9

Evaluation Turned on Itself:

The Vindicatory Circularity Challenge to the Conceptual Ethics of Normativity

Tristram McPherson and David Plunkett

Introduction

We engage in normative and evaluative thought and talk throughout our lives. For example, we make claims about how we should treat other people, which movies are better than others, what kind of social/political institutions are just, and what makes a scientific theory a good one. In such thought and talk, we deploy a range of normative and evaluative concepts: for example, SHOULD, JUSTICE, COURAGEOUS, IMPOLITE, and GOOD.[1] One possible target of normative and evaluative inquiry concerns those very normative and evaluative concepts themselves. For example, we might ask: are some of these concepts that we currently use defective? Could they be improved? More generally: which normative or evaluative concepts should we be using, and why? We call normative or evaluative inquiry with this sort of target the *conceptual ethics of normativity*. (Henceforth, we will generally use 'normative' as a shorthand way to refer to both the normative and the evaluative.)

There are different motivations one can have for engaging in the conceptual ethics of normativity. One natural motivation is to either vindicate or improve one's existing normative concepts. This chapter aims to clarify and address what we take to be one of the deepest challenges to the conceptual ethics of normativity, when the project is motivated in this way. Put roughly, the challenge arises from the fact that in order to evaluate our normative

[1] In this chapter, terms in small caps (e.g. CAT) pick out concepts. Single quotation marks (e.g. 'cat') are used to mention linguistic items. Double quotation marks (e.g. "cat") are used for a variety of tasks including quoting others' words, scare quotes, and mixes of use and mention.

Tristram McPherson and David Plunkett, *Evaluation Turned on Itself: The Vindicatory Circularity Challenge to the Conceptual Ethics of Normativity* In: *Oxford Studies in Metaethics Volume 16*. Edited by: Russ Shafer-Landau, Oxford University Press (2021). © Tristram McPherson and David Plunkett.
DOI: 10.1093/oso/9780192897466.003.0009

concepts, we seemingly need to use some of those very concepts. This might seem objectionably circular, akin to trying to verify the accuracy of a ruler by checking it against itself. We dub this the *vindicatory circularity challenge*.

If this challenge succeeds, it would seem to suggest that *all* of normative inquiry rests on worryingly arbitrary foundations. This would be true both of conceptual ethics and of normative inquiry in domains such as ethics, political philosophy, epistemology, or aesthetics. All such inquiry involves deploying normative concepts. And if the vindicatory circularity challenge cannot be met, then our use of our normative concepts in these contexts could only be vindicated in an objectionably circular way.

We organize the paper as follows. In Section 1, we more carefully introduce the conceptual ethics of normativity. In Section 2, we more carefully introduce the vindicatory circularity challenge. In Sections 3 and 4, we consider two important attempts to answer this challenge. The first appeals to the idea that we can answer it by appealing to a distinctive concept: the concept AUTHORITATIVELY NORMATIVE. The second appeals to the idea that the challenge can be dismissed because it cannot be clearly formulated using our concepts. In Section 5, we build on the insights provided by these replies to offer our preferred reply to the vindicatory circularity challenge. This reply begins by noting a similarity between this challenge and certain radical skeptical arguments in epistemology. With this in hand, we argue that we can answer the challenge by adapting existing anti-skeptical resources from epistemology. In Section 6, we explore the implications that certain pessimistic hypotheses about the concept AUTHORITATIVELY NORMATIVE have for the vindicatory circularity challenge. Finally, in Section 7, we reject a pair of competing attempts to address the challenge, which appeal to metaphysical and metasemantic resources respectively.

1 Introducing the Conceptual Ethics of Normativity

In this section, we introduce the idea of the conceptual ethics of normativity, and contrast it with other philosophical projects concerning normativity.

One useful way to orient ourselves to the conceptual ethics of normativity is to focus on a single normative concept (here: JUSTICE), and contrast three broad sorts of questions that we might ask with respect to that concept.

First, we might *use* this concept to ask questions such as: what is just, and why? We could aim to provide a completely general answer to this question. This would be an instance of "systematic" normative inquiry regarding

justice. Alternatively, we could pursue this question with respect to a specific salient cluster of contexts (e.g.what would be a just form of higher education in contemporary America, and why?). This would be an instance of "applied" normative inquiry regarding justice.

Second, we can ask questions such as the following: are thoughts such as *slavery is unjust* ordinary beliefs, or states that are more like desires? Are there facts about justice? If so, how (if at all) can we come to know about them? These questions can be seen as parts of the following project: to explain how our thought and talk using JUSTICE—and what (if anything) such thought and talk is distinctively about—fit into reality. We take the activity organized around this explanatory project to comprise *metanormative* inquiry about JUSTICE.[2]

Third, we might ask normative questions about the concept JUSTICE itself. For example, we can ask: is this concept good to use in evaluating political arrangements? Is it defective in some way? If so, what would be a better "JUSTICE-like" concept to use in its place? These are questions in *conceptual ethics* about JUSTICE.

We argue elsewhere that these three projects are distinct.[3] Briefly, this is because these projects have distinct success conditions. For example, to discover what is just and why (the core success condition of normative inquiry concerning justice) is a different matter than discovering whether using JUSTICE in our theorizing about politics is good or bad (a core concern of conceptual ethics about JUSTICE). This distinctness does not entail that these three projects are irrelevant to each other, however. For example, if normative inquiry revealed that justice is simply whatever is in the interests of the powerful, this revelation would plausibly have implications for conceptual ethics about JUSTICE. In particular, it would strongly suggest that JUSTICE is a bad concept around which to organize our normative inquiry about politics.

With this broad orientation in hand, we now introduce the conceptual ethics of normativity more carefully. We begin by introducing the idea of *conceptual ethics*.[4] As we understand it, the project of conceptual ethics addresses a range of normative and evaluative questions about *thought and*

[2] Our discussion here of the first two kinds of projects (normative inquiry and metanormative inquiry) draws on McPherson and Plunkett (2017).

[3] See McPherson and Plunkett (2021). For connected discussion, see also McPherson and Plunkett (2017, 2020).

[4] Here and below, our discussion of conceptual ethics draws on Burgess and Plunkett (2013a and 2013b). The project of conceptual ethics is very closely associated with projects that are described as "conceptual engineering" or "conceptual amelioration" (see, for example, Haslanger 2000 and Cappelen 2018). For discussion of a range of different ways of using the labels here, see the essays collected in Burgess et al. (2020).

talk (as well as about connected issues involving "representational" devices or systems more generally). Those include questions about which concepts we should use, what makes concepts better or worse, and what we should mean by our words.

