Received: 6 November 2019 | Revised: 22 February 2020 | Accepted: 29 February 2020

DOI: 10.1111/rati.12265

Check for updates

WILEY

SPECIAL ISSUE ARTICLE

On the relation between conceptual engineering and conceptual ethics

Alexis Burgess[1] | David Plunkett[2]

[1]Alcove Learning Cooperative, 321 1/2 E 1st St #206, Los Angeles, CA 90012, USA

[2]Philosophy Department, Dartmouth College, 6035 Thornton Hall, Hanover, NH 03755, USA

Correspondence
David Plunkett, Philosophy Department, Dartmouth College, 6035 Thornton Hall, Hanover, NH 03755, USA.
Email: david.plunkett@dartmouth.edu

Abstract

In recent years, there has been growing discussion amongst philosophers about "conceptual engineering". Put roughly, conceptual engineering concerns the assessment and improvement of concepts, or of other devices we use in thought and talk (e.g., words). This often involves attempts to modify our existing concepts (or other representational devices), and/or our practices of using them. This paper explores the relation between conceptual engineering and conceptual ethics, where conceptual ethics is taken to encompass normative and evaluative questions about concepts, words, and other broadly "representational" and/or "inferential" devices we use in thought and talk. We take some of the central questions in conceptual ethics to concern which concepts we should use and what words should mean, and why. We put forward a view of conceptual engineering in terms of the following three activities: conceptual ethics, conceptual innovation, and conceptual implementation. On our view, conceptual engineering can be defined in terms of these three activities, but not in a straightforward, Boolean way. Conceptual engineering, we argue, is made up of mereologically complex activities whose parts fall into the categories associated with each of these three different activities.

KEYWORDS

concepts, conceptual engineering, conceptual ethics, philosophical methodology, philosophy of language

 © 2020 John Wiley & Sons Ltd

1 | INTRODUCTION

In recent years, there has been growing discussion amongst philosophers about "conceptual engineering". Put roughly, conceptual engineering concerns the assessment and improvement of concepts, or of other devices we use in thought and talk (e.g., words). Conceptual engineering often involves attempts to modify our existing concepts (or other representational devices), and/or our practices of using them. For example: it might involve an attempt to reform the concepts we have, coupled with an attempt to get others to adopt those reformed concepts in practice.

To get a better handle on what conceptual engineering is, consider what we (and many others who discuss conceptual engineering) take to be two paradigmatic cases of it.

First, consider Kevin Scharp on the concept TRUTH.[1] In *Replacing Truth*, Scharp argues that the ordinary concept TRUTH—which we currently express by the term 'truth'—is inconsistent, which in turn leads us into paradoxes like The Liar.[2] To avoid this, Scharp argues that we should—at least in certain contexts (e.g., when doing philosophy)—replace our ordinary concept TRUTH with two new ones that he has constructed. He argues that using these new concepts instead of our original concept TRUTH will help us make epistemic progress on the issues we care about in key parts of philosophical inquiry. In light of this, Scharp advocates that philosophers take up his proposal and change how we talk about "truth" in the relevant contexts. Scharp explicitly takes this to be a project in "conceptual engineering".

Turning to a different subarea of philosophy, consider Sally Haslanger's work on race and gender. In key parts of her work in this area, Haslanger proposes that we use our existing race and gender terms to express *new* concepts.[3] The concepts are ones that pick out features of our social identities: roughly, properties about how people are either systematically subordinated or privileged along certain dimensions (economic, political, legal, etc.) as a result of being imagined by others to possess particular features that purportedly reveal biological, ancestral, or other socially salient facts. Haslanger thinks that paying attention to these properties can help us better understand our social reality. Moreover, if certain people (e.g., certain activists or certain social theorists) use existing race and gender terminology to refer to these properties, this will not only help us better understand our world but also aid us in the struggle for *social justice*.[4] In light of this, Haslanger aims to get her target audience to change how they use race and gender terms.

There are many differences between Scharp's and Haslanger's projects here, even when described only in broad brushstrokes. For example: Haslanger appeals to the value of social justice, whereas Scharp does not explicitly bring in any moral/political values. Yet there are clearly some striking similarities between their critical projects. The idea of *many* philosophers writing about conceptual engineering is that "conceptual engineering" is a way of talking about the thing—perhaps a series of issues, an approach, or a methodology—that is common between (purportedly) paradigmatic cases of "conceptual engineering", such as Scharp's and Haslanger's projects, or Rudolf Carnap's work on explication.[5]

Given that the very idea of "conceptual engineering" is often understood by reference to specific cases such as Scharp's and Haslanger's, one obvious way to make progress on understanding what conceptual engineering is to carefully examine those cases, and then look at what (if anything) they have in common. Some part of what

[1]In this paper, we use small caps (e.g., CAT) to pick out concepts. We use single quotation marks (e.g., 'cat') to mention linguistic items. We use double quotation marks (e.g., "cat") for a variety of tasks, including quoting other authors, scare quotes, and mixes of use and mention.

[2]See Scharp (2013). For connected discussion, see Eklund (2014).

[3]See Haslanger (2000). Note that we are here (and in the rest of this paragraph) focusing on one influential strand of Haslanger's work on this topic, chiefly as articulated in Haslanger (2000). Importantly, not all of Haslanger's work on race and gender involves this self-understanding in terms of *replacing* our current concepts with new ones. See Haslanger (2005) for a different self-interpretation, and Haslanger (2020) for more recent reflection.

[4]Haslanger (2000).

[5]For some of Carnap's work on explication, see Carnap (1947/1956).

those cases have in common might then be used to illuminate the nature of conceptual engineering.[6] We think that is surely one of the core ways by which one should try to make progress in understanding conceptual engineering. In this paper, however, we aim to supplement this approach with a different tactic. In short, we aim to make progress on understanding conceptual engineering by better understanding its relation to conceptual ethics. For reasons that we explain briefly below, we aim to make progress on this by thinking about how we should use the terms 'conceptual engineering' and 'conceptual ethics'.

