

The New Philosopher-Kings: Conceptual Engineering and Social Authority

Nicholas Smyth

Philosophers are only just beginning to study conceptual ethics. That is, they are beginning to ask: which values and principles ought to guide our assessment and modification of human conceptual schemes?¹ It might be thought that this is just the question of what makes a concept *better* or *worse*. But a conceptual engineer is not just a theorist, they are a person operating in an actual, nonideal social context, constrained by various material and social realities. Rather than simply theorize in an ideal spirit about it means to say that a concept itself is good or valuable, we should ask the nonideal-theoretical question: how should actual conceptual engineers proceed?

In this paper, I defend a firmly deontological answer to this question. I argue that any legitimate engineer or group of engineers must establish their own social authority over the relevant group of concept-users, which in turn requires them to relate their proposed changes to the subjective practical reasons possessed by those concept-users. This position entails the falsity of *consequentialism* about conceptual engineering. That theory, which covertly dominates most discussions of conceptual ethics in the literature, asks us to see the value of concepts in terms of the states of affairs they bring about, and it asks the engineer themselves to aim at producing better concepts in this sense. Deploying the notion of an *identity-constituting concept*, I will show that a wholly consequentialist approach to conceptual

¹ See (Cappelen 2018, Burgess, Cappelen, and Plunkett 2020, Queloz and Bieber 2021, Shields 2021, Marques 2020, Simion 2018, Podosky 2022).

engineering is morally and politically bankrupt, because any engineer also who adopts it cannot explain or justify their own social authority over human agents. We must therefore abandon this approach to concepts and conduct corresponding inquiry into *conceptual rights* and *conceptual authority*.²

Before proceeding to the main critique, however, I will try to take a look at the activity of conceptual engineering itself.

The Nature of Conceptual Engineering

To begin, I think it is useful to follow a suggestion of Mona Simion's and divide conceptual engineering projects into two broad camps, the representational and the practical (Simion 2018, 5). The first is a purely metaphysical form of engineering, one which seeks to produce concepts which, in Ted Sider's words, 'carve nature at its joints' (Sider 2011). There is absolutely nothing new about this kind of conceptual engineering, and it has been practiced and theorized since the inception of inquiry itself. After all, from a certain angle, science itself is primarily representational conceptual engineering, or the development of concepts which can figure in statements that accurately describe the nature of reality. When scientists declared that the concept WHALE fell under the concept MAMMAL and not the concept FISH, they were of course engineering concepts, but not in any new or interesting sense. This paper will not concern itself with that project, which sees joint-carving as the *constitutive* aim of a concept. For these engineers, the question of whether we ought to employ some

² In a recent paper, Matthieu Queloz at least *appears* to tackle the question of conceptual authority, but it turns out that for him, this is just the question of which concepts we have reason to use, and not the question of which actual *persons* have the authority to modify our concepts. (Queloz 2022)

concept or other reduces to the question of whether that concept accurately represents independent reality.³

The version of conceptual engineering that seems to occupy a great deal of the contemporary literature is different, since it embraces the idea that conceptual choice can and should be determined by broadly practical considerations. Rather than conceive of it as a domain which seeks to describe the world accurately, these thinkers take the view that metaphysics is fundamentally a social *activity* with the constitutive aim of producing more valuable, useful or liberatory conceptual schemes. In recent years, this program has been both heavily theorized and actually practiced (Haslanger 2012, Jenkins 2016, Cappelen 2018, Burgess, Cappelen, and Plunkett 2020).

Of course, these two types of normative consideration, the representational and the practical, can conflict (Simion 2018). This paper is not an answer to any question about what to do in such cases. Rather, it is simply a contribution to the question of how practical conceptual engineering ought to proceed.

Before proceeding, however, two more clarifications. First, it is worth noting that conceptual engineering projects can aim at different results.⁴ The first we might call *revisionary*. This project seeks to change existing concepts, by eliminating conditions that currently guide its application, by adding application-conditions, or both. This is what the International Astronomical Union was engaging in when it insisted on adding a condition to the concept PLANET which excluded Pluto. Alternatively, there is an *additive* kind of engineering, which

³ For recent pragmatist critics of this representationalist project, see (Price 2014, Thomasson 2014). See also Nietzsche, *The Gay Science*, 344.

⁴ David Chalmers calls these first two programs “homonymous and heteronymous conceptual engineering”, but I’m using labels which I think are a little clearer. (Chalmers 2020, 8-9)

seeks to introduce an entirely new concept to a language. The people who introduced the concept of SEXUAL HARASSMENT to the English language were these sorts of engineers, since they did not seek to alter or replace any older concept. Finally, the *eliminative* project, which has received less attention in the literature, simply seeks to remove a concept from a language without replacing it (Habgood-Coote 2020). Though there are interesting philosophical questions about this simple three-part division, I'll just take it on board for now.⁵ And while this point will have to wait for later, I will eventually say why the normative framework for each of these projects must be importantly different.

Finally, Cappelen and Plunkett suggest that even under this umbrella, there is still room to merely theorize about the ethics of conceptual engineering in a way that is not meant to guide engineering projects at all (Cappelen and Plunkett 2020, 5; Cappelen 2018, 83-4). It is not clear to me that this distinction between ideal and nonideal theorizing about concepts can be made coherent, particularly when we are talking about actual published work.⁶ That said, I won't rely on this contentious point here, and will just note that my paper concerns the ethics of conceptual engineering as it is meant to actually help guide our additive, revisionary or eliminative practice.

These preliminaries aside, I'll now proceed to outline the normative framework that threatens to make so much trouble for practical conceptual engineers.

⁵ For example: if meaning is holistic, then when we add or eliminate a concept to a language, do we not thereby revise the meanings of other concepts? This would entail that there is no such thing as *purely* additive engineering project. There is also the much-discussed question of how much revision we can make before we are merely "changing the subject" (Sundell 2011, 2017, Prinzing 2016, Thomasson 2019).

⁶ Reflection, writing, and scholarly publishing on the ethics of concepts is itself practical activity. For example, to actively define, in a published work, a concept like GOOD as it applies to conceptual schemes *is itself conceptual engineering*, and to do this is to engage with the practical sphere, inviting others to reflect upon your proposal and potentially change their dispositions to think and speak in certain ways.

Consequentialism about Concepts

Consequentialism about concepts is the view that in evaluating and modifying concepts, we must focus entirely on their effects or on what they produce. It is thus inconsistent with the view I will eventually defend, which is that some acts of engineering are impermissible even if they produce more value than available alternatives.

