



Inquiry

An Interdisciplinary Journal of Philosophy

ISSN: (Print) (Online) Journal homepage: <https://www.tandfonline.com/loi/sinq20>

Taking control: conceptual engineering without (much) metasemantics

Jennifer Nado

To cite this article: Jennifer Nado (2020): Taking control: conceptual engineering without (much) metasemantics, Inquiry, DOI: [10.1080/0020174X.2020.1850342](https://doi.org/10.1080/0020174X.2020.1850342)

To link to this article: <https://doi.org/10.1080/0020174X.2020.1850342>



Published online: 25 Nov 2020.



Submit your article to this journal [↗](#)




View related articles [↗](#)



View Crossmark data [↗](#)



Taking control: conceptual engineering without (much) metasemantics

Jennifer Nado 

University of Hong Kong, Pokfulam, Hong Kong

ABSTRACT

This paper critiques the semantics-driven approach to conceptual engineering presented in Herman Cappelen's recent book, *Fixing Language*. I focus on three core aspects of Cappelen's 'Austerity Framework': first, the claim that conceptual engineering targets the meanings of linguistic expressions (rather than concepts); second, the claim that the linguistic phenomenon of 'samesaying' fixes the limits of permissible revision; and third, the claim that the process of conceptual engineering is largely both outside of our control and epistemically inaccessible to us. I contrast the Austerity Framework with my own proposal, which I call the 'Practical Role Account'. The Practical Role Account disavows a substantial role for metasemantics in a theory of conceptual engineering; I argue that it consequently both offers a more inclusive view on the limits of revision, and places the process of conceptual engineering back within our control.

ARTICLE HISTORY Received 10 April 2020; Accepted 2 June 2020

KEYWORDS Conceptual engineering; metasemantics; metaphilosophy

Is it possible to improve the representational schemes we use to navigate the world? Herman Cappelen's recent book, *Fixing Language*, is an examination of this question. The concepts with which we think and the language with which we converse are omnipresent filters, draped across our every perception, thought, and communication; it stands to reason that we should check them for potential defects. Indeed, Cappelen argues, given the infinite variety of potential concepts we might employ, it would be utterly miraculous if we had stumbled upon the best possible set. So how should we address this nagging worry? By engaging in an attempt to diagnose and eliminate conceptual flaws – by improving our conceptual repertoire, through the activity of *conceptual engineering*.

Fixing Language presents Cappelen's own picture of the nature and limits of conceptual engineering, which he dubs the 'Austerity

CONTACT Jennifer Nado  jennifernado@gmail.com

© 2020 Informa UK Limited, trading as Taylor & Francis Group

Framework'. I think it's fair to say that it presents the most comprehensive, detailed theory of conceptual engineering to date. Moreover, it introduces a new challenge for would-be engineers: one we might call the 'control problem'. Cappelen, like most contemporary philosophers of language, endorses externalism as the correct general flavour of metasemantic theory. But if we take the goal of conceptual engineering to involve changing the *meaning* of a term or concept, then a commitment to externalism potentially brings in tow the rather disheartening conclusion that we have little to no control over the engineering process, and little to no ability to even determine whether it has occurred.

I'm wholly on board the conceptual engineering train, but I'm much more optimistic than Cappelen about our ability to understand and influence the changes conceptual engineers effect. This is largely because I'm much less convinced of the relevance of metasemantics to the engineering project. This might sound absurd – surely we must know how our representations gain meaning before we can hope to successfully change those meanings? On my view, however, changing meanings *isn't* the ultimate goal of conceptual engineering. The classifications drawn by our words and concepts are tools for navigating the world, and the goal of engineering is to improve the efficacy of such tools. In many cases, a change of meaning will plausibly come along for the ride. But the primary goal can be achieved in relative semantic ignorance – and, I'll argue, it is very clearly within our control.

In this paper, I'll sketch an alternative framework for the conceptual engineer which minimizes the role of metasemantic theory – I call it the 'Practical Role Account'. It is, in essence, an attempt to prove false a core tenet of Cappelen's Austerity Framework: his claim that 'at the center of any theory of conceptual engineering is a metasemantic theory' (Cappelen 2018, 7). On the Practical Role Account, metasemantic issues are at best peripheral to conceptual engineering. Engineering success is framed in terms of increased practical utility, and semantic changes are held to be relevant only insofar as they increase or decrease said utility. The resulting theory, I hold, presents a plausible, flexible view of conceptual engineering with rather more optimistic prospects for the engineer than the gloomy implications of Austerity.

Readers who have dipped a few toes into the literature will note that the Practical Role Account bears some resemblance to broadly 'functionalist' views of conceptual engineering that have enjoyed recent popularity. Like these, it focuses on the importance of the various uses and purposes to which our concepts are put. But unlike many of these, it

does not privilege any of these purposes as constituting an essential ‘function’ which the engineer’s revisions must leave untouched. In a quasi-Quinean spirit, the Practical Role Account proposes that *no* aspect of the engineer’s target is immune to revision – provided that appropriate adjustments are made elsewhere in our classificatory system. As we’ll see, this leads the Practical Role Account to at least one point of agreement with the Austerity Framework: there can be no short, tidy, informative list of necessary and sufficient conditions for success in conceptual engineering.

In other respects, however, the Practical Role Account is a near opposite to the Austerity Framework. The Practical Role Account rejects three central features of Cappelen’s Austerity Framework; these features will be my focus in what follows. First, the Practical Role Account denies the Austerity Framework’s claim that linguistic meanings are the targets of conceptual engineering. Instead, it holds that the targets are what I will call ‘classification procedures’. These are, essentially, methods for ‘carving up’ the world into categories – more on them later. Second, the Practical Role Account denies the Austerity Framework’s view on the limits of permissible revision: while Cappelen claims that revisions must preserve sameness of topic, the Practical Role Account cheerfully permits topic-changing engineering proposals. Third, the Practical Role Account holds that the process of conceptual engineering is both epistemically accessible to us and well within our control.

1. The Austerity Framework

As noted above, I’ll be focusing on three aspects of the Austerity Framework: its view on the targets of conceptual engineering, its stance on the boundaries of permissible revision, and its implications regarding the epistemic and practical accessibility of the conceptual engineering exercise. Each of these aspects of the Austerity Framework flows fairly naturally from a ‘semantic’ approach to conceptual engineering such as Cappelen’s; the non-semantic approach I aim to endorse will approach these issues quite differently. In this section, I’ll overview the three issues and the Austerity Framework’s take on them. In the next section, I’ll begin to explore how the Practical Role Account will tackle these same topics in a largely metasemantics-free manner.

First up: the targets of the engineering process. Unsurprisingly, Cappelen takes the entities being ‘operated on’ to be expressions of language – or, more precisely, the meanings of those expressions.¹ For Cappelen,

conceptual engineering is not about concepts at all; indeed, given the bewildering polysemy of the term ‘concept’ within philosophy, he (rightly) takes it to be an advantage that the Austerity Framework side-steps concepts entirely. Moreover, true to its name, the Austerity Framework makes only very general, largely uncontroversial commitments regarding the meanings of linguistic expressions. It holds that expressions have extensions and intensions, where these terms are used as they are standardly understood – a term’s extension is just the set of things it picks out or applies to, and a term’s intension is a function from possible worlds (or world/time pairs, or whatever) to extensions. The process of conceptual engineering is then taken to be the process of producing a change in a term’s extension via a change in its intension.