We think the relation between conceptual engineering and conceptual ethics is worth asking about for the following reason: in recent years, the term 'conceptual ethics' seems to have been used to pick out many of the *same* philosophical projects—or, more precisely, at least core *aspects* of those projects—that the term 'conceptual engineering' seems to pick out. This naturally raises the question of how conceptual engineering and conceptual ethics are related to each other. Moreover, this issue is particularly pertinent for us. This is because, in previous work, we explicitly advocated for *introducing* the term 'conceptual ethics' to describe a set of normative and evaluative issues about thought and talk.[7] We did that at the same time that Scharp and others (e.g., Matti Eklund and Herman Cappelen) were explicitly starting to use the label 'conceptual engineering' to cover a similar cluster of issues, or to refer to a philosophical approach that emphasizes the import of such issues.[8]

One difficulty we face in addressing the relation between conceptual engineering and conceptual ethics is that philosophers employ the terms 'conceptual engineering' and 'conceptual ethics' in a variety of ways.[9] This is tied to the fact that the terms 'conceptual engineering' and 'conceptual ethics' are both relatively new in philosophical discourse. This suggests that there might not be settled linguistic facts about the respective meanings of these terms, nor settled facts about what these terms refer to. Thus, it is not at all clear *which* things we are talking about when we discuss "conceptual engineering" and "conceptual ethics".

Does this mean that the core question we are asking is a bad one, or that our inquiry here is doomed to fail? We think not. To approach the issue, we can first engage in the sort of normative inquiry about concepts that is a key part of both conceptual engineering and conceptual ethics. In particular, we can ask the following normative question about the terms 'conceptual engineering' and 'conceptual ethics': how *should* philosophers use these terms? Some might regard this move as a bait and switch: moving from asking a *descriptive* question about the nature of the *things themselves* to a *normative* one about *the terms* we use for describing them. We agree that this kind of move can *often* be a bait and switch. But it can also be a helpful move for clarifying what we should be talking about when we ask our original question, and why. In our view, it is not that the normative question about terminology here (that is: 'conceptual engineering' and 'conceptual ethics') either *replaces* or *settles* the descriptive question about the nature of the things themselves (that is: conceptual engineering and conceptual ethics). Rather, it helps us get clearer on which things we want to ask questions about (e.g., what their nature is, or what properties they have, etc.), and why.[10] We are optimistic that this move has something to contribute in this context. In short, our working hypothesis is that there is something to be gained by explicitly engaging in some conceptual ethics and/or conceptual engineering about the terms 'conceptual ethics' and 'conceptual engineering' themselves.

In this paper, we approach this task in part by thinking about the following combination of issues: (a) our preferred existing usage of the term 'conceptual ethics', which we have previously argued for (and which we stand behind, for the purposes of this paper), (b) connotations of the term 'engineering' in general, (c) existing

[6]For an example of this approach, see Cappelen (2018). Note that (as Cappelen's work underscores) providing such a theory might, of course, involve rejecting the self-understanding of the practitioners of the projects one identifies (including those, like Scharp's, that are *explicitly* pursued as ones in "conceptual engineering") or reforming our views about which projects really are instances of conceptual engineering and which ones are not.

[7]See Burgess and Plunkett (2013a, 2013b). See also Burgess (2014) and Plunkett and Sundell (2013).

[8]See, for example, Scharp (2013), Eklund (2015), and Cappelen (2018). Their use of 'conceptual engineering' draws on the brief use of the terminology in Blackburn (1999).

[9]For a sampling of some of the diversity here, see the papers in Burgess, Cappelen, and Plunkett (2020).

[10]For connected discussion, see Plunkett (2015).

usage patterns of the term 'conceptual engineering', and (d) key examples of work often described as instances of "conceptual engineering". If nothing else, we think that engaging in this exercise can help draw our attention to important distinctions (about projects, issues, theses, etc.) that arise in cases that many think of as paradigms of "conceptual engineering" and/or "conceptual ethics". One could discuss those cases using any range of proposed labels. More ambitiously, if our proposed usages of 'conceptual engineering' and 'conceptual ethics' are in fact onto something, then their adoption might help facilitate philosophical inquiry and discussion going forward, at least in certain contexts.

It should be underscored from the start that our aim here is *not* to determine how everyone, in all contexts, should use the terms 'conceptual engineering' and 'conceptual ethics'. We think there is almost certainly going to be variation in how different people should proceed in different contexts, given their different aims. Perhaps, for example, it is best in some philosophical discussions to not worry about distinguishing between conceptual ethics and conceptual engineering *at all*. Or perhaps the balance we strike between different desiderata—e.g., the balance between fidelity to current use of the term 'conceptual engineering' versus preserving connotations of the term 'engineering' in general—is good for some philosophical contexts, but bad for others. We are optimistic about the fruitfulness of our proposal for much of the current discussion in philosophy about "conceptual ethics" and "conceptual engineering". But we certainly do not mean to argue that our proposal is always the best one for all contexts in which philosophers (or others) use the terms 'conceptual engineering' and 'conceptual ethics'.

In what follows, we work through a number of proposals about how we should use the term 'conceptual engineering'—in relation to a working proposal for the use of 'conceptual ethics' that we introduce early on—before we arrive at our favorite. In presenting these proposals, we often put things in terms of what the relation *is* between conceptual engineering and conceptual ethics as such (rather than how we *should* use the *terms* 'conceptual engineering' and 'conceptual ethics'). This object-level, seemingly descriptive way of talking should be interpreted in light of our methodological remarks above. At the end of the day, we want to know what is true of the relation between those things picked out by the terms 'conceptual engineering' and 'conceptual ethics' *if* the terms were used in the way we think they should be. Our proposal is ultimately a normative one about how the terms should be used, such that they pick out things where these kinds of relations hold. In this context, we think we can put forward that normative proposal by speaking in a seemingly descriptive, object-level way where we simply *use* the terms in question.[11] This might be misleading for some readers, but we think this way of proceeding is okay as long as one keeps in mind what our ultimate goals here are, and what the appropriate methodology is for reaching them.