More precisely, if we understand consequentialism about Xs to be the theory which says that the normative significance of Xs is exhausted by the consequences which are brought about in virtue of Xs being the way that they are, and if we understand those consequences to be evaluated *agent-neutrally*⁷, then we may define:

Consequentialism about concepts: a concept is *better* to the extent that it produces more agent-neutral value, and *worse* to the extent that it produces the reverse. When seeking to change a conceptual scheme, a conceptual engineer ought to aim at producing *better* concepts in this sense.

The fundamental idea here is that a conceptual engineer is trying to produce more value in the world, and that concepts are simply a convenient or efficient *means* with which to produce this value. This, I think, is the core instinct that drives a great many ameliorative projects in this domain; having noticed that language plays a huge role in maintaining bad states of affairs, theorists seek to improve those states of affairs by improving language.

⁷ Some theorists think that within consequentialist theory there is room for *agent-relative* evaluation (Portmore 2011, Dreier 2011). While the agent-neutralist follows classical consequentialism in thinking that consequences are just good or bad as such, the agent-relative view makes room for the idea that consequences are evaluated by the lights of some person or group of persons. On such a view, of which classical egoism may be an instance, there is a world-ranking for each agent, and the agent's job is to try to get reality to conform to their own highest ranked world. For the purposes of this paper, I am only addressing agent-neutral consequentialism about concepts, and I will follow tradition by continuing to refer to this as "consequentialism".)

Concepts, on this view, are just nodes in a causal nexus, and so we must seek to improve their outputs.

Now, not all engineering projects can or should rely on this basic framework. Unsurprisingly, projects which are purely *representational* in the sense outlined above make almost no substantive ethical commitments at all (except, of course, a commitment to the importance of correct representation itself). For example, Matti Eklund and Kevin Scharp argue that the concept of TRUTH might need to be revised because it is *inconsistent*, and this is not itself a practical argument, it need only rely on the premise that inconsistent concepts cannot represent independent reality (Scharp 2007, Eklund 2014).⁸

But amongst a great many *practical* conceptual engineers, consequentialism about concepts often tends to covertly dominate various discussions. No-one in the contemporary literature has explicitly declared allegiance to consequentialism about concepts (this is partly because almost no-one has declared allegiance to *any* thoroughgoing normative framework in this arena, for a recent exception see Queloz 2022). Nonetheless, broadly consequentialist thinking dominates the literature in subtle and no-so-subtle ways. Some explicitly identify the practical question with the question of which good consequences we wish to produce and which bad consequences we wish to avoid.⁹ When other theorists in this domain describe the normative frameworks that guide their engineering projects, they often use explicitly instrumentalist language, speaking in terms of cost/benefit analyses¹⁰, of ‘increasing’ or

⁸ Though this program would eventually have to answer Nietzsche’s question: why should we seek to represent independent reality? An answer to this question, he believed, would have to “stand on moral ground” (Nietzsche 1882, 344).

⁹ See (Simion 2018):

¹⁰ In the background of the theoretical argument... is a kind of linguistic cost-benefit analysis.... We should abandon a term when the linguistic costs of continuing to use it outweigh the benefits, and the costs of organising to abandon the term do not outweigh the potential gains in linguistic utility.... Some words contribute enormously to bad political projects, and other words contribute more marginally. We should only abandon

‘promoting’ certain values¹¹, or of ensuring that their project has the desired ‘impact’¹². Moreover, most theorists in this domain refer to concepts as things, tools, or ‘representational devices’, and this has led a growing number of them to what is being called the ‘functional turn’ in conceptual engineering. According to this increasingly popular view, concepts are just instruments which ought to fulfill desired functions¹³.

It is important to see that anyone who says that the normative significance of Xs is exhausted by their fulfilling some function is almost automatically committing themselves to consequentialism about Xs. This is true on virtually any theory of functions, but it’s worth illustrating by citing the dominant theory of functions, the *etiological* theory, according to which a thing’s function is just the effect for which it has been selected in the past. As Peter Graham writes, such functions are just the result of ‘consequence etiologies’, where a something regularly produces some effect, and continues to exist *because* it has that effect (Graham 2014, 25).¹⁴ To fulfil a function is to produce an effect, and if the normative significance of X is that it fulfils a function, its normative significance is exhausted by its

words that make significant bad contributions.” (Habgood-Cooté 2020, 21). Or, see (Andow 2020) who is admirably forthcoming about their assumption that conceptual engineering rests on a consequentialist cost-benefit analysis.

¹¹ Sally Haslanger describes the fundamental questions guiding her inquiry into concepts: “What is the point of having these concepts...? Are they effective tools to accomplish our (legitimate) purposes; if not, what concepts would serve these purposes better?” (Haslanger 2000, 33) And with respect to specifically social concepts, she says that the eventual aim is to “[improve] our concepts in ways that will promote greater justice,” or which will improve the “functional role(s)” that concepts play in our social lives. (Haslanger 2020, 230)

¹² Robin Dembroff writes that an engineering project is “one that sets out to elucidate and possibly revise or replace our everyday concepts in light of *the impact we would like them to have*...what is important is that the final concept serves the proposed purposes.” (Dembroff 2016, 3, emphasis added)

¹³ See (Simion and Kelp 2020, Thomasson 2021) and (Queloz 2022).

¹⁴ The only way for a functionalist to avoid this is to specify functional value in an agent-relative fashion, as Mathieu Queloz’s recent concern-satisfaction account does. According to this account, concepts are functional for me when they satisfy my individual “needs, interests, desires, projects, aims, and aspirations” (Queloz 2022, 19), but this means that every concept has (at least) as many functions as there are users of the concept. This is a highly nonstandard use of the word ‘function’ but of course it is open to Queloz to use it, and it remains the case that for every other theorist in the field, “function-first” just means “consequentialist”.

producing effects. This is just consequentialism, and the use of alternative labels like “functionalist” often obscures this important fact about a theorist.

It might be objected here: can't a theorist say that a concept's function is to produce some deontologically specified aim? For example, couldn't they say that concept X has the function of securing people's *rights*? It is important to see that this is still a consequentialist model. After all, if we can modify the concept so that one person loses their rights while another one hundred *gain* rights, this model would tell us to do so, since we would be producing a more functional concept. This is just the old observation that consequentialist models always permit *trade-offs*, whereas deontological models do not. In this paper, I will try to show, conclusively, that many such trade-offs are deeply problematic in the conceptual realm.

