

# Conceptual engineering: when do we need it? How can we do it?

Amie Thomasson

To cite this article: Amie Thomasson (2021): Conceptual engineering: when do we need it? How can we do it?, Inquiry, DOI: [10.1080/0020174X.2021.2000118](https://doi.org/10.1080/0020174X.2021.2000118)

To link to this article: <https://doi.org/10.1080/0020174X.2021.2000118>



Published online: 16 Nov 2021.



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



Article views: 218



View related articles [↗](#)



View Crossmark data [↗](#)



# Conceptual engineering: when do we need it? How can we do it?

Amie Thomasson

Department of Philosophy, Dartmouth College, Hanover, NH, USA

## ABSTRACT

This paper addresses several foundational questions in conceptual engineering: When is conceptual engineering needed? When we engage in conceptual engineering, should we think of ourselves primarily as aiming to change concepts or language – and how should we think of either of those? Finally, how is implementing the changes recommended by conceptual engineering possible? I begin by outlining a number of different sorts of circumstances in which conceptual engineering is desirable, bringing out the commonalities across many classic projects in the history of philosophy and contemporary, socially-conscious, work in conceptual engineering. I then argue that some prominent ways of understanding concepts and meanings are unhelpful for conceptual engineering. We can do better if we think first and foremost of engineering *words*, considering words as a certain kind of abstract cultural artifact, which (like other artifacts) have functions and norms of use. Then we can begin to understand better how words *do* change naturally, and use that as a way to investigate how they (like other cultural artifacts and social norms) can also be changed artificially.

**ARTICLE HISTORY** Received 29 March 2021; Accepted 23 August 2021

**KEYWORDS** Conceptual engineering; words; concepts; conceptual change; abstract artifacts

Conceptual engineering is becoming popular. It's easy to see why. For it has applications all over philosophy – whether we are dealing with questions about how to (re-)engineer the concept of art, of freedom, of number, of probability, of truth, of privacy, of gender, etc. I have taken interest in it because I think that a good many traditional debates in metaphysics can be better seen as implicitly concerning problems in conceptual ethics and conceptual engineering – and that we can demystify metaphysics and make better progress by thinking of them in this light.<sup>1</sup> But conceptual engineering also has applications outside of

---

**CONTACT** Amie Thomasson  [amielthomasson@gmail.com](mailto:amielthomasson@gmail.com)

<sup>1</sup>See my (Thomasson, 2017).

This article has been corrected with minor changes. These changes do not impact the academic content of the article.

philosophy – to resolving problems in medicine, psychology, biology, physics, government and the law; and to serving the cause of social justice. Not only *can* conceptual engineering be applied widely – if people like Plunkett and Sundell (2013), are right, renegotiating our concepts, like speaking prose, is something we have been doing all along without realizing it.<sup>2</sup>

And yet, theoretical work about conceptual engineering is just beginning. Elsewhere (2020a, 2020b), I have sketched a pragmatic method for conceptual engineering – one that doesn't require that we discover obscure 'metaphysical facts' to guide our conceptual choices, and that nonetheless gives us objective criteria for success based on the functions the term is to serve – while being cognizant of the multitude of different functions our terms can serve. Here, I hope to sketch out two more parts of the program: one that comes before, and the other after, the functional assessment. The first is identifying targets for conceptual engineering: what are the signs that it is needed? When should we do it? The second comes after the functional assessment, and I will spend longer on that. That involves what has come to be known as the 'implementation problem'<sup>3</sup>: supposing we have determined how a concept *should* be modified, how is it possible to actually modify our concepts? On some views of concepts, changing concepts is simply impossible. On others, like Cappelen's (2018), while possible, the processes of conceptual change are said to be fundamentally inscrutable and out of our control. I will sketch a somewhat more optimistic picture – arguing that, if we think first and foremost of engineering *words*, and think of words as a certain kind of cultural artifact, governed (like other artifacts) by norms of use, then we can begin to understand better how words *do* change naturally, and use that as a way into seeing how they (like other cultural artifacts and normative practices) can also be changed artificially.

## 1. When do we need conceptual engineering?

One of the earliest explicit discussions of when and why we need to engage in conceptual engineering was Rudolf Carnap's discussion of explication (1950, Chapter 1). On the view he there defends, we need to engage in 'explication' when we have a relatively inexact concept, which we need to make more exact so that it can function better in our

---

<sup>2</sup>Following the old joke from Moliere, *The Middle Class Gentleman*, Act 2 Scene 4.

<sup>3</sup>For discussion of the implementation problem, see Cappelen and Plunkett (2020), Deutsch (2021), and Jorem (2020).

scientific theories. We improve a concept when it can be incorporated into a system of well-connected and precise scientific and logical/mathematical concepts, in a way that is also fruitful for formulating many universal laws (empirical laws or logical theorems) (1950, 3).<sup>4</sup>

Peter Strawson famously was interested in a different role for conceptual engineering – one that has played a more central role in the history of philosophy. Whereas Carnap aims to clarify concepts for use in the sciences, Strawson emphasizes that, however useful the rational reconstruction approach may (or may not) be to get an ‘idea into shape for use in the formal or empirical sciences’, if we are looking for ‘philosophical illumination of essential concepts of non-scientific discourse’ it is totally irrelevant (Strawson 1963, 504–5). As Strawson emphasized, *philosophical* problems about everyday concepts typically cannot be solved by laying down rules of use of exact and fruitful concepts in science (1963, 506). To do so, typically, is not to *solve* the problem but to change the subject – Strawson, in effect, thinks that Carnap doesn’t take traditional *philosophical* problems seriously enough, or give sufficient attention to their solution or dissolution.

But traditional philosophical problems are often symptoms that we need conceptual analysis and re-engineering. When our familiar concepts at least *seem to* get us into trouble – leading to paradoxes or other perplexities – conceptual analysis and re-engineering may be needed. Such paradoxes and puzzles have often driven work in conceptual engineering throughout the history of philosophy. The truth paradoxes, for example, have led to a great deal of work on conceptual analysis and re-engineering of ‘truth’ (most recently, in work by Scharp (2013)). Ryle (1949) was motivated to re-examine the ‘logical geography of mental concepts’ by noting the paradoxes and difficulties their use led us to. Kant (1781) was motivated to re-examine how to understand fundamental concepts such as causation and necessity by attention to the ‘antinomies’. Berkeley (1713) argued for revising our conceptual scheme to do away with ‘matter’ based on the paradoxes and difficulties the use of that concept led to.

Where there appear to be puzzles and paradoxes, we have two options: We might, through a kind of reverse engineering, come to better understand the functions the concept serves and the rules it follows, and thereby untangle it in a way that enables us to understand why the original paradoxes or perplexities arose, and yet that relieves the

---

<sup>4</sup>Similarity to the explicandum, and simplicity are also desiderata Carnap mentions (1950, 7).

puzzlement.<sup>5</sup> Or, we might (if such reverse engineering fails to disentangle the concept but instead reveals deep problems) aim to re-engineer the concept, or even to do away with it.