2 | THREE PARTS OF CONCEPTUAL ENGINEERING, AND A BRIEF PREVIEW OF OUR PROPOSAL

A simple, natural thought is that there is no substantive difference between those things we should describe using the term 'conceptual engineering' and those we should describe using the term 'conceptual ethics'. In what follows, in part simply for ease of exposition, we take conceptual engineering and conceptual ethics to be *activities* one can engage in. So, on this simple view, there is just one kind of philosophical activity here, which we can refer to either using the language of 'conceptual engineering' or the language of 'conceptual ethics'.

There are some passages from recent work on conceptual engineering that suggest this view:

> *Herman Cappelen from* Fixing Language: *"This book is in part about the significance of terminological choices and so it is fitting that the terminology for talking about this field itself presents us with difficult choices. The term 'conceptual engineering' is used by Simon Blackburn (1999) and also Matti Eklund*

[11]One way to see how this might be so is to take us to be explicitly engaging in *metalinguistic negotiation*, in which we seem to use (rather than mention) terms to put forward a view about how those very terms *should* be used. See Plunkett and Sundell (2013).

> *(2014, 2015)… Some use Carnap's term 'explication' for a sub-version of this kind of project. Others again prefer 'conceptual ethics' over 'conceptual engineering' (Burgess & Plunkett, 2013a, 2013b). I've settled on 'conceptual engineering' though it is far from ideal."*[12]

> *Matti Eklund from* Choosing Normative Concepts: *"Whereas philosophers often have been concerned to understand our actual expressions and concepts, and what these expressions and concepts stand for, I think we should be more concerned with the space of possible expressions and concepts pertaining to a given subject, and with questions about what the most apt expressions and concepts are. The latter is what some theorists talk about under the heading of* conceptual engineering, *or* conceptual ethics."[13]

To be clear: we do not think that either of the above philosophers we have quoted means to put forward a worked-out view here about the relation between conceptual engineering and conceptual ethics. Our point is just that these passages suggest the thought about identity between conceptual engineering and conceptual ethics, which is a natural, default starting point.

An initial worry with this default starting point is that paradigm cases of things that philosophers currently call "conceptual engineering" suggest the idea that conceptual engineering is more intimately related with two kinds of activities than conceptual ethics is. We think this idea also fits naturally with associations of the term 'engineering' itself.

The two activities are as follows.

First, there is what we can call *conceptual innovation*. Roughly, this activity involves introducing new concepts and/or amending existing concepts.[14]

Second, there is what we can call attempts at *conceptual implementation*. Roughly, this involves a kind of *advocacy*, in which one tries to get some people (ranging from a lone individual to a large population) to actually take up and use the concepts one is in favor of. In other words, this involves some attempt at actual implementation of the conceptual changes that one thinks should be brought about, e.g., by getting people to employ the concepts one previously developed doing conceptual innovation.[15]

It should be underscored that, in both activities, one's target might not be *concepts*, as opposed to other "representational devices" (e.g., words or lexical items). This follows from the more general point that either some (or all) of conceptual engineering, or some (or all) of conceptual ethics, might ultimately best be understood as not about concepts as such.[16]

[12]Cappelen (2018, p. 4).

[13]Eklund (2017, p. 15).

[14]Some might also want to use the term 'conceptual revision' (drawing on Railton (1989)) or 'conceptual amelioration' (drawing on Haslanger (2020)) or 'conceptual design' to describe this kind of project. As Railton's and Haslanger's projects exemplify, conceptual innovation is standardly undertaken with an eye to creating *better* or *improved* concepts. However, nothing in the idea of conceptual innovation as such necessitates that one has such an aim, and, if one does, it is also obviously a separate question whether the new concepts are in *fact* improvements or not. Note also that another possible idea about the label 'conceptual design' we just mentioned above would be to use it as an alternative label for what we are here calling "conceptual engineering". This would be especially apt if one thought 'engineering' and 'design' were rough synonyms in general.

[15]As our usage of 'advocacy' here suggests, we think 'conceptual advocacy' might be another good label here for the thing we have in mind. A drawback is that perhaps *any* positive normative view about X could be seen as a kind of advocacy on its behalf. There will always be gradations here with how much one really tries to implement a normative proposal about which concepts should be used. So this might not be that big of a problem at the end of the day. We stick with the "implementation" label instead, since it draws attention to the end of the advocacy in a way that we think is more illuminating of the kind of activity we have in mind. Another possible label for the activity is 'conceptual activism'. (See Cappelen and Plunkett (2020) for use of that label in roughly this way). We think that using 'conceptual activism' in this way has a number of drawbacks. For example, a government official might be engaged in a project of attempted conceptual implementation. But we do not generally automatically think of a government official as an "activist" solely in virtue of trying to effect social change, e.g., by trying to carry out or enforce an existing policy or law.

[16]For example, Cappelen develops a theory of conceptual engineering according to which there are not any concepts at all. On his view, conceptual engineering involves work on the intensions and extensions of lexical items, rather than on concepts. See Cappelen (2018).

The prima facie connection of conceptual engineering to the activities of conceptual innovation and conceptual implementation raises a natural question: is conceptual engineering somehow definable in terms of these three activities (conceptual ethics, conceptual innovation, and conceptual implementation)?

One simple thought (which we call *Proposal 1*) is this: conceptual engineering encompasses every instance of each of the three things. This proposal is something we once were attracted to. But we now think this proposal is not the best way of thinking about conceptual engineering, for reasons we explain below. We bring it up here just to illustrate what it would be to try to define conceptual engineering in terms of these three activities.

So *should* one try to define conceptual engineering in terms of these three activities? At the risk of spoiling the suspense, here is a preview of our eventual proposal: conceptual engineering *is* indeed definable in terms of these three activities (conceptual ethics, conceptual innovation, and conceptual implementation). But, we do not think it is so definable in a simple Boolean way. Rather, we think the term 'conceptual engineering' is best used to pick out something that is made up of mereologically complex activities whose parts fall into the categories associated with each of these three different activities. The prototypical/paradigmatic projects in conceptual engineering, we think, are exercises in conceptual innovation, guided by conceptual ethics, with an eye (and likely some effort put toward) effecting actual conceptual change. We'll get more precise about this thought below, as we build up toward our proposal by working through what we find wanting in other proposals.