There is another way of putting all of this. On the sort of account we are discussing, the functional object itself has no intrinsic normative significance, rather, its value is entirely given by a causal effect it has on us and the world. If the sole function of hearts is to pump blood, and if a thing's normative significance is entirely given by its function, then anyone who chose to keep their own heart rather than accept a replacement that pumps blood more efficiently would necessarily be making a mistake. But a moment's reflection reveals that this is not at *all* obvious: it cannot be that a resistant person who says ‘but this is *my* heart, I want to keep it’ is just automatically mistaken because the high-tech mechanical replacement is more efficient and less subject to wear and tear. If ownership is a reason to value my own heart, then functionalism about the value of hearts is false. Similarly, as we will see, people have non-consequentialist reasons to keep concepts that are, in a special sense, theirs.

Yet, the dominance of consequentialist thinking in this arena means that all of this has gone unnoticed, and that the contrary, deontological approach to conceptual engineering not made an appearance. Consequentialism about concepts is not anyone's *explicit* position, but it threatens to constitute a kind of unacknowledged default or starting-point for the contemporary conceptual engineer. This paper is meant to act as a bulwark against this tendency. We should explicitly acknowledge that non-consequentialist constraints must regulate the activity of conceptual engineering, and we must start to think carefully about what those constraints are. We can begin to see why this is so when we begin to think about the relation between concepts and personhood.

Concepts and Persons

It is probably fair to say that the phenomenon of situated personhood is at the heart of most classical objections to consequentialism.¹⁵ To many, it just seems implausible that persons are the sorts of things that should be fed into a moral cost-benefit analysis. Or, from a different angle, it seems implausible to many that persons ought to see their own actions in such terms.

In my view, consequentialist language comes so naturally to the contemporary conceptual engineer partly because concepts are being thought of as *things*, and hence, analogous to objects in a toolbox. Prominent authors refer to them as “tools” and as ‘representational devices’, encouraging us to think of concepts as objects that people pick up and use for

¹⁵ For example, John Rawls believed that consequentialism could not accept “the separateness of persons”, and Bernard Williams accused the utilitarian of being unable to account for the structure of intentional agency (Rawls 1971, Williams 1981).

various purposes (Haslanger 2018, Burgess and Plunkett 2013). Yet, no-one should accept that concepts are *things*, not in this sense anyway. This is something that is immediately apparent once we consult popular views in the ontology of concepts.

Eric Margolis and Stephen Laurence have helpfully divided these views into two camps (Margolis and Laurence 2007). One views a concept as a mental representation of some sort, whereas another sees them as a kind of abstract object. On neither view is a concept anything like a tool.

For those who see concepts in the first, roughly Fregean way, concepts are abstract semantic entities, or the *components* of propositions. They do not exist in the mind, nor do they participate in causal relations. Rather, they are the objects which stand in logical relations to one another and which enable us to think contentful thoughts. A concept, such as “Pegasus”, has a *sense* which, when grasped by a user, enables them to think thoughts about Pegasus, and, most importantly, to express propositions using the concept, even if the concept itself fails to refer (Frege 1892).

On the other hand, many theorists—including most conceptual engineers who write on this topic—embrace the idea that concepts are mental entities of some sort. They may be ‘mental representations’ which are “the constituents of propositional attitudes such as beliefs and desires” (Margolis and Laurence 2007; see also Simion and Kelp 2000, 986, Haslanger 2018) or perhaps a kind of *ability* to distinguish between objects in the world (Dummett 1993, Bennett and Hacker 2008).

It should be clear, I think, that on neither view is a concept akin to a hammer that a user might pick up and wield for some independently specifiable purpose. Indeed, on the first

view, it is not even clear what conceptual *engineering* would amount to at all, since there is no room in this view for the modification of abstract semantic entities. On such a theory, you either grasp the concept or you do not, and there is no sense in which we can modify it (Peacocke 1992).

On the second view, which describes concepts as mental entities embedded in the causal nexus, it certainly makes sense to say that concepts can be modified through intentional action. However, on this view, once again, a concept is not actually a tool in any ethically significant sense. It is a constituent part of a person's *mind*. Thus, on the only ontology which enables us to speak in any clear way about intentionally modifying a concept, a concept is not a tool, it is part of what makes up a *person*. Consequentialism about such entities is therefore on rocky footing from the very outset, since, again, persons are not typically the sorts of things that fit, intuitively, into a consequentialist calculus.

Now, of course many conceptual engineers accept what is called *semantic externalism*, or the idea that part of what gives a concept its semantic *value* are its relations, perhaps causal, to a reality outside the concept-user's mind (Putnam 1975, Haslanger 2012). I do not deny this view. My claim is about the *ontology* of concepts. Thus, particularly for the purposes of conceptual engineering, it seems more plausible to say that whatever they are, concepts are not mere impersonal things in the world, they are at least partly constituted by mental states or abilities.¹⁶ This is perfectly consistent with their semantic *meaning* being

¹⁶ Note that this is even true under the most extreme forms of the *extended mind thesis*, which hold that the mind itself is "extended" into the world beyond the body and brain itself. On this view, the mind itself is still *partly* located in the brain (Clark and Chalmers 1998).

determined by external facts, just as the meaning of a word is given by much more than the ink it is printed with.

But if this is correct, then when we are inquiring into the value of concepts, with an eye to modifying them, we are certainly not just inquiring into the goodness of algorithms or tools, we are talking about persons and their mental and social lives. It is therefore very important to see that unqualified consequentialism about concepts is already deeply contestable at the outset, and not something that should enjoy any sort of default status.

Moreover, even if there were some value in distinguishing between concepts and persons at the ontological level, there is still the question of how agents must *see* certain important representations. Certainly, there are concepts on one end of a spectrum, such as HOUSE, where the line between concept and person can seem relatively sharp to the user of those concepts. They are, we might say, not essentially person-involving, at least in the sense that they refer to things that are not persons. However, for *identity-constituting* concepts like JAPANESE or LESBIAN, this line blurs to the point of nonexistence. People *are* what such concepts represent, since they identify with such concepts, and this identification is laden with emotional and narrativial meaning. As Kwame Anthony Appiah says of racial concepts:

Once the racial label is applied to people, ideas about what it refers to come to have their social effects. But they have not only social effects but psychological ones as well; and they shape the ways people conceive of themselves and their projects. In particular, the labels can operate to shape what I want to call “identification”: the process through which an individual intentionally shapes her projects—including her plans for her own life and her conception of the good—by reference to available labels, available identities (Appiah and Gutmann 1998, 78).

This process, as Ian Hacking has persuasively argued, is ubiquitous: all human identities are ‘made up’ in this sense. Since all intentional human action is action under a description, in

carrying out our most personally significant actions and plans, each of us must draw from culturally laden descriptions that existed before we started to act (Hacking 1986, 166, Anscombe 1957). These descriptions provide us with norms, ideals and values that we adopt as central parts of our practical lives. This is one way to see how consequentialism about identity-constituting concepts is particularly odd: theorists are asking which of our values the concepts might serve or promote, but some of these concepts quite literally *are* our values.

