The Anti-Conceptual Engineering Argument and the Problem of Implementation

Steffen Koch

Bielefeld University

Steffen.koch@uni-bielefeld.de

ABSTRACT

Conceptual engineering concerns the assessment and improvement of our concepts. But how can proposals to engineer concepts be implemented in the real world? This is known as the *implementation challenge* to conceptual engineering. In this paper, I am concerned with the meta-philosophical implications of the implementation challenge. Specifically, must we overcome the implementation challenge prior to undertaking conceptual engineering? Some critics have recently answered this question affirmatively. I intend to show that they are mistaken. I argue as follows. First, successful implementation is not an integral part of conceptual engineering. Second, the idea that the value of conceptual engineering relies on successful implementation is in tension with widespread assumptions about normative theorizing.

1. INTRODUCTION

The central idea behind conceptual engineering is that we should not be complacent with the concepts that we have at our disposal. Rather, we should critically assess how we can improve them. But, what does it take to actually implement proposals for engineering concepts? Is this even feasible? This question is often referred to as the *implementation challenge* to conceptual engineering. A particularly troubling version of the implementation challenge arises for conceptual engineers who aim to change the meaning of words in a shared language (such as English). According to Herman Cappelen (2018), word meanings are typically determined by rather messy and complicated arrays of factors, factors that are for the most part inscrutable. Some of the best extant metasemantic theories tell us that these factors are largely beyond our control (see, for instance, Deutsch 2020a). The implementation challenge has been discussed in the literature, but it remains a controversial issue (Deutsch 2021; Jorem 2021; Koch 2021a, 2021b; Nimtz 2021; Pinder 2021; Thomasson 2021).

In this paper, I address the concerns relating to the *meta-philosophical implications* of the implementation challenge. I am specifically concerned with whether or not the implementation challenge must be answered before we can pursue the project of conceptual engineering any further. Some scholars have recently suggested that there is a strong connection between the prospects of implementation and whether or not conceptual engineering is worthwhile. Amie Thomasson, for instance, states that if

> there is no way to *implement it* – to engage in conceptual activism that would involve genuinely changing the linguistic or conceptual scheme we use…then conceptual engineering might seem like a pointless intellectual exercise. (2021, p. 6; italics in the original)

James Andow is concerned that "the success rate of conceptual engineering may be low enough to make conceptual engineering not worth bothering with" (2021, p. 218). Max Deutsch

thinks that an evident infeasibility in implementing conceptual engineering implies that "conceptual engineering…doesn't belong in the philosopher's tool box" (2020a, p. 22).[1]

There is some outward plausibility to these claims. I shall however argue that they are ultimately mistaken. On the view I shall defend here, concerns about the feasibility of conceptual engineering provide insufficient grounds for rejecting conceptual engineering as a philosophical enterprise.

In making my argument, I shall reconstruct and rebut two versions of the so-called *anti-conceptual engineering argument* (ACEA). (ACEA) aims to demonstrate the futility of conceptual engineering based on concerns about the feasibility of its implementation. There are two versions of (ACEA): (1) a *direct* version and (2) an *indirect* version.

1.   According to the direct version of (ACEA), concerns about implementation should prevent us from engaging in conceptual engineering. This is because implementation is an essential ingredient in conceptual engineering.

2.   According to the indirect version of (ACEA), implementation may not be an essential ingredient in conceptual engineering, but it is still a core motive for engaging in conceptual engineering.

I aim to demonstrate how both the direct and the indirect version of (ACEA) are unsound.

Distinguishing, reconstructing, and ultimately rebutting these two arguments yields two novel insights:

1.   It shows which premises (aside from skepticism about the feasibility of implementation) are required to arrive at an anti-conceptual engineering conclusion. Interestingly, these premises have not been explicitly stated yet.

2.   If successful, my arguments pave the way for an ecumenical defense of conceptual engineering. Such a defense is premised neither on controversial metasemantic assumptions nor on a specific theory of what conceptual engineers are

supposed to be engineering. I consider this to be a significant advantage over extant attempts to save conceptual engineering from the implementation challenge.

2. STAGE SETTING

In this section, I develop a schema for (ACEA). I shall then discuss and ultimately rebut two versions of this schema: the direct version and the indirect version. However, the first thing to do is to explicate some assumptions about conceptual engineering. These assumptions will guide the discussion that follows. There are three assumptions about conceptual engineering that are important for our purposes. Let us call these (1) (activity), (2) (re-engineering), and (3) (word meanings).