Notice that, given this understanding of conceptual ethics, a philosopher can engage in conceptual ethics even if they are skeptical about the existence of *concepts* per se.[5]

Similarly, the term 'ethics' in 'conceptual ethics' is intended as a cover-all term for normative and evaluative theorizing, and should *not* be read as privileging moral and political evaluation of concepts. Indeed, a central question in conceptual ethics is the following:

Standards Question What are the *standards* we ought to use to evaluate the concepts for use in the relevant context?

Consider a sample of the wide variety of ways that a conceptual ethicist might answer this question. She might appeal to:

- *Conceptual* standards, such as the absence of incoherence, emptiness, or other conceptual defects;
- *Metaphysical* standards, such as joint-carvingness (e.g. perhaps relativistic MASS AT A VELOCITY is more joint-carving than MASS);
- *Epistemic* standards, such as being understandable by us, or having an extension that is apt for successful investigation by us;
- *Political* standards, such as usefulness to the pursuit of justice;
- *Ethical* standards, such as having morally good effects;
- Other normative standards (e.g., fit with etiquette, the law, clubhouse rules . . .).

And of course, the conceptual ethicist might appeal to a combination of such types of standards.[6] Note that the aim of these examples is to convey the diversity of possible standards, not to legislate how those standards should be grouped or labeled.

Notice that not all of these standards themselves involve moral or political criteria. Indeed, they don't even all involve specifically normative criteria

[5] For example, see Cappelen (2018).

[6] For related discussion of the range of standards a conceptual ethicist might use here, see Burgess and Plunkett (2013b), Cappelen (2018), and Cappelen and Plunkett (2020).

(e.g. contrast the idea of a concept being joint-carving with it being useful for the pursuit of justice). However, any standard is a norm in at least one thin ("formal" or "generic") sense of the term 'norm': it is something which things can either conform to or fail to conform to.[7]

Work in conceptual ethics happens throughout philosophy, from the philosophy of race to fundamental metaphysics to epistemology.[8] The conceptual ethics of normativity focuses on normative thought and talk in particular. It includes the evaluation of actually circulating normative concepts, as well as of possible *alternatives* to those concepts (roughly, normative concepts that we might consider using to augment or replace the stock of normative concepts that we currently use). We can ask of a given normative concept: is it (or its current or anticipated use) defective in some way? Is it better or worse to use than some relevant alternative? Ought we to use it, or avoid using it, in our thought and talk? Parallel questions can be asked about normative words, or other representational devices. In many cases in conceptual ethics, we are interested in questions about what particular agents, in particular circumstances should do. For example: which concepts we (as opposed to idealized agents) should use in our present circumstances, or which concepts Max should use for the purposes of doing sociology.

In short, the conceptual ethics of normativity aims to identify and explain which normative concepts, words, or other "representational (or inferential) devices" have certain normative properties, and to explain why. Drawing on our discussion of normative inquiry above, we can pursue this project either targeting specific salient clusters of contexts (the core project of "applied conceptual ethics" of normativity) or with an aim to maximal explanatory generality (the core project of "systematic conceptual ethics" of normativity). Much more could be said about the conceptual ethics of normativity, but this brief orientation puts us in a position to more carefully state the vindicatory circularity challenge that is our focus in this chapter.

2 The Vindicatory Circulatory Challenge

Given the characterization of the conceptual ethics of normativity offered in the previous section, we can now ask: what rationale is there for engaging in

[7] For connected discussion, see McPherson (2011) and McPherson and Plunkett (2017).

[8] For further discussion, see Burgess and Plunkett (2013a), Cappelen (2018), and Cappelen and Plunkett (2020).

this sort of inquiry? One flat-footed answer is simple curiosity about this topic, or some instance of it.[9] However, the distinctive significance of certain of our normative concepts suggests a more ambitious motive for engaging in the conceptual ethics of normativity. Normative thought and talk play many important roles in our lives. For example, they can play a central role in planning and deliberation, in the meanings we ascribe to our lives and choices, and in structuring our orientation to persons, political institutions, cultures, and the natural environment. The significance of these roles naturally motivates inquiry that aims either to *vindicate* the use of our existing normative concepts in these roles, or to improve on these existing concepts. We call this the *ambitious* aim for the conceptual ethics of normativity.

This ambitious aim can be further motivated by two observations about our existing stock of normative concepts. First, many of the normative concepts we use (e.g. our ethical ones) appear to be folk concepts: concepts that have their home, as it were, in everyday life as opposed to systematic inquiry. It might seem like a bit of a miracle if those folk concepts just happened to be the *best* tools for the various things we do with normative concepts, including using them within systematic normative inquiry. To see this, consider the many cases in which conceptual innovation in the natural and social sciences (and in philosophy) has provided us with better conceptual tools to understand and engage with various parts of our world.[10]

Second, like every other part of our cultural inheritance, our normative concepts have a history. That history plausibly plays a significant role in explaining why we possess the normative concepts that we do, rather than some alternative concepts. Importantly, that history is characterized both by what may appear to be normatively arbitrary forces (e.g. natural selection), and normatively objectionable ones (such as ideological forces). Reflecting on this history might lead us to doubt that the normative concepts that we currently use are ideal for all of the ways that we use them.[11]

[9] We don't mean to suggest here that such curiosity should not be interrogated. For example, part of enculturation into philosophy involves the shaping of one's sense of which questions are interesting or "philosophical," and the prevalent patterns of such enculturation could certainly be criticized in multiple ways.

[10] We don't mean to imply here that there are no reasons to stick with using folk concepts in normative inquiry. In McPherson and Plunkett (2020), we discuss both potential advantages and disadvantages of orienting normative inquiry around either folk or theoretical concepts, respectively, and further discuss the distinction between those types of concepts.

[11] For further discussion of these issues, see McPherson and Plunkett (2020). For connected discussion, see Cappelen (2020).

In the remainder of this section, we introduce a challenge to the conceptual ethics of normativity, when it is guided by the ambitious aim we have just motivated. As we characterized it, this aim is disjunctive: we aim to either vindicate or improve our existing stock of normative concepts. The challenge is simplest to introduce by focusing on the vindicatory element of this aim, so we begin our discussion there.

The general form of the challenge arises because we must (explicitly or implicitly) deploy *normative concepts* in doing conceptual ethics. For example, we deploy a normative concept when asking the Standards Question introduced in Section 1. You can't use a concept you don't possess, so we must use some of our *own* normative concepts to answer this question. This fact might seem to undercut the ambitious motivations for the conceptual ethics of normativity. To see this, consider an analogy: suppose that you became concerned about the accuracy of the ruler you use to measure things. It seems straightforwardly misguided to address this concern by using that very ruler to "check" itself. This would be an objectionably circular way to attempt to vindicate the accuracy of one's ruler. It might seem misguided in just the same way to seek to vindicate one of your normative concepts by testing it against itself. We call this the *vindicatory circularity challenge*.

