Erkenntnis (2023) 88:2123–2143
https://doi.org/10.1007/s10670-021-00446-1

ORIGINAL RESEARCH

Why Conceptual Engineers Should Not Worry About Topics

Steffen Koch[1]

Received: 7 September 2020 / Accepted: 8 August 2021 / Published online: 21 September 2021

Abstract
This paper argues for explanatory eliminativism about topics (and cognates, such as subject matters) relative to the domain of conceptual engineering. It has become usual to think that topics serve an important explanatory role in theories of conceptual engineering, namely, to determine the limits of revision. I argue, first, that such limits can be understood either as the *normative limits* pertaining to the justification of conceptual engineering, as the *metaphysical limits* pertaining to the identity of the concepts in question, or as the *terminological limits* pertaining to usage of the original terminology. Second, I argue that the metaphysical reading is disputable as a theory of concepts and inconsequential for conceptual engineers, and that neither of the two leading accounts of topics that have been presented in the literature—the samesaying account and functionalism—determine the limits of revision in either of the two remaining senses. In the absence of more promising competitors, I conclude that there is no theoretical role for topics to play in theories of conceptual engineering. An upshot of my argument is that conceptual engineers should stop worrying about things like topic (dis)continuity, and instead shift their attention to the issues that really matter for justifying conceptual revisions or replacements, making terminological choices, and underpinning conceptual engineering with a theory of concepts.

1 Eliminativism, Conceptual Engineering, and Topics

Conceptual engineers have a lot to worry about: How does one identify good concepts? How does one change existing concepts? According to the received view, conceptual engineers should also worry about topic continuity—the question how the original topic of a concept can be preserved through the process of re-engineering it. In this paper, I argue that whereas the former worries are real, the latter one is a red herring. Even more radical than that, I will argue for *explanatory eliminativism*

✉ Steffen Koch
steffen.koch@rub.de

[1] Institut für Philosophie II, Ruhr-Universität Bochum, Universitätsstraße 150, 44801 Bochum, Germany

about the concept of topics relative to the domain of conceptual engineering. Explanatory eliminativism is the thesis that a given concept plays no explanatory role in a given domain and that it should therefore be eliminated from it. In this vein, I will argue that topics—or as they are sometimes called, 'subject matters'[1]—do not play any explanatory role in theories about conceptual engineering, and that philosophers should therefore eliminate these concepts from their theorizing about conceptual engineering.

This thesis has two important upshots. The first is that conceptual engineers should stop worrying about what topics are, or how they can preserve them during conceptual revision. This by itself is a surprising insight, for, following Strawson (1991/1963), many philosophers take topic preservation to be among the most important challenges for the project of conceptual engineering (Cappelen, 2018, 2020; Haslanger, 2000, 2021; Nado, 2021; Pinder, 2020; Prinzing, 2018; Sawyer, 2018, 2020).[2] The second upshot is that, once we have removed topics from their central place in debates about conceptual engineering, we will get a clearer view of the issues that really matter. These include how conceptual engineering can be justified, how we ought to make terminological choices, and how conceptual engineering can be underpinned by a theory of concept individuation.

In order to understand these claims better, let me provide some background. A particularly prominent objection to the very idea of conceptual engineering is that it threatens to change the topics of interest instead of providing us with better ways to think and talk about them. An early precursor of this objection can be found in Strawson's (1991/1963) reply to Carnap's (1971/1950) suggested method of explication, regarded by many as a form of conceptual re-engineering (Brun, 2016). In the introduction to his *Logical Foundations of Probability,* Carnap defines the method of explication as the attempt to "transform […] a given more or less inexact concept into an exact one or, rather, [to replace] the first by the second" (Carnap, 1971/1950, p. 3). Very roughly, Strawson contests the usefulness of this method in philosophy by arguing that what philosophers are interested in are illuminations or analyses of our everyday concepts—a project to which the invention of a refined terminology is deemed "utterly irrelevant" (Strawson, 1991/1963, p. 504), for it does "not…solve the typical philosophical problem, but…change[s] the subject" (ibid., p. 505).

What many philosophers have taken on board from Strawson's objection is, first, the idea that revision *can* go too far (if we are not careful enough, then we might end up producing irrelevant results) and, second, the idea that whether a given re-engineering proposal *does* go too far depends on whether the original topic has been preserved. In light of this, many conceptual engineers have sought to make progress on normative questions about when re-engineering proposals are legitimate by providing accounts of topics (Cappelen, 2018, 2020; Prinzing, 2018; Sawyer, 2018, 2020; Thomasson, 2020). The thought is that analyzing or explicating the notion of

[1] I will use these terms interchangeably.

[2] See Belleri (2021) and Knoll (2020) for recent critical discussion. Their conclusions are largely compatible, albeit less radical, than the thesis I defend in this paper. Both papers appeared after this one had been submitted.

a 'topic' will help us to set the boundary between good and bad cases of conceptual re-engineering, i.e., between those that help us to theorize about important issues and those that are, in Strawson's words, 'utterly irrelevant'.