Regarding (activity), I construe conceptual engineering as an activity. Conceptual engineering is currently described as many things: a method, a meta-philosophical position, a field of research, or even a movement. Indeed, it might partly be all these things at once. However, the activity of conceptual engineering appears to be more basic. Starting from the activity of conceptual engineering, we can, for example, define its associated method as the method of engaging in the activity of conceptual engineering. The corresponding meta-philosophical position is one stipulating that philosophers do or should engage in the activity of conceptual engineering. I shall refer to planful kinds of conceptual engineering as conceptual engineering projects. I take this to be consistent with the view that conceptual engineering is, at heart, an activity.

Next is (re-engineering). Following what appears to be the convention, I take conceptual engineering to be concerned with the assessment and improvement of our concepts (where the relevant sense of "concept" is yet to be specified). In principle, sets of concepts can be improved by adding new members, eliminating old members, or changing existing members. Strictly speaking, there are then three forms of conceptual engineering: *elimination*, *introduction*, and *conceptual re-engineering*. My focus in this paper is on conceptual re-engineering. When I use the term "conceptual engineering," I am therefore referring to conceptual re-engineering.

Regarding (word meanings), I shall follow writers such as Cappelen (2018) and Thomasson (2021) in focusing on a semantic kind of conceptual engineering concerned with word meanings. This contrasts with concepts in a psychological sense (Isaac 2021), classification procedures (Nado 2021), or speaker-meanings (Pinder 2021). It is controversial both whether or not the semantic kind of conceptual engineering represents paradigm cases, and whether or not it is the most fruitful approach to pursue. I focus on the semantic kind of conceptual engineering for the simple dialectical reason that it is targeted by the objection I address below. I return to other kinds of conceptual engineering in my concluding remarks.

Given (activity), (re-engineering), and (word meanings), the central question in this paper can be more precisely formulated as follows:

> *Should concerns about the feasibility of implementing changes to word meanings stop us from engaging in the activity of conceptual engineering?*

While critics of conceptual engineering say "yes," I shall argue that we should say "no." The challenge for critics is to find a plausible candidate for (P3) in the following argument schema.

> The *anti-conceptual engineering argument* (ACEA):
>
> > (P1)   Conceptual engineering is the activity of assessing and improving the meanings of words in a shared language.
> >
> > (P2)   Changing what words mean in a shared language is infeasible.
> >
> > (P3)   […]
> >
> > (C)   Therefore: Conceptual engineering is futile.

(P1) and (P2) are not uncontroversial. Regarding (P1), some have questioned whether conceptual engineering actually concerns the meaning of words in a shared language (Jorem 2021; Nado 2021; Pinder 2021). Regarding (P2), I have argued in other works that language speakers possess a kind of collective long-range control over word meanings (Koch 2021a, 2021b). This long-range control renders intentional meaning-change feasible after all. Christian

Nimtz (2021) argues likewise that would-be conceptual engineers can implement their proposals via manipulations of social norms governing the relevant discourse.

I shall nonetheless grant opponents of conceptual engineering the truth of (P1) and (P2). (P1) springs out of (activity), (word meanings) and (re-engineering), and I am subscribing to these assumptions for the purposes of this paper. I shall also grant that there are, at least, some readings of "infeasible" on which (P2) is true. My concern from here on will be with what to insert for (P3). I shall argue that there is no insertion for (P3) that renders (ACEA) sound.

It is, of course, impossible to consider each and every possible insertion for (P3). I shall therefore focus on two general insertion strategies.

The first strategy proceeds as follows. According to (P1), an integral part of the activity of conceptual engineering is to improve the meaning a word has in a shared language. To do so, it is however insufficient to merely suggest a better meaning. One must also actually implement this meaning. But, (P2) tells us that doing so is infeasible, and so doing conceptual engineering is infeasible. From here, we merely need the extra premise stating that it is futile to engage in activities that are infeasible for arriving at (C). This strategy yields what I shall call the *direct anti-conceptual engineering argument* (D-ACEA).

The second strategy does not follow the assumption that implementation is an integral part of conceptual engineering. That said, it does holds that the value of assessing and designing word meanings is closely tied to the feasibility of implementing the ensuing proposals. Should it turn out that implementing conceptual engineering is infeasible, then conceptual engineering would be futile. This strategy yields what I shall call the *indirect anti-conceptual engineering argument* (I-ACEA).

The goal of this paper is to show that neither (D-ACEA) nor (I-ACEA) are sound. The absence of other antecedently plausible insertion strategies leads me to conclude that (ACEA) fails.

3. THE DIRECT ANTI-CONCEPTUAL ENGINEERING ARGUMENT

On the first proposal I shall discuss, (ACEA) can be completed as follows:

The *direct anti-conceptual engineering argument* (D-ACEA):

> (P1) Conceptual engineering is the activity of assessing and improving the meanings of words in a shared language.
>
> (P2) Changing what words mean in a shared language is infeasible.
>
> (P3i) Changing what words mean in a shared language is an integral part of the activity of conceptual engineering.
>
> (P4) Engaging in activities that are infeasible is futile.
>
> (C) Therefore, conceptual engineering is futile.

