

Can conceptual engineering actually promote social justice?¹

Paul-Mikhail Catapang Podosky
New York University (Shanghai)

Abstract

This paper explores the question: What would conceptual engineering have to be in order to promote social justice? Specifically, it argues that to promote social justice, conceptual engineering must deliver the following: (i) it needs to be possible to deliberately implement a conceptual engineering proposal in large communities; (ii) it needs to be possible for a conceptual engineering proposal to bring about change to extant social categories; (iii) it needs to be possible to bring a population to adopt a conceptual engineering proposal for the right reasons; and (iv) it needs to be possible to do (i) – (iii) without producing harmful consequences. I show that, of the three dominant approaches to conceptual engineering in the literature, only one of them seems amenable to the idea that it is possible and legitimate to promote social justice in accordance with (i) – (iv).

Introduction

Conceptual engineering is a heavily contested notion. In particular, there is widespread difference in what theorists argue, or at least assume, should be the relevant subject matter of conceptual engineering; or what aspects of the general representational terrain should be revised (i.e. concepts, meanings, linguistic practices, etc.)² In recent times, many projects in conceptual engineering have aimed to advance pursuits in the project of social justice. I have in mind strategies such as Sally Haslanger's (2000) ameliorative definitions of 'race' and 'gender,' Katharine Jenkins's (2016) trans* inclusive account of 'woman,' Kate Manne's (2018) revisionary analysis of 'misogyny,' and Robin Dembroff's (2016) reappraisal of 'sexual orientation.' Other examples abound.³ While each account offers a way to make our existing representational terrain better, certain questions can be asked. Looking *backward*, we might wonder: How plausible are such suggestions? And, looking *forward* we might think: Should engineers continue to submit proposals?

With this in mind, this paper explores the question: What would conceptual engineering have to be in order to promote social justice? Specifically, I will assess the relative merits of competing approaches to conceptual engineering with respect to how amenable they are to advancing our social justice goals.

Now, this does not mean that whatever approach comes out best is true. My goal is only to discern which approach to conceptual engineering makes sense of justice-oriented ameliorative projects. After all, without a

¹ I owe a lot of thanks to several people for their help and guidance in developing this paper: Karen Jones, Laura Schroeter, the philosophy postgraduate cohort at the University of Melbourne, and the attendees of my talks at the Australasian Association of Philosophy Conference and the Words Workshop. In particular, I am greatly indebted to wonderful reviewers at *Synthese*. Each of them took time to carefully work with me in shaping and framing this project. I am incredibly grateful for their insights and expertise.

² Löhr (2021) has recently argued that we needn't think of conceptual engineering within a representationalist framework. Like most, I will assume representationalism in this paper.

³ For instance, Barnes (2016, 2020) and Jenkins (2018).

plausible story of how to implement improvements to aspects of our representational terrain, it is unclear how useful ameliorative projects are in the fight against injustice. However, if there is independent reason to reject the truth of the theory that comes out best, or there are clear reasons that outweigh moral, social, or political aims, then it might turn out that *no* approach to conceptual engineering is suited to social justice purposes, even if one is ‘best.’ Given this, my aim is not to argue that one approach stands out as best all-things-considered. This paper is a project, relative to the available facts, that examines which account of conceptual engineering, if it were true, is most able to promote social justice.⁴

In what’s to come, I will argue that to deliver on the promise of promoting social justice, conceptual engineering must account for the following: (i) it needs to be possible to deliberately implement a conceptual engineering proposal in large communities; (ii) it needs to be possible for a conceptual engineering proposal to bring about change to extant social categories; (iii) it needs to be possible to bring a population to adopt a conceptual engineering proposal for the right reasons; and (iv) it needs to be possible to do (i) – (iii) without producing harmful consequences.

I show that, of the three dominant approaches to conceptual engineering in the literature, only one of them seems amenable to the idea that it is possible and legitimate to promote social justice through conceptual engineering. What are these approaches? Each differs in what it takes to be the relevant subject matter of conceptual engineering. Quickly, one approach, endorsed by Herman Cappelen (2018), takes conceptual engineering to be a matter of changing not simply our linguistic practices, but object-level facts. In contrast to Cappelen, Sarah Sawyer (2020, 2021) argues that conceptual engineering is about aligning the extension of a term with the relevant subject matter fixed by a concept, where the change in question does not concern object-level facts, but rather a community’s understanding of particular aspects of the world. Finally, theorists such as Sally Haslanger (2020), Amie Thomasson (2020), and Jennifer Nado (2019), advocate for a view of conceptual engineering that focuses on changing representational practices in line with certain functions of such practices. I will argue that this latter approach to conceptual engineering, what I will call The Functional Account, gets closest to what conceptual engineering would have to be in order to promote social justice according to (i) – (iv).⁵

1. What Would Conceptual Engineering Have to Be?

There are certain things that we might expect from conceptual engineering to advance our moral, social, and political causes. It’s not unreasonable to assume that *if* conceptual engineering is to be put in the service of justice, then it must be useful to us as a tool for disruption and change. But what would make conceptual engineering a useful tool in the fight against injustice? Put differently: What does conceptual engineering need to be in order to contribute to achieving more just social relations?

In recent literature, some theorists of conceptual engineering have, more or less explicitly, been interested in questions of *feasibility* (e.g., Podosky 2018, Fischer 2020, Andow 2021).⁶ Like moral and political philosophy, the questions center on whether normative theorizing about a representational practice should be constrained by reasonable expectation. If conceptual engineering is impossible, then it is not unreasonable for

⁴The reader will notice an absence of *internalist* theories of conceptual engineering. By ‘internalist’, I don’t mean to limit myself to theories of meaning, but instead conceptual engineering approaches that discuss higher-order cognition more broadly. It is only recently that such theories have been introduced into the fold (Pollock 2020, Fischer 2020, Isaac 2020, and to some extent Podosky 2018). The reason for this omission is only due to limiting my discussion to approaches of conceptual engineering that have been more prominently examined. It may turn out that internalist theories are better suited to social justice theorizing. However, I have reservations about this given that much injustice is material.

⁵I want to be clear that this paper does not assume that change to the representational terrain is necessary or sufficient for real-world change for justice. It assumes something weaker: conceptual engineering can *promote* social justice.

⁶Podosky (2018) is concerned with the means by which looping effects can be controlled; Fischer (2020) focuses on feasibility with respect to our ability to reason with engineers’ concepts; Andow (2021) argues that we should assess the likelihood of success against the potential value of success.

one to judge that it is pointless; ‘the success rate of conceptual engineering may be low enough to make conceptual engineering not worth bothering with’ (Andow 2021, p. 218).

One might see this paper as a project that attempts to assess the feasibility of competing conceptual engineering frameworks. However, it is perhaps not possible, or at the very least it is perhaps too hard, to articulate a set of desiderata to assess the feasibility of a conceptual engineering framework in a way that is independent of the project of *giving* a framework of conceptual engineering. This is to say that there may be no way to assess the feasibility of a conceptual engineering framework without at the same time assuming a (partial) theory of conceptual engineering. Thus, I sidestep questions or concerns about feasibility by distancing myself from projects of this kind.

