual ethics and conceptual engineering. We then offer a more detailed motive for our foil. Finally, we distinguish our strategy in this paper from alternative ways of casting doubt on the epistemic tractability of conceptual ethics.

⁴In this paper, we use single quotation marks (e.g. ‘bicycle’) to mention linguistic items. We use double quotation marks (e.g. “bicycle”) for a variety of tasks including quoting others’ words, scare quotes, and mixes of use and mention. We use small caps (e.g. BICYCLE) to pick out concepts.

The schematic characterization of conceptual ethics and conceptual engineering that we work with in this paper draws on our previous work.⁵ No characterization of these projects is uncontroversial.⁶ In putting forward an account of what these projects are, we are taking a stand on some of these controversies – e.g. how conceptual ethics relates to conceptual engineering. However, our aim here is not to settle significant controversies on how the projects we discuss should best be carried out. Rather, we seek to provide a framework that is useful for talking about a range of connected issues that philosophers are interested in, across a wide range of different subareas of philosophy, and with a wide range of different philosophical commitments.

Start with “conceptual ethics”.⁷ We take conceptual ethics to be a branch of normative and evaluative inquiry, focused on certain kinds of issues about thought and talk. For example, if Mirabai wonders whether it is ethical to routinely classify persons using the concept *WOMAN*, they are engaged in conceptual ethics. This example involves a question about a *concept*: namely, the concept *WOMAN*. But we understand conceptual ethics broadly, to include questions about words, conceptions, or other things that theorists hold play connected inferential and representational roles in our thought and talk. Those involved in conceptual ethics might see those things (concepts, words, conceptions, etc.) in a range of different ways (including accepting or denying the analytic/synthetic distinction, the idea that the conceptual is sharply distinguished from the non-conceptual, or the idea that words and concepts should be understood mostly in terms of a “representational” role). Indeed, one could engage in conceptual ethics even if one is suspicious of the theoretical utility of concepts.⁸

Just as the term ‘conceptual’ in ‘conceptual ethics’ can be misleading (in implying a narrower focus than we intend) so too is this true of the term ‘ethics’ as used in ‘conceptual ethics’. We aim to use ‘ethics’ in ‘conceptual ethics’ as a useful shorthand to pick out a range of possible types of normative and evaluative inquiry, including, but not limited to, the narrowly “ethical”. For example, some theorists think that moral and political

⁵See especially (McPherson and Plunkett 2021d), (McPherson and Plunkett 2021c), (McPherson and Plunkett 2021b), and (McPherson and Plunkett 2021a), which in turn draws on (Burgess and Plunkett 2013a), (Burgess and Plunkett 2013b), (Burgess and Plunkett 2020), and (Cappelen and Plunkett 2020).

⁶For a range of different approaches, see the essays collected in (Burgess, Cappelen, and Plunkett 2020).

⁷Our gloss on conceptual ethics below draws mostly from (Burgess and Plunkett 2013a) and (Burgess and Plunkett 2013b).

⁸For example, Herman Cappelen focuses solely on the meanings of lexical items in his discussion of “conceptual engineering” (which we take to involve “conceptual ethics”). See (Cappelen 2018).

norms and values are central to conceptual ethics (at least when done properly, at least for certain parts of conceptual ethics), while others place more emphasis on broadly “epistemological” ones (such as facilitating fruitful inquiry) or broadly “metaphysical” ones (such as “carving nature at its joints”).⁹ Nothing in our use of ‘ethics’ in ‘conceptual ethics’ is meant to take a stand on such issues.

Now turn to “conceptual engineering”.¹⁰ We take paradigmatic conceptual engineering projects to involve three different activities. The first is conceptual ethics, understood along the lines sketched above. The second is “conceptual innovation”, which involves introducing new or modified concepts (or other representational devices, such as words).¹¹ The third is “conceptual implementation”, which involves efforts to actually have some relevant group of people employ the linguistic or conceptual changes proposed as conceptual innovations.

Notice that, because paradigmatic instances of conceptual engineering also involve conceptual innovation and conceptual implementation, conceptual engineering projects paradigmatically extend beyond work in conceptual ethics. To illustrate: it is no part of a conceptual ethics inquiry per se – even when that inquiry concludes with endorsing certain conceptual reforms or replacements – to attempt to actually *bring about* those changes. Consider, by analogy, the contrast between an ethical theory project that aims to better understand or know about some ethical topic, and the practical project of seeking to make the world better in some relevant respect.

This paper aims to evaluate an epistemological hypothesis that could serve as a motive for engaging in conceptual ethics (whether as part of a conceptual engineering project or not). This hypothesis is that, in a central range of cases, conceptual ethics projects are both easier and less epistemically mysterious than “traditional” (non-conceptual ethics) projects in philosophy. The qualification “in a central range of cases” is essential, because we simply want to grant that some questions that philosophers are interested in can be easily answered in unmysterious ways.

⁹For an example of work in (what we take to be) “conceptual ethics” that gives a central role to broadly “moral” and “political” considerations, see (Haslanger 2000). For work that gives a central role to broadly “epistemological” ones, see (Pérez Carballo 2020). And for work that gives a central role to broadly “metaphysical” ones, see (Sider 2011).

¹⁰Our gloss on “conceptual engineering” draws from (Burgess and Plunkett 2020) and (Cappelen and Plunkett 2020). We are skating over some of the relatively subtle differences between those two accounts, which don’t matter to our discussion here.

¹¹The question of what distinguishes “reforming” from “replacing” a concept (or a word, etc.) is an interesting one for work in conceptual ethics and conceptual engineering. But it is not one that matters for our core line of argument in what follows. So we leave it to the side in our discussion here.

We propose to grant to our foil, however, that in a wide range of central cases, the epistemology of philosophy is *very challenging*, and in a way that goes beyond the paradigmatic kinds of challenges that arise in inquiry in the natural and social sciences. Consider some examples of this sentiment, from the very local to the very general. Karen Bennett argues that metaphysical debates about material constitution and the special composition question cannot be settled by “local” philosophical argument; that if there are reasons to favor one side or the other, they reside in a controversial collection of “broader theoretical and methodological” considerations.¹² In discussing the epistemology of ethics, one of us (McPherson) suggests that there is a lot to be said for the view that “adequately justified ethical belief is possible, but we are in general poorly equipped to get it”.¹³ And more generally still, Kit Fine suggests epistemological modesty about philosophy in claiming that “in this age of post-Moorean modesty, many of us are inclined to doubt that philosophy is in possession of arguments that might genuinely serve to undermine what we ordinarily believe”.¹⁴

Consider two related ways in which this modesty can be motivated. The first is broadly inductive: time and again, we find that careful, charitable, and fair-minded philosophers develop theories that turn out, on further reflection, to fail. The second concerns the broad scope and seeming intractability of many philosophical disagreements. This might seem to suggest that it is at least very challenging to correctly answer philosophical questions, given that many (most?) well-trained, committed, and thoughtful inquirers appear to be getting it wrong on central questions.

A related issue is that much philosophical inquiry can seem deeply *mysterious*. The mystery might be understood in different ways. For example, it might be based in the idea that much philosophy is, in Ted Sider’s Ayer-esque terminology, “epistemically metaphysical”: that is, amenable neither to empirical investigation nor to conceptual analysis.¹⁵ Or it might be based in Richard Rorty’s “pragmatist” diagnosis: that much philosophy involves the “attempt to step outside of our own skin – the traditions, linguistic and other, within which we do our thinking and self-criticism – and compare ourselves with something absolute”.¹⁶

¹²(Bennett 2009, 72–74).

¹³(McPherson 2020, 29). For a more nuanced version of the claim, see (McPherson 2018b).

¹⁴(Fine 2001, 2).

¹⁵(Sider 2011, 187).

In what follows, we will simply grant that (much, central) “traditional” philosophical inquiry is very challenging. We will not, however, grant that it is *impossible*.¹⁷ For if it were, this could make the *comparative* epistemic challenge for conceptual ethics projects rather trivial to meet.