To set the stage for this work, we should clarify something up front about how we understand "conceptual ethics" in this paper. In previous work, we introduced the term 'conceptual ethics' to refer to a branch of normative and evaluative inquiry.[17] It is a branch that, in short, deals with normative and evaluative questions about concepts, words, and other broadly "representational" and/or "inferential" devices we use in thought and talk.[18] Put in terms of concepts and words: we take some of the central questions in conceptual ethics to concern which concepts we should use and what words should mean, and why. On our way of using the term 'conceptual ethics', conceptual ethics is really a shorthand way of referring to normative and evaluative inquiry about certain kinds of questions about thought and talk. It need not focus on *concepts* per se. And, despite how we often use the term 'ethics', the normative and evaluative issues on the table here need not be ones in broadly moral/political philosophy. For example: nothing in how we understand conceptual ethics suggests that it needs to concern things such as the promotion of social justice or happiness, or respecting the moral rights of agents. One might argue that we should use those concepts that carve reality at its joints, irrespective of whether our use of those concepts makes our lives go better or not, and irrespective of any other broadly "moral" or "political" considerations. As we understand it, that is an argument in "conceptual ethics".

For the purposes of this paper, we proceed on the assumption that this is a good way of regimenting the term 'conceptual ethics'. That is: we take it that to engage in the activity of *doing* conceptual ethics is to engage in this branch of normative and evaluative inquiry. This might, of course, be incorrect. Our previous usage of 'conceptual ethics' is obviously not sacrosanct. Perhaps, for example, another label would be better to use here for the kind of normative and evaluative inquiry we have in mind. After all, it would be good to have a label that didn't seem to suggest it was really all about *concepts* as such and/or that didn't suggest that *practical* values/norms were not the most important ones for the normative/evaluative issues on the table. However, we have not yet found a label that we think is obviously better, all-things-considered. For example: the label 'normative and evaluative work on concepts and/or other related devices in thought and talk' is obviously a seriously cumbersome phrase to use in many contexts, and the label 'conceptual evaluation' makes it seem as if the kind of inquiry we want to discuss is all about evaluative issues (e.g., about what is good or bad) rather than normative ones (e.g., about what agents should do, think, or feel). Thus, for the purposes of our paper, we stick with the term 'conceptual ethics', and take it to refer to the activity of engaging in the kind of normative and evaluative inquiry we glossed above. This choice is important

[17]See Burgess and Plunkett (2013a, 2013b).

[18]As this gloss suggests, we do not take it as given what those devices are. Many philosophers will put questions about *concepts* on the list. But someone who is a nihilist about concepts will not.

to underscore given that, on one way of proceeding, much of our work in this paper could be focused on different proposals about how we should use the term 'conceptual ethics'. Instead, we proceed by largely taking our earlier account as a good starting point and then focusing on how to think about conceptual engineering in light of it.

3 | WORKING TOWARD A PROPOSAL

In this section, we consider and reject four different proposals, all of which aim to define conceptual engineering in terms of the three different activities we identified above (conceptual ethics, conceptual innovation, and conceptual implementation). We can build a better proposal, we think, by reflecting on the relative merits of each of these four proposals.

3.1 | Proposal 1

Recall Proposal 1, which we introduced in the last section. According to this proposal, conceptual engineering is the union of all three activities (conceptual ethics, conceptual innovation, and conceptual implementation). This proposal faces the following problem: you can do any one of these three activities independently without doing either of the others, and there are examples of this in each of the three cases that do not look like a good thing to call "conceptual engineering". Thus, taking the union simply counts too much stuff as conceptual engineering.

Let's start with the conceptual ethics case. Recall that, on the gloss of conceptual ethics we introduced above, conceptual ethics is essentially the activity of engaging in a branch of normative and evaluative inquiry that deals with issues concerning thought and talk, such as which concepts we should use, and why. There is not going to be a theory-neutral way of demarcating exactly which questions are covered by conceptual ethics. For example, if one is skeptical about concepts as such, then the core questions concern something else (e.g., perhaps lexical items, perhaps conceptions, etc.). Another issue concerns whether a given concept can change its intension and extension and remain the *same* concept. Someone might think yes, whereas someone else might think that any such adjustment necessarily involves switching to a nearby, related concept. How one thinks about this issue will affect how one thinks of the possibility of revising or reforming a given concept, as opposed to simply replacing it. Whatever exact topics are encompassed by conceptual ethics, however, there is something important that follows from the basic idea of thinking of "conceptual ethics" as a kind of *normative and evaluative inquiry*. This is that, just as with normative and evaluative inquiry in other subareas of philosophy—e.g., in moral philosophy, political philosophy, aesthetics, or epistemology—one can engage in that inquiry *without* aiming to make a practical difference to how things are in the world, or without an eye toward actually changing anyone's behavior. For example, consider the question of whether the normative facts about which concepts we should use are ultimately determined solely by facts about what best promotes the best consequences. One might think that it is simply *intellectually interesting* to figure out whether this kind of consequentialism in conceptual ethics is true, just as one might think that questions about parallel kinds of consequentialism in ethics and epistemology are interesting. One might try to figure out whether this kind of consequentialism in conceptual ethics is true without trying to change anyone's actual conceptual behavior, and without engaging in any conceptual innovation. The same is true of working on more "applied" questions in conceptual ethics. For example, one might try to figure out which race and gender concepts a specific group of people should use in a given social-historical context, but without trying to change anyone's actual conceptual behavior, and without engaging in any conceptual innovation. To us, it seems a bit unhelpful to count these kinds of activities—ones that simply try to make epistemic progress on a theoretical question in conceptual ethics—as an instance of conceptual engineering.

Similar points apply to examples of activities solely involving conceptual innovation, or solely involving conceptual implementation. For example, consider someone engaged in the project of building new concepts just for

fun, or just as a casual art project.[19] Such a person might engage in this activity of conceptual innovation without ever consciously engaging in any sort of normative and evaluative theorizing about which concepts they should create, or why, and might have no eye toward trying to get anyone to actually employ their creations. Or consider someone promoting, solely for the sake of personal gain, the use of a concept that someone else created. Such a person might engage in this kind of conceptual implementation without doing any amount of conceptual ethics or conceptual innovation. These examples are, again, ones that seem a bit strange to count as examples of "conceptual engineering". Or at least so we submit.

