21

A Pragmatic Method for Normative Conceptual Work

Amie L. Thomasson

I have argued elsewhere (2016), that we should see metaphysics as fundamentally involved in conceptual work, where this conceptual work includes not only descriptive conceptual work like that undertaken by such figures as Gilbert Ryle (1949), Peter Strawson (1963, 1992), or Frank Jackson (1998), but also—and often more interestingly—*normative* conceptual work.[1]

The claim that metaphysics centrally involves both descriptive and normative conceptual work can itself be read as either a descriptive or normative claim. Taken in the descriptive sense, the claim is that a great deal of historical and contemporary work in metaphysics can be aptly *interpreted* as engaging in (descriptive and/or normative) conceptual work. Often this must be a matter of 'deep' interpretation, since many metaphysicians think of themselves as discovering worldly facts when they do metaphysics. Nonetheless, many classic debates in metaphysics can be understood as—at least *implicitly*—involved in conceptual analysis or in negotiating what terms or concepts we should use. As I have argued elsewhere (2016, 2017a), seeing many past debates as implicitly involved in negotiating for what terms or concepts we should use often makes better sense of what competing metaphysicians are *doing* in their debates, and gives us new tools with which to assess and make progress on those debates. But I don't aim to make the case here that a great deal of historical and contemporary work in metaphysics can be so interpreted. While the job has been started elsewhere,[2] continuing work on it would require detailed case studies that must be left for other occasions.

---

[1] This echoes suggestions by others—classically, Carnap; more recently, Simon Blackburn, who writes that "just as the engineer studies the structure of material things, so the philosopher studies the structure of thought" (1999: 2). More recently still, similar views have been defended by David Plunkett (2015) and Matti Eklund, who suggests more broadly that "Philosophy should . . . be thought of as *conceptual engineering*" (2015: 36). Eklund argues that philosophy could be seen as conceptual analysis, but that it would be better to take it as conceptual engineering in the sense of "a study of what concept best plays the theoretical role of our concept of [e.g.] truth and what features this concept has . . ." (2015: 376).

[2] A. J. Ayer makes the case that Locke, Berkeley, and Hume can be seen as engaged in (descriptive) conceptual analysis (1946: chapter 2). In my view, not only ordinary language philosophy but also phenomenology can easily be understood as engaged in descriptive conceptual work (see my 2007)—including under this heading transcendental conceptual work that examines what concepts we *must* have

Taken in the normative sense, the claim is that we *should* come to understand the legitimate work of metaphysics as (descriptive and/or normative) conceptual work. If the case *can* be made that much historical work in metaphysics involves conceptual work, so much the better for the proposal that *we should* come to understand the proper work of metaphysics in that way. For that would lend plausibility to the idea that, even if this doesn't match the explicit self-conception of many contemporary metaphysicians, it maintains enough continuity with what both historical and contemporary metaphysicians have been *doing* to deserve the title. As I have argued elsewhere (2016) understanding the *legitimate* work of metaphysics in this way brings other advantages, too, such as avoiding epistemic mysteries and apparent rivalries with science.

But whether we take the claim that metaphysics centrally involves normative conceptual work as an *interpretation* of what is implicitly at stake, or as a *proposal* for how to understand what metaphysics legitimately *can* do (or both) an important question immediately arises. If metaphysics centrally involves normative conceptual work, how ought we to be doing it? What methods and standards should we employ?

While the idea that work in metaphysics may involve normative conceptual work has begun to gain some traction, far less has been said about *how* that work is to be done. Herman Cappelen aptly notes that we should be interested not just in how conceptual revision has been done, but "We should also be interested in how [conceptual revision] *should* be done—what *should* be considered successful" (2018: 119). Answering the methodological question of how it should be done is important. As Alexis Burgess and David Plunkett put it:

if certain views about metalinguistic disputes are on the right track, we actually already engage in normative argument about representational choices much more often than one might realize. We would therefore do better to pursue these issues wittingly, overtly, and with greater care.   (2013b: 1091−2)

And again:

If we already practice conceptual ethics, let's do it well.   (2013b: 1097)

But what sort of care should we take—what is it to do normative conceptual work well? Presumably there must be constraints, better and worse ways of engaging in this sort of conceptual choice. We do often find *critiques* of certain concepts plausible and powerful—whether critiques like Foucault's (2006) that show the uses of certain

---

to be able to reason or experience the world at all. Some more recent work on what it is to be a person, a work of art, or artifact can also be seen as engaged in descriptive conceptual work. Much revisionary work in recent metaphysics can be seen as engaged in normative conceptual work. Mark Richard (forthcoming) argues that many philosophical debates, for example, concerning knowledge, free action, truth . . . can be seen as engaged in something like conceptual engineering. David Plunkett makes the case that many philosophical disputes can be interpreted as not canonical disputes but rather cases of metalinguistic negotiations (2015: 861). Some even make it explicit that they are advocating for adopting a new conceptual scheme rather than analyzing the old conceptual scheme or discovering some worldly 'essence'. See, for example, Haslanger (2012) on race and gender; Millikan (1984) on function; Bell (1914/1977); and Davies (2004) on art. This also lies just beneath the surface in the debate between Price (2013) and Brandom (2013) about how to use the term "representation(alism)".

concepts to be driven inappropriately by power relations, or by feminist philosophers of science in showing how certain concept choices in the sciences have been driven by sexist assumptions, or critiques of past or current concepts of race (Appiah 1992). To the extent that concepts are open to critique, we must presuppose that there are at least some standards for evaluating concepts: standards that are not lived up to in these cases.

Comparatively little (that I have been able to find) has been said about how this sort of normative conceptual work should proceed. Burgess and Plunkett (2013b) formulate the question, list a range of goods (clarity, consistency, naturalness, social justice . . .) and goals (cooking, playing soccer, doing mathematics) that might be relevant in doing what they call 'conceptual ethics'. They also note that goals may play a central role in determining what goods matter, and what norms we should follow in conceptual ethics—"it's natural to think that different *goals* give rise to different norms" (p. 1105). But—as is appropriate in their agenda-setting paper— they leave most of these questions open, and don't propose anything like a unified approach to addressing problems in conceptual ethics. Even Peter Ludlow, who argues that normative negotiations of meanings are pervasive in everyday speech, says little about how these *should* be conducted, writing:

Application conditions for terms and phrases like 'murder' 'life' 'family values', and 'good character' must be fleshed out and precisified, and it would seem to be a mistake to just blindly follow our neighbors or the powerful for their precisifications. Here is a place where we want to insist on deliberation and good reasons for a choice.   (2006: xx)

What might such good reasons consist in? Here we get little guidance. Matti Eklund, similarly, having argued that philosophy is best seen as conceptual engineering, admits, "Obviously some big questions in the vicinity remain entirely unresolved, such as what the proper methodology for conceptual engineering is" (2015: 382).[3]

1. The Metaphysical Approach and the Pragmatic Approach

In this chapter I aim to sketch a method for approaching normative conceptual work, with an ulterior motive in view. My own interest in the idea that metaphysics can be understood as (to be) engaged in conceptual work—including normative conceptual work—comes from seeing this position as a way of adding strength and depth to the deflationary metametaphysical position I have argued for elsewhere (2015). Call a metametaphysical approach 'deflationary' if it relies on nothing more mysterious than (normative and descriptive) conceptual work—perhaps combined with straightforward

[3] Some readers may have noticed that some of the relevant works mentioned above speak of concepts, while others speak of language. There is as of yet no consensus on whether one should speak (primarily) of engineering concepts or language. I am inclined to think that it is language that is primary—for (parts of) language we can clearly see as historical and changeable and as having functions. Nonetheless, speaking of concepts is a useful way of abstracting from phonetic and orthographic features of words that are typically irrelevant to the engineering at issue in philosophical work. But I will not address these issues here, and will continue to speak of both below.

empirical work—in answering those metaphysical questions that are well formed and fit to answer. Call a metametaphysical approach 'heavyweight' if it presents metaphysics as involved in discovering deep facts about the world that are not knowable by employing straightforward empirical and/or conceptual means—but are rather (in the words of Theodore Sider) 'epistemically metaphysical' (2011: 187). By acknowledging the role of *normative* (as well as descriptive) conceptual work, I have argued (2016), we can retain the epistemic clarity of a deflationary metametaphysical approach and yet also respect the depth, worldly relevance, and difficulty of many metaphysical disputes.