Rather than assess the relative merits of competing frameworks of conceptual engineering against standards of feasibility, I will instead simply discuss what we might expect conceptual engineering to do for us – or what conceptual engineering needs to be in order to promote social justice. Whether or not such expectations are feasible goes beyond the scope of this paper. This is an additional concern. I will simply offer a list of requirements on conceptual engineering that must be met *if* it is to be useful to us in the fight against injustice.

What’s on this list?

1.1. Control

(i) It must be possible to deliberately implement a conceptual engineering proposal in large communities.

Part of the challenge of conceptual engineering concerns our ability to successfully implement an ameliorative strategy; to see a proposal executed out in the wild. I needn’t commit myself to any set of success conditions. Whatever these conditions are, there must be some sensitivity to our ability to *control* the character and direction of the representational terrain towards a proposed improvement. That is, there must some means by which proposed change, developed by an engineer, can be intentionally propagated within a community of thinkers and speakers. Thus, the problem of ‘control’ refers to the extent to which the propagation of a conceptual engineering proposal is possible. This formulation is neutral regarding which facts we must have control over (e.g., intension, extension, reference, definition, etc.).⁷

One might think that control isn’t necessary. In the absence of control, we should just aim to be disruptive in the hope that we will do some good overall. However, there are a few reasons to think this might not be right.⁸

First, it seems that control is important for the purposes of delineating the theoretical boundaries of justice-oriented conceptual engineering. If control is impossible, and we should simply throw everything we can into disruption, then *all* conceptual engineering proposals that purport to promote social justice seem to be theoretically on a par, or unable to be compared from a social justice point of view. This appears to render justice-oriented conceptual engineering a rather uninteresting theoretical endeavor. However, my project considers the question of what conceptual engineering would have to look like *for* it to be an interesting theoretical endeavor.

⁷ Matthieu Queloz and Friedemann Bieber (forthcoming) have recently argued that lack of control in the implementation of engineered concepts might actually be a good thing given our prior commitments to liberal democratic ideals. If control over conceptual uptake were institutionalized, then this power might (i) end up in the wrong hands and (ii) limit freedom of thought which undermines consent. This is a genuine concern, but I don’t have time to explore it in this paper.

⁸ This is in addition to the point disruption in and of itself may not promote social justice in any case.

To fill this out a bit more, the problem is that theories need only *purport* to promote social justice because, in the absence of control, *we cannot know or reasonably assume* that a proposal will in fact promote social justice. This seems to entail that the justification for making a conceptual engineering proposal can be rather thin or minimal. To reiterate, this appears to make justice-oriented conceptual engineering an uninteresting form of philosophical theorizing.

A second problem, related to the first, is that in the absence of control, we lose a means of *comparing* conceptual engineering projects. Why? One reason why we might prefer a particular conceptual engineering proposal is that, *ceteris paribus*, it is more plausibly implemented – which can only be judged assuming control. Another reason is that without control, conceptual engineers will just be shooting in the dark, as it were. So, comparing the relative merits of competing theories, according to which social justice goals they will in fact, or likely, advance, is epistemically impossible. We can only compare such theories according to which goals they purport to advance, which again, is not very interesting since, in the absence of knowing or reasonably assuming whether a proposal will actually promote social justice, anything goes (or at least most things).⁹

1.2. Looping Effects

(ii) It must be possible for a conceptual engineering proposal to bring about change to extant social categories.

My interest is in the resolution of unjust social relations. *Sometimes* injustice can only be resolved through the amendment, and even creation, of social categories (Dembroff 2020). For example, if the social category *woman* is unjust insofar as it excludes trans* women, then justice requires amending this category to be more inclusive. Part of what is required to achieve this is the introduction of new linguistic practices, such as classification and application dispositions, which construct, or reconstruct, categories through *looping effects* (Hacking 1999, Mallon 2016). This is the mutual causal feedback between our linguistic practices (i.e. classification) and social kinds. As Robin Dembroff puts it, ‘by developing classification practices, we create social kinds, which in turn impact classification practices, and so on’ (2020, p. 7). Thus, where ‘control’ concerns the propagation of a conceptual engineering proposal, this requirement concerns the possibility of successful looping effects, brought about with an ameliorative strategy, that create or amend *social categories* (Podosky 2018, p. 13).

1.3. Right Kind of Reasons

(iii) It must be possible to bring a population to adopt a conceptual engineering proposal for the right reasons.

Conceptual engineering projects are often motivated by particular reasons. Yet, it is entirely possible that a conceptual engineering proposal, were it to be taken up, is adopted by members of a thinking and speaking community for reasons other than those that motivated the engineer. After all, a proposal might be propagated by preference for conformity, or a desire to guarantee efficient communication, or social pressure, etc. However, some conceptual engineering projects, for the purposes of social justice, might require that members of a community adopt a proposal for the *right kind of reasons*. For example, Joey Pollock (2020) argues that with respect to the amelioration of ‘marriage,’ conceptual engineers have moral motivations – such as ‘contributing to the dismantling of oppressive social structures, institutions and systems of belief, and replacing them with those that will promote and sustain social equality’ (2020, p. 88). Given this, such reasons must be ‘transmitted’ from the engineer to the individuals intended to take up the revised practice. This is required to ‘overwrite specific existing problematic beliefs and dispositions [involving ‘marriage’] with those that will contribute to promoting and sustaining social equality for the relevant group’ (2020, p. 90). Put differently, for some projects in conceptual engineering, the reasons that motivate changing a practice are *constitutive* of the

⁹ I am indebted to an anonymous reviewer for pressing me on this point.

ameliorative project itself. Thus, this requirement states that, for at least *some* conceptual engineering projects, it must be possible for a population to adopt a proposal for the right kind of reasons – or the reasons that motivated the engineer.¹⁰

1.4. Harmful Consequences and Unfitting Application Dispositions

(iv) It must be possible to do (i) – (iii) without producing harmful consequences.

Sometimes intended revisions to our representational terrain are morally or politically *illegitimate*. Rather than promoting social justice, such revisions instead contribute to creating or sustaining structural inequalities, or producing deleterious effects for certain members of a population. Concerns of this kind have been referred to as linguistic, conceptual, or meaning ‘degeneration’ in contrast to ‘amelioration’ (Cappelen 2018, p. 59). Broadly speaking, the worry of illegitimate revisions is that they can produce harmful consequences for particular social groups.¹¹ Here’s an example. Through repeated association, ‘urban youth’, in the American context, is used in a way that invokes and entrenches pernicious stereotypes or generic characterizations of young Black men as dangerous (Stanley 2016, Ch. 4), leading to reprehensible social practices such as police brutality and even institution-backed murder.

Given the possibility of illegitimate revisions to our representational practices, a requirement is that successful conceptual engineering must not carry with it a *high degree of risk* of having harmful consequences. This concern is closely related to control. If we don’t have control over a representational practice, then of course a conceptual engineering proposal is at risk of becoming illegitimate – say, through being hijacked or appropriated by dominant oppressive powers. However, it is important to assess the risk of harmful consequences within a community even when control is possible, albeit difficult.

1.5. The Problem of Perversion

Teresa Marques (2020) has brought to light a distinctive and problematic phenomenon that faces conceptual engineering. Rather than amelioration, our word meanings can be *perverted*. Marques’ notion of meaning perversion has a precise and technical application. It occurs when a word that is typically used to refer to a particular kind K_1 , and the use of which recruits certain norms, values, or rights in a context, becomes used to refer to another kind K_2 , often through varied forms of implicit or overt manipulation.