Beyond this, the Austerity Framework builds in few substantive commitments about matters linguistic. Its primary such commitment is to the claim that the correct metasemantic theory is externalist – and as such, the intensions upon which the engineer operates will be at least partly determined by external factors. The Austerity Framework is not, however, committed to any particular externalist theory. Another fairly innocuous commitment is to the claim that changes to an expression’s intension are *possible*. One could conceivably hold that expressions should be individuated by their meanings, in which case meaning change would bring identity change in tow – making the process of conceptual engineering as understood by the Austerity framework incoherent. Cappelen dismisses this worry by simply noting that on most externalist views, changes of intension/extension are plausibly fairly commonplace, and that the degree of change may be at least as great as that exhibited by classic examples such as ‘Madagascar’.²

So the Austerity Framework implies that meaning changes are *possible*. But, of course, not all of these potential changes will be beneficial ones. The Austerity Framework doesn’t give us much of an account of what it is to engineer *well*; as noted in the introduction, Cappelen doubts that any simple set of conditions will do the trick. But Cappelen does address one well known, general question which faces all revisionary

¹Cappelen simultaneously holds that conceptual engineering is ‘about the world’, in the following sense. Suppose I successfully alter the meaning of ‘marriage’. Cappelen then argues that I have made a worldly change; I have changed what marriage is. But this is really just a somewhat unexpected implication of some of Cappelen’s metalinguistic commitments – on his account, it will be true to assert ‘what marriage is has changed’. The truth of that claim carries no metaphysical weight. This wrinkle in the account can be set aside in what follows.

²See Evans (1973). According to Evans, ‘Madagascar’ originally named a portion of the African mainland; due to a confusion by Marco Polo, it began to be used by Europeans when speaking of the island off Africa’s eastern coast. Over time, the term ‘Madagascar’ shifted reference to denote the island.

projects in philosophy: how much revision is *too much*? Even if fairly large shifts in intension/extension are *possible*, some such changes might have the consequence that we are simply no longer talking about the ‘same thing’ as we were before. Thus, though ‘Madagascar’ illustrates the possibility of meaning change, it is also very clearly a case where the change in meaning resulted in a change of subject. And to many philosophers interested in conceptual engineering, both advocates and opponents, a change of subject is a Very Bad Thing.

It’s become nearly obligatory to quote the following passage from Strawson, who arguably first pressed the issue, to illustrate why changes of subject are purported to be so problematic:

... to offer formal explanations of key terms of scientific theories to one who seeks philosophical illumination of essential concepts of non-scientific discourse, is to do something utterly irrelevant – is a sheer misunderstanding, like offering a textbook on physiology to someone who says (with a sigh) that he wished he understood the workings of the human heart. (Strawson 1963, 505)

Strawson is here targeting Carnapian explication, which we might (crudely) gloss as the attempt to conceptually engineer more precise, scientifically-respectable versions of everyday concepts. Strawson argues that this sort of conceptual tinkering results in a change of subject, with the consequence that our original philosophical questions simply go unanswered. A clever philosopher might propose some very neat and tidy, very carefully engineered modification to the concept of free will; for Strawson, all this does is evade our original questions regarding the real nature of freedom. Cappelen echoes Strawson’s worry by claiming that changing the subject disrupts the ‘continuity of inquiry’, and adds that subject change opens the door to potential miscommunications and verbal disputes.

This, then, is our second issue: at what point does revision become so extreme as to derail our original inquiries entirely, thereby doing more harm than good? What are the boundaries of permissible revision? Cappelen’s answer invokes the linguistic phenomenon of ‘samesaying’, which can occur even when two speakers use terms that differ somewhat in intension/extension. Suppose, to use a slightly modified version of one of Cappelen’s examples, that two speakers both utter the sentence ‘Externalism is an interesting theory’. If the speakers make these utterances in slightly different conversational contexts, the extensions of their token utterances of ‘interesting’ may differ slightly – the salient

comparison class might be different, for instance. But there's a clear sense in which we can say that the two speakers have 'said the same thing'. Samesaying, moreover, licenses indirect quotation; thus I can say 'Susan said that she's cold' even if the extension of 'cold' in my mouth is slightly different from that of Susan's utterance of 'cold'. By contrast, I cannot correctly report a pre-Marco Polo inhabitant of a certain portion of the African mainland as having said that he lives in Madagascar. For we two are not 'talking about the same thing' – after the reference shift, there has been a change of topic which prevents samesaying.

As Cappelen notes, and as the above examples make clear, samesaying intuitively seems to track sameness of topic – and thus, Cappelen argues that sameness of topic is more 'coarse-grained' than sameness of standard semantic features like intension and extension. The relevance of all this to Strawson's problem should be clear: Cappelen argues that so long as a user of the pre-engineering concept still counts as a samesayer with respect to a user of the post-engineering concept, engineering has not resulted in the sort of problematic change of subject (that is, topic) that Strawson worried about. And thus, presumably, the engineer cannot be held guilty of neglecting the initial philosophical problems that made the pre-theoretic concept of theoretical interest; continuity of inquiry has been preserved. The limits of permissible revision, then, are set by sameness of topic, which in turn is delineated by the linguistic phenomenon of samesaying – again reflecting Cappelen's overarching 'semantic' approach to a theory of conceptual engineering.

Our third issue is where the most troubling aspect of Cappelen's Austerity Framework comes to the fore. Recall that the Austerity Framework holds the process of conceptual engineering to involve changes to a term's intension. But, since the Austerity Framework involves a commitment to externalism, and since externalism states that external factors enter into the determination of intension, changes of intension will not be particularly straightforward to effect. We simply do not have control over many of the various features that externalist metasemantic theories claim to determine meaning. As Cappelen notes, many such features – such as facts about initial baptisms and the like – are features of the past, and thus are unalterable. Even those features that are not unalterable, such as patterns of use in a linguistic community, are difficult to alter. The upshot: the changes involved in conceptual engineering lie largely outside our control.

It gets worse. The facts that determine intension and extension, Cappelen claims, are largely epistemically inaccessible to us. We may not be able

to know the conditions present at an initial baptism; we may not be able to know enough about the usage patterns, past and present, of the linguistic community. And Cappelen denies that there's a tidy algorithm for determining meaning facts from the underlying external facts, anyway. All this implies that an engineer has no way to verify that her conceptual interventions have succeeded. On Cappelen's view, then, conceptual engineers lack both control over and insight into the changes they are attempting to make. They are, it seems, doomed to work in the dark, with their hands tied behind their backs.

2. The Practical Role Account: the targets

There's nothing inherently wrong with Cappelen's Austerity Framework. It's consistent. It provides answers to many of the central questions conceptual engineering projects naturally generate. And it avoids treading into the quagmires surrounding the nature of concepts. But its implications for our ability to actually *do* engineering sure are depressing. That's not in itself a reason to reject the account, of course. But if there were an alternate theory of conceptual engineering that painted a somewhat more optimistic picture, it might be worth considering its merits. In this section, I'll begin to outline such an alternative: the Practical Role Account.