Next, consider why it might seem that a conceptual ethics project is likely to be in better epistemic shape than a traditional, non-conceptual ethics project addressing the same topic. Just for concreteness, we will return to our example of the concept of FREE WILL, mentioned in the introduction. Consider a conceptual ethics argument in favor of a new “engineered” version of our concept FREE WILL, which we then would express with the same term ‘free will’.¹⁸ Why might the inquiry supporting this argument seem to be in better epistemic shape than traditional philosophical inquiry into free will?

First, on a natural picture, traditional philosophical inquiry is about part of extra-conceptual, extra-social-practical *reality* rather than about our language or practices. By contrast, an attempt to engineer a better version of FREE WILL is focused on our thought and talk. One might think that it will in general be easier to investigate conceptual and linguistic entities than it will be to investigate properties like free will.

It might be noted, in rebuttal, that the track record of philosophers’ attempts to produce adequate conceptual analyses casts doubt on this alleged asymmetry.¹⁹ However, here again the conceptual engineer might seem to have the easier task. For she doesn’t need to develop a fully adequate descriptive theory of the *existing* concept FREE WILL. Rather, she can simply know *enough* about it to support the idea that it could be improved along certain salient dimensions. Indeed, she is even free to simply *stipulate* the concept that she would like to engineer. Here it is relevant that even W.V.O. Quine, the most famous critic of “analyticity” (and forms of standard “conceptual analysis” tied to that idea), seemingly grants that this kind of stipulative activity can be clear and straightforward.²⁰

¹⁶(Rorty 1982, xix).

¹⁷For example, this would be true if local skepticism is true about the relevant area of philosophy. Some pragmatists also seem to think that “traditional” philosophical inquiry is impossible, as (e.g.) Rorty suggests about the alleged effort to “step outside of our own skin” (Rorty 1982, xix). Many philosophical projects might also be “impossible” (in a relevant sense) if, as (Scharp 2020) suggests, they are shot through with deeply defective concepts.

¹⁸For an example of an extended “conceptual engineering” approach to free will, which also includes significant amounts of methodological reflection on those projects that we are here calling “conceptual ethics” and “conceptual engineering”, see (Vargas 2013).

¹⁹For a critical discussion of this track record, see (Williamson 2007). For a more optimistic take on the prospects of conceptual analysis given this track record, see (Jackson 1998).

Next consider the alleged “mysteriousness” of the epistemology of traditional philosophy. Here again, the conceptual ethicist might seem to have the advantage, given that conceptual ethics involves critical reflection on our own conceptual and linguistic practices, rather than properties such as *free will* whose place in reality (especially given a broadly naturalistic account of reality) is hard to pin down. It is in this broad spirit that Amie Thomasson argues that conceptual ethics projects – when conducted in a properly “pragmatic” spirit that she advocates for – enable us to avoid “epistemological mystery”.²¹ In essence, she argues that such projects (carried out along specific lines she argues for) retain the epistemic virtues that she has elsewhere advertised for her “deflationary” approach to metaphysical inquiry.²² (Thomasson’s endorsement of this kind of continuity with “traditional” philosophical inquiry, understood in her “deflationary” way, is part of why we don’t take her to endorse the alleged epistemic contrast that is our foil in this paper).

Notice that we have introduced the alleged epistemic contrast along two dimensions: how relatively *easy* our two sorts of inquiry are, and how much they involve epistemic *mystery*. It is worth emphasizing that these issues can come apart. On the one hand, we might sometimes have a method for studying X that we have good reason to think is reliable, but where we don’t really understand much about *why* it is an accurate method to use. To illustrate: compare claims about the “unreasonable effectiveness” of mathematics in the natural sciences.²³ On the other hand, there are questions for which we have no doubt at all about what *would be* unmysterious good evidence, but where we just lack that evidence. This is true, for example, of trivial details about the past that no one happened to record. While we are most interested in the question of relative ease, we take both dimensions to be relevant to our discussion.

We take it that comparative epistemic ease of a kind of inquiry can motivate engaging in that inquiry, in the following way. Suppose two research projects were such that they would be equally valuable if successfully executed. Here is a reason to prefer engaging in the easier of the projects: its being easier makes success more likely, and frees

²⁰Quine writes: “Here [in the case of stipulation] we have a really transparent case of synonymy created by definition; would that all species of synonymy were as intelligible” (Quine 1951, 26).

²¹(Thomasson 2020, 456).

²²See especially (Thomasson 2015).

²³See (Wigner 1960).

resources to do other worthwhile things. Similarly, if there is value in engaging in activities that we can understand, there is a reason to prefer engaging in the less epistemically mysterious of two otherwise worthwhile projects. At the extreme, engaging in an epistemically mysterious project may seem to involve an objectionable act of faith.

Were the apparent comparative epistemological advantages of inquiry into conceptual ethics vindicated, they would also constitute one side of a powerful two-pronged case for philosophers to shift to engaging in conceptual ethics projects. The case goes like this. It seems quite generally plausible that it makes sense to frame our philosophical inquiry around words and concepts that are *good to use*. And, at least prior to engaging in conceptual ethics investigation, one might think that we have no particular reason to believe that the actual words and concepts that frame philosophical discussion are good to use; or at least that they are as good to use *as they could be*.²⁴ This straightforwardly motivates engaging in conceptual ethics projects, at least to the extent of *testing* how good our existing philosophical concepts are, and comparing them to alternatives. If conceptual ethics projects were epistemically vexed, this motive might be defeated. But if (e.g.) conceptual ethics inquiry into *FREE WILL* is in fact typically epistemically *less mysterious* and *easier* than “traditional” investigation of the nature of free will, then this motive would seem to be strengthened.

Having motivated the epistemic case for engaging in conceptual ethics, we now want to mention two important sorts of responses to that case, which we will simply set aside for the purposes of this paper.

The first sort of response attempts to push back against the alleged epistemic mystery or difficulty of “traditional” philosophical inquiry (i.e. inquiry that doesn’t centrally involve conceptual ethics). Consider three examples. First, many strands of research across many areas of philosophy – from philosophy of language to ethics to metaphysics – involve commitments to forms of “naturalism” that, in different ways, seek in large part to render the relevant inquiry epistemically tractable and unmysterious. Some important broadly “pragmatist” projects, such as Thomasson’s work on “easy ontology”, can be fruitfully understood along these lines.²⁵ So too can work from a range of others, such as, for example, Peter Railton, Frank Jackson, and Steven Stich, whose work doesn’t involve as many explicitly “pragmatist” themes.²⁶

²⁴For connected discussion, see (Cappelen 2020) and (McPherson and Plunkett 2020).

²⁵See (Thomasson 2015).

²⁶See (Railton 2003), (Jackson 1998), and (Stich 2011).

A second route to resisting the idea that traditional philosophy involves epistemic mystery appeals to a kind of broadly “anti-representationalist” form of “pragmatism”, along the lines advanced by Rorty. In short, if one thought of philosophy (roughly) in terms of an ongoing cultural conversation that didn’t aim to correctly represent mind-independent reality, this might arguably undercut the seeming epistemic difficulty of “traditional” philosophical inquiry (i.e. inquiry that doesn’t centrally involve conceptual ethics).²⁷ To take one last example, if one thought (following the lead of the later Ludwig Wittgenstein) that much “traditional” philosophical inquiry consisted of pseudo-problems that could be dissolved by paying close attention to our linguistic or social practices, this would undercut a presupposition of the alleged asymmetry: namely, that there is a legitimate object of study for “traditional” philosophical inquiry.²⁸

The second sort of response is to question how well we in fact understand how to successfully engage in conceptual engineering projects. If, as Herman Cappelen suggests, we lack a good understanding of how to make shifts in the meanings of terms (given the complexity of metasemantics), we may not be in a good position to know what it would take to even achieve success in a conceptual engineering project.²⁹ This kind of idea perhaps most directly raises a challenge for conceptual implementation, but could also be developed to raise a challenge for conceptual ethics.

We take these to be important responses, and we have some sympathy for each of them. However, in this paper, we propose to set them aside, to focus on a different cluster of epistemological challenges. The challenges we focus on arise from the fact that conceptual ethics (whether done as part of conceptual engineering, or not) is a branch of *normative* and *evaluative* inquiry.