The ruler analogy can be used to bring out a further important feature of this challenge. Consider two worries one might have about one's ruler. First, one might be concerned about the internal consistency of the ruler. For example, if the distances between adjacent inch marks vary significantly, the ruler will cause predictable difficulties when used to guide a small-scale carpentry project. Checking the ruler against itself is a reasonable response to this sort of concern, because it could help one to detect and correct for such inconsistencies. Second, one might be concerned with whether the ruler conforms to some *external standard*. This would be important, for example, if you are using the ruler to measure items that are intended to fit with objects designed using other measuring devices. This is the context in which checking the ruler against itself appears unhelpful.

Similarly, the vindicatory circularity challenge only arises if one thinks that there is a relevant standard external to our normative concepts, relative to which our use of those concepts might be vindicated or improved. Not everyone will be moved by the idea that there is (or even could be) such a standard.[12] But here the disanalogy with rulers seems potent: for certain

[12] The idea that there is such an external standard is arguably a kind of *realist* idea about normativity. However, this idea is different from many of the standard ways of understanding

tasks in carpentry, any consistent ruler may be as good as another. But there are surely internally consistent normative concepts that it would be substantively awful to consistently use in central deliberative and evaluative contexts.

So far, we have considered the challenge as it applies to a single normative concept in isolation. But the challenge appears to generalize smoothly to more holistic evaluation. Suppose that the totality of the normative concepts that you deploy fully vindicate each other: you are a paradigm of normative-conceptual harmony. Such *mutual vindication* seemingly fails to answer the challenge. This is because, as with particular concepts, it seems possible that a mutually vindicating set of normative concepts is nonetheless substantively awful to use.

The vindicatory circularity challenge extends smoothly to the use of the conceptual ethics of normativity to *improve* one's normative concepts. The issue is not that such improvement is impossible. Whether it is impossible or not is a separate issue. Moreover, it plausibly is possible. For example, suppose one wants concepts that are precise for doing certain scientific work. One might discover that our concept PRECISE itself fails to be precise, and thus needs to be revised. We might be able to introduce a revised concept PRECISE* that is a better tool to use to evaluate other scientific concepts in the relevant context.

While such conceptual improvement is thus possible, attempted improvement is nonetheless vulnerable to the vindicatory circularity challenge. To see why, again consider the ruler analogy: suppose that one attempts to "improve" one's ruler by testing it against itself, and resolving any internal inconsistencies. There is no guarantee that this makes the ruler better accord with the relevant external standard; indeed, depending on the initial state of the ruler, it might well make the ruler err *more* consistently relative to that standard.

Because the vindicatory circularity challenge concerns the conformity of our normative concepts to an (assumed) external standard, it is a broadly *epistemic* challenge. We have stated it in schematic terms. It can be made more precise in a variety of ways: for example, by clarifying exactly which epistemic properties are at issue (e.g. epistemic justification, reliability,

"normative realism." For example, the existence of normative facts or properties might be neither necessary nor sufficient for there to be such a standard. It is also not clear that such a standard must be wholly independent of our attitudes. Finally, it is also not completely clear that such a standard must be "metaphysical" in the way that some "quietists" about metanormative discourse find objectionable, or in a way which couldn't be incorporated by a "quasi-realist" form of expressivism.

warrant, knowledge, etc.). In light of this, one might think there is, strictly speaking, a *family* of related vindicatory circularity challenges that we have put forward. We think that, in many contexts, clarifying exactly which of these challenges one is focused on really matters (e.g. when working out the details of one of the anti-skeptical responses we consider in this chapter). However, for the purposes of this chapter, we abstract away from the differences among these implementations, and consider what can be said about the family of challenges as a whole, and how one might best respond to them in general terms. As part of this, for simplicity, we will continue to talk about "the" vindicatory circularity challenge in what follows.

We can further illustrate the epistemic character of the vindicatory circularity challenge by noting that it can be formulated in a way analogous to familiar epistemic "bad cases." We can seemingly imagine a thinker whose normative concepts are substantively awful to use, but who cannot come to recognize this, without something like a conceptual conversion experience. Such a being would be in what we call the *normative conceptual bad case*. If we think that *we* might be in the normative conceptual bad case, we seemingly cannot rule this hypothesis out by engaging in the conceptual ethics of normativity. For vindicatory circularity ensures that in the bad case, we will mistakenly conclude that we are in a good case (or at least not a *completely* bad case). This might lead us to worry that the conceptual ethics of normativity ultimately amounts to a circuitous kind of self-congratulatory exercise.

It may be illuminating to distinguish the vindicatory circularity challenge from other challenges that are either familiar, or have received recent attention. To begin, consider the contrast with three familiar sorts of challenge. First, this challenge is not that it is *trivial* that our normative concepts will vindicate themselves. Indeed, they might not. Suppose there is an evil demon who tracks your every move, and will slaughter many innocent people if you use the concept MORALLY BETTER in your conceptual ethics work. In this case it would pretty clearly be morally better *not* to use this concept.[13] Second, the challenge is not that the conceptual ethics of normativity is objectionably *conservative*. For all we have said, the conceptual ethics of normativity might, if properly executed, lead us to abandon our existing normative concepts in favor of substantially different alternatives. Third, the challenge is not simply that the conceptual ethics of normativity is

[13] Compare, e.g. Parfit (1984, §9), for the stronger idea that a true ethical theory might entail that you ought not to believe it.

very hard. It might indeed be very hard, for any number of reasons. For example, if it is possible for epistemic and moral norms that matter for conceptual ethics to genuinely conflict, it may seem very unclear what would be a uniquely sensible way of resolving that conflict. The vindicatory circularity challenge arises even if we imagine that we have a clear grasp on how, by our own normative lights, to resolve such tensions.

Next consider two challenges that have arisen in recent discussions about conceptual ethics. The first is the challenge that certain work in conceptual ethics involves problematic conceptual "hypocrisy." Such hypocrisy arises when one uses a concept to argue against its own use. For example, it would be hypocritical, in this sense, to argue that you shouldn't use the concept SHOULD. There might seem to be something "self-undermining" or "self-defeating" about this sort of argument.[14] Questions about conceptual hypocrisy are also questions about applying a concept to evaluate itself. But the issue is distinct from the vindicatory circularity challenge. In short, the worry about conceptual hypocrisy arises when a concept *fails* to be self-vindicating, while our worry concerns the significance of a concept *succeeding* in being self-vindicating.