Such concepts help to constitute human agency. Inquiry into the value of such concepts is therefore not just inquiry into the way the world is, it is also inquiry into the worth or value of persons. A project which seeks to revise or eliminate identity-constituting concepts is seeking to revise or eliminate portions of the practical identities of real persons. As we will see, the issue isn't that trying to improve people is necessarily problematic, what is problematic is attempting to do so under a set of normative assumptions that renders your social authority null and void. One such assumption is that concepts only have instrumental significance.

Deontology, Concepts, and Social Authority

What if concepts have, in addition to instrumental significance, non-instrumental significance? Specifically, what if deontic facts concerning rights, claims or entitlements apply in the conceptual realm? This is not a possibility that has been canvassed in the literature, but I think that it is highly plausible. For example, what if some linguistic groups have a *right* to their concepts, rights which must be respected in order for an engineering

project to count as legitimate? If this is so, then conceptual engineering simply cannot rest on a cost-benefit analysis, rather, the engineer must conduct far more complex normative inquiry, one which also thinks in terms of the *claims* that people have on us with regard to our conceptual activity.

This is where the notion of an identity-constituting concept becomes particularly crucial. For I think that it is deeply plausible that that people have a strong *prima facie* right to such concepts. This means much more than just that they have the right to *think of themselves* as falling under the concept. They have the right to live in a surrounding community which generally respects the classificatory boundaries under which they have formed their identities and which counts them as falling under those boundaries.

An example here is the English language concept GAY as deployed and self-ascribed by certain communities. The ancestor-concept merely referred to a happy, carefree, or showy person, yet by 1960 it had morphed into a slur for homosexual men. During this second stage, the concept was merely imposed upon an unwilling population, and was not identity-constituting in my sense. However, homosexual persons began to reclaim the term as a locus of pride and struggle, often internalizing a set of norms and values in the process. It is at this point that the notion of a *right* to the concept obviously begins to apply, such that now, anyone who successfully managed to get us all to revert to the original meaning would almost certainly wrong this community. Moreover, it is clear that this conceptual right is explained in large part by the fact that the concept has become identity-constituting: for many, it acquired non-instrumental significance, serving as the practical lens through which they see the world, and not merely as an external element of that world. To modify the concept is to risk wronging such persons.

For example, an ambitious engineer might claim that concepts of sexual orientation themselves enable a huge amount of oppression and create a great deal of unhappiness. From a certain detached viewpoint, it can seem odd that we have decided to divide human populations into distinct categories based roughly on who they tend to sleep with. Such thoughts can lead us to ask highly general questions like: what's the overall point or function of this conceptual division? What is it doing, in general, for humanity? This is the mindset of the conceptual consequentialist, who is not even *asking* about the Gay community's rather obvious *prima facie* right to retain their identity-constituting concept, whether or not it is *accomplishing* anything. I will come back to this point later, when I give a similar argument about racial concepts, but for now, I hope the presumptive case against conceptual consequentialism is clear, and that the notion of a right to a concept is at least starting to make sense.

But what is the nature of this right? It's worth noting that such rights are claims against a *community* and not claims against any given person, unless that person has the power to actually enact the linguistic/conceptual shift all on their own. That is, your right to a concept does not entail that any individual person wrongs you by using the concept differently. Rather, your claim is either against the community of language-users, or against a person (i.e. an effective conceptual engineer) who can enact the conceptual shift.

Moreover, intentionally initiated conceptual change with respect to identity-constituting concepts *can* be justified. It might be thought that the function of this paper is to defend a conservative position on various popular engineering proposals in general, but this is not so. The reader might have been worrying that, for example, my position entails that gender concepts cannot be legitimately engineered, and that this paper is thereby committed to some

problematic or controversial stance. But rights to identity-defining concepts are not absolute or inviolable. However, I argue that persons inhabiting modified identities are owed a certain kind of story about social authority:

Conceptual Authority: concept-users have the right to have their existing concepts managed or modified only by persons with the *authority* to do so.¹⁷ Moreover, this right becomes stronger to the extent that the relevant concept plays an identity-defining role in the concept-user's life.

Now, notice that an *additive* engineer need not concern themselves with this requirement. Since such engineers do not propose to directly modify any existing concept, it is clear that questions about their social authority are far less urgent. When Kate Manne proposes that the novel term ‘himpathy’ usefully picks out a socially important phenomenon—disproportionate sympathy for powerful men—she does not incur the burden of explaining to any of us why she has the authority to enact this shift (Manne 2017). This is because she is not seeking to modify any concept that plays any existing role in a person's life. But revisionary and eliminativist engineers must explain their right to revise or eliminate what is in effect a shared resource.

What does it mean to exercise authority in this social sense? There are of course many possible models here and I lack the space to discuss them all. Instead, for the purposes of discussion I will adopt the following relatively mainstream model of social authority:

Authority: a necessary condition on my exercising social authority over a population with respect to the management of some shared resource is that I can tell a truthful story which strongly relates my ends and my chosen means

¹⁷ This right may be weakly defeasible: it may be that some people, via immorality or participation in injustice, forfeit this right, just as the famous “murderer at the door” has no right to the truth (Kant 1797). Thus, no Kantian implications should be read into these passages.

to strong practical reasons possessed by the members of the population itself.

More concretely, it must be the case that most reflective, well-informed members of that population could be shown that they have strong reasons to endorse the social changes that I am attempting to enact. This completes the required deontological theory of conceptual engineering, by supplying a non-consequentialist *prohibition* against a certain kind of illegitimate conceptual management. This hypothetical consent model is obviously related to social contract theories more generally, but without it I simply cannot see how we can avoid trampling on the conceptual rights which are intuitively possessed by many people. Again, I think the intuition here is clear: anyone who wished to significantly revise or eliminate the concept GAY would, I think, bear the burden of showing how most actual gay persons would have strong reasons to endorse that change.

This moderate condition does not entail that affected persons must have all-things-considered reasons in favor of the change, nor does it mean that the actual agents themselves will endorse the change. Moreover, it doesn't entail that *every* person has a kind of veto over conceptual change; no political theory that concerns itself with the regulation of a shared resource can give any one person that sort of control. The idea is simply that absent any real connection between my activities and the well-informed values of most affected persons, I am proposing to merely manipulate or coerce the population and am not exercising legitimate authority over them.¹⁸

¹⁸ This condition is obviously related to *Public Reason* models of political authority, require that the policies which regulate social life be, in some sense, justifiable to all who must obey or comply with such policies. The critical literature here is large; see (Rawls 1999, Gaus 2010, Vallier 1996, Quong 2013, Enoch 2013).