Let’s take stock. In this section, we have introduced the projects of conceptual ethics and conceptual engineering. In order to introduce our foil, we then made a preliminary case for two sorts of epistemological asymmetries between conceptual ethics inquiry (whether done on its own, or as part of a conceptual engineering project) and “traditional” philosophical inquiry. The first asymmetry is that conceptual ethics inquiry appears

²⁷See (Rorty 1980). Note that we say “arguably” here with a nod to the kind of possibility that we discuss later in this paper: namely, that certain anti-realist views don’t actually sidestep many of the relevant epistemological difficulties in the areas they are adopted in (e.g. metaethics).

²⁸See (Wittgenstein 1991 [1953]).

²⁹(Cappelen 2018, Part II).

to be *easier* than relevant central instances of traditional philosophical inquiry. The second is that it appears to be *less mysterious* than such traditional philosophical inquiry. In the next section we explain why this foil, despite its apparent plausibility, should be rejected.

2. The epistemology of conceptual ethics: a cluster of challenges

In this section, we lay out what we take to be the structure of a central cluster of epistemological challenges to conceptual ethics, tied to its nature as a kind of normative and evaluative inquiry. (Henceforth, for ease of exposition, we will often use ‘normative’ in a broad way to cover both the normative and the evaluative). In this section, we argue that, properly understood, normative inquiry involves at least three distinct epistemological tasks, each of which can be quite challenging. We introduce and explain these challenges in turn.

2.1. The normative standards question

There are a wide variety of normative standards. For example, we could evaluate conceptual ethics proposals using *moral* norms, or *epistemic* norms, or the norms scrawled on the wall of little Anna’s treehouse. However, this arguably understates the variety. Suppose that we think of a normative standard as a function from “ought” claims to truth values. The dominant contemporary semantics for ‘ought’ is context-sensitive.³⁰ This means that features of the context of use can alter the function to truth-values in many ways.

Given this, what sort of normative standard does it make sense to use in conceptual ethics? Intuitively, it is not the case that anything goes here. There are, presumably, standards that support engineering the term ‘electron’ (in the context of doing physics research) to refer to breath mints. But it seems absurd to advocate this sort of conceptual engineering proposal. Here is a very natural thought: in doing conceptual ethics, we should deploy normative standards that *really matter* or ones that are *really good to use*. This might mean using what (in other work) we have called “authoritative” normative standards. These are ones that (put roughly) pick out what normative facts about what we “really and truly” should do, or evaluative facts about what “really and truly” matters.³¹

³⁰See (Kratzer 2012). For connected discussion, see (Finlay 2014) and (Silk 2016).

Alternatively, the “authoritative” normative standards might tell us that we should use *different* standards in a given context (even if the latter standards are not themselves authoritative).³² To illustrate, consider that it might be that someone (authoritatively) should use certain culinary standards when engaged in cooking, or certain legal standards when working as a judge. Notice that, in putting forward this broad line about the kinds of standards we should use, and how they either are or are vindicated by “authoritative” ones, we are *evaluating* normative standards themselves.

On either of these ways of understanding the role of authoritatively normative standards, conceptual ethics work would ideally involve identifying (and deploying) authoritatively normative standards. But (to simplify brutally) seeking such identification seems to require that we *evaluate* normative standards. And this raises an important challenge: how are we to *evaluate* our own normative standards?³³ It would seem bizarre or silly to do so by deploying norms that we do not currently accept. But suppose that we deploy some of our own standards to do the evaluation. Then there is a natural worry that our evaluation is objectionably circular: that it would be like trying to verify the accuracy of a ruler by checking it against itself.

This is already a substantial epistemological challenge. However, given certain further assumptions, things get even more vexed. When we imagine worrying about the accuracy of our ruler, we have an idea of an *external standard* in mind, against which it might turn out to be accurate or not (perhaps, the standard length unit conventions prevalent in one’s society). One way to understand such an “external” standard for normative inquiry is in terms of a broadly “realist” metanormative theory. But, as we discuss in §3 with respect to other epistemological challenges, this epistemological challenge can also arise for broadly “anti-realist” views, insofar as they (as many of them do) seek in some way to accommodate or replicate the idea of “external” standards.³⁴ By contrast, consider the worry about whether the normative standards that satisfy our concept REALLY MATTERS match a relevant external standard. It is unclear how we can even intelligibly think about that external standard, understood as distinct from the standards encoded in any of our own normative and

³¹For discussion, see (McPherson 2018a), (McPherson and Plunkett 2017), and (McPherson and Plunkett 2021b).

³²For further discussion of this idea, see (McPherson and Plunkett 2020).

³³The following discussion briefly introduces issues explored in more depth in (McPherson and Plunkett 2021b).

³⁴For further discussion, see (McPherson and Plunkett 2021b).

evaluative concepts. There can seem to be a challenge here that is at once deep and unnerving, while at the same time appearing *ineffable*.³⁵

In recent work, we have argued for cautious optimism in addressing the “vindictory circularity” challenge just glossed, using resources from anti-skeptical epistemology. In particular, we express sympathy for the idea we are entitled to a form of “self-trust” in the concepts we deploy in doing conceptual ethics, in the same vein as the broader anti-skeptical idea that we are entitled to basic trust in our belief-forming capacities.³⁶ In work in progress, however, we suggest that certain sorts of information we could receive about cognitively superior reasoners (e.g. certain forms of “strong AI”) could make the epistemological challenge here even more intense.³⁷ For our purposes here, the core point is that the “vindictory circularity” challenge is a distinctive epistemological challenge about the foundations of conceptual ethics that remains relatively unexplored. Provisionally, however, this challenge appears both powerfully motivated and quite *difficult*.

2.2. The epistemology of familiar normative concepts

Suppose that the “vindictory circularity” challenge we sketched above can be met. In particular, suppose that we meet it in such a way that the normative concepts that it makes sense to deploy in evaluating conceptual ethics proposals are familiar concepts that we already employ in normative inquiry, like JUST, REASON FOR ACTION, OR MORALLY OUGHT. Here we confront a second layer of epistemological challenges for the conceptual ethicist. This arises from the fact that the epistemology of the normative is paradigmatically vexed. We can illustrate this in several ways.

To begin, imagine that your favored news source proclaimed that moral theorists had discovered that abortion is morally permissible in the first trimester. Whatever your views about the ethics of abortion, you would likely take the headline as a sign of journalistic incompetence. We do not, ordinarily, take morality to be a topic of reputedly reportable fact.³⁸

Next, consider the familiar fact that there appears to be a great deal of *disagreement* about central normative questions, both among everyday people (“the folk”), and among “expert” inquirers about normative

³⁵For further discussion, see (McPherson and Plunkett 2021b), drawing on discussion in (Eklund 2017).

³⁶(McPherson and Plunkett 2021b).

³⁷(McPherson and Plunkett Manuscript-b).

³⁸Here we follow (McPherson 2020, 6).

topics, such as those that arise in ethics, epistemology, the philosophy of science, political philosophy, and aesthetics.³⁹

Then there are a host of more theoretical and controversial worries about the epistemology of the normative. These include worries that our central normative concepts are infected with ideology, or are generated by problematic psychological processes, or that we cannot explain in a plausible way how our normative beliefs reliably track the normative facts.⁴⁰

When we put these observations together, the epistemology of the normative does not, at first blush, appear easy. In fact, it appears to have some of the same epistemological features of “traditional” philosophical inquiry, which our foil supposed could be used to motivate shifting focus to a conceptual ethics project.⁴¹

2.3. *The epistemology of the normatively approved standard*

Suppose that we have navigated the two classes of epistemological challenges posed in the preceding subsections. We are still not quite out of the epistemological woods, so to speak. To see this, suppose that, in engineering a semantic content for the word ‘free will’, one determines that, in fact, part of what *really matters most* in evaluating potential contents for this term is how well the content *carves nature at its joints*, such that the content picks out a Lewisian “elite” or “highly natural” property.⁴² Such properties are arguably, in Sider’s sense mentioned in §1 above, “epistemically metaphysical”; that is, we arguably cannot ensure that we latch onto them through solely a combination of empirical inquiry and conceptual analysis. If this happened, the epistemically motivated conceptual ethicist would have in effect weathered the epistemological challenges canvassed in the preceding subsections, only to find themselves facing just the sort of epistemic mysteriousness challenge they had hoped to avoid in the first place.