The second nearby challenge is at the heart of Matti Eklund's recent *Choosing Normative Concepts*.[15] To get a feel for the challenge, consider one of Eklund's own main examples. Bad Guy has bad motives and does bad things. We can use our normative terms to criticize Bad Guy (as we just did). But what if Bad Guy has his own distinct, *alternative* normative concepts: concepts that differ from ours, but which nonetheless play a similar role in Bad Guy's reasoning and motivational structure? And what if Bad Guy's motives and actions are wholly endorsed by those concepts? This poses a deep question: in what non-trivial, "objective" sense could our normative concepts be better than Bad Guy's normative concepts?[16]

Eklund uses this sort of question to press a rich series of puzzles and challenges to a previously under-theorized sort of realist commitment concerning normativity. Eklund's animating worry is metaphysical. In our terms, his challenge questions what it could even be for the world to contain a single, privileged sort of external standard that would adjudicate between Bad Guy's normative concepts and ours.[17] The vindicatory circularity challenge instead

[14] See Burgess (2020) and Burgess and Plunkett (2013b) for further discussion.
[15] Eklund (2017). [16] This gloss on Bad Guy draws from McPherson (2018b).
[17] Along with Eklund, we presuppose that a "plenitudinous" response to this metaphysical challenge won't work, as our use of the terminology of "single, privileged" here suggests. See Clarke-Doane (2020) for connected discussion.

grants that such a standard exists, and questions the epistemic credentials of the activity of using our own concepts to evaluate how well our concepts conform to it. The vindicatory circularity challenge would thus remain even on many solutions to Eklund's challenge.[18]

In this section, we introduced an ambitious motivation for engaging in the conceptual ethics of normativity, and argued that it is compelling. We then posed the vindicatory circularity challenge to this activity, so motivated. In short, the challenge is this: insofar as we want to vindicate or improve our normative concepts, it seems wrongheaded to use those very concepts as the metric by which to test for success. But then it is unclear *which* standard we should be using to measure those concepts, and why. Finally, we aimed to clarify the challenge by distinguishing it both from three familiar sorts of worries that might seem to arise for the conceptual ethics of normativity, and from two recent challenges that have arisen in the conceptual ethics literature.

3 Authoritatively Normative Concepts

In this section, we begin our discussion of potential replies to the vindicatory circularity challenge. We consider several such replies in the remainder of the chapter. In this section and the next, we introduce replies that we do not find fully convincing on their own, but which provide important resources for understanding our preferred strategy.

Thus far, we have discussed the vindicatory circularity challenge in terms of self- or mutual *vindication* among normative concepts. But we learn something important if we instead think about normative concepts that are (at least in a context) mutually *discrediting*. Consider an example: when one is dining with slaveholders, politeness and justice might seem to dictate conflicting courses of behavior. This perceived conflict can extend to the question of which normative concepts to *use* in guiding one's behavior in such contexts. Concern to promote or exemplify justice might discourage one from guiding one's behavior too scrupulously with the concept POLITENESS. And similarly, norms of politeness might discourage one from scrupulously guiding one's behavior with the concept JUSTICE.

[18] For example, it would remain if the solution to Eklund's challenge offered in McPherson (2020) succeeded. At the end of this chapter, we suggest that it also remains given the solution to the challenge that Eklund is most sympathetic to.

This perceived conflict illustrates a striking asymmetry. It would be somewhat bizarre to use concepts like POLITE or CHASTE in conceptual ethics work that evaluates our use of concepts like JUSTICE or MORALLY BETTER. "Why should I care," you might reasonably ask, "whether it is *impolite* to be guided by my use of JUSTICE here?" By contrast, it at least seems reasonable to ask whether justice requires you to constrain your use of POLITE in deliberation, or to altogether stop using the concept CHASTITY to evaluate others.

One compelling explanation of this asymmetry is that the concepts JUSTICE and MORALLY BETTER are more "normatively significant," or, as we will put it, more *authoritatively normative* than the concepts CHASTE and POLITE.[19] This suggests the general principle that it is sensible to evaluate use of less authoritatively normative concepts by using more authoritatively normative concepts, but not vice-versa.[20] If this is right, then our previous discussion of the *mutual* vindication of one's concepts is somewhat misleading. For what is crucial is whether use of a normative concept is vindicated by *specific* normative concepts: namely, the most authoritatively normative concepts.

This explanation might also seem to suggest a natural response to the vindicatory circularity challenge. Recall that this challenge requires the idea of an *external standard* against which our normative concepts might fall short. A natural idea is that authoritativeness provides the relevant standard. What we seek in the ambitiously motivated conceptual ethics of normativity is not self- or mutual vindication among our normative concepts, but instead the identification of which normative concepts we *authoritatively* should use. Because we are not seeking self- or mutual vindication on this approach, the specter of vindicatory circularity might seem to simply disappear.

We think it is exactly right to place authoritativeness at the center of conceptual ethics (including the conceptual ethics of normativity). To illustrate the role we think it should play, consider the following, four-stage structure for work in "applied" conceptual ethics. We distinguish these "stages" for analytical clarity, not because we think they need to be

[19] The terminology here builds on McPherson (2018a). In some previous work, we used the term 'robust normativity' to talk about what we here call "authoritative normativity," drawing on McPherson (2011).

[20] On this usage, the concept AUTHORITATIVE OUGHT discussed in McPherson (2018a) is, if non-defective, part of the *extension* of the concept AUTHORITATIVELY NORMATIVE being discussed here.

pursued in isolation from each other, or in any particular order. Here are the stages:

i. Identify a context of use for which we are evaluating concepts.
ii. Identify candidate concepts that might be used in that context.
iii. Learn relevant facts about these candidate concepts and the context.
iv. Evaluate the candidate concepts for use in that context.

It is in the fourth stage that we have just been suggesting one appeal to AUTHORITATIVELY NORMATIVE. In short, the proposal is that AUTHORITATIVELY NORMATIVE is used to answer what we earlier called the "Standards Question": namely, the question of which *standards* we ought to use to evaluate the concepts being used in the relevant context.[21]

We thus think that a well-executed version of the conceptual ethics of normativity should involve appeal to the concept AUTHORITATIVELY NORMATIVE. However, we don't think such an appeal provides a full reply to the vindicatory circularity challenge. To see this, suppose we grant the above gloss on the conceptual ethics of normativity that gives pride of place to the concept AUTHORITATIVELY NORMATIVE. Now consider a characterization of one way for us to be in what, above, we called the *normative conceptual bad case*, given this gloss on conceptual ethics: it would be for our concept AUTHORITATIVELY NORMATIVE to fail to have the relevant external standard as its extension.

You might think that we could not be in such a bad case. After all, we have arguably introduced the concept AUTHORITATIVELY NORMATIVE precisely to talk in an especially transparent way about the relevant external standard, so the concept could not fail to pick out that standard. But we can grant that *in the good case* the concept AUTHORITATIVELY NORMATIVE is the most transparent way of talking about the relevant standard, without thereby granting that this concept *could not fail* to pick out the relevant standard.