The significance of this kind of problem must be stressed. Meaning perversions explain how unjust social relations are, at least partly, created and maintained. Marques refers to moments in history to show this. For example, Nazi propaganda changed what was expressed by ‘fanatical’ to heroic, the consequence of which increased the number of fanatic soldiers and strengthened the Nazi army (2020, p. 263). The lesson that Marques offers is that ‘the self-declared goodness of a revisionist project doesn’t suffice to avoid meaning perversions’ (2020, p. 260).

It’s important to reiterate the Marques’ notion of meaning perversion is particular. It involves the misapplication of a word to an unsuitable referent (2020, p. 276). Given this, meaning perversions are indifferent to whether conceptual engineering is successful. It is entirely possible that meaning perversions can occur and transgress the limits of permissible change to the general representational terrain (i.e. topic

¹⁰ Unfortunately, how exactly to determine when a conceptual engineering project must be adopted for the right kind of reasons won’t be discussed in this paper. All that I am committed to is that, for the purposes of social justice, an approach to conceptual engineering must be able to accommodate this possibility.

¹¹ By harm, I don’t just mean physical. Instead, I take a broadly Feinbergian view of harm that involves a setback to interests that are reasonable and legitimate – which includes psychological injury (Feinberg 1987).

continuity).¹² In contrast, the notion of degeneration, as I will understand it, occurs when conceptual engineering is successful, but nevertheless produces harmful consequences.

Given its insensitivity to the success of conceptual engineering, meaning perversions are an *equal* risk for any ameliorative project, no matter which approach one is working with. As such, I won't explain how each of the approaches in the upcoming discussion risk meaning perversion. Perhaps the risk of meaning perversion is insurmountable, and as such it undermines any chance of conceptual engineering being suitably justice-promoting. If this is the case, then so be it. I will nevertheless explore which theory comes *closest* to what conceptual engineering needs to be in order to promote social justice. In particular, I will only discuss the requirements on conceptual engineering that operate within the bounds of what is permissible according to each theory, such as topic continuity.¹³

1.6. Why (i) – (iv)?

Why these conditions? To my knowledge, such are the concerns that have been raised by those interested in justice-oriented conceptual engineering. However, there is more to be said. Each condition represents a potential barrier to change. If we cannot *control* the adoption of conceptual engineering proposal, then there is not much point in trying for the purposes of social justice; if we cannot amend or create social categories through *looping effects* with a conceptual engineering proposal, then there is not much point in developing proposals to effect looping effects for the purposes of social justice; if we cannot at least sometimes get a population to adopt a conceptual engineering proposal for the *right kind of reasons*, then there is not much point in trying for the purposes of social justice; and if we cannot develop a conceptual engineering proposal without some guarantee that it won't produce *harm*, then there is not much point in trying for the purposes of social justice.

A couple of things to note. First, this list is not exhaustive. It only captures *some* of the things we might expect from justice-oriented conceptual engineering. Nevertheless, I believe it accounts for the major concerns that have been discussed in recent conceptual engineering literature, and gets to the heart of many of the worries regarding the usefulness of conceptual engineering for promoting social justice. Second, merely satisfying only one of (i) – (iv) is often not, perhaps never, enough for conceptual engineering to promote social justice. For instance, having greater control in the implementation process of conceptual engineering, and its looping effects, does not entail that the ameliorative project won't produce harmful consequences. In saying this, having greater control and being motivated by good intentions perhaps renders the risk of harm sufficiently low enough to make an ameliorative project legitimate.

With all of this on the table, the next question is: which approach to conceptual engineering gets closest to meeting (i) – (iv)?

2. Approaches

Now that I have the basic idea of what conceptual engineering needs to be in order to promote social justice, I want to introduce three competing approaches to conceptual engineering that, to my understanding, are the most prominent in the literature: The Worldly Account, The Alignment Account, and The Functional Account. After this, I will assess the relevant merits of each theory against (i) – (iv) above.

¹² In fact, it is plausible to interpret Marques as saying that meaning perversion always involves *topic* disruption. Although, this might depend on what one means by 'topic.'

¹³ It should be noted that some theorists deflate the problem of topic continuity. See Simion and Kelp (2020)

2.1. The Worldly Account

Perhaps the most prominent account of what would constitute a successful conceptual engineering project is offered by Herman Cappelen (2018). He argues that the limits of conceptual engineering should be understood within a moderate externalist metasemantic framework, in the tradition of Kripke, Putnam, Burge, and Williamson. According to externalist theories of this kind, a range of factors ‘outside of the head,’ such as past patterns of usage, communicative chains, and experts within a community, play a substantive role in fixing reference.

Further, Cappelen’s understanding of conceptual engineering is strongly *metaphysical*. The success of conceptual engineering depends on actual changes to the world. That is, the goal of an ameliorative project is to change reference, its intension and extension (2018, p. 138).¹⁴ Put differently, conceptual engineering is not about changing thinking and speaking practices *per se*, but changing them with the effect of changing object-level facts. For instance, an account that aims to ameliorate linguistic practices with ‘woman’ is ultimately an account that aims to change what, or who, women *are*.

2.2. The Alignment Account

In response to Cappelen, Sarah Sawyer (2020, 2021) has argued that the ‘worldly’ construal of conceptual engineering, one that focuses on changing object-level facts, isn’t the best option to take within an externalist framework. Part of the reason for this is that Cappelen cannot account for *topic continuity*. Conceptual engineering must not change the relevant subject matter of inquiry, for this would not solve the problems that were of initial theoretical interest (Strawson 1963).

Some have understood the problem of topic continuity as a matter of whether conceptual engineering can preserve the extension of a term. If the extension of a term has changed, then the topic has changed – or so it goes. Sawyer disagrees with this. She does not think that the meaning of a *term* fixes the relevant subject matter. Instead, it is the reference of a *concept*, expressed by a term, that connects the term to a topic. Thus, what makes Sawyer’s account distinctive is a broadly Burgean distinction between the *meaning* of a term and the *concept* it expresses.

For Sawyer, the meaning or intension of a (non-indexical) term is externally determined. That is, it is determined by social practices; stable patterns of usage and deference within a community. Put differently, the meaning of a term is a characterization of the relevant subject matter that members of a thinking and speaking community would settle on were they to reach reflective equilibrium (2020, p. 383). And, because patterns of usage change throughout time, the intension of a term, and therefore its extension, can change too. The question, then, is how extensions can change without changing the topic?

The answer: concepts, that are expressed by terms, fix the relevant subject matter. Like the meaning of terms, according to Sawyer, concepts are externally determined. However, concepts, in contrast to terms, are constituent elements of thought that representationally connect thinkers to a topic. And, the relationship between concept and the world is referentially direct: fundamental, non-conceptual/descriptive relations to objective properties.

What’s important to note is the relationship between concepts and meanings. The concept expressed by a term fixes the relevant subject matter, and the meaning of a term is determined by a community’s characterization of

¹⁴ Cappelen (2018) states that no success conditions for conceptual engineering can be given since the terms stated in such conditions are subject to revision.

this subject matter. Thus, the extension of a term can change and yet pose no problem for the continuity of a topic. What mustn't change, in order to preserve topic continuity, is the extension of a concept.