Even if one denied that our goal was to use concepts that “carve at the joints” (where those joints are understood in terms of Lewisian ideas about “naturalness”), a similar kind of “revenge” problem might arise in

³⁹For discussion, see e.g. (McGrath 2011) and (Locke 2017).

⁴⁰For some of the recent discussion of these kinds of issues, see (Street 2006), (Joyce 2006), (Schechter 2017), and (Vavova 2018).

⁴¹Note that different metanormative theories have different implications for the epistemology of the normative, including ones that matter for how difficult it is and what problems it faces. For discussion of different dimensions of these implications, see (McPherson 2012), (Darwall 1998), and (Street 2006).

⁴²For discussion, see (Sider 2011), drawing on (Lewis 1983). See also (Dorr and Hawthorne 2013).

other ways. To illustrate, consider the normative question of which concept we should express by the term ‘torture’ in a given context. If one is faced with two rival definitions – one of which has a higher threshold for what counts as “torture” than the other – what kinds of normative arguments should one appeal to when determining which definition we should use? In many cases, it seems natural to think that the relevant normative arguments will be (at least) closely related to the arguments that ethical and political philosophers currently use to support substantive conclusions about the nature of torture, and the nature of the moral wrong it exemplifies. For example, conceptual ethics arguments about the word ‘torture’ (or which range of “torture-ish” concepts to employ) should seemingly be sensitive to questions about what rights individuals have, and what sorts of protections for individuals should be enshrined by our social and political institutions. Suppose, then, that conceptual ethics inquiry about how to use ‘torture’ involves sorting through very similar normative and evaluative issues as “traditional” philosophical arguments in moral and political philosophy about torture. Then it is very unclear whether shifting to a conceptual ethics argument about ‘torture’ is really going to allow us to sidestep the core epistemological difficulties allegedly involved in the “traditional” arguments.⁴³

Now, we shouldn’t overstate these possibilities. For example, in many contexts, it may be true that *what really matters most* for evaluating the contents that we might pair with a representational vehicle includes some degree of *epistemic tractability*. This is especially true if we are engineering concepts to play a central role in certain sorts of inquiry. It would be very odd if the best content for that sort of job had an extension that we were deeply hopeless at tracking. But we should also be careful not to overstate the significance of this point. Recall the joke about the man searching for his lost keys under a streetlamp. The punchline is that the man is searching there *not because that’s where he lost the keys* but because *the light is better there* than where he lost them. Just as light is only useful to the searcher if it shines where his lost keys are, epistemic tractability will only be important to the extent that the property or pattern that is epistemically tractable is one that it is worth investigating, tracking, and attending to in our conceptual ethics project.

⁴³Our points here draw on connected points in (Plunkett and Sundell 2013), (Plunkett 2015), and (Plunkett and Sundell 2021).

2.4. Taking stock

Our foil in this paper has been the idea that there might be a quite general *epistemological* motivation for engaging in conceptual ethics. The idea was that conceptual ethics seems, at first blush, like it might be more epistemically tractable than “traditional” philosophical inquiry on the relevant topics. In this section, we have argued that this appearance is difficult to sustain. Difficult epistemological questions arise at the levels of determining which normative standards to deploy in our conceptual ethics projects, and at the level of how to determine what those standards require. And even if we can navigate those questions, we have just seen that there is no assurance that those standards instruct us to investigate something that will be especially epistemically tractable. It should be noted that this doesn’t mean that to make progress in all questions in conceptual ethics one needs to explicitly address all of these challenges, any more than it means that doing work in other areas of normative inquiry (e.g. ethics or epistemology) always requires one to explicitly engage with parallel challenges. For example, it might well be that we can significant progress on certain “applied” questions in conceptual ethics, e.g. about the relative merit of using a given concept in a given context, without explicitly engaging with these challenges. Rather, the point is that these challenges are evidence of the underlying general difficulties of conceptual ethics and other forms of normative inquiry.

In our discussion, we have mostly focused on the issue of how relatively easy it is to learn about the relevant normative facts in conceptual ethics. But the same basic strategy we have pursued here also extends to the second issue that we discussed in §1: worries about the “mysteriousness” of certain philosophical epistemology. Consider two examples. First, recall the potential “ineffability” of conceptual ethics inquiry into which normative standards are the ones that “really matter”. This is arguably a paradigm of epistemic mystery: how are we to adequately *understand* a kind of inquiry that we cannot even adequately *describe*? Second, note that, even setting this aside, normative inquiry is one of the areas that *many* philosophers take to be epistemically mysterious – especially on certain metanormative theories.

To illustrate this second point, suppose that some form of non-naturalistic metanormative realism is true about the “authoritatively” normative facts.⁴⁴ Put roughly, non-naturalistic realists endorse a kind of

⁴⁴For some recent examples of this kind of view, see (Enoch 2011), (Shafer-Landau 2003), and (Fitzpatrick 2008). This kind of view has its roots in (Moore 1993 [1903]).

“cognitivism” at the level of thought, according to which we form straightforward truth-apt beliefs about normative facts, and a kind of “descriptivism” at the level of talk, according to which we make straightforward claims about those facts in language, via expressing those beliefs. The non-naturalistic realist then combines these ideas with the thought there are some normative facts and that these facts are “of their own kind”, metaphysically speaking. Importantly, this means that these facts are fundamentally different in kind from (in some relevant sense of “different in kind from”) *all* other kinds of facts, including, crucially, the kinds of “naturalistic” facts we study in the natural and social sciences. This is standardly taken by non-naturalists to mean that (among other things) such normative facts play no role in the causal order, and that such facts are not fundamentally constitutively dependent on facts about our attitudes or activities.⁴⁵ If this kind of non-naturalistic metanormative realism is true, it’s far from clear how we learn about normative facts. At the very least, it is far from clear how the epistemological tools that we use in the natural and social sciences for studying *naturalistic* facts would carry over to discovering these (purportedly) radically different kinds of facts.

Nothing in what we have argued in this paper suggests that non-naturalistic realism is true. In fact, both of us are doubtful that it is. But, at the same time, we each give it some credence and think there are strong arguments on its behalf. More importantly, it is certainly one of the “live options” in contemporary metanormative theorizing. As such, it is a “live option” that could turn out to be true about the relevant kinds of normative facts in conceptual ethics. And if that was the result, then it is hardly as if the epistemological foundations of conceptual ethics would be un-mysterious. Indeed, it might instead turn out that they would be significantly *more* mysterious than many parts of philosophical inquiry that aren’t fundamentally about “authoritatively” normative issues. This illustrates a general point: the correct metanormative theory might well yield the result that the foundations of conceptual ethics are highly epistemologically mysterious. Nothing in the very idea of “conceptual ethics” rules out this possibility. And, indeed, we think that focusing on the fact that conceptual ethics is a branch of normative and evaluative inquiry makes it a philosophically salient possibility.

⁴⁵For a more careful discussion of how to understand the non-naturalist’s distinctive metaphysical commitments, see (McPherson and Plunkett [Forthcoming](#)), drawing on (McPherson [2015b](#)).

3. Can we save the motivation?

In this section, we consider four broad strategies for mitigating the epistemic challenges that we have posed for conceptual ethics projects. These involve, respectively, embracing an epistemically tractable goal or normative standard for conceptual ethics work, embracing a *metanormative* theory that has epistemically helpful implications, rejecting the idea that authoritative norms guide conceptual ethics, and embracing background theories about thought and talk that lessen the epistemic burdens. We suggest that while we can learn important lessons from considering these strategies, none of them unproblematically vindicates the asymmetry.

To begin, consider a natural suggestion: that the epistemic challenges can be solved if we simply accept a certain clear normative theory for the purposes of evaluating conceptual ethics proposals. For example, consider the widespread idea that facts about *what promotes our goals or purposes* help explain the norms to apply in conceptual ethics.⁴⁶ Drawing on this, one might be tempted by the more ambitious “instrumentalist” idea that we should evaluate conceptual ethics proposals *solely* by appeal to facts about what best promotes our goals and purposes. We want to make several linked points about this proposal.