But the need for conceptual research and development does not arise only within philosophy. It also arises when concepts fail to fulfill their functions, whether these are extant functions or new or altered functions. Carnap's own task of explication (1950) falls under the latter heading, for in explication terms acquire the function of serving in a precise set of well-connected concepts for a scientific theory, which will be useful in formulating universal laws. Meeting these higher standards requires us to explicate certain of our familiar concepts – such as 'probability', 'confirmation', 'heat' or 'fish'. At times, empirical discoveries may demonstrate the need for conceptual reform. For empirical discoveries may show that the concept *never could do what we asked it to, or that given certain empirical changes in the world, it can no longer do what we (now) need it to*. For example, as Carnap points out, 'fish', as used for 'sea creature' generally, turned out to be inept for forming biological generalizations, and so we did better to replace it with a narrower version (Carnap 1950, 6). As Appiah (1992) has made beautifully clear, 'race', as used by Thomas Jefferson and Matthew Arnold (and many of their contemporaries and followers) turned out to name nothing like a natural kind that could support generalizations and inferences. An example of the second sort (discussed by Gert, Culver, and Clouser 2006) is the concept of *death*, which we need to tell us when we may cease medical interventions, commence inheritance and funeral proceedings, etc. But given the expense and wide use of end-of-life medical procedures (as well as possibilities of artificial respiration and organ transplantation), we now need the concept to be far more precise in order to serve these functions well.

Another type of situation in which conceptual reform is called for is not when a term fails to fulfill a function, but when it fulfills a function we do not endorse, or do not think should be fulfilled. (This may call for eliminating old concepts rather than introducing new ones). Such cases may arise when a concept mediates between descriptive application conditions and exit rules that endorse problematic norms of regard and treatment, as, e.g. racial, disability, and other pejoratives do.<sup>6</sup>

---

<sup>5</sup>Sometimes this untangling may be done by way of work in linguistics as, for example, work in Systemic Functional Linguistics may show that certain terms enter language as grammatical metaphors, and may reveal that certain paradoxes result from harmless shifts in grammatical category. See Halliday (2009, 131). See also discussion in my 'How should we think about linguistic function?' (forthcoming).

<sup>6</sup>Braddon-Mitchell (2020) similarly suggests that we think of concepts as having not only causal inputs, but also direct connections to behaviors, negative affect, and so on.

Other needs for conceptual research and development arise when we aim to serve some new, perhaps scientific, technological, or social, purpose.<sup>7</sup> (For the first, think of Hans Asperger introducing a term for what he took to be a characteristic set of behavioral and ability traits – presumably in order to engage better in psychological explanation, prediction, and intervention. For the second, think of terms such as ‘website’, ‘bug’, ‘server’, ‘app’, etc. that have been introduced to keep up with and enable communication about our new technology. For the third, think of the introduction of terms like ‘homophobia’, ‘genocidal rape’, and ‘human rights’.)

That no doubt is incomplete, but we can at least give a preliminary list of some cases in which conceptual research and development are needed. Note that each leads to its own criteria of success for the engineering project:

- (1) Internal problems: Concepts apparently leading to paradoxes, puzzles, perplexities. Criterion of success: resolution or dissolution of the relevant paradoxes, puzzles, and perplexities. (This has been central to promoting conceptual change in the history of philosophy).<sup>8</sup>
- (2) Failure to serve an extant (or evolving) function: e.g. Concepts needing Carnapian explication for use in the sciences. Criteria of success: success at fulfilling that function (perhaps in the evolving circumstances, or given new empirical knowledge).
- (3) Serving a problematic function: Concepts (such as pejoratives) leading to problematic inferences and behaviors. Criteria of success: interrupting these patterns of inference and behavior.
- (4) Serving a new function(s): (scientific, technological, social/moral ...) Criteria of success: fulfilling the new function(s).

It is useful to notice both the unity and variety in these cases, which corresponds to the difference in work in conceptual engineering that has been done in the history of philosophy, in Carnapian explication, and in more recent work in re-engineering social concepts, each of which may have its own goals and criteria of success. In all of these cases, however, allowing for the possibility of conceptual engineering

---

<sup>7</sup>Philosophical purposes might also be added here. Consider, for example, Edmund Husserl’s introduction of a range of new terminology (‘noesis’, ‘noema’, ‘hyle’, ...) to serve work in phenomenology.

<sup>8</sup>We might add to these cases in which old concepts are thought to be too mired in philosophical controversies and problems – consider Heidegger’s changes here, as he resists using old philosophical concepts such as person, behavior, consciousness, and matter, and instead introduces his own range of terminology and concepts to speak of *Dasein* and its *comportment* in the *World*.

requires allowing that concepts are things that may be *created*, that can *change*, and that can serve *certain functions*.

## 2. Should we think about concepts or words?

How is conceptual engineering possible? A critic might say: perhaps you can work out what improvements to our conceptual scheme you would *want*, but you can't really *change* concepts – you can only replace them, in a way that threatens to change the subject. Herman Cappelen attributes this view to Mark Richard, expressing it as follows:

You can't improve on a concept by changing its intension and extension because that very idea is incoherent. Concepts have their intensions and extensions essentially. So a change in intension or extension always involves abandoning a concept, and can never be an improvement of the old concept. Insofar as conceptual engineering rests on the idea of changing a particular concept's intension and so extension, the entire project is incoherent. (Cappelen 2018, 104)

If we are Platonists about concepts, thinking of them intensionally, as functions from worlds to extensions, it will indeed be hard or impossible to see how a concept could be *changed*. More generally, as Sarah Sawyer puts it, 'so long as we assume that concepts have their intensions and extensions essentially, it will not be possible for them to be revised' (Sawyer, 2020, 391).<sup>9</sup>

Others argue that even if we put that worry aside (as Cappelen does, by rejecting the need to talk of concepts at all), and even if we could in principle simply aim to find 'better ways to talk about topics', there is no way to *implement it* – to engage in conceptual activism that would involve genuinely changing the linguistic or conceptual scheme we use. If that's the case, then conceptual engineering might seem like a pointless intellectual exercise. This is a major issue for Cappelen. Given the externalist view of language he presupposes (2018, 58), he allows that meaning change is possible (this amounts to reference change), but he maintains that it is fundamentally 'inscrutable' and 'out of our control' (2018, 72–3). For on this view, meaning supervenes on a base structure involving

---

<sup>9</sup>For this reason, she recommends that we separate out the idea of *concepts* from that of *linguistic meanings*, and treat the primary focus of conceptual engineering to be to change *linguistic meanings*. Concepts, she suggests, we should think of as individuated by relations to 'objective properties in the world', which enables them to fix a stable subject matter. By contrast, she suggests that we understand linguistic meanings as supervening on use and capable of change over time – including change for the better, as we get a better view of the subject matter and engage in conceptual engineering (2020, 387).

introduction of the expressions, communicative change, sources of information, speakers' use and dispositions, etc. – but 'we don't have this information and never will' (2018, 73). And changing the intension and extension of a term would require understanding 'the mechanisms of reference change', but 'These mechanisms are ... not known to any of us and might in effect be unknowable' (2018, 74). As a result, he concludes, it is 'an illusion to think that we can be in a position to effectively predict and implement changes' (2018, 74).<sup>10</sup>

But both the Platonist and externalist assumptions, which lead to such dire problems, are relics of other ways of thinking that don't provide helpful ways of thinking about concepts or meaning, for the purposes of conceptual engineering. We can, in bootstrapping fashion, ask instead how we *should* think of language or concepts, for such purposes. Thinking of concepts in a frozen Platonist fashion may be useful for certain purposes – say, in logic, as it enables us to give an (apparently timeless) way to identify meaning, so that we can assess whether a certain argument or inference form is valid, identify equivocations, etc. But this timeless talk of concepts is highly abstracted from the language we use and live in (intentionally so: it's the messiness of ordinary language, with its homonyms, meaning shifts, vagueness, etc., that provoked the need to shift to talk of concepts abstracted from those contingencies, in order to assess the validity of arguments). Thinking of concepts as eternal, timeless, changeless, with identities fixed by their intension and/or extension may be useful for logic, but is not helpful in thinking about conceptual engineering. For that requires that we be able to both *change* old concepts and *create* new ones.