Identity-conferring concepts ground rights which can only be overridden by those with the social authority to do so. This already entails that consequentialism about concepts is false. This is because for at least these concepts, it doesn't matter whether they can be modified to produce more value; if no existing person has the authority to modify them, then they shouldn't be modified.

Arguably, it is not just identity-conferring concepts that trigger this result, since it is also intuitive that people are entitled to many other kinds of concepts, even if they are not identity-conferring. As an atheist, I might easily be persuaded that the human race would be better off if the concept GOD was intrinsically associated with the idea of a mythic invention. In fact, with all due respect to my religious friends, I think I *do* believe this. But I do not believe that this fact alone could justify the modification of the concept in my linguistic group, or that my beliefs about long-term value give me any authority to enact such a shift. This is not just because the means to this end would be problematic, it is also because people seem to have a right to a concept that plays such a central role in their lives. This is so even if it does not literally *define* them in the way that other concepts do.

I hope it is now clear that deontological considerations appear to be deeply relevant to the ethics of concepts, and that corresponding questions about social authority cannot be avoided. Yet, so far as I can tell, aside from a brief discussion by Haslanger a decade ago (Haslanger 2012, 25-27), recent conversations have almost entirely avoided them.¹⁹ The

¹⁹ Haslanger settles on the idea that a critical revisionist project must be one “that the subordinated and their allies find it illuminating or useful” because “it contributes to their quest for social justice.” I am not discussing this account because it is not an account of conceptual authority ; by focusing on only a small number of concept-users, and by imposing a substantive and contested value (social justice) on to the discussions they must have about such proposals, Haslanger guarantees that most concept-users simply won't feature in the account. But this means that the engineer's authority over *them* is not explained.

question of social authority does not even appear in Cappellen and Plunkett’s recently published list of twenty-four definitive questions for Conceptual Engineering (Burgess, Cappellen, and Plunkett 2020, 2-15).²⁰ This, I suspect, is because (as we have already seen) they may be conceiving of conceptual engineering in an ideal-theoretical way, as an intellectual exercise rather than as a practical activity. Moreover, as I have already stressed, for most authors working in conceptual ethics facts about what might be generally useful or expedient are normally decisive with respect to the question of what ought to be done with our concepts. And when the job looks particularly tough, they puzzle over what they call, using unabashed technobureaucratic language, ‘the implementation problem’.²¹

So let me be blunt: human culture is not properly analogous to an engine that needs a tune-up, because it is partly composed of elements that can easily ground rights and entitlements or constitute lived practical identities. That is not an *implementation* problem, it is an *ethical* problem, one that is only solved by the acquisition of legitimate social authority. Moreover, unlike government agents or actual engineers, academics are not elected or even indirectly certified by the persons whose concepts we are proposing to change.²² Adopting a phrase from Williams, we can see that this activity is in danger of becoming *Government House Conceptual Engineering* (Williams 1985). The reference here is to the residences of colonial masters, and it is not irrelevant that consequentialism began its practical life partly as a

²⁰ As the selection shows, the contemporary literature is primarily focused on (1) the nature of the metasemantic base that generates conceptual meaning and whether we have control over it (Cappellen 2018, Pollock 2020, Jorem 2020, Koch 2021), (2) the question of whether conceptual engineering merely amounts to “changing the subject” (Sundell 2011, 2017, Prinzing 2016, Thomasson 2019).

²¹ See “The Implementation Challenge for Conceptual Engineering (ICCE)”, conference held online Dec 10-11, 2020. Abstracts available at <https://www.st-andrews.ac.uk/arche/event/oow-icce/>. See also (Jorem 2021, Deutsch 2020). For a powerful and important critique of the language of ‘implementation’ here, see (Queloz and Bieber 2021).

²² When most people think of the expertise that is legitimately recognized by academic titles, they do not assume that this includes expertise with respect to *what our words ought to mean*.

systematic justification of British imperial domination (Habibi 2006, Ahmed 2010). To be clear, I am not likening the contemporary conceptual engineer to a colonial overlord in any direct moral sense; that comparison would be preposterous. The point is to identify a shared managerial mindset, one which takes on social problems after deciding, in advance, to ignore such things as rights and entitlements, mainly by deploying an instrumental, cost-benefit analysis.

Let me sum up. Conceptual engineers, like all persons seeking to modify a shared resource, face an authority question. Moreover I have argued that the conceptual *consequentialist* cannot provide a satisfactory answer to that question. Legitimate conceptual change requires that at least a significant portion of actual agents within a group change how they think and speak, and so this change must be something that these agents in general have reasonably strong practical reasons to endorse. Otherwise, the agents are *just* being manipulated, and while this doesn't directly entail that the manipulation is unjustified, it does entail that the engineer lacks social authority in a very important sense. But the stubborn human identification with things other than long-term agent-neutral value virtually guarantees that a consequentialist engineer must lack social authority as I've defined it. People value diverse things, and they don't always just care about values being promoted in some objective sense.²³

Because the reader may not yet be convinced that this is the case, I'll finish by illustrating the point with two examples, an engineering project with social authority and one without it.

²³ On the important distinction between promotion and honoring as it applies directly to this issue, see (Pettit 1989).

Social Authority and Concepts: Two Case Studies

In many societies, the concepts MARRIAGE and MARRIED have undergone significant revision in the past few decades. This change has been swift, and while its causes are debated, it is clear that conceptual innovators and political activists were key engineers, provoking what to many has seemed like an obvious social improvement. Yet, it might be thought that the conditions I am imposing in this paper call this shift into question, since a great many people have complained (and continue to complain) about the change on the grounds that they, as married people, do not wish to recognize non-heterosexual marriage. They claim that this move towards inclusivity is one that cheapens the concept for them, much as my hypothetical artist felt slighted by the inclusion of bricklayers under the concept ARTIST. Moreover, my proposed deontic restriction is clearly triggered here, since this concept is at the very least identity-adjacent: many such people identify very strongly with various family roles that are predicated on traditional ideas about marriage. I think that these people should not oppose the conceptual shift, and I furthermore think that the panic over this shift is entirely ungrounded. But haven't I argued myself into portraying these complaints as legitimate, and aren't I forced to say that this shift has occurred without any engineer possessing the relevant social authority?