On Sawyer's account, then, conceptual engineering is the project of alignment. Specifically, it is about changing a linguistic practice in a way that aligns the extension of a term with the relevant subject matter. In Sawyer's words,

If a revisionary analysis is correct and accepted, the effect is to bring the extension of the linguistic meaning of a term in line with the extension of the concept it expresses (i.e., in line with the relevant subject matter); it moves the linguistic practice closer to truth (2020, p. 391).

From this, we can see that Sawyer's take on conceptual engineering isn't metaphysical. It isn't about changing object-level facts. Instead, it is *epistemic*.¹⁵ It is about the accuracy of a community's characterization or theory of a particular subject matter fixed by a concept.

2.3. The Functional Account

In contrast to both Cappelen and Sawyer, some have argued that conceptual engineering is ultimately about the relationship between function and representational practice. There are different ways of construing what this means.

2.3.1. Functional Roles

One way to think about function is to consider the goal or purpose behind a representational practice within a social milieu. In the language of concepts, this understanding is common. Friedrich Steinle argues that '[o]ne of the fundamental characteristics of concepts is their directedness towards a specific goal' (2012, p. 105); Ingo Brigandt (2010), claims that concepts have 'epistemic goals' that figure in explanations and inferences in scientific theorizing; P.F. Strawson (1963) talks about the 'purpose' of a concept, which is echoed by Frank Jackson (2011); and Manuel Vargas (2013) suggests that concepts do particular 'work'.¹⁶ Moreover, Sally Haslanger takes a similar line in her understanding of what constitutes conceptual engineering:

The task is not simply to explicate the normal concept of X; nor is it to discover what things we normally take to fall under the concept we have in common; instead we ask what purpose is served in having the concept X, whether this purpose is well-conceived and what concept (or concepts) would serve our well conceived purposes(s)... (1999, p. 352).

Taking stock, Michael Prinzing argues that philosophical theorizing about concepts has a common theme:

All of these philosophers seem to have converged... on the same thought. The driving idea is that there are 'intentions' or 'goals' behind concepts; they have 'jobs'; they are supposed to do certain 'work'; they have a 'point', 'purpose', 'role', or 'function' (2017, p. 14).

On this way of thinking, a representational practice is engineered when there are changes to our use of a term or concept, but the goal or purpose behind such usage is preserved (Haslanger 2020).¹⁷ This guarantees topic continuity. Moreover, understanding conceptual engineering in this way makes topic continuity less of an issue since it is less demanding than, say, extension preservation. We can change representational practices that result in extensional changes without disrupting the goal or purpose behind such practices. For example, if the

¹⁵ Cf. Haslanger on *epistemic amelioration* (2020, p. 242).

¹⁶ Other examples include: Haslanger (2000), Thomasson (2020), Nado (2019).

¹⁷ In some cases, we might simply want to trash a representational practice in virtue of having a bad goal or purpose.

function of linguistic practices involving ‘marriage’ is, in part, to organize a nuclear family in line with particular roles, we can maintain this function, and therefore preserve topic continuity, even with revised practices with the term that includes application to same-sex couples.

2.3.2. Etiological Functions

Another way to think about function is to consider the *etiological* or *proper* function of a representational practice. Such functions are those that exist in virtue of being beneficial to our ancestors. For instance, the function of a heart, which is the reason why it exists today, is explained by the fact that it pumped blood in our ancestors, contributing to their survival. Etiological functions are those that exist as a result of natural selection over generations.

This kind of view, offered as one type of conceptual engineering, is endorsed by Mona Simion and Simon Kelp (2020). Concepts, at least some of them, continue to exist and operate in our social and representational milieu in virtue of being beneficial to our ancestors; they are selected-for given environmental pressures. Such concepts have a successful history of positive feedback in a thinking and speaking community. Thus, according to Simion and Kelp, one way of implementing an ameliorative strategy, which is under our control, involves designing an enticing practice that people will be disposed to copy. A concept is successfully engineered when its designed function becomes a function that explains its continued existence (2020, p. 12).

2.3.3. Why Not Both?

Given the foregoing, how should we construe The Functional Account? It’s important to note that these accounts of function aren’t inconsistent. Sometimes we might be interested in changing a representational practice while *preserving* its goal or purpose (or trashing a practice if it has a bad goal or purpose). Other times, we might be interested in designing a function for a practice such that, through selective pressures, it becomes a proper or etiological function. For instance, we might want a designed linguistic practice for ‘marriage’ to be selected-for, and at the same time want linguistic practices with ‘food’ to change while preserving its extant goal or purpose (e.g. structuring our eating practices without harming animals). Thus, I will treat The Functional Account of conceptual engineering as inclusive of these two understandings.

An important note is that on the inclusive account of function that I am working with, there is no connection to meaning or concept identity (Nado 2019, cf. Fisher 2015, Prinz 2018). It is the continuation of a *function* that must be preserved post-engineering, rather than *semantic* continuity. The upshot is that conceptual engineering projects can be much more ambitious (Nado 2019).

3. The Relative Merits

I’m up to the point where I can assess the extent to which each of the approaches to conceptual engineering is able to promote social justice in accordance with (i) – (iv) listed in §1. I’ll begin with The Worldly Account.

3.1. Promoting Social Justice: The Worldly Account

When it comes to *control*, Cappelen is the first to admit that his account isn’t promising. Successfully implementing an ameliorative strategy is likely impossible. After all, a great deal of what fixes reference, the target of engineering on Cappelen’s account, occurs outside of the head. Because of this, our ability to know the reference of our terms is undermined. Therefore, so is our ability to control reference. In other words, according to Cappelen, ‘[t]he process governing particular changes [in reference] is typically incomprehensible and inscrutable’ (2018, p. 53). He recognizes that epistemic agents within a community are not fully, or perhaps not even mostly, informed of the relevant empirical facts that determine reference. To compound this

problem, we might think that power-structures obscure the channels through which we could gain such knowledge. For example, we might be mistaken about who counts as an ‘expert’ in virtue of the fact that people occupy social positions that signal epistemic competence and trustworthiness, when in fact they are *neither* competent nor trustworthy (Fricker 1999). Thus, any attempt to actively bring about referential changes to our words is bound to be highly unreliable.¹⁸

Perhaps a benefit of Cappelen’s account, however, is that it focuses on changing reference, which allows for the possibility of *looping effects*. Through changes to our thinking and speaking practices, we can affect which social categories exist in the world. This accommodates social justice projects that aim at amending or creating social kinds for the purposes of correcting unjust social relations (e.g., trans* inclusive feminism, marriage equality, etc.). However, without a means of deliberately bringing about changes to object-level facts, it is unclear whether this is much of a benefit to The Worldly Account.

What about the *right kind of reasons*? It doesn’t seem that Cappelen’s account requires any kind of motivation to adopt a conceptual engineering proposal. This might make the chances of change much more likely. After all, control over thinking and speaking practices, and therefore object-level facts, is hard enough without requiring that people take up a practice for specific reasons. However, if we think, like Joey Pollock (2020), that some reasons are *constitutive* of certain ameliorative projects, such as the reasons that ought to govern linguistic practices with ‘marriage,’ then this is yet another stumbling block for Cappelen’s account. It seems that it will be too hard to bring a population to adopt a conceptual engineering proposal for reasons that motivated the engineer.