But accepting that a great deal of the interesting work in metaphysics has involved and should involve normative conceptual work does not *commit* one to a kind of metaontological deflationism, as Plunkett makes clear. Having argued for the related point that many philosophical disputes (including disputes about ground, supervenience, and essence) turn out to be cases of metalinguistic negotiation (2015: 859−60), he goes on to emphasize that:

Everything I have said is consistent with thinking that there are important issues [in metaphysics, concerning ground, supervenience, real definition, essence, etc.] here to investigate. To see this, consider the following. Suppose one argued . . . that an important part of communication among biologists involves metalinguistic negotiation. (The different meanings of 'species' is a good place to start with such a proposal, as is the different meanings of 'intelligence'). Would that mean that there aren't facts about animals and their behavior to investigate, and that all biological argument is just about normative issues about word and concept choices? Clearly not.   (2015: 860)

For, as he (aptly) emphasizes, our views in conceptual ethics are generally tied up with commitments on a range of first-order normative and descriptive claims about the world.

Once we embrace the idea that much metaphysical work involves normative conceptual work, the question of whether we should adopt a deflationary or heavyweight metametaphysical approach comes down in large part to the question of what methods we can and should use for undertaking normative conceptual work.

One might take either of two broadly different orientations to answering this methodological question—one of which is tied to a heavyweight approach, and the other of which is consistent with a deflationary approach. Either approach could no doubt be fleshed out in many ways, and there may still be other approaches to consider.[4] Call the first of these approaches the 'metaphysical' approach. Someone who takes the metaphysical approach to conceptual ethics, in broad strokes, is someone who accepts that many debates in metaphysics may be seen as negotiating which terms, concepts, or conceptual scheme we ought to use, but who thinks that it

---

[4] One might also take a transcendental or Kantian approach, for example, to the question of what *basic* concepts we must have—say, in order to have cognition of objects at all. (Thanks to Lanier Anderson for bringing this option to my consideration, and to Jessica Leach for further discussion.) One might be able to consider a transcendental approach as part of the pragmatic approach—but a part on which the choice of certain concepts is determined by certain requirements that are non-optional (and so not 'merely pragmatic', in the everyday sense). This possibility suggests that, in taking a pragmatic rather than metaphysical approach, one need not be committed to the view that all conceptual choices are *merely conventional*, or built on contingent practical needs.

is the metaphysical facts of the world that provide the primary standard, for example, suggesting that we should aim to choose concepts that correspond to things that exist, carve the world at its joints, or correctly reflect the essences of the things described.[5] Such views remain meta-metaphysically heavyweight insofar as, even while acknowledging that some of the central work of metaphysics involves normative conceptual work, they hold that this ought to be driven by the metaphysical facts about the world—where these are not discoverable by straightforward empirical work or conceptual analysis. I have argued against such metaphysical approaches to conceptual choice elsewhere (2017b), on grounds that they leave us with epistemological mysteries that are hard to resolve.

In any case, it is clear that the metaontological deflationist cannot say that our normative conceptual choices ought to be determined by the metaphysical facts of the world. Nor is the metaphysical approach open (on any topic) to those who are deflationary *about that topic*. For example, Sally Haslanger notes that on her approach "the world by itself can't tell us what gender is, or what race is" (2012: 224). Similarly, metaethical deflationists cannot treat our conceptual choices in ethics as requiring guidance from heavyweight ethical facts.[6] Thus—given the difficulties faced by the metaphysical approach, as well as the need many will have for a non-metaphysical approach—there is good reason to investigate the prospects for a non-metaphysical approach to conceptual choice.

Those with heavyweight metametaphysical inclinations are prone to see a threat here: if the metaphysical facts of the world cannot determine which of our concepts or terms we ought to use, one might worry that we will be unable to account for intuitions that the world has *structure*, and that accordingly some conceptual choices (e.g., lithium rather than *lithium on earth*; fish and mammal rather than sea creature and land creature; green rather than grue) are *objectively better* than others (Sider 2011: 18–19). But, as Theodore Sider puts it, "It is really, really hard to believe that the fact that electrons go together, in a way that electrons-or-cows do not, is merely a reflection of something about us" (2011: 18). Not only might we be left unable to account for intuitions that the world has structure; a broader worry arises. That worry is that if we do not take a metaphysical approach to conceptual choice, our conceptual choices will be merely arbitrary, not constrained by the world. As David Plunkett describes the worry (without endorsing it), this view might be thought to suggest "that normative issues about what concepts we should use can be settled by voluntary choices that we ourselves make" (2015: 860–1). Finally, the deflationist might seem unable to make sense of the plausible idea discussed above, that there are standards governing conceptual choice, which justify us in critiquing certain concepts.

In short, then, abandoning a metaphysical approach to conceptual choice leads to three interrelated worries: (1) That we will be unable to account for intuitions of structure; (2) That we will have to treat conceptual choices as merely arbitrary, not

---

[5] Peter van Inwagen (2016: 17) suggests something like the view that the metaphysician's conceptual choices ought to be driven by facts about what exists, while Theodore Sider (2011) develops and defends the view that (at least when doing fundamental metaphysics) our conceptual choices should be held to the standard of carving the world at its joints or mapping 'structure'.

[6] Thanks to David Plunkett for emphasizing this point.

worldly; and (3) That we will be unable to critique conceptual choices. If the deflationist is really saddled with these problems, it threatens to leave the deflationist's view of metaphysics (as engaged centrally in normative conceptual work) inherently uninteresting, or even suspect—with each practitioner simply attempting to wield her power in imposing her own perhaps idiosyncratic conceptual choices, incapable of worldly validation.

But these fears are misplaced. There is a far better option available for the deflationist to take: namely, adopting a pragmatic (rather than metaphysical) methodology for normative conceptual work.

We can think of normative conceptual work as involving two projects. One is conceptual *engineering*: that is, holding in place some function or functions the concept is to serve, we may aim to redesign old concepts to serve that function better, or to engineer new concepts that can serve a function that was done imperfectly or not at all by our prior conceptual scheme. But we can also engage in work in conceptual *ethics* proper. Burgess and Plunkett take conceptual ethics broadly, as involving 'normative and evaluative issues about representation' (2013a: 1091), including deeper normative questions regarding what functions our concepts *should* serve, which functions we should pursue and abandon, and more generally "which concepts [we should] use to think and talk about the world" (Burgess and Plunkett 2013a: 1091). This work in conceptual ethics can also be undertaken at two levels. First, keeping fixed some goals we wish to fulfill or see as desirable, we can ask what functions our concepts should serve, to help fulfill these goals. This will issue in hypothetical imperatives: if you have these goals or purposes, then you should employ this range of concepts (using them in this way). Second, there are deeper questions one can raise about which goals we should adopt, and so about which concepts should we employ, all things considered (not just given some presupposed range of common goals and interests). I will leave these questions to the side here. For one thing, addressing such questions would bring us into deep metanormative issues that are beyond the scope of this chapter. For another thing, we (fortunately) do not need to appeal to such deep normative claims to do the work needed here. For I will argue that, *presupposing various widely shared goals that are generally presupposed in debates about what concepts to use*, we can fully account for ordinary intuitions about structure, as well as for the non-arbitrariness, worldliness, and openness to critique of our conceptual choices. In that way (even without commitment to any deep normative claims), we can account perfectly well for the metaphysician's driving intuitions that that some conceptual choices are better than others, and give a picture on which (given those shared interests) conceptual choice is not at all arbitrary, and is open to critique.

The key to developing a plausible pragmatic approach, as I will argue below, is to take the functions of our (ranges of) concepts as playing a central role. As I shall argue below, when we aim to engage in normative conceptual work, we must pay close attention to the purpose or function that is (to be) served by the relevant term, concept, or conceptual scheme.[7] To engineer a concept well, we must attend to its

---

[7] In asking these functional questions, we need not take a stand on whether the concepts and terms we are investigating are natural products of evolution, explicit artifactual creations, or something in between

function. To do first-level conceptual ethics well, and determine what concepts we should adopt, to meet goals that we have, we must attend to function.