What we need is some guarantee that our concept AUTHORITATIVELY NORMATIVE picks out the relevant external standard. One way to do so would be to stipulate that the relevant external standard is simply whichever standard we pick out with our concept AUTHORITATIVELY NORMATIVE. But

[21] Note that, when one uses that concept, those standards might then (for example) turn out be any of the standards we glossed earlier. And those standards might suggest using any range of normative concepts in a given context (including, for example, using ones that are not very authoritative).

this proposal would in effect stipulate away what seems like the genuine worry about our concepts expressed by the vindicatory circularity challenge. This is because it would threaten the externality of the relevant standard in the following way. Given the stipulation, when we ask about whether our concept AUTHORITATIVELY NORMATIVE conforms to the external standard, this is simply a matter of applying that very concept to itself. We would, in effect, be back to checking our ruler against itself.

What is needed is an argument that preserves the externality of the standard, but guarantees that our concept's extension will match it. Perhaps some kind of metasemantic story could provide the basis for such an argument (we return to this possibility in Section 7). But absent such a story, vindicatory circularity looms again. For we will need to use our own normative concepts to assess whether our concept AUTHORITATIVELY NORMATIVE is the one to give the sort of central role we have suggested in this section, in the ambitious conceptual ethics of normativity.

If this is right, then even if we can most helpfully understand the ambitious conceptual ethics of normativity in terms of concepts like AUTHORITATIVELY NORMATIVE, we cannot answer the vindicatory circularity challenge simply by appealing to such concepts.

4 Ineffability

As we now explain, the significance of concepts such as AUTHORITATIVELY NORMATIVE raises another issue. This is that if the authoritatively normative concepts reply to the challenge fails, the vindicatory circularity challenge might be *ineffable* in the following sense: we might lack the conceptual resources to state it. In this section, we explain why this might be so. We then consider (and cast doubt on) the ideas that this ineffability would either undercut the challenge, or render it insoluble.[22]

Thus far, we have tended to formulate the vindicatory circularity challenge in one of two ways: either we have talked elliptically about the "relevant external standard," or we have talked about the possibility of certain norms being (e.g.) "substantively awful to use." As we now explain, neither of these formulations is satisfactory as a rigorous formulation of the challenge.

[22] We are indebted to Eklund's (2017) discussion of ineffability in what follows.

Consider first the "substantively awful" formulation. The problem with this formulation follows from two facts. First, SUBSTANTIVELY AWFUL is one of our own normative concepts. If we suppose that our normative concepts are mutually vindicating, we can thus rule out that any of them are substantively awful to use. Second, we are assuming precisely that we could be in a normative conceptual bad case even if our normative concepts were mutually vindicating. Compare: we could be mis-measuring a space relative to an external standard, even if all of the rulers we used vindicated each other.

Next consider the "relevant external standard" formulation. The problem is that the term 'relevant' here functions as a theoretical placeholder; what makes a standard relevant for the purposes of the conceptual ethics of normativity? In Section 3, we considered the idea that the relevant standard is the standard of being authoritatively normative. But, as we emphasized earlier, we can seemingly still coherently worry about being in the normative conceptual bad case, even when deploying our concept AUTHORITATIVELY NORMATIVE. And this means that we cannot simply identify the external standard with being authoritatively normative. Generalizing, one might suspect that this will be true of *any* concept we might use to try to formulate the external standard. And if this is so, it would seem that there is no way to use our concepts to specify what the "relevant" external standard is. If this turns out to be correct, then the vindicatory circularity challenge would be *ineffable*: we would not be able to precisely express the challenge using our concepts.

This suggests a second natural response to the vindicatory circularity challenge. The thought is simply that if we cannot formulate the challenge with our concepts, its apparent significance must simply be a matter of confusion. If this is right, the vindicatory circularity challenge is not so much answered as undercut.[23]

We think this reaction is too quick. To see this, suppose for the moment that our normative concepts are all excellent, and consider someone *else* in the normative conceptual bad case. Suppose that they engage in the conceptual ethics of normativity. They find that their (in fact substantively awful) normative concepts are mutually vindicating. They worry that (despite being mutually vindicating), their normative concepts might be bad

[23] This reaction is similar to what Eklund calls the "complacent" reaction (Eklund 2017, 6) to his related set of concerns that we mentioned in Section 2.

relative to some ineffable-but-relevantly-important standard. Two observations here. First, against the confusion diagnosis suggested above, the worry here does not seem confused, even if in some sense it cannot be coherently expressed using the worrier's concepts.[24] Second, in the case we are imagining, this worry is, intuitively, *correct*. This is important because (but for some assumptions about conceptual repertoire), the vindicatory circularity challenge is in part motivated by the possibility that we might be in a parallel bad case. While we lack the conceptual resources to describe our case just like we can describe this one, intuitively, they appear precisely parallel. If this is right, it shows that the vindicatory circularity worry is potent even if it is ineffable.

A second reaction to the alleged ineffability of the vindicatory circularity challenge is to conclude that we cannot hope to address this challenge. The rationale for this reaction is straightforward. Arguably, reasoning necessarily operates on articulable contents. So if we confront a genuine worry that cannot be properly articulated using our concepts, our powers of reasoning might simply be powerless to either vindicate or address that worry. What would follow if this reaction is correct? We would not have a *reason* to abandon the ambitious motivation for the applied conceptual ethics of normativity. But—if we chose to engage in this project—we would face a worry about our project that we are simply incapable of addressing. We think the best way of addressing this reaction is by offering a constructive response to the challenge. We do this in Section 5.

5 Skepticism and Self-Trust

Suppose that despite being ineffable, the vindicatory circularity challenge constitutes a genuine epistemic challenge to the ambitious conceptual ethics

[24] Compare also Nagel's response to Putman's semantic argument against epistemological skepticism (from Putnam 1981). Nagel writes:

> If I accept the argument, I must conclude that a brain in a vat can't think truly that it is a brain in a vat, even though others can think this about it. What follows? Only that I can't express my skepticism by saying "Perhaps I'm a brain in a vat." Instead I must say "Perhaps I can't even think the truth about what I am, because I lack the necessary concepts and my circumstances make it impossible for me to acquire them!" If that doesn't qualify as skepticism, I don't know what does. (Nagel 1986, 73)

As this suggests, one might take the inexpressibility of a significant worry to exacerbate that worry.

of normativity. How might we seek to address this challenge? Our reply appeals to an analogy to certain global skeptical challenges in epistemology.