Cappelen is also aware that in virtue of lacking control over semantic facts, conceptual engineering, if attempted to be put into action, runs the risk of *harmful consequences*:

The fact that conceptual engineering is inscrutable and out of our control means that it is also possible (sometimes I think even likely) that those who try to achieve good ends through conceptual engineering will end up causing harms they didn’t intend. We have no prima facie reason to think the process is typically one that leads to amelioration rather than degeneration (2018, p. 159).

Degeneration is clearly an issue. We want to make sure that our ameliorative projects are precisely that – ameliorative. And we can think of the problem more broadly. Attempts to put into effect a conceptual engineering proposal may not result in the improvement of our language, thought, *or* social facts. It might turn out that intended changes to our representational practices produce harmful consequences. For example, in trying to revise our use of ‘food’ so that it does not involve application to non-human animals, we might produce the unintended effect of pushing farmers into poverty as a result of losing their means of living; and perhaps this is something we want, or ought, to avoid.

Of course, we should expect that bad things might happen when we try to implement an ameliorative strategy. However, this expectation must be weighed against our chances of success. On Cappelen’s account, we have no idea what the chances of success are, and thus all we are left with is the risk of degeneration.

¹⁸ Laura and Francois Schroeter (2020) disagree with this. They argue that while we might not know all of the relevant facts that determine reference, this does not mean we are completely in the dark. And Steffen Koch says control is possible, but we need to recognize that it is a group-level project:

‘...even though reference change on causal theories of reference turns out to require a collective long-term effort, it is nevertheless something that we, as a linguistic community, can bring about willingly’ (2021, p. 347)

Even with this in mind, it remains an open question as to whether we have *enough* knowledge or understanding of the determinants of reference for this approach to be suitable for the deliberate promotion or advancement of social justice.

3.2. Promoting Social Justice: The Alignment Account

It seems that The Worldly Account is not suited to promoting social justice. What about The Alignment Account?

It's hard to know how much better Sawyer's approach to conceptual engineering is regarding control. In one sense, it is better than Cappelen's insofar as the goal is not to change object-level facts. This is one less required step. Instead, conceptual engineering is a matter of alignment: the extension of a term must come to align with the relevant subject matter fixed by a concept that the term expresses. However, it is unclear how much control we have in this process of alignment – especially when it comes to social kinds.

Like Cappelen, Sawyer's preferred account of metasemantics makes it hard to know the reference of our terms. As such, it faces similar problems. It may be too much for us to fully grasp all of the linguistic activity that plays a role in determining meaning, such as past patterns of usage, or who the relevant experts are, and this undermines our ability to control semantic facts.¹⁹ However, there is a plausibly deeper issue with Sawyer's account.

Sawyer accommodates the possibility of communal error in understanding aspects of the world. Of course, this does not entail the impossibility of communal knowledge and discovery. We need only think of our better understanding of natural kinds (e.g. gold) to see how communal error can be corrected through empirical investigation. To put it differently, we can knowingly align the extension of a term with the extension of the concept it expresses.

However, things get difficult when we start thinking about the social world. Often, it seems, we cannot (easily) determine when a communal error has been corrected with respect to our understanding of social facts. For example, how can we tell when the extension of 'race' aligns with the extension of the concept that the term expresses? We would need to investigate objective properties, such as the property of having a race. But, besides moving past biological essentialism, and despite a rich history of investigation spanning hundreds of years, we haven't really gotten close to settling on what race is. Kwame Anthony Appiah expresses this concern plainly when he says, 'there is nothing in the world that can do all we ask 'race' to do for us' (1986, p. 35).

So, in the absence of a means of determining when a communal error has been corrected, can we still engage in conceptual engineering on Sawyer's account? It doesn't seem so. After all, the name of the game *is* to correct communal error. This appears to rule out several justice-oriented conceptual engineering projects that are about social kinds, such as race, gender, misogyny, and the like. What we need, then, is a way of working out when we have discovered genuine facts about the social world before we engage in conceptual engineering *qua* alignment.²⁰

This is a problem of 'control' not simply in the sense that we are unable to reliably manufacture change to semantic facts. Instead, it is a problem insofar as sometimes we do not have control over semantic facts *with the aim of aligning such facts with the extension of a concept*. We are often unable to know when intended change to the extension of a term aligns with the extension of a concept that it expresses.

¹⁹ The benefit of Sawyer's account, however, is that even if we do manage to change a linguistic practice, in whatever direction, if the extension changes then this will not disrupt topic continuity.

²⁰ In saying this, sometimes we can know when communal error has been corrected with respect to social kinds, such as our better understanding of sexual harassment or gaslighting (although, both are still contested). So, to be clear, my claim is not that Sawyer's account faces epistemic limitations for all justice-oriented conceptual engineering projects, but rather for many important ones.

Sawyer's approach doesn't allow for much control. What about the possibility of looping effects? Unfortunately not. Remember, for Sawyer, it is the world, connected to a term via a concept, that guarantees topic continuity. And in order to maintain the legitimacy of a conceptual engineering project, we must not change the topic. Thus, change to the realm of (social) facts runs the risk of changing the relevant subject matter that an engineering project is supposed to be about. However, it is the changing of social facts that some conceptual engineers aim at in their attempts to achieve social justice.

For instance, suppose that a conceptual engineer is interested in changing social facts about what women are, or who counts as a woman. As such, she proposes an account that is trans* inclusive, aiming to revise linguistic practices such that an amended category of *woman* is 'looped' into existence. Apart from saying that it is 'intuitively false' that conceptual engineering changes the nature of women (or any other properties), Sawyer argues that such proposals transgress the limits of conceptual engineering (2021, p. 10). After all, if there have been changes to object-level facts that connect a term to a subject matter, then the subject matter has changed – there is no single topic that persists through changes to the extension of a term. Thus, we have not preserved topic continuity. This is recognized by Sawyer, who says:

If the process of conceptual engineering changes, for example, what belief is, and does so, as Cappelen says, by changing the conditions that have to be satisfied in order for an object to fall into the extension of the term 'belief', then the property of being a belief is not a stable objective property that exists independently of the conditions that we associate with the meaning of the term 'belief' (2021, p. 10).

Thus, looping effects, which change object-level facts, are not possible on The Alignment Account. This goes against many conceptual engineering projects that aim at changing social, extralinguistic facts.

Does Sawyer's approach allow for the possibility of bringing a population to adopt a conceptual engineering proposal for the right kind of reasons? It seems unlikely. If achieving alignment between the reference of our concepts and terms is already a tall order, casting doubt on the plausibility of conceptual engineering, then factoring in the added constraint that members of a thinking and speaking community take up a proposal for reasons that motivated the engineer is only going to make conceptual engineering *less* plausible.²¹

Finally, The Alignment Account makes it difficult to assess whether a conceptual engineering proposal will have harmful consequences. Of course, this is due to being 'out of our control' in the sense that we cannot reliably manufacture semantic change. But it also runs the risk of harmful consequences insofar as the relevant subject matter is not always epistemically available to members of a community. So there will always be some doubt as to whether an engineering project has been successful, which might mean we under- or over-shoot our target, or miss it altogether. To give a more situated analysis, if the power to 'decide' whether change has occurred, as a matter of who decides whether we have gotten at the relevant subject matter, is distributed unjustly across a community, then degeneration is a genuine concern.