By appealing to the idea of function, I will argue, the pragmatic approach can respect our ordinary intuitions about structure, give constraints for conceptual choice that ensure that conceptual choice should not be seen as merely arbitrary or subjective, and must be responsive to worldly constraints (though it is responsiveness to empirical facts, not special metaphysical facts, that is needed), and leave room for critique. As a result, I will argue, those inclined to think that the work of metaphysics centrally involves normative conceptual work may combine that with a plausible methodology consistent with a deflationary metametaphysics. The prospects for a pragmatic approach to conceptual choice are far better than its critics have suspected.

2. A Defense of Function

But although it seems natural to think that our concepts (or perhaps better: ranges of concepts) serve certain functions, and that these are relevant to the projects of conceptual engineering and conceptual ethics,[8] Herman Cappelen rejects this function-driven view, denying that concepts or words have functions: "I don't think concepts have purposes and certainly not words (or extensions or intensions)" (2018: 180). Cappelen argues that the appeal to the function or purpose of a concept doesn't "do any work"—in particular, that it fails to provide an adequate answer to Strawson's challenge against conceptual engineering. Thus, before going on to develop the function-driven pragmatic approach below, we should confront these doubts.

Carnap aimed to replace certain everyday concepts with "exact and fruitful" concepts for use in the formal or empirical sciences. Strawson's challenge against Carnap is that any attempt to solve traditional philosophical problems involving concepts used in non-scientific discourse by engaging in conceptually engineering exact scientific concepts would not be "to solve the typical philosophical problem, but to change the subject" (1963: 506). For "the clarification of philosophically puzzling concepts is not achieved by the introduction of related scientific concepts" (p. 506). For, Strawson insists, shifting to the new scientific concepts may simply fail to address the old philosophical perplexities (pp. 504−5).

Cappelen interprets Strawson as presenting a general challenge to the tenability of conceptual engineering, generalizing it as follows:

Generalized challenge: Change of extension and intension ... is a change of topic, so revisionary projects [in conceptual engineering] are bound to fail. Even if the revisions succeed, they do not provide us with a better way to talk about what we were talking about; they simply change the topic.   (2018: 100)

(as with the nests of birds or cities of ants). In any of these cases, we can profitably ask questions about the function of the heart, the toaster, or the nest.

[8] This idea has also been suggested by others working on conceptual engineering, including Brigandt (2010) and Haslanger (2012).

For if we change the conditions that need to be satisfied in order for something to fall into the extension of our concept as we engineer it, the objection goes, we can't claim to have *improved* the concept—we will have changed the concept and thereby changed the topic in such a way that we can't even be answering the questions posed using the original concept. Mark Richard (forthcoming) argues similarly that one can't improve on a concept by changing its intension and extension, since concepts possess these essentially, ensuring that any such changes leave us with new concepts rather than improvements of the old concepts. The challenge, then, is to find a way to understand sameness of topic in such a way that one can allow that conceptual engineering enables us to improve our concepts without changing the subject.[9]

Cappelen takes the question "does conceptual engineering always involve topic revising...or can conceptual engineering in some cases preserve topic" as "a central question for anyone interested in conceptual engineering and its foundations" (2018: 97). One of his recurrent objections against the view that words or concepts have function is that, he claims, appeal to a function won't help answer this question.

The first thing to say in response to the generalized challenge is that we shouldn't take the challenge (as generalized by Cappelen) with high seriousness and feel pressed to search for a univocal answer—to whether we "really" have the same concept or topic as before. The generalized challenge presses us to say when terms are the same, when concepts are the same, when a topic of conversation is the same. But it is entirely coherent with the conceptual engineer's point of view to think that the key terms here: 'term', 'concept', and 'topic', and the like, are themselves underspecified in ordinary language and up for metalinguistic negotiation and re-engineering. In general, as I have argued elsewhere (2017a) (following Plunkett and Sundell 2013), debates about what is/is not essential to Ts are often debates (disguised in the object language) about how we ought to use the relevant terms ('T'). The deepest, though not most direct, response to the generalized challenge is to urge that we not presume that there is an objectively correct 'discovery' of what does/does not count as sameness of topic, concept, or term. What we count as sameness of concept or term may aptly be engineered or negotiated differently depending on the purposes we have. Sometimes (say, in doing etymology) we may wish to track historical continuity; sometimes we might need track sameness of extension, inferential role, or even phonetics or lexical effects...In other cases, including many of those centrally at issue for conceptual engineering, we may want to track sameness of function. In short, the best response seems to me to be a bootstrapping response that begins by asking what function we want the relevant terms (including 'same word', 'same concept', or 'same topic') to serve and presses for a view that will do that job well.

The purpose that is at issue in Strawson's challenge is to give us a way of understanding *concept* that can preserve the sense in which people are 'talking about the same subject' over time—not just 'changing the subject'. To do this, it seems we might do better to look to function and historical continuity in individuating concepts than to

---

[9] Cappelen himself does this by an appeal to same-saying (2018).

rely on precise intensions and extensions. For appealing to function provides a promising way of giving a sense in which we remain on topic across changes in intension and extension—a sense in which we aim to solve the same problem, or to pursue the same goals.[10] Consider, for example, recent revision of our concept of marriage to include same-sex couples: how can we consider it an improvement to our old concept of marriage rather than a simple change of subject? One way to do so is by appealing to the continuous function (or functions) the concept of marriage was to serve. Why it is useful to have a concept like *marriage*: what legitimate role(s) might it play (perhaps along with other social concepts) in our overall conceptual system, and what would we be missing if we lacked such a concept? If we suppose that one legitimate and desired function of a concept of marriage is to mark a range of close relationships that we would help protect by affording a special legal and social status (tied up with some 3,000 relevant legal obligations and entitlements in the U.S.), then one can see that function as served—and served even better—by extending the criteria to include same-sex relationships that otherwise are similar in character to those previously included in the extension. In that way, we can see the change as a conceptual improvement, rather than a mere change of topic. So thinking of concepts or words in functional terms provides ways of legitimating the feeling that we haven't simply 'changed the subject' when we engage in conceptual engineering.

Nonetheless, Cappelen opposes appealing to function to do this work. He initially considers and rejects two versions of a functional view:[11] Haslanger's appeal to the 'central functions of a concept' (2000: 35), and Brigandt's appeal to a concept's epistemic goals (the kinds of inferences and explanations) a concept was intentionally introduced to serve (2010).[12] Haslanger argues that shifts in the meaning of a term are semantically justified "if central functions of the term remain the same, e.g., if it helps to organize or explain a core set of phenomena that the ordinary terms are used to identify or describe" (2000: 35). Against this, Cappelen argues that the only non-controversial approach to identifying the relevant 'core set of phenomena' is disquotationally: for example, that the concept of salmon is to identify or describe salmon (2018: 183). But this clearly won't help with answering the generalized challenge, of saying how we could still be talking about the same things, after we have revised a concept and changed its extension. Otherwise, Cappelen suggests, we simply need more guidance about how to identify the relevant phenomena a concept was to identify or describe—and Cappelen adds a suspicion that "there simply *isn't* a good way to identify 'the phenomenon' except disquotationally and the disquotational identification is unresponsive to the challenge of articulating the limits of revision" (2018: 184).[13] The other approach to identifying conceptual function that

---

[10] Plunkett and Sundell (2013), and Warren (2015), provide a good account of this for the moral case.

[11] Cappelen also considers contextualist approaches that hold that the function of a concept varies from concept to concept, but since this is clearly a non-starter for solving Strawson's challenge of accounting for sameness of topic, I will leave it to the side here.

[12] Cappelen attributes this view to Brigandt, though noting that Brigandt makes only the limited claim that some scientific concepts have such epistemic goals—not that all concepts, or even all scientific concepts do.