I mention (P4) only because it is needed to render the argument valid. I shall not take issue with it. The important new element is (P3i).

Is implementation an integral part of the activity of conceptual engineering? The characterization of conceptual engineering given above does not settle this question. This is because of an ambiguity in what it can mean to *improve* a concept or *improve* meaning.[2] We might think of concepts or meanings as abstract entities, and consider them improved once we have formulated new application conditions that are superior to their actual ones. But, we might also look at concepts or meanings as concrete social practices. Or, we might look at them as the psychological dispositions people have. Understood this way, a concept or meaning is not improved until the relevant practice or relevant dispositions have changed. We have to move beyond the popular characterization of conceptual engineering as "assessing and improving concepts" if we want to discern whether implementation belongs to conceptual engineering.

It may be useful to begin with three activities that are often distinguished in the literature:

1. The assessment or evaluation of existing concepts.

2. The design or construction of new concepts (unlike the "improvement" terminology, speaking of "design" or "construction" makes it clear that this does not involve implementation).

3. The implementation of these new concepts (e.g., by making them the meanings of relevant pre-existing terms).

Once we have distinguished these three activities, we may ask which of them belongs to the overall activity of conceptual engineering.

Alexis Burgess and David Plunkett (2021) call engaging in the assessment or evaluation of concepts "conceptual ethics." They argue that we can engage in conceptual ethics without constructing new concepts or implementing the outcomes of concept assessments. As in various other domains of philosophical inquiry, we can aim to establish normative or evaluative facts about our concepts purely for the sake of intellectual progress.

The same holds for the subsequent construction of new concepts. An application of conceptual assessment may yield a partly negative verdict (e.g., that a certain concept should be reformed). When this happens, we can engage in the construction of an improved successor concept simply out of intellectual curiosity. One might argue for a particular revisionary definition while refraining from advertising its use to one's broader community. Assessing and designing concepts are intellectual activities that can be pursued in separation from the activism required for implementation. If this is so, then the project of assessing and suggesting improvements to our conceptual apparatus becomes viable.

I now argue that there are good reasons to use the label "conceptual engineering" to refer to the project I have outlined. This is more suitable than using the label to refer to the larger project of assessing, suggesting and implementing improvements to our conceptual apparatus.

Recall that, according to the view currently under consideration, conceptual engineering concerns word meanings in shared languages (e.g., the meaning of English such as "to know" or "to believe"). Consequently, implementing conceptual engineering means actually changing the meaning a word has in some shared language (e.g., to make it the case that "to know" means something different from what it currently means).

There are, at least, two reasons for not counting implementation (in the sense described above) as an integral part of conceptual engineering.

The first reason is that conceptual engineering is often described as a *philosophical* (or at least intellectual) activity. It is somewhat analogous to conceptual analysis. Some have suggested that conceptual engineering can fill the methodological hole left by the demise of conceptual analysis (Machery 2017; Nado 2019; Schupbach 2017; Shepherd and Justus 2015). However, while assessing and designing concepts are philosophical activities, successful implementation requires a kind of hands-on approach. Such an approach falls outside the domain of philosophy (and academic disciplines in general). Recall that to implement a conceptual improvement, one must change the meaning of a word in a shared language. The intellectual endeavors philosophers undertake do not come close to achieving this. Philosophers attempt to convince each other of their views, and they disseminate their arguments through publications and talks. Dissemination is at best, though, a very small step toward implementation. Word meanings in shared languages are hardly influenced by publications and talks to any significant degree. To be clear, I am not suggesting that implementation is impossible. My claim is that, to the extent that implementation is possible, it requires the kind of tangible actions that fall outside the remit of philosophy *qua* academic discipline. It is often suggested that conceptual engineering is a philosophical project. If so, then implementation is not part of it.

What if we scale down from *successful* implementation to something like *making a serious attempt* at implementation? Perhaps, philosophers' typical means of dissemination – giving talks and publishing papers – rarely suffice to successfully implement a given conceptual engineering proposal. They may nonetheless be regarded as serious attempts to do so. I agree that this modification would make the view more plausible. However, it makes things worse, rather than better, for defenders of (D-ACEA). If we swap "implementation" for "serious attempt at implementation," and if we construe "implementation" as "changing what words mean," then (D-ACEA) should be reformulated as follows:

The *direct anti-conceptual engineering argument** (D-ACEA*):

(P1) Conceptual engineering is the activity of assessing and improving the meanings of words in a shared language.

(P2*) Making a serious attempt at changing what words mean in a shared language is infeasible.

(P3i*) Making a serious attempt at changing what words mean in a shared language is an integral part of the activity of conceptual engineering.

(P4) Engaging in activities that are infeasible is futile.