Some of what we are thinking here draws on the contrast between the above cases and paradigm examples of conceptual engineering, such as Scharp's work on replacing the ordinary truth predicate and Haslanger's ameliorative accounts of race and gender concepts. Those are cases where it seems that there is some amount of all three activities going on (that is: some amount of conceptual ethics, conceptual innovation, and conceptual implementation). Some of our thinking here also stems from associations of the term 'engineering'. For example: simply creating something without *any* aim to make it useful for some activity, or without any aim at implementation, does not strike us as something naturally described as an instance of "engineering". Nor does abstract normative theorizing on its own, or simply trying to get others to do something for personal gain.

Proposal 1 faces another issue as well, which is tied to what we have argued above. This is a problem about subservience. The problem is that it is at best misleading, and at worst false, to say that conceptual ethics (or one of the other two activities) is *just* "a part of" conceptual engineering. This problem has more to do with misunderstanding conceptual ethics (or the other two activities) and less to do with misunderstanding conceptual engineering.

To appreciate this problem, consider the idea that normative political philosophy is just "a part" of a broader project, which also includes trying to create and implement better (or more just, etc.) institutions. Or the claim that normative aesthetics is just "a part" of a project that involves creating better art. As we emphasized above, someone might engage in normative and evaluative theorizing for any number of reasons. Some of those reasons might have nothing to do with actually trying to bring about change in the real world (e.g., in terms of what actual social/political institutions are like, or what art is actually created). For example, as we discussed above, one might engage in philosophical theorizing—including about normative issues—simply because one finds it intellectually interesting. Or one might find oneself engaged in thinking about these issues because of how they connect to *other* philosophical topics that one finds interesting. This would be parallel to how some philosophers of art might engage with certain issues of philosophy of language because of the relevance of those issues to their work. Given all of this, even if conceptual ethics is a part of conceptual engineering *in some sense*, it would be misleading to give an account of what conceptual ethics *is* in terms of this relation (e.g., in terms of giving something akin to a real definition of conceptual ethics). This would be parallel to how it would be misleading to give an account of normative political philosophy just in terms of how it fits into a broader project that also involves trying to actually create and implement better social/political institutions, or to give an account of normative aesthetics solely in terms of its connection to a broader project that also involves the practice of trying to create better art.

There is also a further, connected point about subservience. Suppose we assume for now that conceptual ethics *is* a part of conceptual engineering (and even that this connection to conceptual engineering is a *part* of the real definition of conceptual ethics). Conceptual ethics can still be a part of lots of other things too. Some of those other part-whole relations might be more illuminating about the nature of conceptual ethics. For example, following our original take on conceptual ethics in our earlier work, we think it is helpful to think of conceptual ethics as a branch of normative and evaluative inquiry (just as, for example, normative political philosophy is such a branch). This way of defining the field need not mention conceptual engineering at all.[20]

[19]As noted earlier, if one thinks new concepts can't literally be created, then the activity here can be reformulated accordingly.

[20]The above points about subservience draw on the introductory discussion of the nature of conceptual ethics by one of us (Plunkett) in Cappelen and Plunkett (2020).

3.2 | Proposal 2

Given some of the arguments we provided against Proposal 1, a natural thought is that perhaps conceptual engineering is not all instances of all three activities (conceptual ethics, conceptual innovation, and conceptual implementation), but rather *some* instances of all three activities. Call this *Proposal 2*. The basic idea here would be that we can ditch the examples like those above, in which simply engaging with some part of conceptual ethics, conceptual innovation, or conceptual implementation did not seem to really be a part of conceptual engineering. The idea, in short, might be that these examples created issues because of the *specific parts* of those activities that were being engaged in. This picture is a view on which conceptual ethics and conceptual engineering partially overlap in terms of what they cover. On this proposal, the overlapping part is something that would be identified at a given social/historical context, and thus could potentially vary in the future. Depending on one's ambitions, the overlapping part could be used to illuminate *both* conceptual ethics and conceptual engineering.

The basic issue with Proposal 2 is that it is not at all clear which subpart of conceptual ethics should be included in conceptual engineering. The same issue arises in thinking about which subparts of conceptual innovation or conceptual implementation should be included. In all cases, you want to avoid the result that it is just *whatever* activities people who are called (or call themselves) "conceptual engineers" end up engaged in, or drawing on. To appreciate the issue here, consider that issues in epistemology sometimes come up in discussions about the philosophy of mind. But that does not mean that all the issues in epistemology that philosophers of mind discuss are a part of the philosophy of mind. Moreover, there is not a well-defined part of epistemology that deserves to be thought of as a proper part of the philosophy of mind. Similarly, issues in conceptual ethics sometimes come up in discussions about the philosophy of race. But it is not as if *all* such issues that come up thereby become part of the philosophy of race. After all, shouldn't we leave room for the idea that inquiry can move (often quite smoothly) from discussion about one subarea of philosophy to another? Furthermore, in parallel to the previous case concerning the philosophy of mind and epistemology, there is not a well-defined part of conceptual ethics that is best thought of as a subset of philosophy of race. There are of course paradigm examples of topics from conceptual ethics that are discussed in the philosophy of race: e.g., whether we should abandon our ordinary racial concepts and instead employ new, ameliorated ones, or perhaps none at all. But different philosophers, with different philosophical approaches, different views in auxiliary parts of philosophy, etc., will end up engaged with *different* issues in conceptual ethics while trying to make progress on an issue in philosophy of race. For example: some philosophers of race engaged in conceptual engineering might well end up thinking a lot about whether the kind of consequentialism about conceptual ethics we introduced earlier is correct or not. But it is hard to see how this would be part of a list of questions in conceptual ethics that are properly a part of the philosophy of race.

This is not to say that we cannot identify some general issues in conceptual ethics, conceptual innovation, and conceptual implementation that will *often* be relevant to thinking about the overall project of conceptual engineering, or to thinking about specific instances of it. Indeed, we think we can; or at least we can do so relative to our epistemic situation in our current social/historical context.[21] It is just that we do not think there is a well-defined list here that can be used in giving an illuminating definition of what conceptual engineering *is*.