The externalist view of meanings, as Cappelen expresses it, relies on a picture where there are 'meaning facts' that may be deeply hidden, but perhaps determined (in ways that follow some complex and mysterious algorithm) by facts about worldly information sources, word introduction events, use patterns, and communicative chains (2018, 74), 'to figure out the current intension of a term, you would need information about the past, about introductory events, and communicative chains. It is

---

<sup>10</sup>This claim has provoked a number of responses. Riggs (2019) argues that externalist meaning can't be what conceptual engineers care about changing anyway. Koch (2021) argues that externalism does allow for a kind of collective long-range control of meaning change. This sort of control, he argues, is sufficient to make projects in conceptual engineering worthwhile. Pinder (2021) argues that conceptual engineering can work largely at the level of speaker-meaning. Jorem (2020) argues that feasible options for implementation remain, even in the face of externalist concerns. For even if it's not feasible to change the *standing* meaning of a term, implementation may not only take the form of changed speaker meaning, but also of meaning modulation or introduction of a new, constructed language.

indisputable that we don't have this information and never will' (2018, 74). This metaphysicalized way of thinking about what 'meaning facts are', is itself implicitly a *proposal* for a new way to think about 'meaning' – one quite different from the everyday sense in which we think we know what some words mean ('desk') and don't know what other words mean ('muniment'); where we think some people don't know the meanings of certain terms (a four year old who requests toast, and then complains that it is toasted), but may be properly instructed and corrected; where we think linguists may know the meanings of words in an explicit and theoretical way ('knowing that') while competent speakers may simply know them in a mastery sense of 'knowing how', and when we think that historical linguists may track changes in meaning over time. The externalist is giving us an implicit *proposal* for how to think about meaning, but this proposal is unnecessarily mystifying and unhelpful.

In short, the Platonist view of concepts and the extreme externalist view of linguistic meanings that led to the troubles were both proposals of their own, but they are ones we should reject if we aim to understand and engage in conceptual engineering. Both a Platonist approach to concepts, and an externalist approach to linguistic meaning, present barriers to thinking about the kind of change that conceptual engineers aim to make.<sup>11</sup> But these are not the only ways to think about what we can aim to change in conceptual engineering. In bootstrapping fashion, we can ask how we *should* think about linguistic or conceptual meaning, for our purposes in engaging in conceptual engineering.<sup>12</sup>

A first question is whether we should think and speak in terms of *language* or *concepts*. As Manuel Gustavo Isaac argues (2021, [forthcoming](#)), there are other ways of thinking about concepts beyond the Platonist philosophical view. We can, instead, make use of a psychological understanding of concepts, which treats concepts as 'bodies of information about some category of referent' (2021, 10), which are activated to be used in categorization, action-planning, induction, and deduction (2021, 11). Such a view of concepts, Isaac argues, has better prospects for allowing concepts to change over time, and enables us to make use of empirical work in psychology in understanding concepts.

---

<sup>11</sup>See also Isaac (2021) and ([forthcoming](#)) for discussion of the difficulties conceptual engineers face if they adopt a Platonist (or, as he calls it 'semantic' or 'philosophical') understanding of concepts. Isaac argues that we should instead adopt a psychological understanding of concepts.

<sup>12</sup>As Riggs (2019) puts it, it is always open to conceptual engineers to *re-engineer* a better sense of 'meaning' that captures what we care about changing – and what it *is* possible to change.

While I agree that such an understanding of concepts may be a helpful improvement over the Platonist approach, and may be useful for various projects in conceptual engineering, here I am going to work instead towards developing an approach that focuses on *language* rather than *concepts*. Why?

Words are what we live with, communicate with; talk of ‘concepts’ is often introduced as a way of abstracting meaning from words – enabling us to talk about what ‘woman’, ‘Frau’ and ‘mujer’ have in common, why one is an apt translation of the other, and so forth. When we talk of concepts (the concept of woman, the concept of truth ...), we do so using words.

I have come at the problem of conceptual engineering as a promising way of recharacterizing what we *have done* and *can do* in metaphysics (see my Thomasson, 2017). As a result, I am interested not just in the popular and important projects of working out how we should engineer concepts such as our race or gender concepts. I am also interested in questions about how we can come to understand and better handle metaphysical debates about properties, numbers, modality, morality, and other perennially problematic philosophical concepts.

The limitations of the psychological view of concepts, at least as it is generally employed, arise from the fact that it treats concepts as ‘bodies of information about some category of referents’ (Isaac forthcoming, 14). For some of the most interesting and promising treatments of modality, morality, numbers, etc. are treatments that *do not* involve treating the relevant terms or concepts ‘representationally’, as way of carrying bodies of information *about* certain worldly referents.<sup>13</sup> If we simply think of concepts as mental representations of certain referents that carry bodies of information about those referents, then we leave out of view a range of interesting and promising neo-pragmatist approaches to our moral, modal, or number concepts. Moreover, some of these concepts, in their philosophically interesting (problematic) form, rely on being expressed in a certain *grammatical* form or category. For example, Quine notoriously claimed to have no problem with assertions that houses and sunsets *are red*, but rejected the ‘ontological commitments’ that seemed to come with talk of *the property of redness*. Yet, if we just speak of *the concept of redness*, we elide this distinction, and may also overlook important functional

---

<sup>13</sup>For general discussion of the problems of the so-called ‘Representationalist’ assumption, see Price (2011). For particular problems with this approach to morality, see Blackburn (1993); for modality, see my (2020b).

differences between adjectives and their nominalizations. (Much the same goes for determiner uses of number terms (e.g. in ‘there are two geese’) versus nominal uses of number terms (e.g. in ‘two is an even number’)).<sup>14</sup>

Working in the first instance with language enables us to preserve these differences, and to keep in view potential neo-pragmatist (and non-representational) options for addressing a range of topics. It also enables us to make use of empirical work in linguistics, in understanding the notions of *function* and of *change* that conceptual engineering relies on. For, as I argue elsewhere (forthcoming), work in Systemic Functional Linguistics can provide helpful inroads to understanding the functions of various linguistic structures (including property nominalizations and modal discourse, to choose just two). We can also draw on work in historical linguistics to better understand how language *does change* over time, and what the natural processes of linguistic change are, so that we can better make use of these in determining what sorts of change in conceptual engineering are more and less likely to succeed, and how to make use of these natural processes of change in implementing the recommendations of conceptual engineering (see Koslow forthcoming).