(C) Therefore, conceptual engineering is futile.

(D-ACEA*) has an obvious problem: (P2*) is false. If what philosophers such as Sally Haslanger or Kate Manne do qualifies as making a serious attempt at changing what words mean (as per the present proposal), then clearly making such an attempt is feasible. So, while the above rejoinder can save the plausibility of (P3i), it does so at the cost of (P2). *Mutatis mutandis*, the same goes for other ways of watering down the implementation requirement in (D-ACEA). For instance, if we say that it is sufficient for philosophers to *contribute* (even a little) to changing what words mean, then (P2) remains false.

The second reason for not counting implementation as an integral part of conceptual engineering relates to the nature of paradigm instances of conceptual engineering. The cases that are typically cited as paradigm instances of conceptual engineering include Andy Clark's and David Chalmers' (1998) extension of our concept of belief, Haslanger's (2012) revisionary analysis of race and gender concepts, or Manne's (2017) account of misogyny (cf. Cappelen 2018; Sawyer 2020). These cases arguably involve the assessment of old concepts, and then the design of new ones. Yet, none of them changed what the relevant expressions mean in a shared language (e.g., what "belief" means in English).

Importantly, such paradigm cases do not only serve to illustrate the general idea behind conceptual engineering. Arguably, they also *fix* the very meaning of "conceptual engineering." Cappelen considers conceptual engineering to be "the kind of activity instantiated by [these]

examples" (2018, p. 9). He also makes clear that these examples fix the subject matter of his general theory:

> there is a kind of activity that all of these philosophers are engaged in […] I will try to develop an improved framework for understanding what it is that these philosophers are doing. (Cappelen 2018, p. 10)

So, there is evidence that (1) the paradigm cases of conceptual engineering do not involve implementation and (2) "conceptual engineering" refers to the kind of activity exhibited by those paradigms.

Simon Blackburn may have been among the earliest users – if not the inventor – of the term "conceptual engineering." My above point is amplified if we look at how he conceived of it. He writes,

> I would prefer to introduce myself as doing conceptual engineering. For just as the engineer studies the structure of material things, so the philosopher *studies* the structure of thought. *Understanding* the structure involves seeing how parts function and how they interconnect. It means knowing *what would happen* for better or worse if changes *were* made. This is what we aim at when we investigate the structures that shape our view of the world. Our concepts or ideas form the mental housing in which we live. We may end up proud of the structures we have built. Or we may believe that they need dismantling and starting afresh. But first, we have to know what they are. (Blackburn 1999, pp. 1–2; my italics)

The activities mentioned by Blackburn are part of an *epistemic* project, that is, the epistemic project of "studying," "understanding," and "investigating" the structures of thought, and of "knowing" what "would happen" if changes were made. Blackburn does not mention that implementing such conceptual changes (e.g., through activism leading to changes in natural language semantics) is part of conceptual engineering.

Critics may point to the open-textured nature of the discourse around conceptual engineering. Conceptual engineering is a young discipline, and many things are still being discussed and

clarified (including the very meaning of "conceptual engineering"). People use the term in different ways, including in a theoretical sense and in a more practical sense. It is therefore possible that "conceptual engineering" might pick out the more complex and tangible activity of implementing newly constructed concepts (even if the paradigms cases do not). I have two responses.

First, the present discussion is in *meta-philosophy*. It is about what kind of activities *philosophers* should engage in. As such, the relevant sense of "conceptual engineering" should be one that describes a "philosophical activity." If there is both a narrow sense of "conceptual engineering" (where it describes a philosophical activity) and a wide sense (where it describes something larger), then the first sense is plainly the relevant one.

Second, it is insufficient for critics to point out that there is (or could be) *some* sense of "conceptual engineering" on which it is infeasible. This would have to be the *dominant* or *most relevant* sense if it is going to seriously undermine conceptual engineering *qua* philosophical activity.

Note that none of this suggests that implementation is not important. Nor does it suggest that implementation is not what might motivate people to engage in conceptual engineering (more on this in the next section).

4. THE INDIRECT ANTI-CONCEPTUAL ENGINEERING ARGUMENT

Where does the discussion of (D-ACEA) leave us? If implementation is not integral to conceptual engineering, then there is no direct route from concerns about the feasibility of implementing conceptual engineering to the conclusion that we should stop doing it. This does not leave conceptual engineers of the hook, however. It does, though, imply the following for critics who want to halt conceptual engineering owing to concerns around implementation. Critics need a different insertion for (P3), one that connects concerns about the feasibility of implementation to the proclaimed pointlessness of conceptual engineering. Critics must demonstrate that concerns about implementation should stop us from engaging in conceptual engineering, even when implementation is not an integral part of conceptual engineering. What

could establish the required connection between the assessment and design of concepts on the one hand and implementation on the other?