3.3 | Proposal 3

Another obvious proposal is to define conceptual engineering as the union of two of the three activities. Call this *Proposal 3*. This gives us three distinct options. Conceptual engineering could be defined as the union of conceptual ethics and conceptual innovation. Or it could be defined as the union of conceptual ethics and conceptual implementation. Or, finally, as the union of conceptual innovation and conceptual implementation.

[21] See, for example, the clusters of issues discussed in Cappelen and Plunkett (2020).

Whichever version of Proposal 3 one chooses, one nice feature of it is that it can mitigate (even if not fully address) the problem of "too much stuff" we discussed for Proposal 1. This is because, by cutting out 1/3 of the activities, one will cut out at least *some* of the cases that do not seem to be aptly described as cases of "conceptual engineering". Moreover, there are more specific benefits of Proposal 3 depending on which variant of it one goes with (that is: depending on which two of the three activities one combines). For example, consider the combination of conceptual innovation and conceptual implementation. This combination sets up conceptual ethics and conceptual engineering as separate but complementary parts of an overarching (as yet unnamed) project.[22] This version of Proposal 3 thus avoids what is perhaps the most pressing of the subservience problems that arose for Proposal 1, which concerns the account of conceptual ethics in particular (as opposed to one of either conceptual innovation or conceptual implementation). On this version of Proposal 3, conceptual ethics is one thing, and then conceptual engineering is something one can do *informed by* the results of one's normative theorizing in conceptual ethics.

We think there is something going for Proposal 3 in general, and for this version of it in particular. But there is also a basic worry about it that seems hard to shake. This is that it is difficult to find good grounds for singling out one of the three activities for exclusion. To illustrate this problem, suppose one is engaged in a project of conceptual innovation and one finds oneself, along the way, thinking about some normative issues in conceptual ethics. When one thinks about those issues, has one really stopped being engaged in a project of conceptual engineering? It seems more apt to describe what one is doing as simply continuously engaging in a single project in conceptual engineering, in which some issues in conceptual ethics have arisen.

3.4 | Proposal 4

We have been canvassing the possibility of defining conceptual engineering in terms of the activities of conceptual ethics, conceptual innovation, and conceptual implementation. A proposal left open by the logical space we have canvased thus far is to just focus on *one* of the three activities in question. Call this *Proposal 4*.

Earlier in the paper, we already set aside the idea that conceptual engineering is identical to conceptual ethics. So that leaves the possibility that it is identical to conceptual innovation, or that it is identical to conceptual implementation.

Of these two options, the identification of conceptual engineering with conceptual innovation is the more plausible one. After all, building, making, or reforming things seems, in general, to be something engineers of various stripes paradigmatically do, in a way that advocacy does not. Indeed, given the way the term 'engineering' is often used, one might think there is a natural division of labor here: engineers make things (many of which are not used beyond the prototype phase), and then other people (e.g., those involved in advertising, sales, management, etc.) try to get people to use certain of those things. This strand of usage of 'engineering' counts in favor of the identification of conceptual engineering with conceptual innovation. Moreover, so does the attractive simplicity of this proposal.

So what should we make of this version of Proposal 4, on which conceptual engineering is just identified with conceptual innovation? We have a lot of sympathy for it, for the reasons given above. But we have two main worries. First, it is not clear how much mileage one should get from the above strand of usage of 'engineering'. After all, creating or reforming something (e.g., whether a concept or a car) just for art's sake, or as a prank, is not exactly a paradigm case of engineering, at least when we consider how the term 'engineering' is often used. Instead, as the term 'engineering' is often used, engineering is more closely tied to creating, adjusting, etc. *for the sake of* creating something with an eye toward actual implementation. Second, and more importantly, this proposal leaves

[22]One idea might be to call this project "conceptual design", opting for a wider definition of 'design'. This would be in contrast to the idea (discussed in footnote 14) of using 'conceptual design' to refer to what we call "conceptual innovation" or to refer to what we call "conceptual engineering".

out the intimate connections to conceptual ethics and conceptual implementation—and especially to conceptual ethics—that seem to be at the heart of core examples of conceptual engineering that anchor much of the current philosophical discussion of "conceptual engineering" (e.g., the examples of Scharp and Haslanger that we started with, the examples that Cappelen provides at the start of *Fixing Language*, or Eklund's work in *Choosing Normative Concepts*). This might not be that huge of a cost. But we think there is an option available that maintains a focus on these intimate connections to conceptual ethics and conceptual implementation that we see in paradigm cases, and which avoids the problems we have discussed for Proposals 1–4. We now turn to spelling out that positive proposal.

4 | OUR POSITIVE PROPOSAL

In our discussion thus far, we've been looking at unions of sets of activities. Notice that conceptual engineering cannot be the intersection of two or three of these sets, because that intersection is empty or at least very small. But now imagine that you count complex activities that have parts belonging to one of the three as being in the "intersection". This idea is a way of backing into our positive proposal, according to which paradigm cases of conceptual engineering involve some conceptual innovation, guided by some conceptual ethics, and in the service of some conceptual implementation.

A question that immediately arises: could "some amount" be zero? Well, it can't be zero for all three activities. And in light of the "too much stuff" objection, we think it won't be zero for two simultaneously either. But, we think, it can be zero for any one of the three.[23] Or at least that is the proposal we want to float here.

Given this answer, one might wonder: is not the proposal we just sketched simply a version of Proposal 3? Recall that, on that proposal, conceptual engineering is defined in terms of two of the three activities we have been discussing (conceptual ethics, conceptual innovation, and conceptual implementation). The answer is no. Proposal 3 concerns unions. Now we are talking about *complex activities* built up from the extensions we previously discussed.