First, this sort of response is compatible with the possibility that our current goals make reference to some epistemically challenging standard. For example, to return to the example discussed in §2.3, perhaps with certain of our scientific vocabulary, our goal is to carve “nature at its joints”.

Second, we should not simply assume that our current goals are normatively ideal. Individually, most of us have discovered at some points in our lives that central guiding goals that we have held (either individually or as parts of larger groups) were confused or substantively objectionable. Interpersonally, we should be familiar with the possibility that others do not share our goals, and that we disagree with the goals they have. It is thus again unappealing to stop short at the mere possession of a goal; we will want to be able to explain how our goal is *substantively attractive*. In light of these points, it is at least *prima facie* unappealing to simply treat our goals *as such* as determining the success conditions of conceptual ethics projects. It might of course be that some purely goal-based view

⁴⁶See (Burgess and Plunkett 2013a) and (Burgess and Plunkett 2013b) for discussion. For some examples of philosophers who appeal centrally to goals in doing (at least some key parts) of conceptual ethics, see (Haslanger 2000) and (Thomasson 2020).

here could be correct in conceptual ethics. But the key point is that such a purely “instrumentalist” view of conceptual ethics should be just as controversial as a purely instrumentalist view in other sorts of normative inquiry, such as ethics.⁴⁷

We might seek to address this problem by proposing a non-instrumentalist norm to guide our conceptual ethics proposal. This non-instrumentalist proposal could take a range of different forms, from those that focus on the operative norms in a social context to those that are context-independent. For example, we might propose to evaluate conceptual ethics proposals by a simple utilitarian standard, according to which a conceptual ethics proposal is right just in case no other available proposal would produce more net pleasure. This solves two problems. First, if this were the correct standard, then one could use utilitarian reasons in seeking to convince others to accept our conceptual ethics proposal, regardless of the aims they have. And second, this sort of proposal might seem to make the epistemology of ethics more tractable: it is just a matter of investigating the consequences of various options. That obviously involves complicated empirical reasoning, but not of a kind that seems either totally intractable or fundamentally mysterious.⁴⁸

Despite these apparent virtues, this strategy badly misses the point, as a response to the relevant epistemic challenge. To see this, imagine, analogously, proposing to “solve” the epistemic problems in ethics by embracing act-utilitarianism as a theory of ethically right action. This is misguided, because the central epistemic difficulties in ethics arise *exactly* in trying to determine whether act-utilitarianism is correct, and if not, which of its competitors might be. If act-utilitarianism is the correct theory about ethical rightness, that is something we need to *discover* rather than choose, at least on most plausible metaethical theories.⁴⁹ (This point generalizes to substantive theories of other normative topics, such as goodness, moral obligations, and authoritative reasons).

⁴⁷We here echo our discussion in (McPherson and Plunkett 2020, 295), drawing on discussion in (Burgess and Plunkett 2013b, 1105). Note that, in saying it should be controversial, we aren’t here committed to denying that instrumentalism is correct. For arguments on its behalf, see (Schroeder 2007) and (Street 2006).

⁴⁸Although see (Lenman 2000) for reasons to be skeptical about how epistemically tractable this kind of investigation really is.

⁴⁹Note that this is true even on most contemporary metaethical theories that give pride of place to our contingent attitudes (such as, for example, (Lewis 1989), (Railton 1986), (Schroeder 2007), and (Street 2006)), insofar as those relevant attitudes aren’t directly chosen by us. The same point applies to most contemporary “quietist” metaethical views (such as (Dworkin 1996) and (Scanlon 2014)), which, put roughly, claim that most apparently metaethical claims can only be understood as further internal normative claims.

A second strategy for seeking to avoid the epistemic difficulties we have raised is to embrace a *metanormative* theory, which has as a consequence that the epistemology of the relevant normative standard is reasonably epistemically tractable. One principled reason for taking this seriously is that there are significant reasons to think that different metanormative theories can have strikingly different implications for the epistemology of the normative.⁵⁰ For example, in §2.4, above, we saw that non-naturalistic realism might make the epistemology of the normative appear more mysterious than it would on some other metanormative theories. The same point applies to how easy it would be to learn about the relevant normative facts according to different theories. Given that, one might think the way to avoid the epistemic difficulties we have been raising about conceptual ethics is simple: just pick the right kind of metanormative theory that avoids these issues.

Our first point in reply to this strategy is very similar to our reply to the previous strategy. Just as we don't get to simply choose whether act-utilitarianism is the correct substantive ethical theory, so too we don't get to simply choose a metanormative theory. Rather, again, the correct metanormative theory is something we need to discover, rather than choose: in this case, using metanormative inquiry. And metanormative inquiry is, again, not exactly a paradigm of epistemic ease.

We can drive home the complexities here by pointing to an important pattern in contemporary metanormative inquiry. Often, the simplest versions of certain sorts of metanormative theory (such as subjectivism, expressivism, or naturalistic realism) appear to imply a highly tractable epistemology of the normative.⁵¹ However, when we turn to the more plausible versions of these views that are currently prominent, their epistemic consequences become much more complicated. Consider two brief examples.⁵²

First, many contemporary "subjectivist" metanormative theories – ones that, put broadly, explain the normative facts in terms of the attitudes we have (either the person making a normative judgment, or the person whose judgments are about) – include some sort of "idealizing" function on our actual attitudes.⁵³ At the limit, in the influential account offered

⁵⁰For discussion, see (McPherson 2012) and (McPherson 2018b).

⁵¹In the case of the simplest form of anti-realist expressivism (such as (Ayer 1952 [1936])), it is not even clear that epistemological questions arise.

⁵²For a discussion of the epistemology of one sophisticated version of naturalistic realism, see (McPherson 2018b).

⁵³For example, see (Railton 1986), (Lewis 1989), and (Street 2012).

by Michael Smith, the idealizing function appears to simply screen off the idiosyncratic attitudes of individuals from grounding the correct normative principles.⁵⁴ And this raises difficult questions about how exactly we can epistemically access the relevant idealized contents.⁵⁵ Importantly, not all contemporary “subjectivist” theories involve extensive idealizing functions on our actual attitudes. But those that don’t, such as Mark Schroeder’s form of the Humean theory of reasons, often involve other complications (such as issues about how different reasons should be “weighed” against each other) that make the overall epistemology of the normative significantly more complicated than one might initially expect.⁵⁶

Second, consider expressivist theories, on which, roughly, normative judgments ultimately consist of “non-cognitive” attitudes (such as desires, plans, etc.) that we express in making normative statements. The first thing to note here is that most expressivists (of many different stripes) want to insist that expressivism should be sharply distinguished from subjectivism. The basic reason for this is not hard to see: the (purported) fact that *what it is* to make a normative judgment involves expressing a non-cognitive attitude does *not* entail that the correctness-conditions for those judgments are whether or not the speaker has those attitudes or not.⁵⁷ If this thought is right (which we think it is), this means that expressivists can’t rely on the epistemology one gets from subjectivism, whether in “straightforward” or more complicated forms of it.

Things get even more complicated for expressivism when we look closer at the details of contemporary expressivist views. Many contemporary expressivists embrace a “quasi-realist” program, according to which expressivism is compatible with felicitous talk of normative “truths”, “facts”, and “epistemically justified normative beliefs”.⁵⁸ However, the

⁵⁴(Smith 1994).

⁵⁵In (Smith 1994), Smith argues that we can simply use the “method of reflective equilibrium” here. But it is not at all clear why, on Smith’s metaethical account, someone should start her normative inquiry with her contingent and potentially idiosyncratic normative opinions, as his gloss on that method suggests. In later work (such as (Smith 2012) and (Smith 2013)) Smith has begun to develop important constitutivist forms of argument that seem a better match for the background metaethics. For general critical discussion about how much guidance we can get from the “method of reflective equilibrium” for the epistemology of the normative, see (McPherson 2015a).