The core claim of the Practical Role Account is that success in conceptual engineering is a matter of devising a tool that will effectively fulfil an intended practical role. I'll elaborate on the notion of a 'practical role' in the following section, but for now a reasonable approximation would be 'the set of purposes for which the tool is used'. A conceptual engineer examines her target concept (to speak loosely for the moment – as with Austerity, the Practical Role Account does not hold that conceptual engineering targets concepts) and asks herself whether there are any changes that could be made that would allow it to perform its role better. Insofar as a conceptual engineer aims to eliminate 'semantic defects' such as vagueness or inconsistency, this is held to be merely instrumental to the ultimate aim of designing a better tool. Insofar as 'changing the subject' is viewed as problematic, this is only in virtue of its tendency to bring in tow an inability to successfully fill the original concept's role.

The Practical Role Account aims for complete neutrality on metase-mantic questions. Thus, the purposes that make up a practical role are not held to constitute or determine any semantic features whatsoever

(either of words or concepts).³ There is no claim here that a concept's content is determined by its use, for instance. Moreover, the Practical Role Account holds that there is no reason to assume that subject-changing semantic changes entail that an engineer's invention fails to successfully serve the role of its predecessor. Thus, there is no reason to take semantic continuity (even in Cappelen's weak sense of 'samesayability') to constrain permissible revision; rather, permissible revision is constrained directly by degree of suitability to a role. Finally, since successful conceptual engineering does not essentially consist in bringing about semantic change, externalism neither blindfolds nor shackles us – conceptual engineering is well within our control, and its outcomes are within our epistemic reach.

That will do for now by way of an outline; I'll elaborate on the nature of a 'practical role' in the next section when we return to the limits of permissible revision, and I'll defend my claims regarding the control problem in the section following. But first, we must dispense with the pesky 'concept-talk. I'm with Cappelen here: the term 'concept' is simply too elusive, too slippery, to ground a good theory of conceptual engineering. But the Practical Role Account denies that conceptual engineering is fundamentally a semantic activity; and therefore linguistic expressions and their meanings cannot be its chosen targets, either. So what's left?

For Cappelen, the target of conceptual engineering intervention is an expression's intension. Though the notion of an intension can be formulated in multiple ways (and Cappelen stays neutral on these), the most basic formulation casts an intension as a function from possible worlds to extensions. Now, consider the fact that we can make sense of such functions divorced from questions of meaning – a function from worlds to sets of entities is a wholly respectable function regardless of whether it's the intension of any given linguistic item. So here's a potential first-pass answer to the target question: the engineer produces (or discovers, or selects) 'intension-candidates', which are simply functions from worlds to sets of entities, and which may or may not be currently 'attached' to an existing expression. These intension-candidates are, more or less, 'carvings' of logical space – or even more simply, categories. Different 'carvings' will have varying degrees of suitability for various purposes, and

³Though there will presumably be at least some degree of correlation – the exact degree will depend on the correct metasemantic theory. Nevertheless, since the engineering process is more epistemically transparent than meaning change (I'll argue), there is no reason for the engineer to fuss over pinning down the correct theory of meaning before embarking on an engineering project.

the core task of the engineer is finding the best category for the practical role she aims to fill.

There's a lot to unpack in the idea of fitting a category to a role; but first, we need to amend the first-pass answer above. Functions from worlds to sets of entities are ultimately too coarse-grained to be the targets of conceptual engineering, for essentially the same reason that they are too coarse-grained to be the meanings of expressions. When intensions are understood as functions from worlds to sets of entities, all necessarily co-extensive expressions are *co-intensive*. Thus, for example, 'equiangular triangle' and 'equilateral triangle' have the same intension. But they certainly don't seem to have the same *meaning*. The parallel of this problem threatens the use of intension-candidates as targets for the Practical Role Account, because it is intuitively quite plausible that co-intensive concepts (if you'll pardon the quick relapse into concept-talk) might differ in their suitability for a given practical role. Perhaps, for instance, we might find distinct philosophical uses for a variety of necessarily uninstantiated concepts (perhaps, like 'omnipotent', their utility is tied up with explaining *why* they cannot be instantiated). Similarly, the concepts 'even prime' and 'successor of 1' surely have different roles to play within mathematics.

For semantic purposes, we can refine the grain of intensions in various ways – for instance, by moving to a two-dimensional semantics. But it's not immediately clear how to apply such a strategy to secure a hyperintensional target for a non-semantic approach to conceptual engineering. I suggest, instead, that we take the primary targets of conceptual engineering to be what I will call *classification procedures*. A classification procedure is a set of steps or rules which, when followed, determines an intension-candidate – that is, a function from worlds to sets of entities, which may or may not be the intension of any actual linguistic expression. It is a method for 'getting at' the intended function, and there may be many such procedures that get at the same function – just as the arithmetical procedures 'add three then subtract one' and 'add two' get at the same function. The notion of a classification procedure is intended to encompass any series of instructions or rules that enables a user to determine, for any given x in any given world w , whether or not x belongs in the output set of the associated function. The notion I have in mind here is very much like that of an algorithm; however, 'algorithm' is too narrow for current purposes due to the existence of uncomputable functions.

Linguistic definitions are one way of expressing classification procedures – they provide instructions, in the form of a set of necessary

and sufficient conditions, for determining whether or not a given entity belongs in the category being defined. But the fact that definitions purport to ‘attach’ such necessary and sufficient conditions to a linguistic item is, for the Practical Role Account, irrelevant. After all, I might communicate my desired classification procedure by offering up a definiens without a definiendum – e.g. by saying ‘in this book, our subject will be the medical condition defined by the following symptoms ...’, or even just ‘bring me all and only the books on that shelf which are authored by Dostoevsky’. I might even communicate my desired classification procedure non-linguistically – perhaps I demonstrate to a pre-verbal child, by means of repeated example, that the blue blocks should be selected from the displayed toys and placed in the basket.⁴

A classification procedure is an abstract object. It is something like a set of rules, rather than e.g. a set of psychological dispositions to sort in a certain way.⁵ Multiple subjects may share a classification procedure, in the sense that they may each possess sorting dispositions that (likely approximately) adhere to the procedure – much as multiple people can follow the same recipe. As such, the notion of a classification procedure resembles that of a ‘mode of presentation’; and therefore, insofar as ‘concept’ is sometimes used to mean an abstract Fregean sense, the notion I have in mind may start to sound quite like a concept. But ‘sense’, ‘concept’, and ‘mode of presentation’ come with quite a bit of baggage, often semantic baggage, which I would prefer to avoid. Perhaps concepts are senses, and perhaps senses are classification procedures. Perhaps not. Crucially, we don’t need to know the answer before engaging in conceptual engineering. It should be non-controversial that procedures for making classifications exist, and the Practical Role Account claims that those procedures are what the engineer aims to

⁴Plausibly, the dispositions that the child comes to adopt by such means would only approximately implement the classification procedure; the child is likely to make performance errors in non-optimal conditions, and there are likely to be cases where the child’s dispositions fail to result in a classification (perhaps cases of blue–green blocks, or other non-paradigmatic blue blocks). Indeed, I as an instructor may have intentions that underdetermine a precise classification procedure. I think, likewise, most actual mental dispositions in actual agents will only approximate (and will likely underdetermine) a classification procedure. Nonetheless, I think the precision of classification procedures (which are, as I note below, abstracta rather than sets of actual dispositions) makes them suitable as an ideal to which conceptual engineers ought to strive. Necessary and sufficient conditions fare poorly in capturing our actual concepts, but I see no reason to abandon them as the gold standard in conceptual engineering.