A reasonable starting point is the following undeniably plausible assumption:

> (A)   The benefits that derive from conceptual engineering's implementation make it worthwhile (where conceptual engineering is the assessment and design of concepts).

According to (A), philosophers do not (or should not) assess and design concepts simply for the sake of doing so. Rather, they do so because they expect that certain benefits will follow from implementing their proposals. If so, then practical hurdles obstructing the implementation of conceptual engineering proposals may render conceptual engineering futile, even while implementation is not part of conceptual engineering. This allows critics of conceptual engineering to present the following argument:

> The *indirect anti-conceptual engineering argument* (I-ACEA):

>> (P1) Conceptual engineering is the activity of assessing and improving the meanings of words in a shared language.

>> (P2) Changing what words mean in a shared language is infeasible.

>> (P3ii) If changing what words mean in a shared language is infeasible, then conceptual engineering is futile.

>> (C) Therefore, conceptual engineering is futile.

I now argue that (I-ACEA) is unsound. Yet, unlike the case of (D-ACEA), this is not simply because (P3ii) is false. Rather, as will become clear, (P3ii) does not interact well with (P2), in the following sense: There is a sense of "feasible" in which (P3ii) is true, but if we apply this sense to (P2), then (P2) comes out false. There is also a sense of "feasible" in which (P2) is true, but if we apply this sense to (P3ii), then (P3ii) comes out false. There is no single sense of "feasible" in which they are both true.

Let us begin by taking a closer look at (P2). Is it feasible to change what a word means in a shared language? Yes and no. Bracketing anomalous cases, no individual can decide whether

and how any given word changes its meaning. Yet, this does not imply that word meanings are independent of us. We are not entirely unable to change them. As David Kaplan points out, words are "objects of the created realm, created by language makers" (1990, p. 117). What words mean is a product of what we do with them and how we act on them. We can capture this relationship with the following principle:

> *Meaning follows use* (MFU): Changes in how participants in a linguistic community use a word can change that word's semantic content.

Not all changes in usage yield changes in meaning, however. Sometimes new usages are consistent with a term's meaning. Even when they are not, word meanings allow for misapplications. Moreover, metasemantic views provide different explanations for the truth of (MFU). For instance, on Wittgensteinian theories of meaning, meaning follows use simply because meaning *is* use. On causal theories of reference, use serves to anchor a term to a referent. That said, the following two claims seem highly plausible:

> 1.   Words do, and therefore can, change their meaning or reference.

> 2.   Changes in word meaning or reference somehow result from changes in how those words are used.

If (MFU) is as plausible as I believe, then it is feasible for linguistic communities to change word meanings. If we change the way we use words, then their meanings will follow suit. We can capture this result by drawing a distinction between two notions of feasibility: (1) *feasible-narrow* and (2) *feasible-wide.*

> 1.   Something x is feasible-narrow if x can be brought about by individuals or small groups of individuals over short periods of time.

> 2.   Something y is feasible-wide if y can be brought about by large groups of individuals over long periods of time.

These notions are, of course, far from rigorously defined, but they will do for present purposes. Employing these notions, we can now say that changing what words mean is feasible-wide but

not feasible-narrow. (P2) in (I-ACEA) appears to be true as long as we read "feasible" as "feasible-narrow" rather than "feasible-wide."

What about (P3ii)? This premise is crucial for deriving an anti-conceptual engineering conclusion from the implementation challenge. As we have seen, not any version of (P3ii) will do. "Feasible" must be read as "feasible-narrow." As I now argue, this reading of (P3ii) is highly implausible. We can see why if we generalize (P3ii) to the following principle:

> *Narrow feasibility principle* (NFP): Assessing and suggesting improvements of x is futile if it is not feasible-narrow to implement a normative proposal for x.

(NFP) states that it is not worth making any normative proposal that prescribes a situation that cannot be realized by an individual or small group of individuals over a short period of time. The problem is that this flies in the face of many inquiries conducted in normative ethics, political theory, and other normative domains. For instance, abandoning democracy (as suggested in Brennan 2016) is not something an individual or small group of individuals can do over a short period of time. In this case, (NFP) implies that we should stop thinking about whether democracy is our best political system.

The same holds, *mutatis mutandis*, for debates over police-abolitionism, incarceration, and taxation. It also holds for normative views such as extreme longtermism (championed by effective altruists), and for many views about the ethics of consumption, the ethics of climate change, and population ethics. In these domains of inquiry, philosophers often argue for proposals whose implementation is not feasible-narrow (they cannot be done by individuals or small groups of individuals over short periods of time). (NFP) prescribes that such discussions are futile and should therefore be abandoned. I doubt that most critics of conceptual engineering are ready to endorse this radically pessimistic conclusion about philosophical practice.