In forming beliefs, we must make use of our capacity to reason. One sort of global skeptical challenge in epistemology begins by challenging the reliability of this capacity. Any non-skeptical epistemology must, in a natural sense, beg the question against this challenge. This is because we could only hope to answer this challenge by using our own capacity to reason, the reliability of which is precisely what is questioned by this sort of skeptic. In doing so, the epistemologist faces a vindicatory circularity challenge: there is an external standard (reliability), and one is using one's capacity to reason in order to check whether that very capacity conforms to that standard.

Suppose (at least of the sake of argument) that the epistemologist proceeds correctly when she rejects this global skeptical challenge. This would show that—at least in one case—a kind of vindicatory circularity is acceptable. Our motivating thought is that the similar structure of these two challenges suggests that it should be appropriate to address the vindicatory circularity challenge in conceptual ethics in whatever way it is appropriate for the non-skeptical epistemologist to address the global skeptic.

There are many competing ways of responding to the global epistemic skeptic. Here, we briefly sketch just one possibility that we find attractive in the context of both global skepticism and our vindicatory circularity challenge. This response is to suggest that we are entitled to a kind of epistemic *self-trust*: roughly, the trust that one's reasoning capacities are not root and branch unreliable, and are, in fact, acceptable starting points for inquiry.[25] In our view, the most promising reply to the ineffable version of the vindicatory circularity challenge is analogous. The idea is that we are entitled to *normative-conceptual self-trust*: that is, that we are entitled to treat our own normative concepts as acceptable starting points for the conceptual ethics of normativity.

One way to motivate the analogy begins by noting the general scope of the vindicatory circularity challenge with respect to our reasoning. Suppose that the central assumption underlying the project of conceptual ethics is correct: that there can be better or worse words or concepts to use in representing the world. If our normative and evaluative thinking about our own words and concepts is put in question by the vindicatory circularity

[25] Keith Lehrer (1997, 6) attributes this idea to Thomas Reid. Note that those partial to a less individualistic approach to epistemology can substitute (or add) trust that *our* reasoning capacities are not hopeless.

challenge, then so too is the representational scheme that structures our thought and talk quite generally. If the vindicatory circularity challenge does have such broad implications, this arguably reinforces the appeal of our self-trust reply. For many philosophers, part of what underlies the self-trust reply to global epistemic skepticism is the generality of the capacities being targeted by that skeptic. The fact that there is a similar generality at play in the vindicatory circulatory challenge does not by itself show that this self-trust reply can be extended to this case. But it does give us initial reason to take self-trust seriously as a strategy for responding to the vindicatory circularity challenge.

Normative-conceptual self-trust means that, even if we grant that there is an important worry that using our own concepts might fail to help us to track the (perhaps ineffable) external standard, this does not undermine the project of the conceptual ethics of normativity. We are entitled, on this view, to use our own normative conceptual starting points in engaging in this project. This reply does not address the vindicatory circularity challenge by showing that we can *avoid* (at least some amount of) vindicatory circularity. Rather, it claims that vindicatory circularity is epistemically acceptable here, as it is in replying to the global skeptic.

It is worth emphasizing another reason that we are arguably entitled to normative-conceptual self-trust. The conceptual ethics of normativity is just *one* context where we need to trust these concepts. We also need to trust them in *all* of our other normative reasoning: including, for example, when engaged in normative inquiry in ethics, political philosophy, and epistemology, as well as in prosaic normative reasoning in one's everyday life. We suggest that there is parity between the normative-conceptual self-trust required to sensibly engage in conceptual ethics in hopes of improvement, and the same self-trust required to sensibly use normative concepts in other contexts. If such self-trust were illicit, that would vitiate normative reasoning quite generally.

It is important to emphasize that reasonable self-trust is not *uncritical*. One can investigate whether some of one's reasoning capacities are reliable in certain contexts. One might discover, for example, that certain of one's reasoning capacities are prone to be influenced by unreliable framing effects in some contexts, thereby defeating one's trust in those capacities in those contexts. Similarly, normative-conceptual self-trust does not entail that we must uncritically embrace our normative-concepts. We can imagine that a well-executed conceptual ethics might serve to discredit one or more of our central normative concepts.

A pressing task for any reply to the vindicatory circularity challenge is to explain what went wrong with the ruler analogy. We are now in a position to offer this diagnosis. The details here will be controversial, but it is plausible that a crucial factor is the *generality* of the capacities being targeted. As we have just seen, reasonable self-trust treats one's own reasoning capacities in a critical way. And part of appropriate critique is bringing to bear one's critical capacities to the best of one's abilities. So, suppose that one was concerned about the value of *one* of one's normative concepts. Then it would be very odd to use only *that* concept to vindicate itself. One should, one might think, evaluate that concept using (appropriate elements of) one's full set of normative concepts. On this diagnosis, it is only when the challenge is extended to target one's use of normative concepts *quite generally* that it fails.

It is worth stepping back at this stage. The appeal to self-trust is only one of the standard ways of diagnosing what goes wrong with radical skeptical arguments in epistemology. And while we find it compelling, we are not going to defend the idea that it is the only or best approach. We suspect that several of the other leading anti-skeptical strategies could be extended by analogy to defend the conceptual ethics of normativity.[26] But others may not be as friendly to this extension. We thus take our argument here to be provisional: we think there is reason to be optimistic that the vindicatory circularity challenge can be answered, even if any particular way of arguing for such an answer will likely itself be controversial.

The anti-skeptical response we have just floated does not need to be given on its own. In particular, it might be supplemented with resources we introduced when discussing other strategies for responding to the vindicatory circularity challenge. To see this, let's spell out the relationships among the three main strategies for responding to this challenge that we have discussed thus far in this chapter.

The appeal to the concept AUTHORITATIVELY NORMATIVE (which we discussed in Section 3) is, in our view, the most natural way of treating the vindicatory circularity challenge as stateable using our concepts: the idea is that we can formulate the relevant standard as whatever is picked out by this concept. Notice that, if the challenge is stateable in this way, we would have no need for normative-conceptual self-trust. For our use of the concept

[26] So too might resources from the adjacent discussion in epistemology about "debunking" arguments (which are sometimes seen as distinct from "skeptical" arguments). On that front, one possible example is Katia Vavova's (2018) insistence that debunking challenges to your belief that P need to provide good independent reason to believe that you are mistaken about P, rather than showing that you have no independent reason not to believe that not-P.