[13] To be fair, Haslanger does point to ways of identifying relevant phenomena in the case of race and gender, including the need to "identify and explain persistent inequalities between females and males, and

Cappelen considers is from Brigandt, who appeals to the idea that conceptual change in science can be understood as rational by appealing to its epistemic goals: the kinds of inferences and explanations the concept is intended to support. This of course, as Cappelen notes (2018: 185), is too narrow to apply generally as a view of the function of a concept. One might attempt to develop this view more broadly by identifying the function of a concept with whatever function it is intended to serve. But this will clearly be problematic as well, since few concepts are intentionally designed at all: any that are innate or that gradually evolve in a community rather than being explicitly designed and introduced for a purpose will not be intentionally designed to serve any function.

This needn't be a worry, however, for those inclined to appeal to function in laying out a method for conceptual ethics. We need only attend to the recent philosophical work on function to see that there are more plausible options available, which don't identify function with intended function, and aren't left with a mere disquotational story about the function of words or concepts. There have been two large trends in understanding the notion of function in post-Darwinian biology: in terms of evolutionary/selection history (a historical story about what the ancestors of such things did that accounts for their reproduction and survival) ('proper function'), and in terms of a thing's current capacities/dispositions, with particular attention paid to the role such things play in the overall system in which it is embedded (what Beth Preston calls 'system function') (1998: 221). In neither case do we need to identify function with what anyone intends the function to be.

The notion of proper function has been most famously developed by Ruth Millikan, who is not concerned merely with biological functions, but rather explicitly aims to identify functions of 'language devices' that are "not found either by averaging over idiolects or by examining speaker intentions" (1984: 4). On Millikan's view (roughly), a member of a 'reproductively established family' has as its proper function whatever its ancestors did that contributed to the reproductive success of the family, which contributes to explaining the existence of that member (1984: 28). Millikan explicitly applies the view to cultural products, including language, as much as to biological entities such as hearts and lungs. Meaningful linguistic devices, on her view, are also members of "first-order reproductively established families" (p. 29), and Millikan argues that "language devices must have direct proper functions at some level or levels. It must be because they correlate with functions that they proliferate" (p. 31). This is clearly a view on which the (direct) proper function of a concept or term need not be identified with anyone's intentions or beliefs about what the function is, nor with actual dispositions of speakers to use the term in certain ways, or an average over the relevant ways in which, or purposes for which, it is used. It is thus a view that avoids Cappelen's objections.

The other dominant approach to function is to adopt a 'system function' view like that defended by Robert Cummins. On this model, the function of an item is whatever it *does* within the system as a whole—whatever its current capacities

between people of different 'colors'", and to be "sensitive to both the similarities and differences" among people considered to be male/female or of different 'colors' (2000: 36). So, it seems further discussion should focus on these suggestions, and whether and if so how more general guidance can be given.

contribute to the capacities of the whole, so that we can give a compositional analysis "of the capacities of containing systems in terms of their component parts" (Preston 1998: 225). As Preston argues, these two notions of function—as proper function and as system function—must not be conflated, but we also need not consider them rivals. Instead, we may need to recognize them as "equally important for a viable general theory of function"—whether we are concerned with the functions of biological entities or of artifacts (p. 226). Typically new proper functions begin life as *system* functions: it is because these entities can do something for the system that they (and their later copies) tend to be kept around; that is why identifying what a range of vocabulary *can* do (its system function) can be a useful tool in determining its *proper* function.

Cappelen, however, also rejects the idea that we can usefully identify function in these ways.[14] He argues, first, that we can't identify the functions of vocabulary, considered as *words*, by asking what it "enables us to do that we couldn't do (or couldn't do as effectively or efficiently) without it". For we could clearly do the same work by exchanging one symbol ('1') with another ('2') (2018: 187). But this misses the overall point. The point is to ask (e.g.) what *nominative number terms* (or ethical terms, or a truth predicate) do for us that we couldn't do, or couldn't do as effectively or efficiently without them. The question clearly isn't a matter of what this typographical shape type can do for us. Nor are functional analyses that identify what hearts or forks do for us, which might account for their being reproduced, undermined by imagining that the same work might be done by an artificial heart or a differently shaped piece of silverware. The relevant counterfactual asks us to evaluate what we couldn't do as effectively without hearts or forks, *holding other aspects of the background context in place*—not while we vary the background by providing a substitute to do the job. In any case, we can clearly avoid the above worry by rephrasing, asking: what a range of vocabulary "enables us to do that we couldn't do (or couldn't do as effectively or efficiently without it, *or an apt translation*)".[15]

While Cappelen expresses some willingness to accept that we can identify functions "by looking for what makes [terminologies] useful for us (and hence perpetuated in our culture)" (2018: 187), such functions, he suggests will be nothing more than disquotationally specified functions, such as "the reason 'salmon' is useful for us is that it can be used to talk about salmons (or denote salmons)" (p. 187). For "beyond these disquotationally specified functions, there's variability" in how the terms are used in different speech acts (p. 187).

But this response misses two important points. The first is a functional pluralist point, which (following Huw Price (2011)) I have emphasized elsewhere: that we shouldn't assume that all terms have as their function (or even among their central functions) to track or denote entities of a certain kind.[16] As I discuss in section 3 below, many of our most philosophically interesting terms (such as mathematical,

---

[14] What follows in this section responds to the published material in Cappelen (2018), which was criticizing material found earlier in an earlier draft of this chapter.

[15] Another option is to speak in terms of concepts rather than words here. I aim to remain neutral here regarding which provides the best approach.

[16] The tracking or denoting function is what Price (2011) calls an 'e-representational function'.

moral, and modal terms) may plausibly serve very different functions from this kind of denoting or tracking function that can be given disquotationally.

The second important point is that variability in how something is used does not entail that nothing informative can be said about its function. The parallel argument would never be accepted for holding that biological or artefactual entities can't be said to have functions, in a way that can be given substantively and informatively. What is the function of a dog's mouth? There are a great many things that can be, and have been, done with them—but that does not show that the mouths of dogs can't be said to have a biological proper function, or functions. The case is even clearer for artifacts: there are a great many things that can be, and have been, done with forks, or screwdrivers (they are exapted in all sorts of ways), but that does not show that forks or screwdrivers do not have proper functions (Preston 1998). For a proper function is not identified just by looking to anything that *can be* or *has been* done with the item in question. If we think of a language, and the terms in it, as human creations—as artifacts (abstract artifacts, in the sense I have elsewhere articulated (1999)), it is natural to think that linguistic items, too, may have proper functions, identifiable separately from the diverse uses to which they are put. I have elsewhere[17] distinguished the 'practical significance' or proper function of a range of vocabulary (the function it serves in our overall linguistic apparatus, which explains why it is useful to have vocabulary like that in place), from that of its use in different speech acts (what it is used to do on particular occasions). Even where uses vary, a more stable proper function or functions may be identified.

In any case, I want to leave open what view, precisely, of function should we adopt with respect to the functions of our terms (or ranges of vocabulary) or concepts. That is a major topic for discussion in its own right. A great deal of work remains to be done in determining how we should understand the notion of linguistic function most relevant to normative conceptual work, and how we can best discover the relevant functions of our concepts, terms, or ranges of vocabulary.

The central point for now is that we can legitimately maintain that our terms and concepts have functions, without thereby having to think of all functions as intentionally endowed, and without having to limit them to functions that can be understood disquotationally. I have tried to at least point towards some familiar ways of understanding 'function' that might help us do the job here. With that defense in place, we can return now to our story—and try to utilize this appeal to function as a way of sketching a pragmatic method for undertaking normative conceptual work.

3. A Pragmatic Method for Normative Conceptual Work

Here I aim to sketch the beginnings of a pragmatic approach to normative conceptual work—and to show that such approaches have every prospect of accounting for core

[17] In my *Norms and Necessity* (forthcoming), following Michael Williams (2010), who uses slightly different terminology.

intuitions about structure, and of avoiding accusations that such approaches must leave normative conceptual work arbitrary, subjective, or insusceptible to critique. The key, as I have already suggested, involves beginning from an appeal to function. As Strawson puts it, "The kinds of concept we employ are not independent of the kinds of purpose for which we employ them" (1963: 506).