There are two viable response strategies to the above:

1.   Give an argument for the general truth of (NFP).

2.   Deny that conceptual engineering and other domains of normative theorizing deserve equal treatment.

Let us consider these responses in turn.

Regrading response strategy 1, can critics of conceptual engineering find support for (NFP) in moral and/or political philosophy? Many philosophers hold some version of the ought-implies-can principle (Streumer 2007). There are also ongoing debates around feasibility (Southwood 2018, Lawford-Smith 2013; Räikkä 1998) and the ideal/non-ideal theory distinction (Mills 2005; Valentini 2012) in political and feminist philosophy. Nonetheless, even staunch defenders of feasibility requirements or non-ideal theory constraints shy away from (NFP). It is generally accepted that feasibility should not be conflated with easiness (Tessmann 2010), and that some worthwhile normative projects cannot be accomplished by individuals or small groups. All serious debates about political feasibility, for example, concern feasibility in the wide sense.

A brief look at Nicholas Southwood's (2016) discussion of whether *ought* implies *feasible* helps to bring this point home. For Southwood, there is one core practice-serving ought that implies feasibility and another that does not. What matters here is that, on Southwood's view, even the ought that implies feasibility implies feasibility only in the wide sense (as what an agent or a set of agents can bring themselves to do) (2016, pp. 11–17). Absent some case-specific reasons to the contrary, changing how we use words is clearly something that we can bring ourselves to do. But, if the relevant sense of "feasible" is "feasible-wide," then difficulty changing word meanings does not undermine the truth of normative claims about which meanings should be used. This is the case regardless of whether and which "ought" implies "feasible."

Regarding response strategy 2, is there a way of rejecting (NFP) while holding on to (P3ii)? This requires identifying disanalogies between conceptual engineering and the kinds of normative enterprises just mentioned. Let us briefly discuss two potential candidates: (1) *feasibility* and (2) *below implementation value.*

(1) *Feasibility*. Is the implementation of conceptual engineering harder than the implementation of other normative proposals? One might think so. Conceptual engineering requires changing the meaning of words, which could follow a *threshold model*:

> Either one achieves everything by crossing the relevant threshold or one does not. If one does not, then one achieves nothing.

Other normative projects are not like this. Think of the project of creating a world where people do not eat meat. While one might never convince most people to stop eating meat, every person one does convince constitutes a small step towards a perceived good. Such projects follow a *summative model* rather than a threshold model.

However, many of the normative projects listed above follow a threshold model. One cannot officially abandon democracy, for example, without changing the constitution. And, changing the constitution requires crossing a certain threshold of political power. If one fails to meet the threshold, then the constitution will not be changed and democracy will remain in place. The same holds, *mutatis mutandis*, for revisionary proposals about policing, taxation, and the like. Moreover, it is not clear that semantic change even follows a threshold model. Some have argued that meaning-change comes *gradually*, perhaps through stages of partial reference (Devitt 1981), polysemy (Fischer 2020), or dialectization (Gasparri ms).

(2) *Below implementation value*. Other normative projects generate a variety of intellectually and practically worthwhile results even if they are not (and perhaps never will be) fully implemented in actuality. When discussing the advantages and disadvantages of democracy, we might learn a lot about how democracy functions, about what alternatives there are, and even how to mitigate some of the problems typically associated with democratic decision-making. It is reasonable to think that insights into these and other issues make a critical assessment of democracy worthwhile. This can be the case even if it is infeasible to change whether or not we will live in a democracy in the future. Can the same be said about conceptual engineering projects?

I think so. The critical assessment and design of concepts generates roughly the same kinds of insights as other normative projects. These include:

1. The identification of a concept's function (or our aim in using the concept).

2. What other similarly functioning concepts there are (or could be).

3. How some of a concept's most-pressing problems might be mitigated (even if it is infeasible to revise or replace the concept altogether).

Each of these three points plays a major role in current theorizing about conceptual engineering. There is a growing body of literature on the relevance of conceptual functions for undertaking conceptual engineering (see notably Nado 2019; Simion and Kelp 2020; Thomasson 2020; Queloz 2022). Matti Eklund (2021) and Rachel Rudolph (2021) discuss the importance of *conceptual explorations* and *conceptual evaluations*. Through these activities, we learn about potential alternatives to concepts and about how to mitigate some of the problems related to current concepts.

Let us take stock. The second premise of (I-ACEA) – (P2) – is plausible if rendered in terms of feasible-wide. It is implausible if rendered in terms of feasible-narrow. When it comes to the third premise – (P3ii) – the opposite is true. While normative proposals that cannot be realized by large groups of people over long periods of time might be futile, the narrow feasibility requirement is just too strong. In conclusion, while there are plausible versions of both of (I-ACEA)'s crucial premises, these versions do not jointly support the argument's conclusion.