⁵But see also Haslanger (2020), who offers a similar view which holds intensions to be the contents of concepts, and characterizes possessing a concept in terms of possessing various dispositions with regard to that content. The current proposal differs in that it makes no claims about the semantic features or possession conditions of concepts; in addition, Haslanger’s primary focus is on semantic changes at the intensional level (specifically, changes to the content of a concept), rather than non-semantic changes at the hyperintensional level.

develop.⁶ Whether or not those procedures happen to *also* be invoked in the correct theory of concepts or meanings is an interesting, but peripheral, issue.

3. Practical roles and the limits of revision

Let's return now to the central claim of the Practical Role Account: that conceptual engineering is a matter of devising a tool that successfully fills an intended practical role. Given the discussion above, this can be tentatively elaborated as the claim that conceptual engineering is a matter of devising a *classification procedure* that successfully fills an intended practical role. Next up: what is a practical role?

The notion I have in mind is best framed by comparison to a group of similar views which I will call 'functionalist'. Functionalist views hold the *functions* or *purposes* of a concept to be in one way or another the core of the conceptual engineering enterprise. Cappelen devotes a chapter to arguing against this type of account in *Fixing Language*; he flags Haslanger (2000), Brigandt (2010) and Thomasson (2020) as exemplars of the approach. Other recent variants can be found in Fisher (2015), Prinzing (2018), and Simion and Kelp (2019). The following passage from Haslanger, outlining the sort of questions a conceptual engineer ought to ask, gives a good sense of the general spirit of functionalist accounts: 'What is the point of having these concepts? What cognitive or practical task do they (or should they) enable us to accomplish?' (Haslanger 2000, 33).

Functionalist accounts vary on just what a concept's function *is*. Some accounts tie a concept's function to its semantic features (e.g. Brigandt (2010) and Fisher (2015)). Thomasson (forthcoming) suggests making use of Millikan's notion of 'proper function', while Haslanger (2020) plumps for a version of the 'systems' approach to functions. But a fairly common thread running through the diversity is that an appeal to function can serve to answer Strawson's problem, via the claim that a proposed revision is permissible so long as the concept's *function* is preserved.

⁶Note that the engineer's target is the abstract procedure, rather than a set of dispositions in actual users. This is an important feature of the Practical Role Account, for it emphasizes the separation between 'invention' and 'implementation' – a point to which we'll return later. For now, simply note that on the Practical Role Account successfully engineering is a matter of inventing or discovering an effective/useful procedure; whether or not one's fellows actually come to have the corresponding classification dispositions is a separate concern.

The Practical Role Account departs from the above views in that it holds that the purposes which are relevant to conceptual engineering far exceed those that might comfortably be called ‘functions’. Both human artefacts and biological structures serve purposes other than their ‘function(s)’ – one can use a book as a paperweight, or one’s skin as a canvas for self-adornment. I’ve known several brave souls to use their teeth as bottle openers. The Practical Role Account holds suitability for such non-function purposes to be partially determinative of the success of an engineering proposal.

Moreover, the purposes recognized by the Practical Role Account are not in any sense essential or even intrinsic features of classification procedures. Cappelen expresses doubt about the idea that concepts have functions, but notes that ‘of course, *people* have goals and aims and purposes when they use words on particular occasions’ (Cappelen 2018). The Practical Role Account claims that it’s just *those* sorts of purposes that matter. If a group of subjects *S* uses classification procedure *x* for purpose *P*, then *P* is part of *x*’s practical role. And the success of a conceptual engineer’s proposed revision or replacement of *x* will be at least in some part determined by how well said revision/replacement contributes to *P*.

Emphasis on the ‘in some part’: in nearly all cases there will be a wide array of purposes for which a given *x* is used, and the ‘weight’ that each of those purposes should get when assessing a proposed change will be determined by all sorts of factors. How many people use *x* for purpose *P*? How often is *x* used for purpose *P*? How important is purpose *P* to *S*? Even then we’ve barely scratched the surface, because the Practical Role Account recognizes partial replacements, as well as replacements that are one-to-many or many-to-one. An engineer may offer up multiple successor procedures, each taking on some portion of the original’s practical role; she may suggest that the original be retained for some subset of its uses, and only replaced in certain specialized contexts. Thus, since one might replace procedure *x* with procedure *x*’ in only some contexts of use, often an engineer will only be concerned with a subset of *x*’s practical role, corresponding to the purposes for which it is used in a target set of contexts.⁷ In fact, partial replacement characterizes a large class of real-world instances of conceptual engineering: the introduction of specialized, rigorous definitions for everyday terms within professional

⁷Thanks to an anonymous reviewer for pointing out that purposes would need to be relative to contexts in this way.

contexts. Examples here include the use of ‘depression’ among psychiatrists, ‘weight’ among physicists, and ‘credence’ among philosophers.⁸

Further, purposes can be discarded. If an engineer comes to believe that a given purpose is undesirable or at least unneeded, the successor she devises is not obliged to retain it. For instance, racial concepts used to be implicated in pseudo-scientific explanations of variation in intellectual ability; this is clearly a theoretical purpose best rejected. Relatedly, some purposes will be sufficiently ‘fringe’ to be safely ignored – if a concept is used for a certain purpose by only a small minority of a community, for instance. This will be a matter of case-by-case judgment. But it should be noted that an engineer’s neglect of a ‘fringe’ function really only implies that there may be cases where certain community members retain use of the predecessor concept – similar to, e.g. on-call doctors retaining the use of pagers due to their more reliable signal reception, or aspiring novelists retaining the use of typewriters in order to project a certain air of hipster credibility.

Finally, the Practical Role Account holds that success in conceptual engineering is partially determined by yet *further* practical factors which are not even comfortably called ‘purposes’ – for instance, the procedure’s ease of learnability, or how likely it is to produce errors in use, or its tendency to produce such-and-so emotional response. This last is an instance of what Cappelen terms ‘lexical effects’; I would hold that similar effects can attach to classification procedures. Such effects will not, of course, be intrinsic features of the procedures as abstracta, but instead contingent facts about the effects that following such procedures have upon users.⁹ Roughly, the Practical Role Account holds that the relevant factors for determining engineering success encompass any reasonable answer to the Haslanger-style question ‘why should we

⁸To elaborate on one of these examples: when the technical, DSM-defined use of ‘depression’ was engineered, only the purposes to which ‘depression’ is put within the context of psychiatry were relevant. In common usage, ‘depression’ has several uses involving the flagging of general low mood – to explain behavior, to provide excuses, and so forth (as in, ‘sorry, I’m feeling a bit depressed today, I can’t make it to the party’). Many of these uses are not relevant within a psychiatric context, since they do not signal a need for treatment. The core use of ‘depression’ within psychiatry is to identify treatment needs, and thus this purpose receives the bulk of our ‘weighting’ when evaluating proposals for a DSM definition. Since the everyday purposes persist, the ‘ordinary’ concept of depression is retained in contexts where those uses arise – such as making excuses for absence at a party.