While Carnap was interested in conceptual engineering, primarily in the sense of devising new, technical languages, most work in metaphysics (traditional and contemporary) does not involve devising new terms, but rather working with, and making normative choices regarding, common terms of our long-familiar vocabulary. Thus, if we think of metaphysics as engaged in conceptual negotiation regarding terms such as 'freedom', 'person', 'art', 'good', 'responsible', 'number', 'property', 'species', and the like, then we must acknowledge that these are terms that are *already* part of our shared vocabulary and conceptual scheme—not terms we do or can engineer on a blank slate. So how should we begin?

3.1. Reverse Engineering

Revisionary work has been increasingly popular in metaphysics, sometimes including recommendations that we do away with certain ranges of vocabulary—often in order to 'avoid' 'problematic ontological commitments'. But before removing a piece of a car engine, lines in a software program, or an organ from the body, it is always a good idea to begin with reverse engineering: working out what the part does for the engine, program, or organism as a whole. Thus, in conceptual engineering as it is (to be) practiced in philosophy, we must often begin not with simple constructive conceptual engineering, but rather with 'reverse engineering'.[18]

In some cases, we may get important clues about the functions a range of vocabulary has served by engaging in conceptual genealogy—looking back to when and why the term was introduced, how it has been used, and what functions it served in its original and later historical contexts.[19] As David Plunkett argues, "conceptual history can help us when we engage in conceptual ethics" (2016: 59). Understanding conceptual history can not only help us determine what functions those concepts have served, but also (Plunkett argues) might help in (re-)evaluating our purported justifications for using those concepts, help us in engaging in conceptual analysis, and thereby aid us in determining which concepts will be the most helpful for a given sort of inquiry.

But whether one aims to discover the function of a term or of other artifacts, such historical information is not always available, and gives only defeasible clues about how the relevant item *currently* functions. Even in those cases, however, one may engage in reverse engineering. But how is reverse conceptual engineering to be done? Consider how one approaches reverse-engineering a piece of software—say, a discovered piece of malware. One may have no access to historical information about how and why it came into being. Instead, one proceeds first by investigating what it

---

[18] I am indebted to David Sanford for suggesting this point.
[19] For an excellent overview of the role of conceptual genealogy in philosophy, see Dutilh Novaes (2015). For a general defense of the idea that conceptual history may play a useful role in undertaking work in conceptual ethics, see Plunkett (2016).

does and can do, and thereby gains clues to determine what (system) function(s) it serves and how it serves them. So similarly, in engaging in conceptual engineering one may aim to engage in reverse engineering the concepts in question—aiming to determine what they do or can do—why it is useful to have such (ranges of) concepts at our disposal, what we can do better with such a range of concepts than if we lacked it. For that identification of system function might provide an explanation of why ranges of vocabulary to express these concepts emerged and persisted—providing evidence of their proper function.

It is not hard to find philosophical analyses of functions that fit in these molds. Consider, for example, Stephen Yablo's (2005) analysis of what we can do using mathematical discourse that we couldn't otherwise. On his analysis, introducing noun terms for numbers enables us to simplify our statements of laws in certain effective ways—so that we can state in finite form laws that otherwise would take an infinite series of infinitely long sentences. Or consider the expressivist's analysis of the point of moral discourse, as enabling us to express and coordinate our attitudes in ways that put pressure on certain forms of agreement that thereby enable us to better live together. Paul Horwich's (1999) view of the role that the truth predicate serves as a device of generalization can also be understood in this light. In each case, these analyses purport to identify something that this range of concepts *does* or (better) *enables us to do*, that we couldn't do (or couldn't do as effectively or efficiently) without it (or an apt translation). Analyses like these can serve as clues to proper function analyses: to why it would have been useful to have concepts like this, why terms that express them might have been perpetuated in our culture. Nonetheless, it is important to emphasize that saying that a concept has a function is not to say that it is overall beneficial, aids the general utility, or anything of the sort. Some concepts, or ranges of concepts, may serve a function not *for us*, but rather for *some*: perhaps for those in power, who have the power to keep them in use. Of course, this does not make terms different from other artifacts—say, weapons, poisons, fences, elitist school systems, gendered clubs—which may serve functions for some, without being beneficial overall.

Engaging in conceptual genealogy and/or reverse engineering may yield various important results. On the one hand, we may find a useful function a range of vocabulary has served for us. For example, if Yablo is right about the useful functions of using nominative terms for numbers, and we have purposes that would be well served by being able to state scientific laws of these sorts in finite form, then we should hesitate before we suggest doing away with the vocabulary and 'making do' with some replacement nominalist language on grounds of alleged 'ontological' concerns. (Similar points could be made about mentalistic vocabulary, property talk, truth talk, etc.) On the other hand, we may find that the terms have served an insidious function that we don't think should be served—I will come back to this when we return to discussing critique below.

3.2. Identifying the Function to Be Served

A second crucial step is to undertake work that is more explicitly in conceptual ethics: determining what functions (if any) these concepts *should* serve, and are *to*

serve going forward, given the goals and purposes we have. In some cases, this may involve no change (supposing the original concept served a function or functions that were useful, and that we still desire to be served). In other cases, it may involve proposing changes. In the most radical cases, it may involve proposing that we drop or replace a range of concepts or terms entirely, if we find that they have no hope of serving their function (say, in the case of terms for failed scientific posits such as phlogiston or Vulcan), or if we find that they have served an insidious function that we want to abandon.[20] For example, one might argue for rejecting pejorative terms from our vocabulary on grounds of their serving functions of insult and exclusion, which (given our current and public purposes) we think should be dropped. In other cases, the second step may involve a proposed shift in function, while retaining the old terms. Sally Haslanger's work on race and gender concepts makes this move explicitly. While the old concepts may have served to give artificial pseudo-scientific legitimacy to discriminatory practices, Haslanger argues that concepts in this vicinity *should* be retained to serve a different function: serving as "effective tools in the fight against injustice" (2012: 226). For without race concepts, it is difficult to address questions about the consequences being African American, say, has on one's job opportunities, educational prospects, treatment by the police, and so on—and thus hard to identify, and aim to fight against, racism.

3.3. *Engineering to Serve the Function*

Once we identify what function the relevant concepts are *to serve*, we can do far better at engaging in the third step, of constructive conceptual engineering. It is no good engineering a boat, or deciding which boats should be kept, modified, or replaced without a clear idea of whether the boat is to function in providing a fast and nimble escape from police boats, in transporting masses of heavy cargo across the wide oceans, or in undertaking exploration in the icy arctic. Some features would require repair for any boat—failure to float for example. So similarly some features of a concept would require repair regardless of purpose—such as its propensity to ensnare us in contradictions.[21] But no detailed evaluation can be made without an assessment of the functions that are to be served. Once a purpose (or multiple purposes) is/are identified, we can go on to use that in engaging in conceptual engineering—determining what sorts of rules or constraints would best (or better) enable it to fulfill its function(s), going forward.

---

[20] This also enables us to properly criticize other cases in which one might engage in conceptual negotiation with merely personal goals in mind—as with the case Burgess and Plunkett raise of Karl, a politician who (for purely personal gain of getting elected) employs traditional race and gender concepts rather than ones that would better serve social justice (2013b: 1105–6).

[21] However, Alexis Burgess' (2013) interesting arguments that we may, for example, have reasons to retain our concept of truth 'as is', even if its meaning is given by principles that classically entail a contradiction, should give us pause before assuming that even inconsistency always is sufficient to justify revising or replacing a concept.

4. A Defense of the Pragmatic Approach

Bringing functions into account enables us to develop a pragmatic approach to conceptual ethics (taken broadly as including practical work in conceptual engineering) that clearly avoids the problems often thought to plague deflationary methods.

4.1. *Preserving Core Intuitions about Structure*

First, it enables us to give due respect to the idea that the world is structured into natural kinds to which our concepts should be responsive, and to the idea that some concepts seem 'objectively' better than others. For it is plausible that some words serve something like a joint-tracking function—and are best left serving that function. The prime candidates for these 'joint-tracking' terms are the terms David Lewis and D. M. Armstrong began with in developing the joint-carving idea: predicates that aim to pick out 'natural' properties and relations. These are those that will figure in laws (Armstrong) and in our natural-scientific theories. Since these terms have the function of serving in explanatory and predictive scientific theories, which in turn aim to predict and explain, there are worldly constraints on what concepts we ought to adopt.