For example, let's say that the reference of the concept WOMAN is trans* inclusive. Further, suppose that attempts to redefine the term 'woman' fall short of full trans* inclusion, and stop at those who have had sex-reassignment surgery. This might be due to prominent voices in the community who advocate for this more narrow ameliorative gender project, or because the government, seen as an epistemic authority, has passed legislation about who counts as a woman. While the amelioration of a representational practice has occurred to

²¹ Given that the goal of Sawyer's approach is *just* for the extension of a term, determined by linguistic practices, to align with the relevant subject matter, it doesn't seem to matter what motivates people to adopt a linguistic practice. Alignment can be achieved with people adopting a proposed linguistic practice for different reasons. For example, suppose that the relevant subject matter represented by the concept FOOD excludes meat. The meaning of 'food' can be changed to match this subject matter through the adoption of revised linguistic practices with 'food' for different reasons – e.g., health, moral, environmental, fashion, etc. The benefit is that accommodating different motivations behind the adoption of a linguistic practice makes change more plausible.

some extent, it does not go far enough. Our understanding of what women are fails to align with what in fact women are. And this produces harmful consequences to the extent that certain people are unjustly denied being seen as belonging to a category to which they in fact belong. Put differently, the problem is that if the goal of conceptual engineering is epistemic, as Sawyer's approach suggests, our attempts to improve linguistic practices will come up against (unjust) power that obscures our channels to knowledge and discovery. And such power can render our attempts at amelioration harmful.²²

3.2. Promoting Social Justice: The Functional Account

The Worldly Account and The Alignment Account appear to be inadequate approaches to conceptual engineering for promoting social justice. Can The Functional Account do any better?

With the shift from meaning to function,²³ having control over the processes of conceptual engineering becomes much more plausible. After all, the goal is not to change a representational practice in order to bring about meaning change or preserve concept identity. With respect to promoting social justice, conceptual engineering on The Functional Account is a matter of changing a representational practice with the aim of having an ameliorative impact on social facts. On the different ways of understanding what 'function' means, we might want to change a representational practice but preserve its underlying goal or purpose; or else we might want to introduce a new representational practice, motivated by certain reasons, that will be selected-for.

Controlling changes to a representational practice is, in the first place, much easier on The Functional Account insofar as we are not nearly as epistemically limited. We do not need to know all of the linguistic activity that plays a role in determining meaning. In some cases, all we need to identify is the dominant representational practice in a context, and reveal the relevant goal or purpose of this practice via some process of rational interpretation of what we trying to achieve (e.g., Edward Craig-style (1990) conceptual ethnography). Sally Haslanger (2020) captures this in her example of our use of 'family.'

According to Haslanger, the function of our concept of family is to organize our lives together with respect to the 'coordination of domestic life, for example intimacy, sex, raising of children, economic partnerships, intergenerational transfers of traditions and property' (2020, p. 251). In Western contexts, the typical arrangement that stabilizes coordination in domestic life consists of a husband, a wife, and biological offspring. The issue with this, however, is that the function of family thought and talk is contingently caught up in hetero- and bio-normativity. But this does not exhaust all domestic arrangements that can constitute a family. For instance, there are same-sex couples, adopted children, single or unmarried parents, and extended families. Acceptance of such domestic arrangements has been a result of societal pressure in the form of activism and material change (2020, p. 251).

²² There is a sense in which Sawyer's approach to conceptual engineering does well with respect to degeneration compared to The Worldly Account. This is because it isn't possible to change the extension of a concept through conceptual engineering, despite it being possible to change the extension of a corresponding term. On The Worldly Account, the goal is to change object-level facts. This brings with it the risk of producing a unique harm: inappropriately changing the membership conditions for a social kind. For example, through extensive changes to linguistic practices, and other social facts, it is possible for a society to change what *sexual harassment* is – say, making it more impoverished by narrowing its application to a smaller set of behavior. The downside of The Worldly Account is that, despite being more impoverished, sexual harassment is determined by the conditions we associate with the meaning of 'sexual harassment.' On The Alignment Account, this is impossible. Sexual harassment is determined *independently* from the conditions we associate with the meaning of 'sexual harassment.' Thus, it does not risk inappropriately changing the membership conditions for social kinds. Rather, we just risk *misunderstanding* what the membership conditions are.

²³ Some like Fisher (2015), and perhaps even Prinzing (2018), tie function with meaning – the former offers a teleosemantics; the latter offers a view of concept identity. I follow Nado (2020) in thinking that meaning and function, at least in our effort to preserve continuity, should be treated separately.

This example shows us that the function(s) of certain representational practices is epistemically available to us, and in virtue of this, such practices can be scrutinized and revised. It has been recognized, for some time now, that existing practices with ‘family’ have been unjustly exclusive of other forms of domestic arrangement that can preserve the goal or purpose of such language.

Moreover, an issue with The Worldly Account and The Alignment Account is that there are often no clear-cut clues as to when conceptual engineering has been, more or less, achieved. Sometimes we do not know whether meaning has changed, or whether meaning has aligned with the relevant subject matter. But on The Functional Account, there are obvious signs. We need only ask: Have social facts changed? Have trans* and genderqueer parents been accepted *qua* parents? This we can see from the commonsense point of view, or more clearly at the level of sociological and social psychological inquiry. Thus, not only is it possible for an engineer to recognize a representational practice, but it is also possible for an engineer to make an informed judgment as to whether their project has been, at least somewhat, successful.

So far, I have only spoken about control with respect to one notion of function. Is control possible when it comes to implementing etiological functions? Simion and Kelp (2020) argue that the answer is ‘yes.’ To reiterate, the mark of successful conceptual engineering, on this approach, is when a designed function of a representational practice becomes a function that explains the existence of that practice. The designed function must catch on and be stabilized within a community of thinkers and speakers. And, for Simion and Kelp, a conceptual engineer can achieve this by acting on the environment that will drive adaptive change. The environment needs to be tinkered with in order to drive demand for the designed function. Thus, given that we can change the environment in this way, then we have a non-trivial amount of control over the direction and character of our practices.

How does The Functional Account do with respect to looping effects? It is evident in the foregoing example that looping effects are not only possible but available for us to understand and actively set in motion. After all, once upon a time it might have been inconceivable that our linguistic practices with ‘family’ might include application to same-sex couples. However, with revised classificatory practices, off the back of social movements, same-sex couples now fall comfortably, but not without resistance, within the extension of ‘family’ (in many countries). This example shows that the *amendment* of social kinds, through looping effects, is possible on The Functional Account. But what about the *creation* of social kinds? This is also possible. Suppose that the goal of gender thought and talk is to keep track of people’s self-identification, or the internal maps that they use to navigate social reality (e.g. Jenkins 2016). This goal has been preserved with the introduction of linguistic practices with ‘non-binary.’ And given how looping effects work, the classification of people as ‘non-binary’ has brought into existence a new social kind: *non-binary* people.