A worry about our proposal is that *any* activity that only involves parts of two of our three activities (rather than parts of all three) might not look like something worth calling "conceptual engineering". For example: could you really have conceptual engineering with just some parts of conceptual ethics and conceptual implementation? Some may think that answering "yes" is a problem here. But we think it is not. Consider here papers advocating for eliminativism about a concept; e.g., papers advocating for no longer using (at least in certain contexts) our ordinary concept FREE WILL, or no longer using our ordinary concept BELIEF. In such cases, no new concepts need be introduced, and no existing concepts need be revised. Let us suppose the authors are trying to get us (or others) to actually follow the normative view on the table. This feels like a project involving a kind of conceptual engineering, despite the lack of conceptual innovation. One might want to insist that maybe we are innovating/amending our whole "conceptual scheme" even if we are not creating or revising a specific concept. Perhaps so. But we think it is just as faithful a representation of the project on the table to say that it does not really involve conceptual innovation, despite being a project in conceptual engineering. Other examples of conceptual engineering that do not involve conceptual innovation are ones where people draw on conceptual innovations that others have made in the past. For instance, suppose one is simply advocating for using a concept that someone else already created, perhaps a century or more earlier.

Another worry about our proposal is that, like Proposal 1, it is over-inclusive. Do we really want to say that absolutely *all* cases of (i) conceptual ethics and conceptual implementation, (ii) conceptual ethics and conceptual innovation, and (iii) conceptual innovation and conceptual implementation really count as conceptual engineering? Will we not face embarrassing counterexamples in much the same way that Proposal 1 did? We grant that this

[23]This makes "conceptual engineering" out to be something like Goodman's "imperfect community" as discussed in Goodman (1951/1977).

is indeed a worry. But we think that it can be mitigated by thinking about the following considerations. First, our proposal is *less* over-inclusive than many other proposals. For instance, earlier cases involving just one of the three activities will not be counterexamples to our view here. Second, it would be easy to come up with a proposal that is *under*-inclusive. We think our proposal strikes a nice balance here, including enough of the paradigm examples of conceptual engineering that help ground the current discussion of the topic, but without excluding too many of them.

We want to close by making three final points about our proposal.

First, if our view is right, then different projects in conceptual engineering can vary from each other in a number of significant ways. For example, contrast a project in conceptual engineering that does not involve any serious engagement with the project of conceptual implementation with one that does. Those two projects have quite distinct success-conditions: one is trying to do something that the other is not. Thus, the critiques of how these projects are carried out will vary. For example, issues about the difficulty of conceptual change, or the degree to which we can control conceptual change, will (at least prima facie) have more bite for the second project than the first. This illustrates the following point: critiques of conceptual engineering *as such* need to pay attention to the diversity of projects within conceptual engineering. Tied to this, so too should critics be attentive to the distinction between conceptual engineering and conceptual ethics, since, again, the projects have distinct success-conditions.

The second point stems from the fact that, on our view, conceptual engineering is a kind of *activity*. As part of engaging in that activity, a philosopher might make any number of claims or advance any number of theses. For example, she might make claims about the nature of concepts, or about the internalism/externalism debate in the philosophy of language. Which of those claims are properly thought of as claims "in conceptual engineering"? We favor a sort of context-sensitive answer here.[24] There are different relations a claim can have to a given project of conceptual engineering, or to the project of conceptual engineering as such. In asking whether a given thesis (e.g., content externalism in the philosophy of language) is a claim "in conceptual engineering", we might be interested in whether the claim plays an important role in an existing theory in conceptual engineering as such.[25] Or perhaps it is simply crucial to an attempt to carry out the engineering of a particular concept, or group of concepts. Alternatively, we might be interested in claims that *might* play an important role in debates about the nature of conceptual engineering as such, or a given contemporary debate about this. We will be focused on different such relations that a claim can have to a given project of conceptual engineering (or to the project of conceptual engineering as such) in different contexts. Which relations are important, we think, should be reflected when asking about whether individual claims are ones in conceptual engineering. Individual claims will count as ones "in conceptual engineering" when we focus (either explicitly or implicitly) on certain relations, and not when we focus on others. At the end of the day, we think, the status of *claims* (or questions, theses, etc.) matters less than the *projects* to which these claims are tied. Keeping those projects in mind—whether it is the overall project of conceptual engineering as such, or the conceptual engineering of some concept or group of concepts—is the key thing.

The above point that different claims—about a range of different topics—will be relevant to different projects in conceptual engineering also ties into thinking about the relations between conceptual engineering and things *other* than conceptual ethics. This leads to our third point. On our view, conceptual engineering has a particularly intimate connection to conceptual ethics, conceptual innovation, and conceptual implementation. These are the three activities we have built into our proposed account of what conceptual engineering is. But this does not mean that conceptual engineering is not also tightly connected (even if not as intimately so) to other activities.

[24]Our account below parallels the account of what counts as a "claim" in a given area in McPherson and Plunkett (2017) (focusing on the case of metaethics) and Plunkett and Shapiro (2017) (focusing on the case of general jurisprudence).

[25]For example: whether it plays an important role in Cappelen's theory of conceptual engineering in Cappelen (2018), or in Amie Thomasson's account of it in Thomasson (2020), etc.

For example, consider the activity of trying to understand, or simply describe, the concepts we currently employ. This sort of descriptive project might be understood as a kind of *conceptual analysis*. The kind of descriptive information one gets from doing conceptual analysis will *often* be very helpful for engaging in a range of projects in conceptual engineering (just as it will be for many projects in conceptual ethics). After all, many projects in conceptual engineering take off from the idea that a current concept is defective in some way, or at least that it could be improved. It would be hard to justify such an evaluative claim without any understanding of what the current concept involves, not to mention an understanding of what the proposed alternative concept involves. Consider the analogy: if one wants to fix a broken car by replacing one old part with a new one, one will need to understand some *descriptive* information about both car parts.

Another kind of activity that might often be of use to those engaged in conceptual engineering is the activity of better understanding the history of our concepts; e.g., where they came from and what people have done with them in the past.[26] Uncovering such facts of *conceptual history* might well matter to a range of projects in conceptual engineering. For example, looking at facts about the history of a concept might help us better understand features of that concept, which in turn might lead us to appreciate ways in which it could be improved.[27]

There are numerous other activities we could list here as well. For instance, engaging in general normative issues in epistemology or ethics will be helpful for many projects in conceptual engineering (and many in conceptual ethics). Our account gives room for the import of information and ideas gleaned from such other activities, despite not granting them pride of place in our proposed account of conceptual engineering.