AUTHORITATIVELY NORMATIVE would be justified by the fact that this is the very concept with which the relevant challenge is stated. In Section 4, we argued that the ineffability of the vindicatory circularity challenge need not undermine its significance. In this section, we argued by parity with general anti-skeptical replies that we could appeal to normative-conceptual self-trust to entitle us to use our own concepts in addressing such an ineffable challenge in an epistemically legitimate way. One important such concept is AUTHORITATIVELY NORMATIVE. As we argued in Section 3, we cannot formulate the vindicatory circularity challenge using that concept to pick out the external standard, which undercuts the idea that the challenge can be met simply by appeal to that concept. Despite this, we still think that the concept AUTHORITATIVELY NORMATIVE has a central role to play in the ambitious conceptual ethics of normativity. This is because the arguments of Section 3 suggest that this concept has a privileged place in our thinking about normativity, and normative-conceptual self-trust suggests that we should provisionally accept this privileged role.

6 Doubts about the Concept AUTHORITATIVELY NORMATIVE

We have just suggested that the concept AUTHORITATIVELY NORMATIVE provisionally has a privileged role within our preferred reply to the vindicatory circularity challenge. In this section we explore the consequences for the challenge of three sorts of pessimistic hypotheses about that concept.

One kind of worry we might have concerns our ability to learn about which normative concepts are more authoritative than others using the concept AUTHORITATIVELY NORMATIVE. For example, perhaps creatures like us are just very bad at learning about such facts using this concept, or at least are so in our current social/historical circumstances. Suppose a radical version of this worry is correct and we are *very* bad at learning about which normative concepts are more authoritative. If this is right, then the ambitious kind of conceptual ethics that we have been considering in this chapter may be hopeless, for reasons wholly distinct from the vindicatory circularity challenge. In short, unless there is some other normative concept that we can introduce that would provide us with the relevant way to learn about which normative concepts (or facts, etc.) are more authoritative than others, the ambitious form of conceptual ethics would be an in-principle sensible project, but one that we cannot reliably pursue.

Consider a second set of worries: that the concept AUTHORITATIVELY NORMATIVE has an empty extension, or that it is incoherent, or that it is defective in some other sort of way.[27] If one of these hypotheses is correct, it might cast doubt on the idea that there is any external standard that is relevant for evaluation or improvement of our normative concepts, given the ambitious motivation for the conceptual ethics of normativity. After all, as we have emphasized, not just any external standard will do: it has to be a standard that is distinctively suitable to use in evaluating normative concepts. Some proposals for why AUTHORITATIVELY NORMATIVE is incoherent, or has an empty extension, could effectively cast doubt on the very idea that there is a distinctive external standard of this kind.

Finally, it might turn out that there is not *one* "authoritativeness" role for normative concepts, but a variety of "authoritative-ish" ones. For example, some of these roles might be tied to action-guidance, others to advice, and yet others to speaker-endorsement.[28] The upshot of this result for the conceptual ethics of normativity is not straightforward. Much depends upon whether and *how* and *to what extent* the competing concepts overlap; for example, a partial ordering of normative concept improvement might still be possible.

It is worth emphasizing that in spelling out the upshots of these possibilities about our concept AUTHORITATIVELY NORMATIVE, we have been assuming that we are in some sense stuck with that concept. But even if one of the hypotheses just mentioned is correct, it is possible that we might be able to introduce a novel concept that suitably represents the relevant external standard for the conceptual ethics of normativity, even if none of our existing concepts do. Recall, by analogy, the possibility of introducing a more precise alternative to the concept PRECISE.

7 Alternative Responses to the Vindicatory Circularity Challenge

In this chapter, we have considered three general strategies for answering the vindicatory circularity challenge: one which appealed to the concept

[27] For relevant discussion of the emptiness possibility, see McPherson (2018a). For the defectiveness idea, see Baker (2018) and Tiffany (2007). For replies to the defectiveness worries, see McPherson (2018a) and Wodak (2018).

[28] This idea is briefly floated in McPherson and Plunkett (2017) and discussed further in Plunkett (2020). See also Finlay (2019).

AUTHORITATIVELY NORMATIVE, one which dismissed the challenge due to its (purported) ineffability, and one which appealed to anti-skeptical resources in epistemology (e.g. self-trust). We take each of these replies to warrant further exploration, although we are most optimistic about a version of the third response, coupled with resources drawn from the discussion of the first two. In this section, we briefly criticize two broad types of competing reply to the challenge, which we find less promising. These are instances of metaphysical and metasemantic strategies for addressing the challenge, respectively.

The first strategy takes inspiration from the fact that the heart of the vindicatory circularity challenge is to identify an *external standard*. This strategy proposes to escape the challenge by investigating this standard directly: to find out *what really matters*, and evaluate our concepts by using what we learn about normative reality.

The core problem for this metaphysics-first strategy is that there is a *lot* of normative reality. If Section 6 is on the right track, we would need to be investigating a specific bit of that reality: whichever bit constitutes the relevant external standard. It is hard to believe that such an investigation will not need to be guided by our concept AUTHORITATIVELY NORMATIVE (or, given the discussion that concluded Section 6, a related normative concept). And if it is, the vindicatory circularity challenge can simply target that guidance, and the self-trust reply (or something like it) appears needed to explain why we are entitled to orient our investigation of normative reality by deploying *this* concept.

The second strategy is inspired by a related discussion by Matti Eklund.[29] This strategy begins by asking what would need to be true for us to worry about vindicatory circularity. One seemingly plausible answer is this: there would have to be a concept with a different extension from that of our concept AUTHORITATIVELY NORMATIVE, that nonetheless plays the same role in thought as that concept. The second strategy suggests that the vindicatory circularity challenge can be avoided if this very role in thought is reference-fixing; that is, if *every* concept that plays that role has the same extension.

Here is a reason to think that this, by itself, fails to answer the vindicatory circularity challenge. Suppose that you are worrying about whether it is good to use the concept PRACTICAL OUGHT in certain central deliberative

[29] See Eklund (2017).

contexts. Then you are told that *any* concept that you used in the same way in those central deliberative contexts would have the same extension as PRACTICAL OUGHT. Does this *fully* answer your worry? We don't think so. This is because this hypothesis seems compatible with the possibility that the relevant extension might be substantively awful to be guided by. It might be a *tragedy* that insofar as we deliberate, we cannot help but guide our actions by using a concept that picks out *that* extension, etc. Compare again the initial analogy with the ruler. If you learn that you are under a spell, so that you cannot help but use a given ruler, that doesn't give you any assurance at all that the ruler matches up to the relevant external standard.[30]

Here we only considered one instance each of a metaphysical and metasemantic strategy for addressing the challenge. However, we suspect that the general form of our replies will generalize to other instances of these strategies. For example, we suspect that any sort of metaphysical strategy will at some point need to at least tacitly utilize our normative concepts in a way that regenerates the challenge. And we think that, in general, metasemantic hypotheses are poor candidates to fully address evaluative anxieties. For, in short, facts about how we think and talk (or even how we must do so) aren't the right kind of facts to address these worries.[31]

Conclusion

In this chapter, we have explored what we take to be one of the most pressing foundational questions for the conceptual ethics of normativity: namely, the vindicatory circularity challenge. We have aimed to clearly explain and motivate this challenge, and to identify the most promising strategies for addressing it. As we have argued, one of the central difficulties we face in *characterizing* this challenge is that the most general form of the challenge appears likely to be ineffable, which may cast doubt on how well we can understand or grapple with it.