The Functional Account also does well with respect to accommodating the possibility of bringing a population to adopt a representational practice for the right kind of reasons. It’s worth emphasizing that on the *etiological* function approach to conceptual engineering, it appears that the name-of-the-game *is* to get people to adopt a representational practice for the reasons that motivated the engineer.

On the etiological approach to conceptual engineering, the goal is not *merely* to change a representational practice, but for the practice to be adopted for the reason it was designed. That is, in the language of concepts, when an ameliorative strategy is ‘launched on a competitive market of concepts’, the hope is that the designed concept will ‘*be used to do what [it is] designed to do, and [it] will continue to be used in virtue of the fact that the way of thinking about the world [it] made available was beneficial to users*’ (Simion and Kelp 2020, p. 6 my emphasis). Thus, representational change must be brought about and sustained in a particular way. One natural interpretation of this, at least when it comes to social justice concerns, is that a representational practice must be adopted for the right kind of reasons. Joey Pollock (2020) argues for this way of thinking when it comes to concepts for social equality, such as MARRIAGE:

...in ameliorative projects of the sort under consideration, engineers are motivated by the pursuit of moral goods – they are interested in contributing to the dismantling of oppressive social structures, institutions and systems of belief, and replacing them with those that will promote and sustain social equality. When, in implementing an ameliorated concept, engineers seek to transmit their own motivating reasons to the individuals who adopt the new concept, it is because acceptance of these reasons is partly constitutive of the sorts of changes that they wish to effect – individuals who adopt the new concept ought (if the project is considered to be successful) to be appropriately motivated by specific moral considerations (2020, p. 88).

Pollock believes that an ameliorative strategy is often not solely about the adoption of a concept, but includes, as a matter of *constitution*, a certain set of motivating reasons.

One might think: Requiring that people adopt a conceptual engineering proposal for the right kind of reason is simply implausible. It would be wonderful if, say, everyone revised application practices with ‘food’ because they were motivated by a moral concern for non-human animals. But it seems that we should be happy with different motivations *so long as pernicious representational practices, and thus oppressive social facts, change*. Given that the world is unjust, in which many people are suffering as a result of being situated in systems of dominance and subordination, adopting a revised practice for the wrong reasons may be *good enough*. It will go a long way toward resolving oppressive social relations, even if it doesn’t go *all* the way – at least, not immediately.

I accept that this might be true *most of the time*. However, my goal in this section is only to examine whether The Functional Account can accommodate the possibility that, for at least *some* ameliorative projects, people adopt a representational practice for the right kind of reasons.

As discussed above, Simion and Kelp (2020) show how control is possible on their functional approach to conceptual engineering. What is interesting to note is that control, on their way of thinking, *just is* getting people to take up an ameliorative proposal for reasons that motivated the engineer. Thus, given that control is possible, by acting on the environment that will drive adaptive change, it is therefore possible to get people to adopt a representational practice for the right kind of reasons.

Can the approach that understands the notion of function as the *goal or purpose* behind a representational practice accommodate the possibility of bringing a population to adopt the right kind of reasons?

It seems that if we combine the requirement that the goal or purpose of a representational practice must be preserved post-engineering *and* the requirement that people take up a representational practice for the right kind of reasons (at least for some conceptual engineering projects), then we are not too far away from the etiological functional account of conceptual engineering. Of course, there are differences – for example, how topic continuity is preserved. But when it comes to the possibility of bringing a population to adopt a representational practice, it seems that we simply have to manipulate the environment in a way that will drive people to adopt the practice for reasons that motivated the engineer (while preserving its goal or purpose, of course).

For instance, suppose the goal or purpose of practices with ‘food’ is to organize a community around a set of eating resources. And, suppose that an engineer, motivated by moral reasons, wants to revise practices with ‘food’ such that it no longer applies to non-human animals *and* that people adopt such practices for the same moral reasons. It seems that the conceptual engineer can achieve²⁴ this by making changes to the social and material environment in which the practice will take place. For example, the engineer might create and engage

²⁴ Not by themselves, of course, but as a member of a community of activists.

in widespread, high-profile endorsement of certain documentaries (e.g. *Dominion*, *Cowspiracy*, etc.) that target individuals at the level of moral affect, such as triggering feelings of empathy, in the hope that those who watch such documentaries will suggest it to others.

So, if we already accept that bringing a population to adopt a representational practice is possible on the etiological functional approach to conceptual engineering, then it must also be possible on the approach that understands the notion of function as the goal or purpose behind a practice. Both approaches need only accept that control and the adoption of the right kind of reasons is a matter of being able to manipulate the environment in which thinkers and speakers are embedded. This is something we *can* do.

What about the risk of harmful consequences? On The Functional Account, degeneration is a genuine concern. Given that control is possible, then the environment can be manipulated by bad eggs in a way that produces harm in a community. And this is something that we have already seen: just think of figures such as Donald Trump and practices with ‘fake news’ to undermine credible news sources. It seems reasonable to think that the effect of this usage by such a high-profile figure is that a vast number of the U.S. population have come to trust, or trust more, news sources that are in fact untrustworthy – news sources that promote pernicious ideas about members of marginalized groups, thus cultivating or solidifying discriminatory, bigoted, and even violent attitudes and behavior.

Despite this, The Functional Account does not suffer the same problem of degeneration as The Worldly Account and The Alignment Account. For the latter accounts, the problem of degeneration is a matter of shooting in the dark, as it were. Since we do not know when meaning change occurs, we do not know when meaning change will produce harmful effects. We cannot predict the effects of our attempt at revision, and thus we risk degeneration in a serious sense: we have no reliable means of foreseeing what will happen when we promote a particular representational practice. On The Functional Account, we do not need to know when meaning has changed. So, we can avoid this problem of prediction. This is not to say that we can be certain of the effects of promoting a particular practice. Instead, it is to say that we can make more reliable judgments about what will happen in a community if an ameliorative strategy were to be adopted. After all, the only thing one needs to judge is whether the purpose or goal behind a representational practice will be preserved, or whether it will be selected-for, and whether one’s aim for bringing about such change will be achieved. And this can be empirically discerned (e.g. from the point of view of commonsense or sociological inquiry).²⁵ Thus, on The Functional Account, we can think more critically about whether to promote a representational practice given that we have, more or less, the ability to reliably judge whether that practice will have positive or negative effects for a community. This is clearly something important for social justice theorizing: we want to be able to make assessments about what will make things better, not worse.

One final note about degeneration. Even though The Functional Account allows for the possibility of bad eggs having control over our representational practices, the visibility of our changing practices means that we can scrutinize people for pushing the practice in one way rather than another. Of course, we can scrutinize people for *attempting* to change meaning. However, we cannot judge them for having a substantive effect on meaning change since we do not know how much their contributions count. On The Functional Account, we can judge people, like Donald Trump, for having a substantive impact on our representational practices because we can see it first hand – we can see its widespread usage in our everyday lives. Thus, The Functional Account offers another benefit: we are able to make moral, social, and political evaluative and normative assessments about someone’s thinking and speaking behavior. We can tell people that their promotion of a conceptual engineering proposal is right or wrong *because* of its effect on representational practices.²⁶

²⁵ We can even update implementation strategies if our attempts at bringing about representational change are unsuccessful.