Some might share the worry raised by Isaac (2021, 6) that if we see the target of conceptual engineering as language rather than concepts, that makes ‘conceptual engineering’ a misleading misnomer – and one that undermines its own goals if it leads people astray (2021, 6). There is a point to this, but I do not see it as a decisive worry. For conceptual engineering has long been presented as a way to understand, evaluate, and modify how we *think and talk*. Both of those projects remain – particularly since, at least for a great many of the scientific, social, and philosophical concepts of interest, we tend to think *in* and *using* language. And when we take interest in conceptually engineering philosophically interesting terms such as ‘person’, ‘freedom’, ‘cause’, etc., we are not interested in their phonetic or typographic structure,<sup>15</sup> but rather (as we might naturally put it) in the *concept they express*. In short, in thinking about words and how to engineer them, we *are* thinking about concepts, so it is not misleading. We can even aim to shift up from any conclusions we reach regarding particular terms of English, German, Spanish, Chinese or Tagalog, to ask about to what extent

---

<sup>14</sup>I don’t mean to suggest that there is no way one could *develop* a psychological view of concepts that could handle these phenomena. I simply don’t know of one yet; and linguistics provides helpful and clear inroads.

<sup>15</sup>Unless this impacts implementation issues – see below.

those conclusions carry over to good translations of those terms in other languages (acknowledging that we can abstract away from phonetic or typographical issues, which aren't of primary interest).<sup>16</sup> When we reach conclusions for how or whether we should re-engineer 'woman', that may tell us something about what we should do with 'Frau' or 'mujer', too.<sup>17</sup> In short, talk about concepts, in this way, can be arrived at by starting from linguistic analysis, and then seeing where it is legitimate to abstract from accidental features of the linguistic formulation in question to apply the same analysis and results to appropriate translations.

So, for purposes of this paper, at any rate, I will speak first and foremost about understanding, engineering, and changing, *words*, rather than *concepts*. I'll also suggest a particular view of language that fits very well with conceptual engineering – and is fruitful, natural, and coherent with linguistics as an empirical science concerned with language. It is not entirely new, and I am not going to argue that we *have to* take this view over the alternatives (and I certainly won't argue that it's the picture that matches the metaphysical realities). But I will try to show how very natural, fruitful, and unifying this view is, as applied to issues in conceptual engineering. While Cappelen takes his book as showing what view of conceptual engineering can go with an externalist view of language, take this as suggesting a view that would go with a rather different – artificialist and inferentialist – view of language.<sup>18</sup>

### 3. How should we think about words?

It is natural to think of words as human creations: as things that come into existence and can change. This is entirely cohesive with work in historical linguistics, particularly etymology, which aims to trace out 'how new words enter a language' (Durkin 2009, 43) and how they change over time – including not just changes in form or sound, but changes in meaning. As Philip Durkin (the chief etymologist for the *Oxford English Dictionary*) puts it,

The meaning of a word [such as 'friar'] can also be traced historically. We can see how the meaning broadened in Latin and French, but how the English

<sup>16</sup>Except perhaps with regard to implementation – as new words that are hard to pronounce for speakers of a given language may be less likely to catch on, and new words that sound like an old word may create confusion (see Koslow (forthcoming)).

<sup>17</sup>Though, as one anonymous referee aptly suggested, we can't just *assume* that the relevant suggestions will be equally apt across languages, given their different sets of inferential connections and the varying social contexts in which they are used.

<sup>18</sup>Cappelen invites the development of other accounts that might assume, for example, an inferentialist framework (2018, 187–8).

borrowing showed only a very narrow component of the donor word's meaning. (Durkin 2009, 6–7)

Or, in the simpler terms of linguist John McWhorter, 'Changes in meaning are as natural to words as changes of pitch are to music' (2016, 85). Once we think in terms of the change and creation of *words*, we can make use of the knowledge in linguistics to help us understand what conceptual (or linguistic) engineering would be, and how we might be able to do it.

So how should we think of words? This is also a matter of some controversy (one which I don't intend to enter into (much) here). But again, if we are to think of words in a way that is consistent with historical linguistics, and a way that might be useful to understanding conceptual engineering, we must think of words as entities with a *history*, as things which can be *created*, and which can *change*. Nurbay Irmak argues that preserving the idea that words are *created* and *can change* are, 'two indispensable desiderata ... for any successful ontology of words' (2019). We can best fulfill these, he argues, by treating words as a kind of 'abstract artifact' (in my (1999) sense): as 'artifacts that are created to fulfill various kinds of purposes' and as abstract in the sense that they 'are not located in space', though they may have a temporal beginning and ending. Irmak's work follows up on the work of Kaplan (1990), who also insists, against Platonists<sup>19</sup> about words, that words can be created and undergo change. As Kaplan puts it, 'there is no metaphysically fixed form in either speech or spelling ... There are spelling variations, there are pronunciation variations, there are all kinds of variations that take place over time' (1990, 100). On Kaplan's view, words (including names) are 'objects of the created realm, created by language makers' (1990, 117).

Once we think of words in this way, it is clear that words are things that not only *can* change but that *do* change. Moreover, while it is not entirely predictable what meaning changes a given word will undergo over time, work in historical linguistics has given us some known principles regarding how new words can be introduced, and how old words come to change in their meanings. New words may be introduced via a variety of patterns involving affixation, compounding, borrowing, conversion of a word of one grammatical class to one of another: ('knife'/'to knife') (Durkin 2009, 114), blending, onomatopoeia, etc. (Durkin 2009, Chapter 4). There are similarly typical patterns to ways in which word meanings may change over time, including broadenings or narrowings, step-wise associations ('merry' – see McWhorter 2016, 72), 'metaphorical variations'

---

<sup>19</sup>Irmak (2019) identifies Katz and Wetzel as Platonists about words.

(McWhorter 2016, 73), and what McWhorter calls shifts from ‘impersonal’ to ‘personal’ use: from a denotative use, to use as ‘tools for reinforcing and marking various shades of interpersonal understanding’ (McWhorter 2016, 85).<sup>20</sup>

Allison Koslow (forthcoming) makes the case that understanding the natural process of meaning-change in words can also help us, if what we aim to do is to change the meaning of a term in common use (rather than a technical term or a term to be used by a like-minded cohort). For it can help us distinguish which sorts of meaning change can be induced, and which sorts of proposed change would be highly unlikely to succeed, given the full range of our communicative desires and the features of the extant language (what homonyms or similar-sounding words there are, if it calls to mind anything offensive or distracting ...). Moreover, once we understand the natural processes of meaning change, we may be able to exploit that knowledge to increase our chances of success at artificially introducing changes of meaning as we aim to implement proposals in conceptual engineering.<sup>21</sup>

Thinking of words in this way, as historical abstract artifacts, also brings clear advantages for thinking about conceptual engineering. For if we think of words as (abstract) artifacts,<sup>22</sup> we can make use of what else we know about artifacts. Two features of artifacts are especially relevant here. First, (as is widely agreed (see, e.g. Baker 2008; Hilpinen 2018)) artifacts characteristically *serve some purpose or function* (where that function may be intentionally imposed or may evolve over time, and may change). Second, as I have argued elsewhere (2014), instances of public artifact types are subject to certain *norms of use*. Part of what it is for something *to be a fork* rather than a comb or a backscratcher, is for it *to be appropriately used* in some ways (and for some purposes and by some sorts of individuals) and not others.