One objection to our view is that these two other activities that we focused on above (conceptual analysis and conceptual history) should really be *included* in the definition of conceptual engineering as such, and not just noted as closely connected activities. Our basic response is simply that these activities are not as explanatorily central in helping illuminate the nature of the distinctive activity that is common to core cases of "conceptual engineering". To make that case in full would require much more engagement with particular case studies than we have room for in this paper. So, instead of doing that here, we want to offer two additional thoughts that matter in assessing our proposal. First, our proposal (involving three activities) is simpler than a proposal involving five activities, which is arguably *some* sort of virtue. Second, it is not clear that engagement with these two activities—conceptual analysis and conceptual history—is *more* crucial to defining conceptual engineering than, say, the activity of engaging in research on metasemantics, or on normative epistemology.

5 | CONCLUSION

In this paper, we have put forward a view about how we should use the terms 'conceptual engineering' and 'conceptual ethics'. We put forward this proposal with the hope that it can help illuminate key philosophical distinctions and projects that are closely related to each other, and help us make better progress in thinking about them in the years ahead. At the end of the day, the sorts of distinctions which our paper highlights matter regardless of the labels that are used to discuss them. We think the labels 'conceptual engineering' and 'conceptual ethics' are currently good ones to use for the relevant activities, but we very much take it as an open question whether better labels will emerge in the coming years. Indeed, if future work in conceptual engineering and conceptual ethics is successful, people might well discover better labels for the activities that we have discussed. If nothing else, we hope our work in this paper can serve as a helpful reference point for those who will carry out such work.

[26]Note that, as the wording here suggests, the information here might be best understood (strictly speaking, given one's theory of concepts) not as the history of certain concepts as such, so much as the history of people's use of those concepts.

[27]For more on conceptual history, see Plunkett (2016) and Dutilh Novaes (2015).

ACKNOWLEDGEMENTS

An earlier version of this paper was presented at a workshop on conceptual engineering at NYU in September 2018. Thanks to everyone there for helpful questions and feedback. Thanks also to David Braddon-Mitchell, Herman Cappelen, Esa Diaz-Leon, Matti Eklund, Sally Haslanger, Tristram McPherson, Eliot Michaelson, Adrian Russian, Kevin Scharp, Tim Sundell, Amie Thomasson, and an anonymous referee for *Ratio*.

ORCID

David Plunkett https://orcid.org/0000-0003-1293-5770

REFERENCES

Blackburn, S. (1999). *Think: A compelling introduction to philosophy*. Oxford, UK: Oxford University Press.

Burgess, A. (2014). Keeping 'true': A case study in conceptual ethics. *Inquiry*, *57*(5–6), 580–606.

Burgess, A., & Plunkett, D. (2013a). Conceptual ethics I. *Philosophy Compass*, *8*(12), 1091–1101.

Burgess, A., & Plunkett, D. (2013b). Conceptual ethics II. *Philosophy Compass*, *8*(12), 1102–1110.

Burgess, A., Cappelen, H., & Plunkett, D. (2020). *Conceptual engineering and conceptual ethics*. Oxford, UK: Oxford University Press.

Cappelen, H. (2018). *Fixing language: An essay on conceptual engineering*. Oxford, UK: Oxford University Press.

Cappelen, H., & Plunkett, D. (2020). A guided tour of conceptual engineering and conceptual ethics. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 1–26). Oxford, UK: Oxford University Press.

Carnap, R. (1947/1956). *Meaning and necessity: A study in semantics and modal logic*. Chicago, IL: University of Chicago Press.

Dutilh Novaes, C. (2015). Conceptual genealogy for analytic philosophy. In A. C. J. Bell & P. M. Livingston (Eds.), *Beyond the analytic-continental divide: Pluralist philosophy in the twenty-first century (Routledge studies in contemporary philosophy), 2015* (pp. 77–108). New York, NY: Routledge.

Eklund, M. (2014). Replacing truth? In A. Burgess & B. Sherman (Eds.), *New essays in metasemantics* (pp. 293–310). Oxford, UK: Oxford University Press.

Eklund, M. (2015). Intuitions, conceptual engineering, and conceptual fixed points. In C. Daly (Ed.), *The Palgrave handbook of philosophical methods* (pp. 363–385). New York, NY: Palgrave MacMillan.

Eklund, M. (2017). *Choosing normative concepts*. Oxford, UK: Oxford University Press.

Goodman, N. (1951/1977). *The structure of appearance*. Dordrecht, the Netherlands: D. Reidel Publishing Company.

Haslanger, S. (2000). Gender and race: (What) are they? (What) do we want them to be? *Nous*, *34*(1), 31–55.

Haslanger, S. (2005). What are we talking about? The semantics and politics of social kinds. *Hypatia*, *20*(4), 10–26.

Haslanger, S. (2020). Going on, not in the same way. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 230–260). Oxford, UK: Oxford University Press.

McPherson, T., & Plunkett, D. (2017). The nature and explanatory ambitions of metaethics. In T. McPherson & D. Plunkett (Eds.), *The Routledge handbook of metaethics* (pp. 1–25). New York, NY: Routledge.

Plunkett, D. (2015). Which concepts should we use?: Metalinguistic negotiations and the methodology of philosophy. *Inquiry*, *58*(7–8), 828–874.

Plunkett, D. (2016). Conceptual history, conceptual ethics, and the aims of inquiry: A framework for thinking about the relevance of the history/genealogy of concepts to normative inquiry. *Ergo*, *3*(2), 27–64.

Plunkett, D., & Shapiro, S. (2017). Law, morality, and everything else: General jurisprudence as a branch of metanormative inquiry. *Ethics*, *128*(1), 37–68.

Plunkett, D., & Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint*, *13*(23), 1–37.

Railton, P. (1989). Naturalism and prescriptivity. *Social Philosophy and Policy*, *7*(1), 151–174.

Scharp, K. (2013). *Replacing truth*. Oxford, UK: Oxford University Press.

Thomasson, A. L. (2020). A pragmatic method for normative conceptual work. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual engineering and conceptual ethics* (pp. 435–458). Oxford, UK: Oxford University Press.

How to cite this article: Burgess A, Plunkett D. On the relation between conceptual engineering and conceptual ethics. *Ratio*. 2020;33:281–294. https://doi.org/10.1111/rati.12265