²⁶ This is clear in the case of high-profile figures, such as Donald Trump. But it is perhaps unclear how much we can judge lay people for their promotion of a representational practice.

4. Taking Stock

So, what would conceptual engineering have to be in order to promote social justice? The foregoing reasons lean clearly in favor of The Functional Account. Neither The Worldly Account nor The Alignment Account seem to do well with respect to control, looping effects, and the risk of harm – though, The Alignment Account does better than The World Account. In contrast, The Functional account does relatively well with respect to all of (i) – (iv), however when it comes to the right kind of reasons, it might turn out that functions are simply too hard to manufacture, at least for morally motivated representational practices.

I want to emphasize again that this assessment of the relative merits of each approach to conceptual engineering does not settle the dispute as to which is correct. If it turns out that there is no way of plausibly defending the idea of the function of a representational practice, then of course The Functional Account should be abandoned.²⁷ All I have established is that The Functional Account gets closest to what conceptual engineering needs to be in order to promote social justice.

References

- Andow, James. 2021. Conceptual Engineering is Unlikely to Work. So What? *Inquiry* 64 (1 – 2): 212 – 226.
- Appiah, Kwame Anthony. 1986. The Uncompleted Argument: Du Bois and the Illusion of Race in “*Race, Writing, and Difference*,” Henry Louis Gates (ed.). Chicago: University of Chicago Press.
- Barnes, Elizabeth. 2016. *The Minority Body: A Theory of Disability*. Oxford: Oxford University Press.
- Barnes, Elizabeth. 2020. Gender and Gender Terms. *Nous* 54 (3): 704 – 730.
- Brigandt, Ingo. 2010. The Epistemic Goal of Concept: Accounting for the Rationality of Semantic Change and Variation. *Synthese* 177: 19 – 40.
- Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Craig, Edward. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- Dembroff, Robin. 2016. What is Sexual Orientation? *Philosophers’ Imprint* 16.
- Dembroff, Robin. Forthcoming a. Real Talk on the Metaphysics of Gender. *Philosophical Topics*.
- Feinberg, Joel. 1987. *Harm to Others*. Oxford University Press.
- Fischer, Eugen. 2020. Conceptual Control: On the Feasibility of Conceptual Engineering. *Inquiry* DOI: 10.1080/0020174x.2020.1773309
- Fisher, J.C. 2015. Pragmatic Experimental Philosophy. *Philosophical Psychology* 28 (3), pp. 412 – 433.
- Fricke, Miranda. 1998. Rational Authority and Social Power: Towards a Truly Social Epistemology. *Proceedings of the Aristotelian Society* 98 (2): 159 – 177.
- Hacking, Ian. 1999. *The Social Construction of What?* MA: Harvard University Press.
- Haslanger, Sally. 1999. What Knowledge Is and What It Ought to Be: Feminist Values and Normative Epistemology. *Philosophical Perspectives* 13: 31 – 55.

²⁷ Jared Riggs (2021) has recently argued that we should deflate the functional turn in conceptual engineering.

- Haslanger, Sally. 2000. Gender and Race: (What) Are They? (What) Do We Want Them to Be? *Nous* 34 (1): 31 – 55.
- Haslanger, Sally. 2020. Going On, Not the Same Way in *Conceptual Engineering and Conceptual Ethics*, Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), pp. 230 – 260. Oxford: Oxford University Press.
- Isaac, Manuel Gustavo. How to Conceptual Engineer Conceptual Engineering. *Inquiry* DOI: <https://doi.org/10.1080/0020174X.2020.1719881>
- Jackson, Frank. 2011. On Gettier Holdouts. *Mind & Language* 26 (4): 468 – 481.
- Jenkins, Katharine. 2016. Amelioration and Inclusion: Gender Identity and the Concept of Woman. *Ethics* 126 (2): 394 – 421.
- Jenkins, Katharine. 2018. Toward an Account of Gender Identity. *Ergo* 5.
- Koch, Steffen. 2021. The Externalist Challenge to Conceptual Engineering. *Synthese* 198: 327 – 348.
- Löhr, Guido. 2021. Commitment Engineering: Conceptual Engineering without Representations. *Synthese*. DOI: <https://doi.org/10.1007/s11229-021-03365-4>.
- Mallon, Ron. 2016. *The Construction of Human Kinds*. Oxford: Oxford University Press.
- Manne, Kate. 2018. *Down Girl: The Logic of Misogyny*. Oxford: Oxford University Press.
- Marques, Teresa. 2020. Amelioration vs. Perversion in *Shifting Concepts: The Philosophy and Psychology of Conceptual Variability*, Teresa Marques and Asa Wikforss (eds.), pp. 260 – 284. Oxford: Oxford University Press.
- Nado, Jennifer. 2019. Conceptual Engineering, Truth, and Efficacy. *Synthese* DOI: <https://doi.org/10.1007/s11229-019-02096-x>
- Podosky, Paul-Mikhail Catapang Podosky. 2018. Ideology and Normativity: Constraints on Conceptual Engineering. *Inquiry* <https://doi.org/10.1080/0020174x.2018.1562374>
- Pollock, Joey. 2020. Conceptual Engineering and Semantic Deference. *Studia Philosophica Estonia* 12: 81 – 98.
- Prinzing, Michael. 2017. The Revisionist Rubric: Conceptual Engineering and the Discontinuity Objection. *Inquiry*, DOI: [10.1080/0020174X.2017.1385522](https://doi.org/10.1080/0020174X.2017.1385522).
- Queloz, Mattheiu and Friedemann Bieber. Forthcoming. Conceptual Engineering and the Politics of Implementation. *Pacific Philosophical Quarterly*.
- Sawyer, Sarah. 2020. Talk and Thought in *Conceptual Engineering and Conceptual Ethics*, Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), pp. 379 – 395. Oxford: Oxford University Press.
- Sawyer, Sarah. 2021. The Role of Concepts in Fixing Language. *Canadian Journal of Philosophy* 50 (5): 555 – 565.
- Schroeter, Laura and Schroeter, Francois. 2020. Inscrutability and its Discontents. *Canadian Journal of Philosophy* 50 (5): 566 – 579.
- Simion, Mona and Kelp, Christopher. 2020. Conceptual Innovation, Function First. *Nous* DOI: [doi: 10.1111/nous.12302](https://doi.org/10.1111/nous.12302)
- Stanley, Jason. 2016. *How Propaganda Works*. NJ: Princeton University Press.

Steinle, Friedrich. 2012. Goals and Fates of Concepts: The Case of Magnetic Poles in *Scientific Concepts and Investigative Practice*, Uljana Feest and Friedrich Steinle (eds.), pp. 105–125. Boston, MA: De Gruyter Mouton.

Strawson, P.F. 1963. Carnap's View on Constructed Systems versus Natural Languages in Analytic Philosophy in *The Philosophy of Rudolf Carnap*, Paul Schilpp (ed.), pp. 502–518. La Salle: Open Court.

Thomasson, Amie. 2020. A Pragmatic Method for Normative Conceptual Work,' in *Conceptual Engineering and Conceptual Ethics*, Alexis Burgess, Herman Cappelen, and David Plunkett (eds.), pp. 435 – 458. Oxford: Oxford University Press.

Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility*. Oxford: Oxford University Press.