The first feature, the idea that words (like most other artifacts) serve a *function*, is useful for our project, since (as I argue (2020a, 2020b)) the appeal to functions gives us a non-arbitrary way of assessing various proposals in conceptual ethics and conceptual engineering – according to the success they may or may not have in serving the relevant function (s). Engineering words (whether creatively or reparatively) may, like

---

<sup>20</sup>That is, they gradually convert into what linguists call ‘modal pragmatic markers’ such as ‘actually’, ‘rather’ ... ‘Throughout any language, words of all kinds are always going personal to a certain extent: the subjective exerts a gravity-style pull on words’ meanings’ (McWhorter 2016, 16).

<sup>21</sup>See Koslow (forthcoming) for details of this work.

<sup>22</sup>I introduced the idea that various cultural creations can be understood as ‘abstract artifacts’ in my work on fiction and literature (e.g. in my 1999).

creation or repair of other artifacts, be assessed for its success by reference to how well these efforts result in an artifact that can serve the relevant functions.<sup>23</sup>

But in saying that words, or other parts or aspects of language, serve functions, we must be careful to clarify what is not being said. First, we should not presuppose that each word has a single function, but rather allow that *many* functions may be served. We also should not presuppose that the right unit of analysis is (always) be to ask about the functions of individual *words*; we may sometimes do better to ask after the functions of categories of terms, of grammatical formulations, or of other linguistic structures – or even to begin by considering the functions served by a linguistic system as a whole, and then moving back to ask what the relevant parts and aspects of that system contribute to those functions. Finally, we should not presuppose that these functions are beneficial to *everyone* or overall: like weapons and pesticides, words may have functions even if they are not beneficial to all concerned or *all things considered*.

In saying that parts or aspects of language serve certain functions, we also must be clear what sense of ‘function’ is relevant here (on this see also my 2020a, 2020b). First, it needn’t be understood as an *intended* or ‘design’ function, since much of language is not intentionally designed by anyone, and at any rate the functions served may drift from any functions that are intentionally envisioned. Second, *pace* Cappelen (2018, 180) we are *not* restricted to thinking of function disquotationally (that ‘the function of “salmon” is to pick out the salmon’, and no more can be said (see my 2020a, 2020b)). Nor need the function of a part or aspect of language be identified with the *uses* to which it is put.

From the point of view of conceptual engineering, a better way to think of the question is: how *should* we think about functions here? One prominent option (building on well-known work by Millikan (1984)) is to begin by asking what the relevant term or linguistic structure has done or does for us (or some of us) that we couldn’t do as effectively without it, and which that has led to its perpetuation, or by asking what we couldn’t do (so well) if we lacked the relevant linguistic structures (or what speakers who haven’t mastered them can’t do as well). We can also look to work in linguistics – particularly in systemic functional linguistics – for empirical work in identifying functions (in this sense) served by various kinds of linguistic structure (see, e.g. Halliday (2009)). While it is important

---

<sup>23</sup>Of course, we may need other standards to assess such efforts for their moral standing or all-things-considered value.

to add the above clarifications to avoid misunderstandings or easy dismissals at the outset, detailed work on how we should understand linguistic function will have to wait for another occasion (see my forthcoming).

The second feature, the idea that words (like other public artifacts) have certain *norms of use*, is also quite helpful here. For (while it doesn't *entail* it) it fits quite naturally with an inferentialist account of meaning, on which the meaning of a word may be understood in terms of the (broad) norms governing its use. (One reason it fits naturally is because inferentialist views begin by looking to the function these terms serve in our lives, and move from there to specifying 'content'. Thinking of words artifactually fosters looking first to what we do with parts or aspects of language, and to what they do for us.) And that, in turn, can help us understand *what* meaning-change for a word would consist in, *why* it may be so important, and *how* it can be done. Lynne Tirrell makes roughly the same claim – that modes of discourse are 'specified in two ways: structurally, by their inferential networks, and functionally by their goals and practices' (1999, 56).

The sort of inferentialism I have in mind is roughly the 'broad' inferentialism developed by Brandom (2007), according to which the content of concepts (or, here, words) is determined by their broad inferential role, including their role in:

- Logical/formal inferences (where correctness is settled just by logical form)
- Material inferences (which articulate the contents of non-logical concepts). These include inferences from 'x is green' to 'x is colored'; from 'A is East of B' to 'B is West of A', and so on (Brandom 2007, 657–8). They also include material incompatibilities (e.g. that the applicability of 'triangular' precludes the applicability of 'square' to the same thing at the same time.)
- Inferences between 'non-inferential circumstances and consequences of (appropriate) application': That entitle us to introduce or apply terms from perception/world, or that articulate exit rules to forms of action, attitude, or behavior (Brandom 2007, 658)

Now this can help us understand what sorts of change we might hope to get from conceptual engineering. In changing words in this way, we are not looking to change the pronunciation or spelling, but rather to change the meaning – *and that is (on this conception) to change the total inferential profile of the term, changing its norms of use*. Re-

engineering a term, as Tirrell puts it (for the special case of a targeted group reclaiming pejoratives), requires: ‘a reorganization of the inferential structure associated with the term’ (1999, 60). And that is no easy matter, given the holism that governs our inferential structures.<sup>24</sup>

This kind of change in inferential role may be needed so that a term can better fulfill its function. So, for example, if we agree with Carnap that ‘fish’ was re-engineered, we can see that re-engineering the term in a way that no longer permits an entry rule simply from ‘swimming creature that lives in the water’ to ‘fish’; given empirical investigations and the new entry rules (which require a cold-blooded aquatic vertebrate with gills, etc.), we also get new entitlements to make inferences about how the creature will reproduce, respire, etc. These inferential entitlements were unavailable given the old use of the term, and so better enable it to fulfill its function as a scientific term, by enabling it to figure in a wider variety of laws.

The inferential role approach can also enable us to see why conceptual change (elimination or revision) is sometimes crucial to avoid fulfilling undesirable functions. As Tirrell makes beautifully clear, an inferential role approach ‘offers a powerful conceptual framework for analyzing the social problems reflected in and the linguistic problems created by derogatory terms’ (1999, 46), and thereby makes evident why re-engineering terms may be so *socially* important. For certain terms, including pejoratives, carry among their exit rules certain norms of how we are to regard, treat, etc. those who are so named, and how they are to behave, regard themselves, etc. ‘The assertional commitments of these terms tell members of the target group how they *ought* to be, under the guise of describing how they *are*’ (1999, 53). Once that is noted, we can either try to ‘reclaim’ the term (at least for use within a certain group, in which we aim to reject the inference to the relevant norms of regard and treatment), or we can reject the use of the term entirely (see Tirrell 1999 for an excellent discussion of the options here).