5. CONCLUDING REMARKS

The anti-conceptual engineering argument uses concerns about the feasibility of implementing conceptual engineering proposals as a premise for rejecting conceptual engineering as a worthwhile pursuit. While this strategy initially seems plausible, I have argued that it ultimately fails. My reasoning proceeded as follows.

An extra premise is required to make the anti-conceptual engineering argument valid. This premise should connect the infeasibility of implementing conceptual engineering with the futility

of the overall activity. However, both antecedently plausible premises are indefensible upon scrutiny. In the absence of any further candidate premise, it appears that there is no sound version of the anti-conceptual engineering argument.

Both versions of the anti-conceptual engineering argument (and my rebuttal of them) were targeted at the semantic type of conceptual engineering. However, much of the discussion also applies, *mutatis mutandis*, to other types of conceptual engineering. My suggested defense applies to ways of implementing conceptual engineering that satisfy the standard of feasibility-wide. Plausibly, this standard is met by many (perhaps all) of the leading proposals on what it means to successfully re-engineer a concept. With enough time it is surely within the power of linguistic communities to change the following:

1. Linguistic conventions and norms (Thomasson 2021).

2. At least some of the information retrieved by default from long-term memory (Isaac 2021; Machery 2017).

3. What people speaker-mean by their utterances (Pinder 2021).

4. The classification procedures used (Nado 2021).

What now of the implementation challenge? Conceptual engineering is often thought to be capable of generating significant real-world changes. The argument given in this paper is neutral with respect to whether such ambitions are warranted. I have argued that successful implementation of conceptual engineering proposals is neither likely nor unlikely. This means that finding ways to address the implementation challenge (e.g., by developing feasible ways to change semantic meanings) is a live issue. Yet, if my argument is sound, philosophers need not wait until this research succeeds before getting stuck into assessing and designing concepts. Contra the critics, concerns about implementation should not prevent us from placing conceptual engineering firmly on the philosophical agenda.

Acknowledgments

Notes

1. Note that neither Thomasson nor Andow ultimately endorse skepticism about conceptual engineering.

2. This is not an idiosyncrasy of my formulation. The same formulation can be found in numerous other topical works. In Cappelen's seminal book *Fixing Language*, conceptual engineering is characterized as "the process of assessing and improving our representational devices" (2018, p. 3).

References

Andow, James. 2021. "Conceptual Engineering is Extremely Unlikely to Work. So What?" *Inquiry: An Interdisciplinary Journal of Philosophy,* vol. 64, no. 1–2, 212–226. https://doi.org/10.1080/0020174x.2020.1850343

Brennan, Jason. 2016. *Against Democracy.* Princeton, NJ: Princeton University Press.

Burgess, Alexis, and David Plunkett. 2021. "On the Relation Between Conceptual Engineering and Conceptual Ethics," *Ratio*, vol. 33, no. 4, 281–294. https://doi.org/10.1111/rati.12265

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

Deutsch, Max. 2020a. "Speaker's Reference, Stipulation, and a Dilemma for Conceptual Engineers," *Philosophical Studies,* vol. 177, no. 12, 3935–3957. https://doi.org/10.1007/s11098-020-01416-z

------ 2020b. "Trivializing Conceptual Engineering," *Inquiry: An Interdisciplinary Journal of Philosophy,* online-first. https://doi.org/10.1080/0020174X.2020.185334

------ 2021. "Still the Same Dilemma for Conceptual Engineers: Reply to Koch," *Philosophical Studies*, vol. 178, no. 11, 3659–3670.

Devitt, Michael. 1981. *Designation.* New York: Columbia University Press.

Eklund, Matti. 2021. "Conceptual Engineering in Philosophy," in *The Routledge Handbook of Social and Political Philosophy of Language,* ed. Justin Khoo and Rachel Sterken, 15–30. New York: Routledge.

Fischer, Eugen. 2020. "Conceptual Control: On the Feasibility of Conceptual Engineering," *Inquiry: An Interdisciplinary Journal of Philosophy*, online-first. https://doi.org/10.1080/0020174X.2020.1773309

Gasparri, Luca. *Unpublished Manuscript.*

Haslanger, Sally. 2000. "Gender and Race: (What) Are They? (What) Do We Want Them to Be?" *Nous,* vol. 34, no. 1, 31–55. https://doi.org/10.1111/0029-4624.00201

Isaac, Manuel Gustavo. 2021. "Which Concept of Concept for Conceptual Engineering?" *Erkenntnis*, online-first. https://doi.org/10.1007/s10670-021-00447-0

Jorem, Sigurd. 2021. "Conceptual Engineering and the Implementation Problem," *Inquiry: An Interdisciplinary Journal of Philosophy,* vol. 64, no. 1–2, 186–211. https://doi.org/10.1080/0020174X.2020.1809514