As you might have guessed, I am not forced to say any such thing. For it is not the case that the actual normative standpoint of any existing concept-user must be respected if my condition is to be met. My account of social authority, like all such accounts, *idealizes* away from the real perspectives of affected people. There are many difficult questions about just

how far this idealization can go, but it is clear that in order to legitimately dissent from a proposed conceptual innovation, a person must not be operating under any obvious bias, nor can their opposition be driven by mistaken empirical facts. And in the case of opposition to gay marriage, I am perfectly prepared to defend the idea that both of these distorting factors are normally at play, though of course that is a substantive argument that remains to be made. This argument might begin by noting that at least as of the mid-90s, nearly one-half of Americans opposed to gay marriage reported feelings of disgust towards gay and lesbian persons, and roughly the same proportion reported believing, falsely, that homosexuality is a choice (Nussbaum 2006, Rimmerman and Wilcox 2007, 223). It is unsurprising that opposition to same-sex marriage has declined as precisely these distortions and falsehoods have started to fade away.

I am therefore confident that an empirically informed normative argument could show that most or possibly all moderately idealized persons will have strong reasons to approve of the expansion of the concept MARRIAGE, and so the authority condition is met.²⁴ Thus, if a conceptual activist is questioned about their social authority with respect to this expansion, they can legitimately say to their opponent: ‘a more rational version of your present self would have no significant problem with this shift.’ They will not just be saying: ‘the world will be better after this shift’, they will respect the population’s right to see this shift as an improvement. However, other engineering projects seem forced into silence at this crucial

²⁴ Indeed, I think that anyone who wishes to portray this opposition to gay marriage as a moral error must be willing to make the same kinds of claims: you cannot say that someone else is systematically mistaken on any matter without being prepared to offer some explanation for how their beliefs have been distorted (CITATION REMOVED).

juncture, precisely because they focus so heavily on that possible-future *just world*. This is easiest to see when we examine certain revisionary proposals in the Philosophy of Race.

Haslanger has provided an oft-cited and dramatically revisionary engineering proposal in this domain, one which perfectly illustrates the tension between consequentialism about concepts and principles of social authority. In what follows, I will try to show that her proposal has content that virtually guarantees that it cannot be carried out with the requisite authority. This, I should stress, is not because of who Haslanger is, but because of the consequentialist normative perspective that guides the proposal.

While not a straightforward eliminativist about race, Haslanger has proposed that racial concepts should nonetheless acquire a very different meaning, such that they pick out members of a group that are “socially positioned as subordinate or privileged” and “*marked* as a target for this treatment by observed or imagined bodily features presumed to be evidence of ancestral links to a certain geographical region.” (Haslanger 2000, 44) The goal here, as she said at the time, was to eventually move towards a post-racial society, one which has successfully “undermine[d] those forces that make being a member of a racialized group possible”, such that we can “refuse to be raced.” (Haslanger 2000, 48) And this proposal was meant to be judged by whether it “serve[s] both the goal of understanding racial... oppression, and of achieving.. racial equality.” (Haslanger 2000, 47) The reader hopefully will, by this point, recognize the hallmarks of a consequentialist normative framework. Here, it is valuable future-possible states of affairs (social understanding, social equality) which are said to justify the conceptual shift.

Haslanger has since qualified this proposal, stating that it is mainly meant to provoke discussion, and of course not anything she wished to actively impose on anyone (Haslanger 2019). Yet, as part of this proposed discussion, we should ask: what would it mean to try to produce this conceptual shift?

Note that the racial concept BLACK, on the revisionary analysis, should ideally be used to talk about what “really matters”, i.e. the systemic subordination of people according to presumed recent geographic origin (Africa). Blackness, according to this revised conception, is something that is done *to* a person, a set of disadvantages they accrue as a result of how other people see them. But BLACK is an identity-constituting concept, one which plays a vital role in the lives of so many people (Haslanger 2013, 18). So my proposed deontic restriction on this engineering project is certainly triggered. Given these facts, do most or even many existing African-American persons have strong practical reasons to support this proposed shift?

Of proposals like Haslanger’s, Chike Jeffers writes:

What it means to be a Black person, for many of us, including myself, can never be exhausted through reference to problems of stigmatization, discrimination, marginalization, and disadvantage, as real and as large-looming as these factors are in the racial landscape as we know it. There is also joy in blackness, a joy shaped by culturally distinctive situations, expressions, and interactions, by stylizations of the distinctive features of the black body, by forms of linguistic and extralinguistic communication, by artistic traditions, by religious and secular rituals, and by any number of other modes of cultural existence. There is also pride in the way black people have helped to shape Western culture... directly through cultural contributions, most prominently in music and dance. (Jeffers 2013, 422)

In a similar spirit, Kathryn Belle writes:

The idea that we can simply make racial categories disappear is a naive one which does not take into account the historical significance of race. Race is not just a

negative category used for the purpose of oppression and exploitation or for the purpose of establishing a sense of supremacy over others. Race has also come to represent a more positive category that encompasses a sense of membership or belonging, remembrance of struggle and overcoming, and the motivation to press forward and endeavor towards new ideals and achievements (Belle 2003, 56, cited in Jeffers 2013). |

Such thoughts are not in any sense idiosyncratic: they are echoed by many people from many walks of life when prompted to speak about the meaning that the racial and ethnic concepts have for them.²⁵ Of course I do not mean to suggest that Haslanger is unaware of all of this, but I'm not sure that its significance has been recognized in the conceptual engineering literature at all.²⁶ Positive thoughts about achievement, struggle and overcoming are quite obviously central to other racial (and ethnic) identities as well. Such thoughts can therefore ground what moral philosophers call an *integrity objection* to consequentialist conceptual engineering: why should actual embodied people think about their own identity-grounding concepts in the terms set by a theorist who is only focused on long-term agent-neutral value (Smart and Williams 1973)?

Importantly, Haslanger has made it clear that she is not overly interested in what ordinary race talk is “really” about, specifically in the intuitive conditions that constitute its so-called “manifest” content. “My priority in this inquiry,” she writes, “is not to capture what we do mean, but how we might usefully revise what we mean for certain theoretical and political

²⁵ See for example the wide array of perspectives collected in Touré's *Who's Afraid of Post-Blackness? What It Means to be Black Now* (Touré 2011, Womack 2010). For criticism, see (Baker Jr and Simmons 2015). See also *Black Cool: One Thousand Streams of Blackness* (Gates Jr 2012).