⁹Just as hearing a lexical item can trigger various emotional associations, so too can merely tokening a concept; thinking of rape generates anger just as hearing the term ‘rape’ does. Thus, if following a classification procedure involves tokening the concept RAPE, it will potentially trigger emotional effects. Note that while classification procedures are not themselves concepts in the psychological sense (they are instead abstracta), the psychological process that a user goes through to follow the procedure will involve the tokening of concepts. It is for this reason that, insofar as concepts can trigger associations and emotions, classification procedures can be said to produce an analogue to lexical effects.

employ this classification procedure?’ And reasonable answers to that question run the gamut from ‘it aids me in ascribing moral responsibility’, to ‘it factors into a systematic explanation of the tides’, to ‘it’s an easier-to-learn alternative to a more technical concept’, to ‘it puts a smile on my face when I think about it’.

One might worry about the liberality of this approach. Consider, for instance, Cappelen’s response to Brigandt’s (2010) functionalist account. Brigandt appeals to the ‘epistemic goal’ of a concept, where this is understood as involving the ‘inferences and explanations that the concept is intended to support’ (Brigandt 2010, 24). Cappelen argues that ‘there are too many speakers with too many and too varied intentions ... appealing to actual intentions of real speakers to fix the functions of concepts is unpromising as a general account of what “the core function” of a concept is’ (Cappelen 2018, 185). But the Practical Role Account doesn’t claim that classification procedures have a single ‘core function’ – the practical role of a classification procedure is constituted by *all* the various purposes that its users employ it for, and which of those purposes are relevant to a proposed engineering replacement (and to what degree) will vary by case. Is this sort of permissiveness a problem for the conceptual engineer? It would be, if complete preservation of practical role were a requirement of successful revision. But it is not. Here too, the Practical Role Account departs from the bulk of its functionalist cousins.

As noted earlier, functions are often invoked to constrain permissible revision. This can be done in a variety of ways. One might subscribe to some variety of metasemantic theory upon which functions are at least partially determinative of meaning, and then conclude from this that preservation of function will guarantee sufficient semantic similarity to avoid a change of subject. Or, one might simply directly identify a concept’s function as That Which May Not Be Altered, claiming that preservation of function, *rather* than preservation of subject, sets the boundaries of permissible revision.

The Practical Role Account handles things a bit differently. First, it is worth noting that the Practical Role Account implies that conceptual engineers don’t really *revise* at all; all instances of conceptual engineering are replacements. Procedures, presumably, have their ‘steps’ essentially – any alteration of such steps results in a different procedure. The categories (intension-candidates) which classification procedures pick out are functions, and therefore sets of ordered pairs; any change of membership results in a different category. An engineering proposal, then, is a proposal to replace, within some or all contexts of use, a given

classification procedure with an improved successor, and thereby (in most cases) replace a given category with another. Typically, use of a classification procedure is *correlated*, perhaps strongly, with the use of a word or the tokening of a concept; and thus a recommendation to replace a procedure will often amount to breaking one such correlation and implementing another. But the Practical Role Account doesn't claim that such correlations constitute or determine meaning/content. They might, if the correct metasemantic view turns out to be some sort of pseudo-Fregean, quasi-descriptivist view. But I wouldn't bet on it.

Let's look at a concrete example: the recommendation to include same-sex partnerships within the category of marriage. The Practical Role Account characterizes this as a proposal to replace a classification procedure roughly expressed by 'the category whose members are pairs consisting of one man and one woman who have entered into such-and-so formal recognition of their relationship' with a similar classification procedure which drops the requirement that the pair be of different genders. The classification procedure to be replaced is (or at least was) associated by many speakers with the term 'marriage', and the procedure associated with 'marriage' inherits an extensive practical role – legally, it is used to determine who receives various financial benefits, who is granted guardianship over children, and so forth, and it possesses numerous social and religious uses as well. The primary task of the conceptual engineer is to evaluate whether the proposed replacement procedure would be better suited to this practical role – or to some subset of this role, if the proposed replacement is partial.

Note that the above casts *procedures* as the bearers of practical roles. In fact, there's a tricky issue lurking here. Cappelen offers the following complaint against functionalist accounts generally – the complaint that the only function he sees for a concept/term is in enabling us to talk about its referent. 'The reason "salmon" is useful for us is that it can be used to talk about salmons' (Cappelen 2018, 187). True enough, but surely this just shows that there is a further, more pertinent question: why do we want to talk about *salmons*? And in this case the answer is plausibly something like, 'Because we take them to instantiate a natural subdivision of ray-finned fishes – and thus they form part of a scientific taxonomy that enables us to better understand the biological world'. Maybe. Or maybe the answer is, 'Because they taste good, and I need to tell the fishmonger what sort of thing I want for dinner'. Either answer provides a good reason to carve reality salmon-wise, rather than salmon-or-pufferfish-wise.

But perhaps an analogous worry applies to procedures; perhaps the only purpose a given classificatory procedure has is to enable us to pick out its corresponding intension-candidate function. So ought we to say that the bearer of a practical role is the procedure, or the function – or perhaps the entities selected by the function? There's an element of truth in all of these. One purpose to which salmon (the entities) are put is eating; and one purpose to which the function which picks out salmon is put is in selecting what to eat. We use that function for that purpose *via* deployment of a classification procedure. I think, then, that it is reasonably apt to say that 'selecting what to eat' is one purpose for which the classification procedure we associate with 'salmon' is used. Moreover, some factors which the Practical Role view holds to be relevant to engineering success plausibly accrue to the procedure itself rather than to the category or its members – for instance, ease of use. I propose, then, to retain procedures as the loci of the relevant practical roles; admittedly, though, this is to some degree an oversimplification.

Let's now more directly discuss the boundaries of permissible revision. First: note that, given all that has been said thus far, a proposed successor procedure need retain no *particular* element of the predecessor's practical role. We might make an analogy with the hire of a new employee – a new hire might take over a previous job role entirely, or multiple hires might be made to fill the role, or the previous employee might remain in a reduced role, or the tasks corresponding to that role might be redistributed to others, or some of the tasks might be eliminated from use ... and so forth. So long as all needed tasks are filled by *some* employee who is well suited to them, no arrangement ought to be deemed impermissible. *Mutatis mutandis* for the practical roles which classificatory procedures fill.

Roughly, then, the Practical Role Account permits a replacement of procedure x with procedure x' in contexts C_1 – C_n provided that, post-replacement, all purposes to which x is put in C_1 – C_n are either (1) sufficiently well-served by x' , or (2) justifiably abandoned. This leaves open what it is for x' to serve a purpose sufficiently well, and under what conditions abandoning a purpose is justified. But here the answers will be far too complex for any pithy guidelines. If x' serves purpose P better than x , then presumably that will be sufficient; but in other cases we may accept trade-offs, such as when the use of x' in context C results in a drop in effectiveness with regard to purpose P_1 , but a correspondingly greater increase in effectiveness with regard to purpose P_2 . The conditions under which such a tradeoff is worthwhile will vary case by

case. The same applies to the circumstances under which we may abandon purposes. As with Cappelen, I doubt any simple list of conditions will do the trick here. To attempt one would essentially be to attempt a complete theory of practical reasoning.