#### 4. How can we implement the changes we need?

I have so far argued that, in discussing conceptual engineering, we should work first and foremost with a conception of *words*, considering these as

---

<sup>24</sup>This also connects up with work by Shields (2018) who sees speakers engaged in ‘deep disagreements’ as implicitly *stipulating* how certain key terms are to be used, where what they are doing in making these stipulations is giving both speaker and audience entitlement to make certain inferential connections.

abstract artifacts capable of being introduced and changed (in their meaning as well as in other ways). This conception, I have argued so far, not only coheres well with historical linguistics, but also enables us to see words as entities with a *function* and with *norms for their proper employment*. The former enables us to adopt a method for conceptual engineering, and for evaluating concepts, in a manner that does not rely on matches with mysterious ‘metaphysical facts’, and yet nonetheless gives an objective way of evaluating various proposed changes with respect to their ability to fulfill the relevant function(s). The latter – thinking of words as having norms for their use – fits well with a broad inferentialist conception of meaning, which enables us to get a clear picture of what meaning change is and why it can be so important – whether for our scientific or our social and moral projects.

There is also another important consequence of thinking of words in this way. For thinking of word meanings as a matter of *social norms regarding how the relevant terms are to be used* (including entry rules that may appeal to the world, and exit rules that may license linguistic or non-linguistic behavior) enables us to understand how we might create linguistic change, by turning to the rich and growing literature on how social norms may be changed. If we can see how social norms may be changed, we might not be stuck with Cappelen’s pessimistic perspective, that conceptual engineering is fundamentally inscrutable and out of our control.

To understand social norm change, we can turn to Bicchieri and Mercier (2014), who distinguish social norms from mere conventions: Social norms are behavioral rules supported by a combination of empirical expectations (that others will behave this way) and normative expectations (that others believe I ought to behave this way, and may punish me if I transgress) (2014, 61). Mere conventions, by contrast, ‘are supported by empirical expectations of compliance, and a preference to follow the convention provided one expects most ... others to comply with it’, but are not supported by *normative* expectations.

Linguistic rules, I would suggest, should count as norms in Bicchieri’s and Mercier’s sense. For they involve not only expectations that others will use terms in the relevant way, but also normative expectations that others believe I ought to speak this way and will punish me if I transgress. The sanctions for violation of (local or more global, informal or formal) linguistic rules may vary. As Sterken (2020) points out, there is always the risk of misunderstanding and of failures of communication when speakers engage in linguistic interventions aimed at changing the meaning of an

extant term. Even where our interlocutors figure out our speaker meaning nonetheless, there are sanctions for violating linguistic norms: whether undergoing embarrassing public corrections and shaming, downgrading our work or rejecting our articles or giving less credence to our public statements, or ostracism and social out-grouping – for, as Herbert and Kukla make clear ‘negotiating and constituting community boundaries are among the important pragmatic functions of language’ (2017, 4).

How, in general, can we work to change social norms? As Bicchieri and Mercier put it

We can ... devise specific interventions to effect norm change by acting upon the expectations that support the norm we wish to eradicate or, when it is a new norm we want to establish, work at creating new expectations, and focus on those factors that will bridge expectations and behavior. (2014, 62)

When it is norms, not merely conventions, that we aim to change, we must do so by changing the relevant normative expectations – to gain assurance that we won’t suffer negative consequences if we change. Bicchieri and Mercier identify some internal and external factors that tend to increase the chance of successful norm change:

Regarding the norm itself:

- The change must not be too distant from existing social norms.
- There should be a viable alternative to the norm that is to be changed. (2014, 71)

I will come back shortly to discuss changes in linguistic norms that are ‘too distant’ from existing norms to catch on.<sup>25</sup>

We can also add one element to this. As Koslow (*forthcoming*) argues, once we understand, from work on historical linguistics, how linguistic norms (meanings) naturally change, we can also rely on and exploit those processes to help certain proposed changes (in ‘artificially induced’ conceptual engineering) catch on. An introduced term that follows one of the standard processes of language introduction is much more likely to be understood and to catch on than is an idiosyncratic neologism.

---

<sup>25</sup>As Sally Haslanger aptly reminded me, we also should take care to acknowledge the role of other important factors (beyond empirical and normative expectations) in norm change, including the surrounding material conditions in the extant social structures and their role in maintaining norms. What material conditions might play a role in norm change regarding language? Such factors such as access to education, communication technology, to (consumption and control of) printed media, etc. might play a central role. I will leave this for future investigation.

Think of the process of introducing the term ‘homophobic’: a term introduced in 1969 by psychologist George Weinberg to refer to a heterosexual man’s fear that others might think he was gay, and adopted in 1971 by Kenneth Smith with the related meaning of a personality profile that involves a psychological aversion to homosexuality. The term has a nicely transparent meaning, involving a blend of ‘homosexual’ and the Greek root ‘phobos’ for fear, following standard word-formation processes that make it easy for others to understand, and easy for it to catch on.<sup>26</sup> It also instituted other inferential connections: ‘phobia’ suggesting that the aversion is irrational, and the medical sound of the term suggesting that the aversion is a medical problem – as indeed Weinberg describes it. The introduction of this term, and its success at ‘catching on’, as well as the American Psychiatric Association in 1963 ceasing to classify homosexuality as a mental disorder, helped reorient our conceptual scheme in ways that saw *aversion to homosexuality* (rather than homosexuality itself) as the problem that needs to be ‘fixed’. This is a fine example of a case of conceptual engineering exploiting natural processes of word formation to improve the chances of success, and instituting new inferential connections that helped change not only linguistic but non-linguistic norms and behaviors, in service of a social function.

But for social norms to change, given the dependence of norms on expectations, there also needs to be a change in expectations about the behavior of others, and about what we think others think we should do (Bicchieri and Mercier 2014, 74). How can we achieve those sorts of change? Bicchieri and Mercier mention several methods:

- (1) Public argumentation or deliberation
- (2) Government injunctions
- (3) Trendsetters

The primary method of change they identify is a process of public and collective argumentation or deliberation, ‘For most beliefs the most effective way to change them, and thus eventually change the practices they support, is through argumentation’ (2014, 73). But it must be an argument that rests on beliefs and values in the listener that are (explicitly or implicitly) shown to be inconsistent with the norm to be changed.

---

<sup>26</sup>As Durkin writes ‘the lexicon of any language will be extended by speakers in an ad hoc way, as new words are formed by productive word-forming processes such as derivation or compounding. These will normally be understood very easily by other users of the language from their transparent composition and from clues in the context of the utterance which help to explain the meaning’ (2009, 58).

Sometimes (so that the listener doesn't feel manipulated) the argument must be left implicit – leaving people to discover the inconsistency on their own (2014, 70–71). That 'public argumentation' appears first on the list should give some comfort to philosophers who aim to change the norms for using our terms – for argumentation is always what we have done best. Consider, for example, the ways hard determinists have aimed to show the inconsistency between thinking of our actions as following laws of causal determinism, and thinking of individuals as free and responsible in a sense that entitles us to retributively punish them. Appiah (1992) shows the inconsistency in our culture's concept of race. These kinds of arguments, if made publicly and known to be public in a given community, can give powerful reasons for conceptual change.