²⁶ Her gloss on this is to say: “In fact, I believe that many forms of racial identity are important, valuable, and in some cases even inevitable responses to racial hierarchy. As I see it, a racial identity is a kind of know-how for navigating one's position in racialized social space. The apt content for a racial identity, then, may be positive, affirming, and empowering, even if the racialized social position one occupies is oppressive.” It is worth noting, however, that such identity is still portrayed as fundamentally reactive, as a kind of protective mechanism shielding an agent from the worst of oppression, and not as an *independently* valuable, life-giving force. Racial identity is valuable *because* it mitigates injustice. Moreover, tellingly, she goes on to question the idea that this protective value ought to receive any priority “in the long run”. This, I believe, is just the consequentialist's famously problematic “long run”. (Haslanger 2019, 30)

purposes” (Haslanger 2012, 224). However, the problem of social authority *requires* that we inquire into the manifest content of identity-conferring concepts, because the shape of that ethical problem is given precisely by the emotional and existential attachments that existing people have to that content. To block this question at *this* juncture is to entirely give up on the project of securing social authority. Since BLACK is an identity-constituting concept, the question of social authority is urgent, and its contours are delineated by the professed perspectives of Black persons who reflect on what Blackness means for them. Jeffers and Belle are plainly not just worried about externalist semantics or conceptual analysis, rather, they are offering personal and political objections arising from a shared practical perspective. The reasons that partly constitute this perspective are grounded in such things as joy, hope, community, defiance, and as Anthony Appiah recognizes, from the need to live an *authentic* existence in the face of more dominant or hegemonic cultural forces (Appiah 1996).

In making this proposal, Haslanger envisioned a future in which racial categories simply mark a form of oppression which we then seek to eliminate. We may even grant, for the sake of argument, that this future is “better” in some agent-neutral sense (perhaps the world will be happier, or more just, or something like that). The only way to move from *this* fact to the conclusion that we ought to create that world is to embrace some form of consequentialism about concepts, and this is why I say that the proposal is surely consequentialist in nature. Yet, right now, existing people identify with the cultural norms and values associated with racial concepts, and it is clear that such distinctive norms and values could not survive in a post-racial world. There would be no such thing as distinctively “Black” music, language, dance, religions, or modes of personal appearance, because, as Jeffers notes, these elements are “cultural contributions whose significance can only be fully understood when they are

placed in proper context as emerging from a racialized people.” (Jeffers 2013, 422). It is not hard to see why Jeffers and Belle are nervous about all of this: their real-world commitment to a shared identity and a way of life means they cannot actually prioritize this possible-future ‘better’ conceptual scheme.

Of course, their actual nervousness is not decisive, for *if* the idealized counterpart of a Belle or a Jeffers would drop their resistance, then Haslanger’s proposal would pass the authority test. But in order to establish this counterfactual, we’d need to portray such attachments and resistances as *distortions* or *prejudices*, grounded in some kind of moral bias or false belief, as homophobia is. But this, to put it mildly, is an uphill battle. Unlike the case of homophobic resistance to gay marriage, it is just completely unclear how personal reasons to embrace and promulgate the joy, hope, struggle and shared community that comes along with certain forms of racial identity could ever be reasonably seen as bias or prejudice in any interesting sense. If they could be, then I would happily concede that the project might go ahead, but importantly, it would be going ahead on *non-consequentialist grounds*, as a project which has properly posed and answered deontic questions about rights and authority before trying to improve the world. My basic argument in this paper would be unaffected by any such result.

Yet, because the prospects for describing Jeffers and Belle as biased or misinformed seem dim, I conclude that any such revisionist proposal almost certainly cannot secure the social authority it needs. That is, while the modified conceptual scheme might be ‘better’ in some impersonal sense, such that an objectively real problem has been ameliorated, the project itself could only be pursued by ignoring perfectly justified resistance to it on the part of affected groups. This, I suggest, is *because* the normative justification is itself

consequentialist. The essential idea is that existing racial concepts produce on-balance harm or injustice, and so they must be revised in order to eliminate this harm. This is just the consequentialist model I outlined earlier, whereby theorists evaluate concepts instrumentally, in terms of their causal effects, and seek to produce better overall states of affairs by modifying the concepts themselves. But this normative perspective, we can now see, is virtually guaranteed to lack social authority in the social world we actually inhabit.

What About Indirect Consequentialism?

Finally, I want to note that at this point in the normative ethical dialectic, a different position usually rears its head. This is *indirect* consequentialism, theory which restricts itself only to a definition of right action. Crucially, this means that the theory does not recommend that any agent should actually *think* like a consequentialist, because the theory does not tell anyone how to think or act (Bales 1971, Parfit 1984, Pettit 1997, Wiland 2007). If a conceptual engineer were to adopt this framework, they might claim that they are only saying something very simple: the best conceptual scheme is the one which produces the most agent-neutral value, and we ought to do whatever will promote that scheme. If people in general will most efficiently move towards this scheme by (for example) endorsing conceptual rights and holding on to identity-constituting concepts, then so be it.

Such a position does not help at all with questions of social authority. Recall that this paper is entirely focused on the nonideal question of what ought to actually guide a conceptual engineering project. The indirect theory, if it is to be put into practise, must be effective through some agent or group of agents who seek to discover and implement an optimific strategy. And should those agents attempt to deploy consequentialist reasoning, the problem of authority will return to haunt those same agents and their activity. Moreover, the fact that the indirect consequentialist is proposing to *hide* the decisive normative facts from the population at large guarantees that they cannot actually explain their own social authority to the population itself. Thus, I conclude that consequentialist reasoning should not guide any actual conceptual engineering project, and, to repeat, this entails that consequentialism about concepts is false.

Conclusion

This is consequentialism about conceptual engineering: a normative perspective that can have no story about its own authority over our linguistic activity. The cure here is to reject the abstract, totalizing perspective that is being implicitly imported into recent discussions. This means that we can insist that individuals do *not* have to be haunted by the mere possibility of an agent-neutrally ‘better’ conceptual scheme floating around in logical space, that there are limits on the kind of conceptual change we can demand of communities, and that we should respect a linguistic community’s agent-relative reasons for holding on to various concepts, most prominently identity-constituting ones. Such moves are familiar in normative ethics; I suggest that they need to be made here as well, if the normative perspective of the conceptual engineer is to have any hope at all of legislating to humanity.