Now, given this picture of the boundaries of revision, topic changes certainly seem to be on the menu – procedure x and procedure x' might well delineate categories that are not plausibly the ‘same thing’. Is this problematic? I’d argue that it is not. ‘Changing the subject’ isn’t somehow *inherently* bad; if it is problematic, this must be because of its consequences. Strawson was concerned that Carnapian explication resulted in abandonment of the puzzles philosophers were concerned with; Cappelen’s helpful generalization of this worry is that topic changes threaten to disrupt inquiry. A further worry Cappelen suggests for topic-changing revisions is that they produce verbal disputes and miscommunication.

Let’s deal with these latter worries first. Verbal disputes are peripheral to the process of conceptual engineering as envisioned by the Practical Role Account; they result from retaining a linguistic item in contexts where a replacement of classification procedure has occurred. In situations where this is likely, the problem can be averted by substituting a neologism. Similar gambits should avert miscommunication, if carefully enough applied. This is a reasonably common practice already, but a move away from ‘semantic’ approaches to conceptual engineering ought to leave us with an even greater willingness to employ it. Cappelen’s Austerity Framework doesn’t cover introductions of new terminology – he writes that he doesn’t ‘have much new to say about introducing terminology from scratch’ (Cappelen 2018, 37). The Practical Role Account, by contrast, doesn’t hold such cases to be different in type from the cases Austerity does cover.

Continuity of inquiry is, at least *prima facie*, a more serious worry. If I propose replacing the classification procedure currently used in contexts associated with the term ‘woman’ with a classification procedure which picks out all and only can-openers, I can rightly be accused of abandoning the questions that, say, feminist philosophers are concerned with. And there is certainly a sense in which changing the topic from F to G invariably disrupts any inquiry into the nature of F .

But there is inquiry, and there is *inquiry*. In many cases, ceasing inquiry into F in order to inquire into G instead is conducive to progress in some *higher-level* inquiry. To progress in the understanding of combustion, inquiry into phlogiston had to be abandoned. To progress in the study

of innate mental faculties, inquiry into various phrenological protrusions and had to be abandoned. Numerous biological taxa have been abandoned since Linnaeus, as have many medical and psychiatric classifications. Along the way, all sorts of once-central questions regarding those categories have ceased to hold interest. What is the mass of phlogiston? Is the brachycephalic skull type characteristic of darker-skinned races? What is the most effective treatment for female hysteria?

The moral of the story is twofold. First, topic-changing replacements are consistent with continuity of inquiry, when inquiries are not individuated too narrowly. Second, when inquiries *are* individuated narrowly, disruption of inquiry may well be a good thing: we may simply be focused on the wrong topic, and therefore *asking the wrong questions*. Strawson worried about leaving time-honoured philosophical puzzles unresolved. But some topics just may not be worth puzzling over – we are mortal, and must be selective about which carvings of logical space warrant study.

So suppose an engineered replacement for a concept – let's say, *free will* – serves all the needed purposes to which its predecessor was put, and does so more effectively to boot. And suppose a philosopher then objects that the proposed replacement abandons our initial concern, which was to investigate the nature of *free will*. On the Practical Role Account, the engineer's reply should simply be: why exactly should we care about investigating *that*? And what possible answer could be given to this question which did *not* advert to some practical or theoretical role that the concept plays in our lives – which, by hypothesis, the replacement better serves? Of course, perhaps the philosopher simply has a brute interest in the particular portion of reality picked out by the English expression 'free will'; in which case she, like the typewriter-wielding hipster, may feel free to retain the predecessor for her own idiosyncratic inquiries. But she ought not protest when the rest of us shift our inquiries to those portions of reality which further our practical interests.

Both the Austerity Framework and the Practical Role Account, then, have answers to Strawson's problem. Which answer is superior? I don't take the Practical Role Account to be the knock-down, obvious winner here, but its greater flexibility is intuitively an advantage. For instance, this flexibility enables the Practical Role Account to cover certain exemplars of conceptual engineering that threaten to slip through the cracks of the Austerity Framework – such as Kevin Scharp's (2007, 2013) work on the truth-concept. Cappelen notes that Scharp explicitly portrays his project as replacement rather than revision; moreover, I'd argue that his suggested replacement results in a change of topic. In brief, Scharp

proposes replacing the truth-concept (in certain contexts) with two successor concepts, called ‘ascending truth’ and ‘descending truth’. These two concepts have different inferential properties, corresponding to different directions of the T-schema – if p , then ‘ p ’ is ascending true; if ‘ p ’ is descending true, then p . Now, the issue is that samesaying doesn’t seem to occur across these two concepts. If I say “ p ’ is ascending true’, I have not said the same thing as someone who says “ p ’ is descending true’. But suppose that samesaying *does* occur over uses of ‘true’ and ‘ascending true’, and over ‘true’ and ‘descending true’. It would follow that we have a failure of transitivity for samesaying. This would seem to be problematic, for samesaying and sameness of topic presumably involve *identity* of ‘what is said’ and of topic, and identity is a paradigmatically transitive relation. Thus, we should conclude that samesaying does *not* occur over uses of ‘true’ and Scharp’s successor concepts, and therefore that there has been a change of topic.

At worst, then, the Austerity Framework implies that Scharp’s engineering exceeds the limits of revision (I would maintain it clearly does not). At best, Cappelen will have to argue that cases like Scharp’s fall out of the scope of the Framework, along with e.g. *de novo* conceptual inventions. The Practical Role Account, meanwhile, can happily embrace Scharp’s proposals – so long as his conceptual subdivision does not somehow prevent one of the original purposes of ‘truth’ from being served.

The Practical Role Account exhibits this nice flexibility and inclusiveness while retaining a principled response to Strawson’s concern about neglected questions. Though the *exact* questions that prompted investigation of a pre-engineering category may go unanswered, if our engineering work attends carefully to the relevant practical roles then appropriate analogue questions will remain topics of investigation. Conceptual engineering goes hand in hand with engineering ways of framing inquiry; when we attempt to determine which concepts we *should* use, it’s only natural that this will bleed into an investigation of which questions we *should* ask. The Practical Role Account’s approach to both is the same: we should use the concepts, and ask the questions, which best serve our aims – explaining the world around us, predicting future events, enabling social harmony, or simply figuring out what to have for dinner.

4. The Practical Role Account and the question of control

Let’s turn now to the question of control. The Austerity Framework implies conceptual engineering is largely out of our hands; can the

Practical Role Account do better? To answer this, it will help to delve a bit deeper into the relationship between classification procedures and language. As noted earlier, there are correlations between use of a term and use of a classification procedure, in the following sense: for any term *T*, a given speaker will be disposed to classify certain objects as falling under *T*, and those dispositions will approximately, likely underdeterminately, conform to a classification procedure. But there is no claim here that this determines the meaning of a speaker's utterance. If the correct metasemantics is externalist, then there will be innumerable cases where the meaning of a speaker's utterance and the classification procedure that best captures her sorting dispositions come apart: as in, e.g. a Burge-style case where a confused patient tells her doctor that she has arthritis in her thigh.¹⁰

Though there is a fairly strong correlation between words and procedures, conceptual engineering *isn't* about what our words should mean, or even about how we should use our words. It is about *how we should classify*. When a conceptual engineer proposes replacing a procedure *x* with a procedure *x'* in some contexts of use, this is not reducible to a recommendation to change how a correlated term is used. Much of the work our classification practices do occurs at the level of thought, entering into inferences and decision-making without a single syllable being uttered. Consider, as an example, the term 'person'. The classificatory procedure a pregnant woman associates with 'person' may affect whether she chooses to seek an abortion – a decision-making process which may occur without any verbalization at all. Note too that the actual semantic facts about 'person' aren't a direct predictor of her decision here; just as the actual meaning of 'arthritis' does not prevent Burge's imagined patient from seeking medical treatment for a thigh-ache. If we want our conceptual engineering interventions to affect how people infer and behave, then changing the meaning of a term seems a rather inefficient stratagem. Why not target the classificatory practice directly?