⁵⁶See (Schroeder 2007). For discussion of a relevant foil of a kind of “simple subjectivism”, see (McPherson and Plunkett 2017).

⁵⁷For further discussion of this point, see (Gibbard 2003) and (Schroeder 2014).

⁵⁸For example, see (Blackburn 1993) and (Gibbard 2003). For critical discussion, see (McPherson 2022). Note that the relevant part of the “quasi-realist” program we are talking about here is the part that Sebastian Köhler discusses as “accommodationist” expressivism in (Köhler 2021).

contemporary literature shows that once we develop expressivism in these ways, it is no longer clear that the expressivist escapes the sorts of epistemic challenges that face their realist foils.

Consider two examples. First, some philosophers have argued that quasi-realist expressivists face the kind of “reliability” challenge that is one of the largest epistemological challenges to non-naturalistic realism.⁵⁹ It might well be that expressivists have a good response to this challenge, or that they can sidestep it in some way.⁶⁰ The point is just that this issue is much more complicated than one might initially think for contemporary forms of expressivism, especially given just how much of the core claims of non-naturalistic realism expressivists want to embrace in some form. Second, consider Allan Gibbard’s views in *Thinking How to Live*, where he suggests that the epistemology of the normative, on his kind of expressivism, involves a kind of intuitionism similar in key ways to the epistemology that G.E. Moore embraced in defending non-naturalistic realism.⁶¹ While some form of intuitionism might well be the correct epistemology of the normative, many have taken it to be a paradigm of epistemic mysteriousness.

We now turn to a third strategy. On certain substantive background views about language, thought, and normativity, the sorts of challenges that we have been posing for the epistemology of conceptual ethics might well turn out not to be that deep.⁶² One might appeal to such theories in seeking to reject our arguments in this paper. We make two observations about this strategy. First, to defend this strategy, one would need to take on substantive philosophical views in a range of areas of philosophy. If “traditional” philosophical inquiry is epistemically vexed, it will be hard to justify such views. After all, to defend those views, it (at least *prima facie*) seems that one would need to engage in that kind of inquiry about the topics at hand (e.g. issues in the philosophy of mind and philosophy of language). Second, many philosophical views that would have this result will also entail that the epistemology of “traditional” philosophical inquiry is easier than it initially appears as well. For example, suppose that one is drawn to certain kinds of “pragmatist” views that make inquiry in conceptual ethics non-mysterious and tractable.⁶³ The same kinds of considerations that support that kind of view

⁵⁹See (Street 2011) and (Schechter Manuscript).

⁶⁰For arguments that this is so, see (Gibbard 2011) and (Dreier 2012).

⁶¹(Gibbard 2003), discussing the kind of “intuitionist” view of epistemology of the normative found in (Moore 1993 [1903]).

⁶²This might be true, for example, given the commitments argued for in (Jackson 1998) and (Thomasson 2015).

⁶³For example, such as the view argued for in (Thomasson 2020).

in conceptual ethics might well also support views about what is going on in (appropriately conducted) “traditional” philosophical inquiry that makes it relatively non-mysterious and tractable as well.⁶⁴ For all that we have said in this paper, both such views might turn out to be right. But notice that if that turned out to be right, it would undermine the alleged epistemic *asymmetry* between conceptual ethics and traditional philosophical inquiry that is our foil.

A fourth strategy is to abandon the idea we introduced earlier in this paper about the kinds of norms that one should use in doing conceptual ethics. We have suggested that it is plausible that conceptual ethics should be guided either by “authoritative” norms, or by standards endorsed by authoritative norms. We think this is a compelling idea.⁶⁵ However, perhaps in order to sidestep many of the epistemic difficulties for conceptual ethics that we have highlighted in this paper, one might be tempted to reject it.

Based on this, suppose that one decided to engage in conceptual ethics using norms that one came up with oneself, which were designed to be epistemically easy to apply, and which weren’t meant to in fact be authoritatively normative ones or supported by such norms. (For example, appeals to facts about our contingent aims in a given context might play a key role here, or facts about social norms that are operative in a given context). This might well mitigate the sorts of epistemological difficulties we highlighted in the previous section.⁶⁶

The core trouble with this strategy is that, put bluntly, it threatens to make conceptual ethics less interesting and relevant to our underlying philosophical interests. By comparison: one could make ethics or political philosophy easier by just giving up on the idea of discussing issues about how we “really and truly” should live, or how we “really and truly” should organize our social/political institutions, in favor of some allegedly more tractable goal. But the farther we move away from inquiring about the authoritatively normative facts, the more it will seem that we are failing to answer at least some of the core normative questions that we wanted to ask. The same is true in conceptual ethics. Whether there are authoritatively normative facts about which concepts we should use (or about other topics in conceptual ethics) is something for metanormative

⁶⁴For example, such as the view argued for in (Thomasson 2015).

⁶⁵For further sympathetic discussion of this idea, see (McPherson and Plunkett 2020) and (McPherson and Plunkett 2021b).

⁶⁶It also might not, if our ability to understand our own normative or evaluative commitments is limited. Such limits might be suggested by the wide variety of different systematic views on offer in the contemporary metanormative literature, and the large amounts of disagreement in this area.

inquiry to investigate – and thus, as we emphasized earlier, not something we get to choose. If there are such facts, then we should want some part of conceptual ethics to be trying to investigate them. For those drawn to a vision of conceptual ethics in that spirit, the epistemological issues we have raised will loom large.

Conclusion

In this paper, we have investigated a cluster of issues about the epistemology of conceptual ethics. We started by considering a foil: the idea that conceptual ethics projects will tend to be epistemically easier and less mysterious than relevantly related “traditional” philosophical inquiry. We have argued that, taken as a general thesis, this idea is deeply implausible. Conceptual ethics is a form of *normative* inquiry, and insofar as we take seriously the epistemic difficulties that face normative inquiry, we should not expect conceptual ethics to constitute an epistemic panacea. Indeed, given the difficulties in the epistemology of the kinds of normativity that (at least seem to) matter to conceptual ethics, there may be significant reasons to worry that a shift to conceptual ethics *exacerbates*, rather than *mitigates*, our epistemic burdens. In this final section, we emphasize two limitations to our conclusions, and one complication, before sketching the broader significance of the project of this paper.

To begin, we emphasize two ways in which our conclusions should be understood to be limited. First, as we have noted just above, there are possible views about language, thought, and normativity, on which the sorts of challenges that we have been posing for the epistemology of conceptual ethics dissipate. On some of these views – perhaps, in particular, views on which the epistemology of the normative is *distinctively* unproblematic – the contrast that we take as our foil might in fact be vindicated. While this paper has offered some brief reasons to be doubtful about the promise of some such views, we have certainly not decisively refuted them. Hence, the view that is our foil could potentially turn out to be correct.

Second, we intend our conclusions to be understood as *generally* true in a way that is amenable to important exceptions. Partly this is because, while we take the epistemology of the normative to be challenging, we do not take it to be *impossible*. And in particular, we are sympathetic to the idea that there are *some* normative questions where we are in a comparatively strong epistemic position.⁶⁷ Consider one example: the

conceptual ethics conclusion that we should *abandon* the use of certain pejorative terms, in certain contexts.⁶⁸ We think that the substantive case for the pernicious character and effects of some such terms (e.g. certain slur terms) is clear enough that the epistemology involved in evaluating this sort of conclusion is, comparatively, reasonably easy.

Now turn to the complication. For expository simplicity, this paper has been organized around the assumption that there is a clear methodological *contrast* between conceptual ethics projects, on the one hand, and “traditional” philosophical projects on the other. But it might well be that much familiar philosophical work is perhaps best understood as involving significant amounts of conceptual ethics claims or arguments, even if only implicitly.⁶⁹ At least initially, this might seem to undermine the interest of our comparative question in this paper. However, we think it does not.