Given the relevant function and the constraints that come with it, it is easy to see why we do better to have in our chemical theory the current chemical concept of 'lithium' rather than a concept that would apply to lithium on earth, but not to the same chemical kind if found on Mars (cf. Sider 2011: 7). Geographic constraints in themselves are not helpful to chemical explanations and predictions, so the limited 'lithium' concept would not be as useful in a chemical theory. Similarly, it is easy to see, on these grounds, why the concepts of 'fish' and 'mammal' would be more useful than 'sea creature' and 'land creature' to serve the function of figuring in explanations and predictions in biology: more predictions of behavior, internal construction, disease susceptibility, reproduction, and so on will be facilitated by use of the former concepts than the latter concepts.[22]

I have respect for natural kind structure. I have a child with a nut allergy. It is a matter of life and death ("death in seven minutes", her allergist tells us) whether something is biologically a tree nut or is something *called* a 'nut'. It is a matter of life and death because it enables us to *predict* whether ingesting something will cause a life-threatening allergic reaction. It is not just a subjective matter whether 'tree nut' is a better concept than one that includes all and only things *called* 'nut' (including hazelnuts, peanuts, coconuts, nutmeg, and doughnuts (only the first of which is biologically a tree nut), and excluding cashews, pistachios, and almonds). That one concept but not the other is usefully and efficiently *predictive* in this way, which has life-or-death consequences, is all I need to be fully convinced that one set of concepts is objectively better. Moreover, the choice of concept of course has other useful consequences beyond predicting allergic reactions—consequences for its use in biological theory, farming, government regulations, etc. The important thing to

---

[22] On the other hand, as Sundell points out, where the concept of *fish* is employed with a different role in old seafaring contexts (aboard whaling ships, say), it is far less clear that the current biological concepts are better (2011: 14–15).

note here is that there *is* a way to justify the claim that one concept is better than another, for worldly reasons, but that our choice of concepts is vindicated *empirically*, given our shared purposes—it does not require additional *metaphysical* vindication.

The pragmatic approach to conceptual ethics can and does take into account all of these perfectly objective, worldly, empirically-driven reasons for choosing one set of concepts over another, *where these concepts are designed to figure in our explanatory and predictive theories*. And the advantages are not just in the simplicity of stating our biological theories, but also in the living of our lives, communication with others, safety, formulation of laws of state (as well as expression of laws of nature), etc. We can thus take into account our ordinary intuitions that some concepts are 'objectively better' than others, given a widely shared and generally unquestioned set of purposes: in this case, the purpose of designing theories that enable us to better predict and explain. In these cases, clearly our normative conceptual choices must also be world-responsive—in an empirical way. For when concepts (such as natural kind concepts) are designed to be useful in our empirical explanatory and predictive theories, we are thereby committed to being deferential to the world—letting experimental evidence help determine which *do* best serve in our predictions and explanations. Those concepts that function well in these explanatory and predictive roles will tend to be those we think of as 'carving at the joints' in the ordinary sense of marking those similarities and differences that are most relevant to our overall body of predictions and explanations, and so can preserve the everyday sense in which we think of the world as 'structured'.

Sider insists that "Joint-carving thought does not have merely instrumental value" (2011: 61) but rather is a constitutive aim of inquiry. But the instrumental value of employing concepts in our scientific theories that are particularly useful in predicting and explaining is sufficient to account for the ordinary intuitions about structure used to motivate the theory: that we ought to employ concepts like the purely chemical concept of water, the biological concept of mammal, and even the color concept of green (rather than grue). But it does so without positing extra metaphysical facts about 'structure', without invoking epistemic difficulties about how we could know such facts, or about why such metaphysical facts should be theory-guiding.[23] Once we can account for the worldliness and objectivity of criteria for conceptual choice in these cases, it's worth asking how very powerful are any remaining intuitions about 'real structure' that aren't accommodated in this way, and whether we need any further metaphysical vindication of our intuitions that some concepts are objectively better than others.

4.2. *Respecting Non-Arbitrariness*

However, we should not assume that all terms serve the same function: of tracking features of reality that enable us to better explain and predict.[24] Even in cases where

---

[23] On the last point, see Dasgupta (2018).

[24] Sundell accepts that 'metaphysical naturalness' serves as a norm governing the use of our scientific terms, but he too insists that "across a wide range of activities, speakers regulate their usage according to norms that are largely orthogonal to metaphysical naturalness" (2011: 10).

this is not the function, however, we can preserve the idea that our conceptual choices are not merely arbitrary, and often must be responsive to worldly constraints.

As our examples above show, many of the functions attributed to philosophically interesting concepts, including mathematical concepts, moral concepts, modal concepts, the concept of truth, and so on, plausibly serve functions very different from those of natural kind terms. Where functions vary, the criteria for evaluating, retaining, rejecting or rejigging extant concepts will vary accordingly—in Timothy Sundell's terms, there will be different 'measures of appropriateness' (2011: 15). It won't *always* be an apt criticism of a concept to say that it doesn't 'track the joints of reality' or serve in our best scientific theories, any more than it is always an apt criticism of a boat to say that it couldn't carry more than 1,000 tons of cargo. For social and institutional terms like 'married', 'citizen', 'person', or 'voluntary', or philosophically interesting and contested terms like terms for the moral, modal, or mathematical, the proper criteria for evaluation might not be whether the terms or concepts serve well in building explanatory and predictive theories, but whether they properly serve other purposes we have—say, endowing certain close human relationships with legal protections, enabling us to assign legal and moral rights and responsibilities, enabling us to coordinate our plans and attitudes, or enabling us to simplify our expressions of laws.

Even where the function of a concept is not predictive/explanatory, however, the pragmatic approach can nonetheless allow that our choices in engineering the concept are not merely arbitrary or subjective. For such conceptual choices also must be responsive to worldly factors. Consider as an example the concept of death, as examined by Bernard Gert, Charles Culver, and K. Danner Clouser (2006). They argue that the concept of death serves a variety of functions, including to enable us to determine when medical care should cease, funeral preparations should begin, survivors' benefits put into effect, and so on. Yet (they argue), there is no precise joint in nature marked by the concept of death, but rather a continuum of changes that go on in the process. Choosing, precisifying, or engineering the concept of death must be responsive to worldly matters. Certain empirical discoveries might place new pressure on our old vague concept of death. First, the old way of treating cessation of spontaneous breathing and circulation as a criterion for death comes into question with the availability of artificial ventilation, and puts pressure on finding new ways of identifying criteria for death. At the same time, the use of new and increasingly expensive medical technologies, and the critical need for organ transplantation to be done quickly, puts new pressure on determining the time of death more precisely than before, so that expensive treatments can be stopped, and organs harvested with greater chance of success for the recipient. These are empirical factors that put pressure on the old concept, and give reasons for at least precisifying the concept and altering the criteria typically used in applying it, so that it may continue to serve its functions. Where the function of a concept like 'death' involves, at least in part, enabling us to make decisions about when medical care should cease, and that medical care becomes increasingly costly or scarce, we may have a need to precisify the concept of death beyond the vaguer concept that served us well a hundred years ago (Gert et al. 2006: 284). The pragmatic approach to normative conceptual work is certainly worldly in that, to do it well, one must be responsive to worldly

constraints and new empirical situations. In conceptual engineering no less than civil engineering, the question of which design (of concept or bridge) will best fulfill the relevant function, given the requirements, does not leave room for a merely 'arbitrary' or power-driven answer, and must be addressed while being sensitive to a variety of worldly factors.[25]