¹⁰Following Burge's (1979) case more closely, we can further note that the content of a subject's *beliefs* need not correspond to the classification procedure she approximates when tokening a *concept*. If externalism about mental content is true, then the subject may have a false belief *about arthritis* because the classification procedure her dispositions instantiate picks out a function whose outputs include cases of non-joint inflammation. Note that not all false beliefs about arthritis will work this way; employing the correct classification procedure doesn't give one omniscient access to a category's properties without information about which world one occupies. The classification procedure expressed by 'x is arthritis iff x is a case of joint inflammation' likely fixes on the actual intension of 'arthritis'; nonetheless I might correctly associate this procedure with 'arthritis' while still having all sorts of false beliefs about the prevalence of arthritis, how to treat it, and so forth.

Of course, if an engineer wishes to affect how others classify, she will need to convince others of the utility of her preferred classification procedure. And to do that, she will need to communicate it. The most straightforward way to do so is through language, via use of a stipulative definition. So what happens, semantically, when an engineer offers up such a stipulative definition? Plausibly, a single act of stipulation can't alter the meaning of an entrenched term like 'woman' or 'marriage'. But two other scenarios do seem possible. First, the engineer might succeed in introducing, via stipulation, a new term which is a homonym of the original. One might claim, for instance, that when Haslanger (2000) suggests modifying the concept 'woman', what she is really doing (perhaps unbeknownst to her) is stipulating a new term which is a homonym to the English term 'woman'.

In fact, I see no reason for Cappelen to object to this possibility; meaning-fixing via stipulation is not incompatible with externalism. Furthermore, samesaying can quite plausibly occur over at least some cases of stipulative replacement. Samesaying can cover cases of non-type-identical linguistic expressions – an example Cappelen has used in the past is that I may report Galileo as saying that the earth moves, though of course he spoke in Italian. And plausibly, samesaying can occur even over non-identical linguistic items that vary somewhat in intension/extension. The Chinese character 羊 refers to both sheep and goats, but I can correctly report a Chinese-speaking friend when I say that she told me she was born in the Year of the Goat. So I'd suggest that samesaying (and topic preservation) can at least in principle occur even if I stipulate a homonymic replacement for a target of conceptual engineering.

Alternately, perhaps the engineer does not introduce a new homonym, but instead simply produces an utterance where *speaker-meaning* diverges from *semantic-meaning*; the engineer specifies what *she* will mean by term T, and an attentive interlocutor will get the message regardless of what T actually means. This possibility has been championed by Mark Pinder (2019), who notes that it renders the conceptual engineering process not only within our control, but almost trivially easy to achieve. Cappelen, however, is not convinced by the speaker-meaning gambit. He writes:

At best your definition will tell charitable readers how to get at the speaker's meaning, i.e. what you had in mind. However, even that isn't a safe space in which we have control. The content of what's called 'speaker's meaning' is just as externalistically determined as linguistic meaning. (Cappelen 2018, 76)

Even stipulation, it seems, can't free me from the shackles of externalism – telling my readers what I mean doesn't guarantee that I mean it.

The thought here, presumably, is that speaker-meaning depends on what one intends to mean, and the contents of mental states like intentions are externalistically determined. This is technically to add an additional commitment to Austerity – externalism about language doesn't *entail* externalism about mental content. Pinder's response is as follows:

even if mental content is externalistically determined, I have significant control over what I intend. For example, when uttering 'Luddite', I control whether I intend to convey *Luddite* or, alternatively, *one who does not use social media and is opposed to it due to its negative effects*. (Pinder 2019, 15)

This is true enough, but we might reply on Cappelen's behalf that lack of control can still occur over the *internal* elements of the relevant definition. Suppose, for instance, I offer the following rather odd stipulation: 'to be disabled is to have one of the following conditions: muscular dystrophy, blindness, or severe arthritis'. And suppose that I have a Burge-style misunderstanding of the nature of arthritis. In such circumstances, my speaker-meaning would plausibly slip out of my control; for presumably, the content of my communicative intention involves *arthritis*. Though I have, via my stipulation, ceased to defer to experts on the meaning of 'disabled', I continue to defer on the meaning of 'arthritis'. Unbeknownst to me, then, the definition that I speaker-mean in such a case does not cover thigh inflammation.

Fortunately, the Practical Role Account's focus on classification procedures allows us to say that the degree of control we have over our speaker-meaning is ultimately beside the point. On the Practical Role Account, whether or not an instance of conceptual engineering succeeds simply isn't determined by whether one speaker-meant *y* by one's utterance of *x*. As Max Deutsch has recently noted, speaker-meaning something (and perhaps convincing others to do so as well) 'seems like a rather trivial and easy thing to do. Surely it is not the sort of thing the exciting terminology of "conceptual engineering" was designed to describe' (Deutsch 2020, 7). It's plausible, I think, that Pinder is right about what happens when an engineer offers up a stipulative definition. But successfully speaker-meaning a given content with one's stipulation isn't the *goal* of conceptual engineering; it's a way to communicate the classification procedure that one has engineered. On the Practical Role Account, the primary goal of the engineer is to determine which

classification procedure best suits a given practical role. That's not trivially easy – it's extremely hard. But it is just as much within our control as, say, developing an elegant proof or inventing an effective machine; our success is limited only by our own ingenuity.

We should distinguish between the goal of inventing or discovering an effective procedure, and the goal of convincing others to *use* our invention – that is, to adopt actual psychological dispositions to classify in accordance with the procedure. The former is engineering; the latter is advertising. And it is the former that strikes me as the philosopher's primary task. Philosophers do sometimes work to popularize their arguments and theories, but the degree to which they succeed isn't determinative of philosophical success. The fact that Dan Dennett hasn't convinced the entire American public that God does not exist doesn't impugn standard arguments for atheism. *Mutatis mutandis* for pretty much every philosophical argument or theory; and *mutatis mutandis* for conceptual engineering. Philosophers generate all kinds of theories, even very 'worldly' ones having to do with just wars, medical ethics, artistic value, and so forth. They rarely hold out hope that these will become common coin outside of the ivory tower. So conceptual engineers, I'd argue, simply share their lot with the rest of the philosophical community: even the most brilliant conceptual engineer can, at best, aspire to be among the happy few whose philosophical thought will trickle down just enough to seep into the zeitgeist of future eras.