To see this, return to our earlier example of inquiry concerning free will and the concept *FREE WILL*. Suppose that much “traditional” inquiry about free will turns out to involve some degree of conceptual ethics argument about *FREE WILL* or related concepts (or: words, etc.). This is compatible with the thesis that we *could* inquire about free will in a way that carefully eschews conceptual ethics inquiry. We can then ask about the relative epistemic merits of *this* sort of inquiry relative to conceptual ethics inquiry about *FREE WILL*. With this in hand, we can then also contrast philosophical inquiry that doesn’t involve conceptual ethics with inquiry that involves both conceptual ethics inquiry *and* the sorts of inquiry involved in such philosophical inquiry. This is important to keep in mind given that the idea of conceptual ethics as such is fully compatible with the idea that many of the questions that philosophers have been standardly interested in – including, for example, issues in “heavyweight” metaphysics about such issues as ground, essence, and real definition – are well-formulated and important questions worth asking.⁷⁰ If that is right, then those questions don’t go away once we start doing conceptual ethics. Rather, issues in conceptual ethics can be understood as additional ones we can take on, which might then in turn help us with better asking those other questions, or which might be worth asking in their own right.⁷¹

⁶⁷This is a point emphasized in more detail in (McPherson 2018b).

⁶⁸For more on “abandonment” as an option in conceptual ethics, see (Burgess and Plunkett 2013b), (Cappelen *Manuscript*), and (McPherson and Plunkett *Manuscript-a*).

⁶⁹For further discussion of different versions of this idea, see (Cappelen and Plunkett 2020), (Plunkett 2015), (Thomasson 2016), (McPherson and Plunkett 2021a), and (McPherson and Plunkett 2021c).

⁷⁰For connected discussion, see (Plunkett 2015).

We conclude with two broader points about the significance of our project in this paper. First, it is no part of our project in this paper to argue *against* engaging in conceptual ethics. Even if work in conceptual ethics (usually) cannot be motivated on the kinds of comparative epistemic grounds we have discussed here, that does not mean that it cannot be motivated. We are sympathetic to a number of such motivations. Here are four in particular that we think are worth emphasizing. First, there is an *ameliorative* motivation: even if it is difficult, it might be worthwhile to attempt to *improve* our conceptual and linguistic repertoire.⁷² Second, as we have just mentioned, there are reasons to think that some amount of conceptual ethics might well be a large part of actual philosophical practice, if only implicitly. If conceptual ethics is already happening implicitly, we can hope for modest gains by explicitly coming to grips with it and trying to do it well.⁷³ Third, paying attention to issues in conceptual ethics might help us better understand *which* topics we should be investigating in philosophy, and why.⁷⁴ This is because, by reflecting on issues in conceptual ethics, we are forced to confront questions about why we are employing certain concepts rather than others (and, tied to this, then investigating certain properties rather than others, or certain topics rather than others).⁷⁵

Finally, conceptual ethics questions might well be – like many other normative questions – of intrinsic interest to many philosophers. For example, it strikes us as a matter of great interest, as philosophers, whether many of our central philosophical terms and concepts are defective, as Kevin Scharp argues.⁷⁶ Likewise, it strikes us as a matter of great interest *what we should do*, if this sort of hypothesis turns out to be true. For example: should we abandon philosophy's terminological legacy? Seek to engineer better semantic values for it? Or just retain the defective terms and concepts and learn to live better with them?

We take these sorts of questions to be important. The main argument of this paper has suggested that they will not be easy to answer in any

⁷¹In other work, we argue that in the foundations of ethics and epistemology, (i) it is comparatively rare for practitioners to carefully distinguish conceptual ethics projects from what we call "metanormative" projects, and (ii) it is nonetheless very useful to distinguish these sorts of projects from each other, because while each such project can be powerfully motivated, their constitutive success conditions are quite different. See (McPherson and Plunkett 2021a) and (McPherson and Plunkett 2021c).

⁷²For connected discussion, see (Cappelen 2020) and (Haslanger 2020).

⁷³See the end of (Burgess and Plunkett 2013b) for a similar idea. See also (Plunkett 2015).

⁷⁴For discussion of how we think about "topics" in relation to issues about concepts, see (McPherson and Plunkett 2021d).

⁷⁵For further discussion of this idea, see (Plunkett 2015) and (McPherson and Plunkett 2020).

⁷⁶(Scharp 2020).

general way. And this, we think, suggests the most important upshot of this paper. Thus far, there has been comparatively little explicit attention paid to the epistemology of conceptual ethics. If we are right, this is regrettable: it would, for example, be helpful to know what we can know, and how to come by that knowledge, in this area. Relatedly, it would be helpful to better understand what we can *learn*, *have justified beliefs about*, etc. in this area. To make progress on these kinds of epistemological questions, we think, philosophers need to take seriously that conceptual ethics is a branch of normative inquiry, and the complexities that arise from doing so. We hope that this paper helps to spur more attention to these issues, so that the epistemology of conceptual ethics can take its place as a locus of philosophical attention, alongside, for example, the epistemology of morality and that of metaphysics.

Acknowledgements

Thanks to Alexis Burgess, Herman Cappelen, Ian Cruise, Terence Cuneo, Raff Donelson, Tyler Doggett, Andy Egan, Jesse Ferraioli, Natalie Dokken, Céline Henn, Yvonne Hütter-Almerigi, Zöe Johnson King, Zachary Lang, Amanda Li, Jake McNulty, Kate Nolfi, Jonathan Phillips, Björn Ramberg, Timothy Rosenkoetter, Jada Twedt Strabbing, Tim Sundell, and Amie Thomasson for helpful discussion and feedback.

Disclosure statement

No potential conflict of interest was reported by the author(s).

References

- Ayer, A. J. 1952 (1936). *Language, Truth, and Logic*. New York: Dover.
- Bennett, Karen. 2009. "Composition, Colocation, and Metaontology." In *Metametaphysics: New Essays on the Foundations of Ontology*, edited by D. Chalmers, D. Manley, and R. Wasserman, 38–76. Oxford: Oxford University Press.
- Blackburn, Simon. 1993. *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- Burgess, Alexis, Herman Cappelen, and David Plunkett. 2020. *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Burgess, Alexis, and David Plunkett. 2013a. "Conceptual Ethics I." *Philosophy Compass* 8 (12): 1091–1101.
- Burgess, Alexis, and David Plunkett. 2013b. "Conceptual Ethics II." *Philosophy Compass* 8 (12): 1102–1110.
- Burgess, Alexis, and David Plunkett. 2020. "On the Relation Between Conceptual Engineering and Conceptual Ethics." *Ratio* 33 (4): 281–294.
- Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.

- Cappelen, Herman. 2020. "Conceptual Engineering: The Master Argument." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 132–151. Oxford: Oxford University Press.
- Cappelen, Herman. Manuscript. "The Concept of *Democracy*: An Essay on Conceptual Amelioration and Abandonment."
- Cappelen, Herman, and David Plunkett. 2020. "A Guided Tour of Conceptual Engineering and Conceptual Ethics." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 1–26. Oxford: Oxford University Press.
- Darwall, Stephen L. 1998. *Philosophical Ethics*. Boulder: Westview Press.
- Dorr, Cian, and John Hawthorne. 2013. "Naturalness." In *Oxford Studies in Metaphysics: Volume 8*, edited by K. Bennett and D. Zimmerman, 3–77. Oxford: Oxford University Press.
- Dreier, Jamie. 2012. "Quasi-Realism and the Problem of Unexplained Coincidence." *Analytic Philosophy* 53 (3): 269–287.
- Dworkin, Ronald. 1996. "Objectivity and Truth: You'd Better Believe It." *Philosophy and Public Affairs* 25 (2): 87–139.
- Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Fine, Kit. 2001. "The Question of Realism." *Philosophers' Imprint* 1: 1–30.
- Finlay, Stephen. 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- Fitzpatrick, William. 2008. "Robust Ethical Realism, Non-Naturalism, and Normativity." *Oxford Studies in Metaethics* 3: 159–205.
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge: Mass.: Harvard University Press.
- Gibbard, Allan. 2011. "How Much Realism? Evolved Thinkers and Moral Concepts." In *Oxford Studies in Metaethics*, edited by R. Shafer-Landau, 33–51. Oxford: Oxford University Press.
- Haslanger, Sally. 2000. "Gender and Race: (What) Are They? (What) Do We Want Them to Be?" *Noûs* 34 (1): 31–55.
- Haslanger, Sally. 2020. "Going On, Not in the Same Way." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 230–260. Oxford: Oxford University Press.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon.
- Joyce, Richard. 2006. *The Evolution of Morality*. Cambridge: MIT Press.
- Köhler, Sebastian. 2021. "How to Have Your Quasi-Cake and Quasi-Eat It Too." *The Canadian Journal of Philosophy* 51: 204–220.
- Kratzer, Angelika. 2012. *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.
- Lenman, James. 2000. "Consequentialism and Cluelessness." *Philosophy and Public Affairs* 29 (4): 342–370.
- Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343–377.