Bicchieri and Mercier also discuss the role of government injunctions in leading to changes in social norms. They can change the cost/benefit structure of following certain norms (say, if female genital cutting is punished, this may change the incentive structure and help counter-balance the social costs of not cutting) – though there are limits to what can be enacted in that way, as the law must be seen as legitimate, must be enforced, and the change must not be too far from the standing norms (2014, 64).

Where change in linguistic norms is concerned, laws and official edicts from governing bodies may also have effect, as we require legal definitions of 'property', 'minor', 'person', 'voluntary', etc., which go on to play a role in what legal consequences follow from application of a given term.<sup>27</sup> But here, as with changes in other norms, there are limits to what may be accomplished 'by decree'. Cappelen defends his claim that we lack control of conceptual change by citing the U.S. Supreme Court's definition of corporations as persons: 'even an institution as powerful as the US legal system (backed by police power) has failed to change the meaning of "person"' (2018, 77). But the work by Bicchieri and Mercier can show us why this attempt at change failed: this was too great a shift to the meaning of our common word 'person' to create a change in our standard norms of usage. Instead, it just introduced a technical legal definition for 'person' as used in legal contexts, without changing the informal social norms concerning the use of the shared English term. But this does not tell us that the decrees of bodies with

---

<sup>27</sup>These would aim to bring about what Sterken (2020) refers to as 'top down' linguistic change, rather than 'bottom up', or grass roots, linguistic change (2020, 419).

institutional authority (courts of law, the psychiatric Diagnostic and Statistical Manual of Mental Disorders (DSM), other professional bodies, sporting organizations ...) can never be (intentionally) effective in bringing about meaning change, or in introducing terms with new meanings (where there are not old norms to conflict with). It only tells us that in this case, it was poorly done, *if* the effort was to change the meaning of the shared term 'person', rather than just to change the legal standing of corporations (which did succeed).

The more general point is that sometimes authorities are granted institutional power to introduce new terms, re-define old terms, etc. to be used within the institutional realm in which they are granted authority (legal, medical, psychological, governing bodies of sports ...). Making institutionalized conceptual changes in this way is a somewhat easier matter than changing the use of shared terms of natural language. But institutional changes may succeed in yielding the relevant institutional entitlements (to tax benefits, insurance coverage, eligibility for play), and yet it may remain an open question whether they succeed or fail at 'catching on' more broadly, depending on the extent to which the relevant terms are commonly used for purposes outside the realm of the given institution. Consider the DSM definition of 'depression', which sets up an institutional definition for 'depressed' that fails to coincide with, and remains a rival to, everyday uses of 'depressed' (we note the distinction by calling the first 'clinical depression'). And compare this to the DSM introduction of 'bipolar disorder', which had no such common language rival.

But, as Bicchieri and Mercier again make clear, it is not only official institutional authority that is relevant to enacting changes in social norms: 'It has been shown ... that it may take a small number of "trendsetters" who question the standing norm and start behaving differently to effect a major change' (2014, 63). Where changes in linguistic norms are concerned, there are various levels and sorts of informal authority in a variety of speech communities and sub-communities: from the norms governing how professional philosophers speak, to those governing how the hip kids in a particular high school talk. As Herbert and Kukla emphasize, 'Speaker authority may come from an official institutional role ... or it may be more informal. It may come from one's position within a social hierarchy and community or from something more neutral ...' (2017, 16). Matthew Shields emphasizes this, too, noting that

the stipulations that get uptake – especially stipulations of an understanding of a concept – will likely track familiar informal sources of authority that are linked

to a speaker's social and material position. The ability to claim this kind of informal authority depends heavily on one's position in a network of power hierarchies and relationships. (2018, 20)

And this kind of informal linguistic change need not involve merely explicit stipulation. Sterken identifies various sorts of linguistic disruptions, in which a speaker intentionally uses a term with some meaning other than its current public meaning, or intentionally (mis-)interprets others in this way, in order to provoke metalinguistic reflection and change (2020, 421). These may include: '(i) neologisms, protologisms, and semantic introductions; (ii) the reappropriation of slurs and insults; (iii) transgressive uses of definitional or normative generics; (iv) semantic elimination and interpretive uncharity; and (v) blocking and flouting (semantic) presuppositions' (Sterken 2020, 424). If someone is recognized as an authority or trendsetter in a group, what they say *about* how we are to use words, or (more simply) their own behavior in how they *do* use words (including using the disruptive tactics identified above) may send a powerful signal to those who recognize her or him as authoritative, to expect changes in behavior and normative attitudes, which in turn are essential for changes in social norms.

Simple changes in *use*, especially by those who are recognized as (officially or informally) 'authoritative' within a group, can also lead to changes in the norms governing our speech by way of the 'expressive commitments' such speakers exhibit. Lynne Tirrell develops the idea of expressive commitments as follows: 'When we speak, we make expressive commitments, which are commitments to the viability and value of ways of speaking' (1999, 42). Simply in using terms in certain ways, Tirrell argues, speakers undertake expressive commitments, which can spread by an authoritative speaker's use (an implicit endorsement). The commitments of other speakers can also be publicly challenged by anyone or, most effectively, by (formal or informal) authority figures. Thus, we can challenge the use of pejoratives, diminutives ('girl'), or other inappropriate or harmful terminology – and those challenges that come from authority figures will be most likely to have effect on changing norms of how (or whether) to use the relevant terms.<sup>28</sup>

---

<sup>28</sup>The thought that we may be engaging in reinforcing or changing the social norms that govern which terms we are to use and how we ought to use them simply as we speak, and undertake or challenge the relevant expressive commitments, fits well with the idea familiar from Ludlow (2008), and Plunkett and Sundell (2013) that we may often be engaged implicitly 'metalinguistic negotiation', whether or not we acknowledge it and make it explicit. On this see also Shields (2018) who argues that those engaged in a deep disagreement, say, about 'what a person is' are implicitly aiming to shape the language by which we will make sense of the world.

## 5. Conclusion

Where have we been? After outlining some different sorts of occasions or situations in which conceptual engineering is needed, I have suggested that we think first and foremost of engineering *words*, and that we should think of words as historical cultural entities, on the model of *abstract artifacts* that I have developed elsewhere (1999) (and along the lines endorsed by Kaplan (1990) and Irmak (2019)). Doing so, I have argued, fits well with work in historical linguistics, and enables us to rely on some of that work in beginning to understand processes of word formation and meaning change over time – natural processes that we can make use of to increase our chances of success at conceptual engineering. I've also argued that understanding words as abstract artifacts fits naturally with seeing words as having a function (in terms of which we can evaluate the success of various projects in conceptual engineering) and seeing them as governed by norms of appropriate use. The latter in turn enables us to make use of research on change in social norms to begin to get a better picture of how conceptual engineering could be implemented – how we could really enact the sorts of changes in what words we use and how we are to use them. Despite doubts some have raised, we do have reasons to think such changes are possible – particularly where those changes are not too far from our existing practices and build on extant and recognized processes of meaning change and word formation. Changes are also more likely to succeed if they are supported by official edicts of governments or other institutional bodies, or supported by exemplary expressive commitments (or challenges to expressive commitments) made by those in formal or informal positions of authority, in ways that change the expectations of speakers about the behavior of their interlocutors, and that change the standards and sanctions which speakers can expect to be held to.