[30] It should be underscored that we don't take this to be an argument against Eklund's own use of this kind of reference-fixing idea. As we mentioned earlier in this chapter, Eklund is focused on a topic distinct from the vindicatory circularity challenge. For our further discussion of Eklund's use of this idea, see McPherson (2020) and Plunkett (2020).

[31] For further discussion of this idea, see McPherson (2020), Plunkett (2020), and Leary (2020). For connected discussion, see Enoch (2006). The idea that there might be concepts we *must* use in our thinking is tied to issues about the existence of "conceptual fixed points" or "bedrock" concepts, as discussed in Eklund (2015) and Chalmers (2011).

One way of responding to the vindicatory circularity challenge is to deny that we need to worry about ineffable challenges. While the question of how and what to think about any alleged ineffability is inevitably vexed, we have cast doubt on this sanguine conclusion. A second way of responding suggests that the challenge can be both formulated and answered in terms of our concept AUTHORITATIVELY NORMATIVE. Here we noted the need for a clear argument that we could not fail to think about what matters most while using this concept. We suggested that the most promising general strategy for addressing the challenge involves noting an analogy with parallel challenges in the foundations of general epistemology. We suggested that a plausible hypothesis is that we are entitled to the same resources here as we are in addressing the skeptic about the reliability of our capacity to reason. As with the preceding strategies, however, our execution of this strategy was extremely schematic, and leaves many important questions unanswered. We hope that this chapter spurs further work on each of these strategies, in addressing this important challenge.

One reason to investigate these strategies in more depth concerns the broader significance of the possibility that we cannot answer the vindicatory circularity challenge. Consider what this would mean for the project of substantive normative inquiry (whether in epistemology, moral philosophy, political philosophy, or elsewhere). If the challenge cannot be answered, this arguably casts doubt on the use of normative concepts quite generally. For if our normative concepts are not appropriate tools to use in investigating which concepts to deploy, it is hard to see how they could be appropriate tools to use in normative reasoning more generally. As we noted in Section 5, the significance of the challenge is arguably broader still, putting into question the representational scheme that structures our thought and talk quite generally. We hope that we have contributed here to the project of clarifying how to best answer (or at least cope with) such doubts.[32]

[32] Thanks to Mitch Berman, Alexis Burgess, Herman Cappelen, Matti Eklund, Sukaina Hirji, Nadeem Hussain, Zoë Johnson King, Alex King, Victor Moberger, Jonas Olson, Josh Petersen, Adrian Russian, Dan Singer, Tim Sundell, Amie Thomasson, and two referees for this volume, for helpful feedback and discussion. Thanks also to participants in discussions of previous drafts of this chapter at Dartmouth College, University of Michigan, Stanford University, Stockholm University, Pompeu Fabra University, the Frankfurt School of Finance and Management, the Madison Metaethics Workshop, University of Oslo, and the Arizona Ranch Metaphysics Workshop.

Works Cited

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

Burgess, Alexis. 2020. Never Say 'Never Say "Never"'? In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen, and D. Plunkett. Oxford: Oxford University Press.

Burgess, Alexis, and David Plunkett. 2013a. Conceptual Ethics I. *Philosophy Compass* 8 (12): 1091–101.

Burgess, Alexis, and David Plunkett. 2013b. Conceptual Ethics II. *Philosophy Compass* 8 (12): 1102–10.

Burgess, Alexis, Herman Cappelen, and David Plunkett. 2020. *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.

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. Oxford: Oxford University Press.

Cappelen, Herman, and David Plunkett. 2020. A Guided Tour of Conceptual Engineering and Conceptual Ethics, edited by A. Burgess, H. Cappelen, and D. Plunkett. Oxford: Oxford University Press.

Chalmers, David J. 2011. Verbal Disputes. *Philosophical Review* 120 (4): 515–66.

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

Eklund, Matti. 2015. Intuitions, Conceptual Engineering, and Conceptual Fixed Points. In *The Palgrave Handbook of Philosophical Methods*, edited by C. Daly. New York: Palgrave Macmillan.

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

Enoch, David. 2006. Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action. *Philosophical Review* 115 (2): 169–98.

Finlay, Stephen. 2019. Defining Normativity. In *Dimensions of Normativity: New Essays on Metaethics and General Jurisprudence*, edited by D. Plunkett, S. Shapiro, and K. Toh. 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.

Leary, Stephanie. 2020. Choosing Normative Properties: a reply to Eklund's Choosing Normative Concepts. *Inquiry* 63 (5): 455–74.

Lehrer, Keith. 1997. *Self-Trust: A Study of Reason, Knowledge, and Autonomy*. Oxford: Oxford University Press.

McPherson, Tristram. 2011. Against Quietist Normative Realism. *Philosophical Studies* 154 (2): 223–40.

McPherson, Tristram. 2018a. Authoritatively Normative Concepts. In *Oxford Studies in Metaethics Vol. 13*, edited by R. Shafer-Landau. Oxford: Oxford University Press.

McPherson, Tristram. 2018b. Review of Choosing Normative Concepts. *Notre Dame Philosophical Reviews*: available at https://ndpr.nd.edu/news/choosing-normative-concepts/.

McPherson, Tristram. 2020. Ardent Realism without Referential Normativity. *Inquiry* 63 (5): 489–508.

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. 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. Oxford: Oxford University Press.

McPherson, Tristram, and David Plunkett. 2021. Metaethics and the Conceptual Ethics of Normativity. *Inquiry* (online first version available at https://doi.org/10.1080/0020174X.2021.1873177)

Nagel, Thomas. 1986. *The View from Nowhere*. Oxford: Oxford University Press.

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

Plunkett, David. 2020. Normative Roles, Conceptual Variance, and Ardent Realism about Normativity. *Inquiry* 63 (5): 509–34.

Putnam, Hilary. 1981. Brains in a Vat. In *Reason, Truth, and History*. Cambridge: Cambridge University Press.

Tiffany, Evan. 2007. Deflationary Normative Pluralism. *Canadian Journal of Philosophy* 37 (5): 231–62.

Vavova, Katia. 2018. Irrelevant Influences. *Philosophy and Phenomenological Research* 96 (1): 134–52.

Wodak, Daniel. 2018. Fictional Normativity and Normative Authority. *Canadian Journal of Philosophy* 49 (6): 828–50.