Kaplan, David. 1990. "Words," *Aristotelian Society Supplementary Volume,* vol. 64, no. 1, 93–119. https://doi.org/10.1093/aristoteliansupp/64.1.93

Koch, Steffen. 2021a. "The Externalist Challenge to Conceptual Engineering," *Synthese*, vol. 198, no. 1, 327–348. https://doi.org/10.1007/s11229-018-02007-6

------ 2021b. "There is No Dilemma for Conceptual Engineering. Reply to Max Deutsch," *Philosophical Studies*, vol. 178, no. 7, 2279–2291. https://doi.org/10.1007/s11098-020-01546-4

Lawford-Smith, Holly. 2013. "Understanding Political Feasibility," *Journal of Political Philosophy*, vol. 21, no. 3, 243–259. https://doi.org/10.1111/j.1467-9760.2012.00422.x

Machery, Édouard. 2017. *Philosophy Within its Proper Bounds*. Oxford: Oxford University Press.

Manne, Kate. 2017. *Down Girl: The Logic of Misogyny*. Oxford: Oxford University Press.

Nado, Jennifer. 2019. "Conceptual Engineering, Truth, and Efficacy," *Synthese*, vol. 198, suppl. 7, 1507–1527. https://doi.org/10.1007/s11229-019-02096-x

------ 2021. "Classification Procedures as the Targets of Conceptual Engineering," *Philosophy and Phenomenological Research*, online-first. https://doi.org/10.1111/phpr.12843

Nimtz, Christian. 2021. "Engineering Concepts by Engineering Social Norms: Solving the Implementation Challenge," *Inquiry: An Interdisciplinary Journal of Philosophy*, online-first. https://doi.org/10.1080/0020174X.2021.1956368

Pinder, Mark. 2020. "Conceptual Engineering, Speaker-Meaning and Philosophy," *Inquiry: An Interdisciplinary Journal of Philosophy*, online-first. https://doi.org/10.1080/0020174X.2020.1853342

------ 2021. "Conceptual Engineering, Metasemantic Externalism and Speaker-Meaning," *Mind*, vol. 130, no. 517, 141–163. https://doi.org/10.1093/mind/fzz069

Queloz, Matthieu. 2022. "Function-Based Conceptual Engineering and the Authority Problem," *Mind*, vol. 131, no. 524, 1-32. https://doi.org/10.1093/mind/fzac028

Räikkä, Juha. 1998. "The Feasibility Condition in Political Theory," *Journal of Political Philosophy*, vol. 6, no. 1, 27–40. https://doi.org/10.1111/1467-9760.00044

Rudolph, Rachel E. 2021. "Conceptual Exploration," *Inquiry: An Interdisciplinary Journal of Philosophy,* online-first. https://doi.org/10.1080/0020174X.2021.2002053

Scharp, Kevin. 2013. *Replacing Truth.* Oxford: Oxford University Press.

Southwood, Nicholas. 2018. "The Feasibility Issue," *Philosophy Compass*, vol. 13, no. 6, e12509. https://doi.org/10.1111/phc3.12509

Schupbach, Jonah N. 2017. "Experimental Explication," *Philosophy and Phenomenological Research,* vol. 94, no. 3, 672–710. https://doi.org/10.1111/phpr.12207

Shepherd, Joshua, and James Justus. 2015. "X-Phi and Carnapian Explication," *Erkenntnis*, vol. 80, no. 2, 381–402. https://doi.org/10.1007/s10670-014-9648-3

Simion, Mona and Chris Kelp. 2019. "Conceptual Innovation, Function-First," *Nous*, vol. 54, no. 4, 985–1002. https://doi.org/10.1111/nous.12302

Streumer, Bart. 2007. "Reasons and Impossibility," *Philosophical Studies*, vol. 136, no. 1, 351–384. https://doi.org/10.1007/s11098-005-4282-1

Tessman, Lisa. 2010. "Idealizing Morality," *Hypatia,* vol. 25, no. 4, 797–824. https://doi.org/10.1111/j.1527-2001.2010.01125.x

Thomasson, Amie L. 2020. "A Pragmatic Method for Conceptual Ethics," in *Conceptual Engineering and Conceptual Ethics*, ed. Herman Cappelen, David Plunkett, and Alexis Burgess, 435–458. Oxford: Oxford University Press.

------ 2021. "Conceptual Engineering. When Do We Need It? How Can We Do It?" *Inquiry: An Interdisciplinary Journal of Philosophy,* online-first. https://doi.org/10.1080/0020174X.2021.2000118

Valentini, Laura. 2012. "Ideal vs. Non-ideal Theory: A Conceptual Map," *Philosophy Compass,* vol. 7, no. 9, 654–664. https://doi.org/10.1111/j.1747-9991.2012.00500.x