- Lewis, David. 1989. "Dispositional Theories of Value." *Proceedings of the Aristotelian Society* 63: 113–137.
- Locke, Dustin. 2017. "The Epistemic Significant of Moral Disagreement." In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett, 499–518. New York City: Routledge.
- McGrath, Sarah. 2011. "Skepticism about Moral Expertise as a Puzzle for Moral Realism." *Journal of Philosophy* 108 (3): 111–137.
- McPherson, Tristram. 2012. "Unifying Moral Methodology." *Pacific Philosophical Quarterly* 93 (4): 523–549.
- McPherson, Tristram. 2015a. "The Methodological Irrelevance of Reflective Equilibrium." In *The Palgrave Handbook of Philosophical Methods*, edited by C. Daly, 652–674. Houndmills, Basingstoke: Palgrave Macmillan.
- McPherson, Tristram. 2015b. "What is at Stake in Debates among Normative Realists?" *Noûs* 49 (1): 123–146.
- McPherson, Tristram. 2018a. "Authoritatively Normative Concepts." In *Oxford Studies in Metaethics, Vol. 13*, edited by R. Shafer-Landau, 253–277. Oxford: Oxford University Press.
- McPherson, Tristram. 2018b. "Naturalistic Moral Realism, Moral Rationalism, and Non-Fundamental Epistemology." In *The Many Moral Rationalisms*, edited by K. Jones and F. Schroeter, 187–209. Oxford: Oxford University Press.
- McPherson, Tristram. 2020. *Epistemology and Methodology in Ethics*. Cambridge: Cambridge University Press.
- McPherson, Tristram. 2022. "Expressivism Without Minimalism." In *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*, edited by B. Dunaway and D. Plunkett, 147–170. Ann Arbor, MI: Maize Books, University of Michigan Publishing.
- McPherson, Tristram, and David Plunkett. 2017. "The Nature and Explanatory Ambitions of Metaethics." In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett, 1–25. New York: Routledge.
- McPherson, Tristram, and David Plunkett. 2020. "Conceptual Ethics and the Methodology of Normative Inquiry." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 274–303. Oxford: Oxford University Press.
- McPherson, Tristram, and David Plunkett. 2021a. "Conceptual Ethics, Metaepistemology, and Normative Epistemology." *Inquiry*: 1–33.
- McPherson, Tristram, and David Plunkett. 2021b. "Evaluation Turned on Itself: The Vindictory Circularity Challenge to the Conceptual Ethics of Normativity." In *Oxford Studies in Metaethics, Vol. 16*, edited by R. Shafer-Landau, 207–232. Oxford: Oxford University Press.
- McPherson, Tristram, and David Plunkett. 2021c. "Metaethics and the Conceptual Ethics of Normativity." *Inquiry*: 1–34.
- McPherson, Tristram, and David Plunkett. 2021d. "Topic Continuity in Conceptual Engineering and Beyond." *Inquiry*: 1–27.
- McPherson, Tristram, and David Plunkett. *Forthcoming*. "Ground, Essence, and the Metaphysics of Metanormative Non-Naturalism." *Ergo*.

- McPherson, Tristram, and David Plunkett. [Manuscript-a](#). "Conceptual Ethics and the Categories of 'Ideal Theory' and 'Non-Ideal Theory' in Political Philosophy: A Proposal for Abandonment."
- McPherson, Tristram, and David Plunkett. [Manuscript-b](#). "Strong AI and the Foundations of Ethics."
- Moore, G. E. 1993 (1903). *Principia Ethica*. Cambridge: Cambridge University Press.
- Pérez Carballo, Alejandro. 2020. "Conceptual Evaluation: Epistemic." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 304–332. Oxford: Oxford University Press.
- Plunkett, David. 2015. "Which Concepts Should We Use?: Metalinguistic Negotiations and the Methodology of Philosophy." *Inquiry* 58 (7-8): 828–874.
- Plunkett, David, and Tim Sundell. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13 (23): 1–37.
- Plunkett, David, and Tim Sundell. 2021. "Metalinguistic Negotiation and Matters of Language: A Response to Cappelen." *Inquiry*: 1–25.
- Quine, Willard V. O. 1951. "Two Dogmas of Empiricism." *Philosophical Review* 60 (1): 20–43.
- Railton, Peter. 1986. "Moral Realism." *The Philosophical Review* 95: 163–207.
- Railton, Peter. 2003. *Facts, Values, and Norms*. New York: Cambridge University Press.
- Rorty, Richard. 1980. *Philosophy and the Mirror of Nature*. Princeton: Princeton University Press.
- Rorty, Richard. 1982. *Consequences of Pragmatism: Essays 1972–1980*. Minneapolis: University of Minnesota Press.
- Scanlon, T. M. 2014. *Being Realistic about Reasons*. Oxford: Oxford University Press.
- Scharp, Kevin. 2020. "Philosophy as the Study of Defective Concepts." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 396–416. Oxford: Oxford University Press.
- Schechter, Joshua. 2017. "Explanatory Challenges in Metaethics." In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett, 443–458. New York: Routledge.
- Schechter, Joshua. [Manuscript](#). "Does Expressivism Enjoy an Epistemological Advantage Over Realism?"
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, Mark. 2014. "Does Expressivism Have Subjectivist Consequences?" *Philosophical Perspectives* 28 (1): 278–290.
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.
- Silk, Alex. 2016. *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Cambridge: Blackwell.
- Smith, Michael. 2012. "Agents and Patients, or: What We Learn about Reasons for Action by Reflecting on Our Choices in Process of Thought Cases." *Proceedings of the Aristotelian Society* 112 (3): 309–331.
- Smith, Michael. 2013. "A Constitutivist Theory of Reasons: Its Promise and Parts." *LEAP: Law, Ethics, and Philosophy* 1: 9–30.

- Stich, Stephen. 2011. *Collected Philosophical Papers Volume 1: Mind and Language, 1972–2010*. Oxford: Oxford University Press.
- Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value." *Philosophical Studies* 127 (1): 109–166.
- Street, Sharon. 2011. "Mind-Independence Without the Mystery: Why Quasi-Realists Can't Have It Both Ways." In *Oxford Studies in Metaethics, Vol. 6*, edited by R. Shafer-Landau, 1–32. Oxford: Oxford University Press.
- Street, Sharon. 2012. "Coming to Terms with Contingency: Humean Constructivism about Practical Reason." In *Constructivism in Practical Philosophy*, edited by J. Lenman and Y. Shemmer, 40–59. Oxford: Oxford University Press.
- Thomasson, Amie L. 2015. *Ontology Made Easy*. Oxford: Oxford University Press.
- Thomasson, Amie L. 2016. "Metaphysical Disputes and Metalinguistic Negotiation." *Analytic Philosophy* 57 (4): 1–28.
- Thomasson, Amie L. 2020. "A Pragmatic Method for Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett, 435–458. Oxford: Oxford University Press.
- Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility*. Oxford: Oxford University Press.
- Vavova, Katia. 2018. "Irrelevant Influences." *Philosophy and Phenomenological Research* 96 (1): 134–152.
- Wigner, Eugene. 1960. "The Unreasonable Effectiveness of Mathematics in the Natural Sciences." *Communications in Pure and Applied Mathematics* 13: 1–14.
- Williamson, Timothy. 2007. *The Philosophy of Philosophy*. Oxford: Blackwell.
- Wittgenstein, Ludwig. 1991 (1953). *Philosophical Investigations*. Oxford: Wiley-Blackwell.