   This is not the only way, however, in which our conceptual engineering work is subject to constraints that make our choices non-arbitrary. Civil engineering projects must take into account not only the function to be served, but also the constraints of the site: what the relevant land and geography are like, what the constraints are on surrounding extant structures and geographic features, etc. Similarly, when we engage in conceptual engineering, we must engage in descriptive conceptual work so that we can analyze, assess, and go on to be mindful of the multiple inferential connections our concepts bear to other concepts and practices.[26] Gert et al. again emphasize this point for the concept of death, appealing to the conceptual connections between death and a wide range of other social and personal (not merely medical) concepts as a way of criticizing the conceptual revisionism of certain medical doctors, who aimed to (re-)define death in such a way that they would be permitted to harvest organs sooner, when they would have a greater chance of success with transplantation. Such physicians, they argue, make the mistake of noticing only the connections between 'death' and other *medical* terms and practices, not the wider system of concepts and social practices in which 'death' plays a central role. Gert et al. use this example as part of a generalized argument for conservatism in conceptual change: while new circumstances (such as new medical technologies in keeping patients alive using artificial respiration, and in enabling organ transplantation) may require new precisifications of terms like 'death', Gert et al. argue:

When a term plays an important part in social and legal practices, as 'death' does, then the greater the change in the meaning of the term, the greater the likelihood that there will be significant social and legal problems.   (2006: 284)

Given the dangers of introducing confusion, distrust, and other social and legal problems in changing a common term, they defend a strong principle of conservativism regarding meaning change:

It is almost impossible to describe a situation in which it is appropriate to redefine a term with widespread ordinary use in order to change any particular medical (or even social or legal) practice, in which that term plays a significant role.   (2006: 285)

---

[25] This is not to say, however, that there will always be a uniquely best answer. While there may be some bridge designs that are far better than others, there nonetheless may be two or more that do the work (of transversing the chasm, safely supporting the intended vehicles, and staying within budget) equally well. So, similarly, in my view, we should allow the possibility that two or more different conceptual choices may (like different axiomatizations of geometry, or choices of different logical constants) serve equally well, without assuming there must be a 'worldly' fact to determine which of these 'carves at the metaphysical joints'.

[26] I suspect that this is related to the point Eklund makes as he argues that one cannot 'selectively' engineer the quantifier (or, presumably, other concepts) (2015: 380).

I think it is an under-appreciated point that conceptual engineering, no less than civil, does not take place in a vacuum, and that it is extremely important to note and be responsive to the inferential connections between the term in question (which we are considering revising or eliminating) and our other terms and broader practices.

Nonetheless, I think this is better taken as a caution than as an argument for a general principle of conservativeness in conceptual engineering. In civil engineering it may be a good—but defeasible—principle in constructing your new bridge or building to interfere as little as possible with surrounding roads and structures. But when problems get bad enough, or there are overriding social or moral purposes at stake, there are times for a more complete clearing. So similarly, though 'marriage' is connected to a wide range of social practices, those who value equality and happiness had good reason to change the legal definition to not precisify but rather expand the applicability of the term to same-sex partnerships, just as those who do not endorse racism had reasons to bulldoze the whole network of race concepts such as 'octaroon', 'quadroon', and 'mulatto' that played an influential social and legal role in former slaveholding and colonial societies. (As Burgess and Plunkett note, one question in conceptual ethics is "whether we ought to be using a given concept *at all*" (2013a: 1095)).

4.3. *Leaving Room for Critique*

This brings us to the third point: respecting the thought that concepts may be subject to critique. The functionally-driven pragmatic approach to conceptual ethics makes it clear why (and under what conditions) conceptual critique may be in place. One way critique can be appropriate, on this model, is in showing that the function of certain concepts cannot be fulfilled. If, for example, the term 'Vulcan' had the function of tracking a certain heavenly body, supposed to explain eccentricities in Mercury's orbit, then later discoveries led us to see that no term could fulfill this function—and it was time for getting this term out of our astronomical theories.

Another appropriate role for critique (closer to Foucault's work) is to show that the ostensible function of a range of terms or concepts comes apart from what it *really* serves to do and has done. For example, if race terms that were ostensibly introduced as natural kind terms—to explain and predict—have failed in that function, and have served instead to lend pseudo-scientific legitimacy to oppressive social practices, we have grounds to show that, whether our purposes were genuinely scientific or anti-racist (or both), they have gone wrong. Similarly if (as Foucault's work (2006) suggests) terms like 'madness' and 'mental illness' have served not so much to diagnose and treat as to give artificial pseudo-medical authority to practices of ostracism and exclusion, we have reason to reevaluate our attachment to and use of these concepts—provided we presuppose shared interests in transparency, human well-being, and/or inclusiveness.

In other cases, we may have reason to engage in critique once we notice that the function certain terms serve is only for a privileged few. If concepts like Hochdeutsch or 'received' English have served to reinforce and to legitimate regional and class biases, which we now seek to undermine, we will have reason to unmask and rethink

these concepts.[27] The same goes for the case above of eliminating terms such as 'octaroon' and 'quadroon', and revising the concept of marriage to meet the goals of building a fairer society.

Uncovering the functions of various ranges of vocabulary can thus pave the way to the sorts of critiques engaged in by Nietzsche and Foucault. Exposing how our terms function and for whom, where those functions can't be fulfilled or don't fit our current shared values and goals, may give us entirely non-arbitrary reason to reject or revise the concepts at hand.

5. Conclusion

Where have we come? I have aimed above to sketch a blueprint for how normative conceptual work can be done, on a pragmatic model. No doubt it requires a great deal of expansion and revision. But already here we can note some important features of the proposal.

The first is that there is every prospect of adopting a method that does not require appeal to specifically *metaphysical facts* for guidance. On this model, all that is required is both descriptive and normative conceptual work, and also empirical work. To do normative conceptual work (on our extant concepts) explicitly, we engage in reverse engineering to figure out, empirically, what function(s) the concepts *have served* and *do* serve (where these, of course, might differ), and do descriptive conceptual work in figuring out how they work and what the 'site constraints' are: how they are related to other concepts and practices. But we must also do work in conceptual ethics to determine what functions our concepts *are* to serve, going forward, given our shared purposes. Finally, we combine that with empirical work, in doing constructive conceptual engineering: determining whether (given worldly constraints) certain modifications or precisifications would better enable the term to fulfill its function. This gives us a non-mysterious pragmatic approach to conceptual ethics that may not only be defensible against objections, but even be preferable to metaphysical approaches to conceptual ethics. For the latter leave us with familiar epistemic mysteries about how the relevant guiding metaphysical facts may be discovered. Fully evaluating these problems and comparatively evaluating the two approaches requires a more extensive discussion elsewhere.[28] Nonetheless, there is at least *prima facie* reason to think that, by appealing to nothing more than empirical, conceptual, and normative work, the pragmatic approach may

---

[27] And note that to engage in this kind of critique and make it seem non-arbitrary to do away with such concepts, we need not rely on any form of moral realism—the critique may simply involve pointing out what the terms purport to do (what is the legitimation for having such terms), versus what they actually do, and leave it to readers to decide what, from a practical standpoint, we should do given that uncovering. But nor do we, as readers who share a certain outlook, think that—given the relevant uncovering—it would be simply arbitrary to do away with the critiqued concepts or revise them heavily. *Given the moral views we may be presupposed to share*, such moves will not be arbitrary at all.

[28] See my (2017) for a more thorough evaluation of the metaphysical approach to conceptual choice, and Dasgupta (2018) for arguments against the idea that we can appeal to metaphysical naturalness to guide theory choice.

retain the epistemic high ground over metaphysical approaches to conceptual choice.[29]

Most importantly for present purposes, I hope to have shown that a pragmatic approach to conceptual ethics is a viable option for metaontological deflationists who still hope to make some sense of the difficulty, depth, and value of work that has often gone under the heading 'metaphysics'—as well as being a viable option for those who are (merely) deflationary *about a certain topic*. Like civil engineering, conceptual engineering is not a matter for discovery but for invention. But also like civil engineering, that does not mean that the choices we make are arbitrary, unconstrained, merely subject to our will, or 'subjective'. Which boat, or development of a concept, will work best *given our shared goals, purposes, and situation* may often be an objective matter, once all constraints are in. Of course, this is not to suggest that there will always be a *uniquely best* solution to a problem in civil or conceptual engineering. But that is no embarrassment for the deflationist, who may recognize the value in developing a plurality of concepts to serve a plurality of functions, as well as the possibility that two or more concepts could (like different bridge designs) serve a function equally well.