Cited Works

- Andow, James. 2020. "Fully experimental conceptual engineering." *Inquiry*:1-27.
- Anscombe, Gertrude Elizabeth Margaret. 1957. *Intention*: Harvard University Press.
- Appiah, Kwame Anthony, and Amy Gutmann. 1998. *Color conscious*: Princeton University Press.
- Arthur, John. 2009. "World Hunger and moral obligation : The case against Singer." In *Exploring Philosophy: An Introductory Anthology*, edited by Steven M. Cahn. London: Oxford University Press.
- Baker Jr, Houston A, and K Merinda Simmons. 2015. *The trouble with post-blackness*: Columbia University Press.
- Bales, R. Eugene. 1971. "Act-Utilitarianism: Account of Right-Making Characteristics or Decision-Making Procedure?" *American Philosophical Quarterly* 8 (3):257-265.
- Belle, Kathryn Sophia (Kathryn T. Gines before 2014/2017). 2003. "Fanon and Sartre 50 years later: To retain or reject the concept of race." *Sartre Studies International* 9 (2):55-67.
- Burgess, Alexis, Herman Cappelen, and David Plunkett. 2020. *Conceptual engineering and conceptual ethics*: Oxford University Press.
- Cappelen, Herman. 2018. *Fixing language: An essay on conceptual engineering*: Oxford University Press.

- Chalmers, David J. 2020. "What is conceptual engineering and what should it be?"
Inquiry:1-18.
- Dembroff, Robin A. 2016. "What is sexual orientation?".
- Deutsch, Max. 2020. "Speaker's reference, stipulation, and a dilemma for conceptual engineers." *Philosophical studies*:1-23.
- Dreier, James. 2011. "In Defense of Consequentializing." In *Oxford Studies in Normative Ethics, Volume 1*, edited by Mark Timmons. Oxford University Press.
- Egré, Paul, and Cathal O'Madagain. 2019. "Concept Utility." *The Journal of Philosophy* 116 (10):525-554.
- Eklund, Matti. 2014. "Replacing truth." *Metasemantics: New essays on the foundations of meaning*:293-310.
- Enoch, David. 2013. "The Disorder of Public Reason." *Ethics* 124 (1):141-176.
- Frege, Gottlob. 1892. "On sense and reference." oversatt av Max Black, i J. Guitérrez-Rexach (red.): *Semantics: Critical concepts in linguistics* 1:7-25.
- Gates Jr, Henry Louis. 2012. *Black cool: One thousand streams of blackness*: Catapult.
- Gaus, Gerald. 2010. *The order of public reason: A theory of freedom and morality in a diverse and bounded world*: Cambridge University Press.
- Habgood-Coote, Joshua. 2020. "Fake news, conceptual engineering, and linguistic resistance: reply to Pepp, Michaelson and Sterken, and Brown." *Inquiry*:1-29.
- Hacking, Ian. 1986. "Making up people."
- Hardimon, Michael O. 2012. "Minimalist Biological Race." In *The Oxford Handbook of Philosophy and Race*, edited by Naomi Zack. OUP.
- Hardimon, Michael O. 2017. *Rethinking race*: Harvard University Press.

- Haslanger, Sally. 2000. "Gender and race: (What) are they? (What) do we want them to be?" *Noûs* 34 (1):31-55.
- Haslanger, Sally. 2012. *Resisting Reality: Social Construction and Social Critique*: Oxford UP.
- Jeffers, Chike. 2013. "The cultural theory of race: Yet another look at Du Bois's "The Conservation of Races"." *Ethics* 123 (3):403-426.
- Jorem, Sigurd. 2021. "Conceptual engineering and the implementation problem." *Inquiry* 64 (1-2):186-211.
- Kitsik, Eve. 2020. "Conceptual Engineering and Ways of Believing." *Erkenntnis*:1-22.
- Marques, Teresa. 2020. "Amelioration vs. perversion." *Shifting Concepts: The Philosophy and Psychology of Conceptual Variability*:260-84.
- Mill, John Stuart. 1859/1998. *On liberty and other essays*. New York: Oxford University Press, USA.
- Nietzsche, Friedrich, 1882. *The Gay Science*. London: Oxford University Press.
- Nussbaum, Martha C. 2006. *Hiding From Humanity: Disgust, Shame, and the Law*: Princeton University Press.
- Parfit, Derek. 1984. *Reasons and Persons*: Oxford University Press.
- Pettit, Philip. 1989. "Consequentialism and respect for persons." *Ethics* 100 (1):116-126.
- Pettit, Philip. 1997. "The Consequentialist Perspective." In *Three Methods of Ethics*, edited by M. Baron, P. Pettit and M. Slote. Blackwell.
- Podosky, Paul-Mikhail. 2022. "Agency, Power, and Injustice in Metalinguistic Disagreement." *Philosophical Quarterly*:1- 24.

- Portmore, Douglas W. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. London: Oxford University Press.
- Price, Huw. 2014. "Where would we be without counterfactuals?" In *New Directions in the Philosophy of Science*, 589-607. Springer.
- Queloz, Matthieu, and Friedemann Bieber. 2021. "Conceptual Engineering and the Politics of Implementation." *Pacific Philosophical Quarterly*.
- Quong, Jonathan. 2013. "Public Reason." In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta.
- Rawls, John. 1999. *The Law of Peoples: With, The Idea of Public Reason Revisited*. Vol. 20: Harvard University Press.
- Rimmerman, Craig A, and Clyde Wilcox. 2007. *The politics of same-sex marriage*: University of Chicago Press.
- Rosch, Eleanor, and Carolyn B Mervis. 1975. "Family resemblances: Studies in the internal structure of categories." *Cognitive psychology* 7 (4):573-605.
- Scharp, Kevin. 2007. "Replacing truth." *Inquiry* 50 (6):606-621.
- Shields, Matthew. 2021. "Conceptual domination." *Synthese* 199 (5-6):15043-15067.
- Simion, Mona, and Christoph Kelp. 2020. "Conceptual innovation, function first." *Noûs* 54 (4):985-1002.
- Smart, John Jamieson Carswell, and Bernard Williams. 1973. *Utilitarianism: for and Against*: Cambridge University Press.
- Thomasson, Amie. 2021. "Conceptual engineering: when do we need it? How can we do it?" *Inquiry*:1-26.
- Thomasson, Amie L. 2014. *Ontology made easy*: Oxford University Press.

Touré. 2011. *Who's afraid of post-blackness?: what it means to be Black now*: Simon and Schuster.

Vallier, Kevin. 1996. "Public justification."

Wiland, Eric. 2007. "How Indirect Can Indirect Utilitarianism Be?" *Philosophy and Phenomenological Research* 74 (2):275-301.

Williams, Bernard Arthur Owen. 1985. *Ethics and the Limits of Philosophy*. Vol. 83: Harvard University Press.

Womack, Ytasha L. 2010. *Post Black: How a new generation is redefining African American identity*: Chicago Review Press.

Cited Works

Frege, Gottlob. 1892. "On sense and reference." *oversatt av Max Black, i J. Guitérrez-Rexach (red.): Semantics: Critical concepts in linguistics* 1:7-25.