Seen in this light, change in the meanings of our words is neither impossible, uncontrollable, nor inscrutable. It may be difficult, and require much further investigation, to determine how best to change social norms, including those of language use. But sometimes it's just too important to give up.

## Acknowledgements

This paper was originally presented at the NYU conference on The Foundations of Conceptual Engineering, held September 14–15, 2018. Many thanks go to the audience there, as well as to the audiences at further presentations of this paper at the

University of Pennsylvania and the Jowett Society at Oxford University, for many helpful questions and comments. Thanks also to the editors and an anonymous referee for helpful comments on an earlier version of this paper.

## Disclosure statement

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

## References

- Appiah, Anthony. 1992. *In My Father's House*. New York: Oxford UP.
- Baker, Lynne Rudder. 2008. "The Shrinking Difference Between Artifacts and Natural Objects." In *Newsletter on Philosophy and Computers*, edited by Piotr Boltuc, 2–5. American Philosophical Association Newsletters, Newark, DE.
- Berkeley, George. 1713. *Three Dialogues Between Hylas and Philonous*.
- Bicchieri, Christina, and Hugo Mercier. 2014. "Norms and Beliefs: How Change Occurs." *Iyyun: The Jerusalem Philosophical Quarterly* 63: 2014.
- Blackburn, Simon. 1993. *Essays in Quasi-Realism*. New York: Oxford University Press.
- Braddon-Mitchell, David. 2020. "Reactive Concepts." In *Alexis Burgess, Herman Cappelen, and David Plunkett, eds., Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Brandom, Robert. 2007. "Inferentialism and Some of Its Challenges." *Philosophy and Phenomenological Research* LXXIV (3): 651–676. doi:10.1111/j.1933-1592.2007.00044.x.
- Cappelen, Herman. 2018. *Fixing Language: Conceptual Engineering and the Limits of Revision*. Oxford: Oxford University Press.
- Cappelen, Herman, and David Plunkett. 2020. "A Guided Tour of Conceptual Engineering and Conceptual Ethics." In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 1–34. Oxford: Oxford University Press.
- Carnap, Rudolf. 1950. "On Explication." In *Logical Foundations of Probability*, 1–18. Chicago: University of Chicago Press.
- Deutsch, Max Emil. 2020. "Speaker's Reference, Stipulation, and a Dilemma for Conceptual Engineers." *Philosophical Studies*, 1–24. doi:10.1007/s11098-020-01416-z.
- Durkin, Philip. 2009. *The Oxford Guide to Etymology*. Oxford: Oxford University Press.
- Gert, Bernard, Charles M. Culver, and K. Danner Clouser. 2006. "Death." In *Bernard Gert, Charles M. Culver, and K. Danner Clouser, eds., Bioethics: A Systematic Approach*, 283–308. Oxford: Oxford University Press.
- Halliday, M. A. K. 2009. *The Essential Halliday*. Ed. Jonathan Webster. London: Continuum.
- Herbert, C., and Rebecca Kukla. 2017. "Ingrouping, Outgrouping, and the Pragmatics of Peripheral Speech." *Journal of the American Philosophical Association* 2 (4): 576–596.
- Hilpinen, Risto. 2018. "Artifact." In *The Stanford Encyclopedia of Philosophy (Summer 2018 Edition)*, edited by Edward N. Zalta. <https://plato.stanford.edu/archives/sum2018/entries/artifact/>
- Irmak, Nurbay. 2019. "An Ontology of Words." *Erkenntnis* 84 (5): 1139–1158.

- Isaac, Manuel Gustavo. 2021. "Broad-Spectrum Conceptual Engineering." *Ratio* 00: 1–17. doi:10.1111/rati.12311.
- Isaac, Manuel Gustavo. forthcoming. "Which Concept of Concept for Conceptual Engineering?"
- Jorem, Sigurd. 2020. "Conceptual Engineering and the Implementation Problem." *Inquiry*, 1–26. doi:10.1080/0020174X.2020.1809514.
- Kant, Immanuel. 1781. *Critique of Pure Reason*.
- Kaplan, David. 1990. "Words." *Aristotelian Society Supplementary Volume* 64 (1): 93–119.
- Koch, Steffen. 2021. "The Externalist Challenge to Conceptual Engineering." *Synthese* 198: 327–348.
- Koslow, Allison. forthcoming. "Meaning Change and Changing Meaning." *Synthese*.
- Ludlow, Peter. 2008. "Cheap Contextualism." *Philosophical Issues* 18.
- McWhorter, John. 2016. *Words on the Move*. New York: Henry Holt and Company.
- Pinder, Mark. 2021. "Conceptual Engineering, Metasemantic Externalism and Speaker-Meaning." *Mind* 130: 141–163.
- Plunkett, David, and Tim Sundell. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosopher's Imprint* 13 (23): 1–37.
- Price, Huw. 2011. *Naturalism Without Mirrors*. New York: Oxford University Press.
- Riggs, Jared. 2019. "Conceptual Engineers Shouldn't Worry About Semantic Externalism." *Inquiry*. doi:10.1080/0020174X.2019.1675534.
- Ryle, Gilbert. 1949. *The Concept of Mind*. New York: Hutchinson.
- Sawyer, Sarah. 2020. "Talk and Thought", in *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 379–95. Oxford University Press.
- Schärp, Kevin. 2013. *Replacing Truth*. Oxford: Oxford University Press.
- Shields, Matthew. 2018. "On the Pragmatics of Deep Disagreement." *Topoi*. doi:10.1007/s11245-018-9602-0.
- Sterken, Rachel Katherine. 2020. "Linguistic Intervention and Transformative Communicative Disruptions." In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 417–434. Oxford: Oxford University Press.
- Strawson, Peter. 1963. "Carnap's Views on Constructed Systems Versus Natural Languages in Analytic Philosophy." In *The Philosophy of Rudolf Carnap, Library of Living Philosophers, Volume XI*, edited by Paul Arthur Schilpp, Open Court: La Salle, Illinois: 503–518.
- Thomason, Amie L. forthcoming. "How Should We Think About Linguistic Function?" (Promised for *Inquiry* Special Issue on Conceptual Engineering and Pragmatism, ed. Celine Henne and Yvonne Hütter Almerigi).
- Thomasson, Amie L. 1999. *Fiction and Metaphysics*. Cambridge: Cambridge University Press.
- Thomasson, Amie L. 2014. "Public Artifacts, Intentions and Norms." In *Artefact Kinds: Ontology and the Human-Made World*, edited by Pieter Vermaas, 45–62. Synthese Library: Springer.
- Thomasson, Amie L. 2017. "Metaphysical Disputes and Metalinguistic Negotiation." *Analytic Philosophy* 58 (1): 1–28.

- Thomasson, Amie L. 2020a. "A Pragmatic Approach to Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 435–458. Oxford: Oxford University Press.
- Thomasson, Amie L. 2020b. *Norms and Necessity*. Oxford: Oxford University Press.
- Tirrell, Lynne. 1999. "Derogatory Terms: Racism, Sexism and the Inferential Role Theory of Meaning." In *Language and Liberation*, edited by Christina Hendricks, and Kelly Oliver, 41–79. Albany, NY: SUNY Press.