The crucial point here is that, once we understand the approach better, we can easily see that the problems thought to plague the pragmatic approach are avoidable. If we can properly develop and understand an approach to normative conceptual work in this way, then even metaontological deflationists will be able to account for the intuitions that have motivated many to take a metaphysical approach to conceptual choice. Given the purposes we commonly assume in the background, we can account for the central intuitions that the world is 'structured' and that there are worldly constraints on conceptual choice, allow that our conceptual choices are non-arbitrary and that some are 'objectively better' than others, and leave room open for the critique of problematic concepts. Most importantly, even if we adopt the deflationary metametaphysical approach, and the pragmatic approach to conceptual choice that comes with it, we will have room to account for the difficulty, depth, and importance of work in metaphysics—and to do so without invoking epistemological mystery.

---

[29] Some might be tempted to worry that any claim to the epistemic high ground is illusory, however. For determining what functions our concepts *should* serve (it might be argued) requires discovering *normative* facts, about what functions our concepts *ought to* have and whether we *ought to* revise them in various ways or reject them, and whether there are sufficient moral or political reasons to justify overriding the usual site constraints. But discovering such deep normative facts (it might be thought) involves epistemic barriers every bit as formidable as discovering metaphysical facts. This epistemic problem is avoidable, however, as long as there is some acceptable and non-mysterious approach to moral epistemology. Certain reductive naturalist approaches, for example, purport to render moral knowledge non-mysterious. Another approach, which goes naturally with the functional pluralism I have advocated above, is to adopt a form of non-descriptivism, for example, seeing our moral statements as expressions of certain kinds of non-cognitive attitudes or plans (as in the work of Blackburn or Gibbard)—thereby eliminating the principled epistemic problem, and leaving us with difficult, but pedestrian, problems of coordinating our plans and attitudes and figuring out what to do. We needn't settle these contested issues here, but only note that as long as some non-mysterious approach to moral epistemology is both tenable and combinable with the deflationary metametaphysical approach, the latter can indeed retain the epistemic high ground, giving it a substantial advantage over the heavyweight metaphysical approach.

Acknowledgments

References

Appiah, Anthony Kwame. 1992. *In my Father's House*. New York: Oxford University Press.

Ayer, Alfred Jules. 1946. *Language, Truth and Logic*. New York: Dover Publications.

Bell, Clive. 1914/977. Art as Significant Form: The Aesthetic Hypothesis. In George Dickie and R. J. Sclafani (eds.), *Aesthetics: A Critical Anthology* (pp. 36–48). New York: St. Martin's Press.

Blackburn, Simon. 1999. *Think: A Compelling Introduction to Philosophy*. Cambridge: Cambridge University Press.

Brandom, Robert. 2013. Global Anti-Representationalism? In Huw Price (ed.), *Expressivism, Pragmatism and Representationalism*. Cambridge: Cambridge University Press.

Brigandt, Ingo. 2010. The Epistemic Goal of a Concept: Accounting for the Rationality of Semantic Change and Variation. *Synthese* 177 (1):19–40.

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

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

Burgess, Alexis. 2013. Keeping 'True': A Case Study in Conceptual Ethics. *Inquiry*. http://dx.doi.org/10.1080/0020174X.2013.851866

Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.

Dasgupta, Shamik. 2018. Realism and the Absence of Value. *Philosophical Review* 127 (3):279–322.

Davies, David. 2004. *Art as Performance*. Oxford: Blackwell.

Dutilh Novaes, Catarina. 2015. Conceptual Genealogy for Analytic Philosophy. In J. Bell, A. Cutrofello, and P. M. Livingston (eds.), *Beyond the Analytic-Continental Divide: Pluralist Philosophy in the Twenty-First Century* (pp. 75–108). London: Routledge.

Eklund, Matti. 2015. Intuitions, Conceptual Engineering, and Conceptual Fixed Points. In Christopher Daly (ed.), *Palgrave Handbook of Philosophical Methods* (pp. 363–85). London: Palgrave Macmillan.

Foucault, Michel. 2006. *A History of Madness*. London: Routledge.

Gert, Bernard, Culver, Charles M., and Danner Clouser, K. 2006. Death. In *Bioethics: A Systematic Approach* (2nd edn) (pp. 283–308). Oxford: Oxford University Press.

Haslanger, Sally. 2000. Gender and Race: (What) Are They? (What) Do We Want Them to Be? *Nous* 34 (1):31–55.

Haslanger, Sally. 2012. *Resisting Reality*. Oxford: Oxford University Press.

Horwich, Paul. 1999. *Truth* (2nd edn). Oxford: Oxford University Press.

Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.

Ludlow, Peter. 2006. The Myth of Human Language. *Croatian Journal of Philosophy* 6 (3):385–400.

Millikan, Ruth Garrett. 1984. *Language, Thought and Other Biological Categories*. Cambridge, Mass.: MIT Press.

Plunkett, David. 2015. Which Concepts Should We Use? Metalinguistic Negotiations and the Methodology of Philosophy. *Inquiry* 58 (7–8):828–74.

Plunkett, David. 2016. Conceptual History, Conceptual Ethics, and the Aims of Inquiry. *Ergo* 3 (2):27–64.

Plunkett, David, and Sundell, Tim. 2013. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosopher's Imprint* 13 (23):1–37.

Preston, Beth. 1998. Why is a Wing like a Spoon? A Pluralist Theory of Function. *The Journal of Philosophy* 95 (5):215–54.

Price, Huw. 2011. *Naturalism without Mirrors*. Oxford: Oxford University Press.

Price, Huw. 2013. *Expressivism, Pragmatism and Representationalism*. Cambridge: Cambridge University Press.

Richard, Mark. forthcoming. *Meanings as Species*.

Ryle, Gilbert. 1949. *The Concept of Mind*. Chicago: University of Chicago Press.

Sider, Theodore. 2011. *Writing the Book of the World*. Oxford: Oxford University Press.

Strawson, Peter. 1963. Carnap's Views on Constructed Systems versus Natural Languages in Analytic Philosophy. In Paul Arthur Schilpp (ed.), *The Philosophy of Rudolf Carnap*, Volume XI: *Library of Living Philosophers* (pp. 503–18). Peru, IL: Open Court Publishing.

Strawson, Peter. 1992. *Analysis and Metaphysics: An Introduction to Philosophy*. Oxford: Oxford University Press.

Sundell, Timothy. 2011. Disagreement, Error, and an Alternative to Reference Magnetism. *Australasian Journal of Philosophy*. DOI: 10.1080/00048402.2011.614266 (pp. 1–17).

Thomasson, Amie L. 1999. *Fiction and Metaphysics*. Cambridge: Cambridge University Press.

Thomasson, Amie L. 2007. Conceptual Analysis in Phenomenology and Ordinary Language Philosophy. In Michael Beaney (ed.), *The Analytic Turn: Analysis in Early Analytic Philosophy and Phenomenology*. London: Routledge.

Thomasson, Amie L. 2015. *Ontology Made Easy*. New York: Oxford University Press.

Thomasson, Amie L. 2016. What Can We Do, When We Do Metaphysics? In Giuseppina D'Oro and Soren Overgaard (eds.), *The Cambridge Companion to Philosophical Methodology*. Cambridge: Cambridge University Press.

Thomasson, Amie L. 2017a. Metaphysical Disputes and Metalinguistic Negotiation. *Analytic Philosophy* 58 (1):1–28.

Thomasson, Amie L. 2017b. Metaphysics and Conceptual Negotiation. *Philosophical Issues* 27:364–82.

Thomasson, Amie L. forthcoming. *Norms and Necessity*. New York: Oxford University Press.

Van Inwagen, Peter. 2016. The Neo-Carnapians. *Synthese*. DOI 10.1007/s11229-016-1110-4

Warren, Mark. 2015. Moral Inferentialism and the Frege-Geach Problem. *Philosophical Studies* 172 (11):2859–85.

Williams, Michael. 2010. Pragmatism, Minimalism, Expressivism. *International Journal of Philosophical Studies* 18 (3):317–30.

Yablo, Stephen. 2005. The Myth of the Seven. In Mark Eli Kalderon (ed.), *Fictionalism in Metaphysics*. Oxford: Oxford University Press.