Nonetheless, most conceptual engineers *are* interested in advertising their proposals – futile as it may be to, say, convince the general population to employ Haslanger's definition of 'woman'. So how much control do we have over disseminating our discoveries? As we've seen, even the speaker-meaning of a stipulative definition can come apart somewhat from the classification procedure the speaker associates with the definiendum. So there is always potential for miscommunication when we promote our proposed replacement. But this applies to all communication whatsoever – conceptual engineering doesn't generate any *special* worry here. In fact, use of a stipulative definition presumably mitigates, as much as linguistic communication allows, the potential for miscommunication.

We have plenty of evidence that the use of stipulative definitions *can* succeed in changing classification behaviour, particularly within a limited group. Technical terminology within various professions is a clear example. There is a widespread practice of employing stipulative definitions in philosophy, the sciences, law, and beyond, and it seems to do pretty well in enabling the sort of classificatory changes the

stipulators intend. Indeed, it's hard to imagine successfully diagnosing illnesses, enforcing legal codes, calculating spaceship trajectories, and so forth *without* such a practice. The messy, workaday concepts of everyday life just don't hold up to the definitions experts have offered up for 'gastroenteritis', 'mens rea', or 'inertial mass'. Getting professionals to adopt a classificatory practice that adheres to a proposed definition may be a bit of a challenge, but the means to doing so are clear – one offers up arguments for the merits of one's classificatory scheme, and hopes for the best. One's control is limited by one's rhetorical skills (or perhaps one's influence and popularity), but nothing more.

How difficult is it to know whether one's engineering proposal has been taken up by others? We don't have direct epistemic access to the classificatory dispositions of our fellows, but their behaviour and verbal reports can give us a pretty good indication of whether the desired change has taken effect. Again, no unique problem for conceptual engineering here – our epistemic access to other minds is always indirect. Capelen suggests that conceptual engineering success is opaque, because we cannot hope to know all the various facts (about usage patterns, baptisms, etc) that determine a given term's meaning. We therefore cannot hope to know when (or if) a term's meaning has changed. But it's much easier to determine whether or not classificatory practices have changed. For instance, we can run a poll to see how many Americans believe that transgender women count as women. This won't tell us whether the meaning of 'woman' has changed; but it will tell us quite a lot about the classificatory procedures Americans use. There is, then, no epistemic barrier to conceptual engineering on the Practical Role Account – or, at least none that exceeds the normal lack of certainty that accompanies inquiry generally.

The Big Picture is as follows. Conceptual engineering proposals are generally offered up in the form of stipulative definitions. But this does not commit us to the view that successful engineering requires changing the meaning of the defined term. Plausibly, the actual semantic upshots of presenting a conceptual engineering proposal will be as Pinder states – the proposer speaker-means a new content for the term at issue. This suffices to communicate her intended classification-procedure, which a receptive hearer will then utilize to make future classifications in appropriate contexts. However, ultimately this is all to do with 'advertising'. One nice feature of the move away from a semantic approach is that advertising can be viewed as secondary to engineering proper – that is, secondary to the actual discovery or creation of a classification procedure that effectively

fills a needed practical role. Since a semantic approach views conceptual engineering as a change of meaning, which presumably requires widespread change in usage patterns, advertising is for such views an intrinsic part of the process. By contrast, the Practical Role theorist can claim that there is no more a control problem for conceptual engineering *proper* than there is for devising a mathematical proof. Moreover, since the alteration of actual dispositions to classify is much more within our control than the alteration of a term's meaning, the Practical Role Account's view on advertising permits substantial control in that arena as well.

5. Conclusion

This paper has contrasted two potential accounts of conceptual engineering – Cappelen's Austerity Framework, and my own Practical Role Account. The Austerity Framework conceives of conceptual engineering as a linguistic process, and teases out various limitations on that process as a consequence of metasemantic theory. The Practical Role Account, by contrast, views conceptual engineering as a process of inventing or improving methods for carving up reality – for categorizing, for classifying, for making distinctions. Such methods, considered abstractly, need not be identified with either the concepts we possess or with the meanings of the words of our languages. Thus a conceptual engineer is not obligated to endorse any particular metasemantic view, nor any particular view on concepts.

The fundamental perspective of the Practical Role Account is pragmatic. We divide up the world in various ways for *reasons* – often to do with explanation and prediction, but not always. Insofar as a given division 'works better', that's all the justification that's needed for its employment. Efficacy in a practical role is, therefore, the criterion of success; semantic features can sometimes be a guide to efficacy, but metasemantic facts place no constraints on the conceptual engineer. We need not wring our hands over questions of meaning change, or change in the contents of our concepts. The Practical Role Account, by rejecting the semantic view of conceptual engineering, puts engineering success wholly within reach – and places us solidly in control.

Acknowledgements

Thanks to Herman Cappelen and Max Deutsch for helpful input on earlier versions of this paper.

Disclosure statement

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

Funding

The work described in this paper was fully supported by a Grant from the Research Grants Council, University Grants Committee Council of the Hong Kong Special Administrative Region, China (Project No. HKU 13603718).

ORCID

Jennifer Nado  <http://orcid.org/0000-0003-1680-2975>

References

- Brigandt, I. 2010. "The Epistemic Goal of a Concept: Accounting for the Rationality of Semantic Change and Variation." *Synthese* 177: 19–40.
- Burge, T.. 1979. "Individualism and the Mental." *Midwest Studies in Philosophy* 4: 73–121.
- Cappelen, H. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Deutsch, M. 2020. "Speaker's Reference, Stipulation, and a Dilemma for Conceptual Engineers." *Philosophical Studies*. doi:10.1007/s11098-020-01416-z.
- Evans, G. 1973. "The Causal Theory of Names." *Proceedings of the Aristotelian Society, Supplementary Volumes* 47: 187–225.
- Fisher, J. 2015. "Pragmatic Experimental Philosophy." *Philosophical Psychology* 28: 412–433.
- Haslanger, S. 2000. "Gender and Race: (What) Are They? (What) Do We Want Them to Be?" *Noûs* 34: 31–55.
- Haslanger, S. 2020. "How not to Change the Subject." In *Shifting Concepts: The Philosophy and Psychology of Conceptual Variation*, edited by Teresa Marques and Åsa Wikforss, 235–249. Oxford: Oxford University Press.
- Pinder, M. 2019. "Conceptual Engineering, Metasemantic Externalism and Speaker-Meaning." *Mind; A Quarterly Review of Psychology and Philosophy*. doi:10.1093/mind/fzz069.
- Prinzing, M. 2018. "The Revisionist's Rubric: Conceptual Engineering and the Discontinuity Objection." *Inquiry* 61: 854–880.
- Schärp, K. 2007. "Replacing Truth." *Inquiry* 50: 606–621.
- Schärp, K. 2013. *Replacing Truth*. Oxford: Oxford University Press.
- Simion, M., and C. Kelp. 2019. "Conceptual Innovation, Function First." *Nous (Detroit, Mich)*, doi:10.1111/nous.12302.

- Strawson, P. 1963. "Carnap's Views on Conceptual Systems Versus Natural Systems in Analytic Philosophy." In *The Philosophy of Rudolf Carnap*, edited by Schilpp Paul Arthur, 503–518. La Salle, ILL: Open Court.
- Thomasson, A. 2020. "A Pragmatic Method for Conceptual Ethics." In *Conceptual Ethics and Conceptual Engineering*, edited by H. Cappelen, D. Plunkett, and A. Burgess, 435–458. Oxford: Oxford University Press.