I oppose this trend, for I think that topics cannot play such a role. As I will argue below, the issues that really matter for evaluating re-engineering proposals have little to do with topics, regardless of how these are construed or explicated. In the next section, I will start by distinguishing three questions regarding conceptual re-engineering. Each of them, I will argue later on, can be used to zoom in on the (alleged) explanatory role of topics. In Sect. 3, I will discuss Strawson's original objection to Carnapian explications in more detail, distinguish it from what I call the 'neo-Strawsonian worry', and argue that the latter lends itself to interpretations in terms of the questions outlined in Sect. 2. In Sect. 4, I dismiss one of these readings right away. Then, in Sect. 5, I discuss two prominent accounts of topics in detail—the samesaying account (Cappelen, 2018, 2020), and the functionalist account (Haslanger, 2000, 2021; Nado, 2021; Prinzing, 2018; Thomasson, 2020) —and argue that neither of them provides us with a notion of topics that can play one of the remaining explanatory roles. It will emerge that the reasons for this failure generalize to other possible accounts of topics. I end in Sect. 6 by considering what is left of the (neo-)Strawsonian worry, and which roles samesaying and conceptual functions might play beyond accounting for topic continuity.

2 Three Questions About Conceptual Engineering

As I will use the term, 'conceptual ethics' (Burgess & Plunkett, 2013a, 2013b) is the activity of *scrutinizing* the quality of representational devices. Understood in this sense, conceptual ethics is neither committed to any particular view of concepts, nor to the idea that conceptual ethics is primarily concerned with them. Conceptual ethics concerns itself with the discovery of representational defects (Cappelen, 2018) or with considerations about the fruitfulness of our representational devices (Carnap, 1971/1950). Conceptual engineering, on the other hand, is the activity that builds on the outcome of conceptual ethics. It is the activity of *improving* our representational devices, e.g., our conceptual apparatuses, once the assessment is done and has reached a (partly) negative verdict. Again, even though the term might suggest this, I do not mean conceptual engineering to be committed to concepts, or to concepts being its primary target. Conceptual engineering can be, and actually is, endorsed by skeptics about concepts (Cappelen, 2018, 2020; Pinder, 2021). Nonetheless, for expository purposes, I will stick with the 'concepts' terminology.

Understood in this broad sense, conceptual engineering can be carried out in either of three ways: (a) by *introducing* new concepts, (b) by *eliminating* existing concepts (either from a given field of interest or tout court), or (c) by *re-engineering*, i.e., replacing or revising (more on this below), a given concept. Which of these strategies works best in a given instance depends on the prior normative assessment. If one's conceptual repertoire is found to be incomplete, such that there are interesting and important phenomena not yet denoted by a (non-complex) concept, then introduction is often the way forward. If, on the other hand, one or more of

the concepts already in place is found wanting, then elimination is to be preferred. Things are more complicated in the case of re-engineering. Typically, re-engineering a given concept is justified by both kinds of assessment: that a given concept is found wanting, and that there is an interesting phenomenon not quite denoted by this concept that deserves special attention and which should thus be denoted by a revised version of the concept or a newly introduced one. Which ameliorative strategy to choose and how to justify it—assuming that any change in the *status quo* demands justification—is an important set of questions that any conceptual engineering project has to engage with.

Now, focusing just on re-engineering proposals, two further questions become pressing: one of a terminological nature, the other primarily metaphysical. Suppose you have successfully argued that a given concept should be re-engineered. Should you now introduce a new term to express the newly engineered concept, or should you continue using the old term? Answering this question is far from trivial. The obvious downside of using the old term is that this might stand in the way of successful communication. If one person uses the term to express the old concept, whereas another person uses it to express the new concept, this might lead to confusion, miscommunication, or verbal disputes (Cappelen, 2018; Chalmers, 2011; Sterken, 2020). On the other hand, if the project is explicitly to re-engineer a certain concept, then using the old term to express the new concept has the positive effect of making explicit which concept has been re-engineered. If I were to use a completely new word to express a concept that is meant to replace another one, then people will easily miss the replacement aspect of this proposal. Further issues that bear on the question of which terms we should use are summarized by Cappelen (2018) as 'lexical effects'. As he convincingly argues, terminological choices matter beyond semantics and pragmatics, as traditionally understood. The mere sound or appearance of spoken or written words often generates associations and expectations that the conceptual engineer may want to either preserve or discard in the process of re-engineering concepts.

Apart from terminological choices, there is also an allegedly deeper metaphysical question that arises for re-engineering projects: Is re-engineering a concept best understood in terms of *replacement*, or can it (at least sometimes) be understood in terms of *revision*? Put differently, can the identity of a given concept be preserved through the process of re-engineering it, or will this automatically amount to replacing it with a new one? I call this a metaphysical question, for it concerns the metaphysics of concepts. On a Fregean view, concepts have their semantic properties essentially. It will then not be possible to change the semantic properties of a concept, just to abandon it. On an alternative understanding of concepts, concepts are four-dimensional entities that persist and change through time (e.g. Prinzing, 2018; Richard, 2019; Sainsbury & Tye, 2013). On these views, it is at least sometimes possible to change a concept without abandoning it.

To sum up, there are three (sets of) questions concerning conceptual re-engineering:

Normative question: What justifies eliminating, introducing, or re-engineering a concept? In particular, what justifies re-engineering instead of introducing a new concept?

Terminological question Suppose we replace one concept by another, or revise it in a certain way. Under what circumstances should we nevertheless retain the old terminology? When should we rather introduce new terminology?

Metaphysical question Can re-engineering sometimes be identity-preserving? Or does re-engineering always amount to replacing a concept rather than revising it? If not, how do we account for the metaphysical difference between replacement and revision?

3 Strawsonian and neo-Strawsonian challenges

As I said in Sect. 1, topics, or subject matters, found their way into current debates about conceptual engineering through an objection that Strawson (1991/1963) raised in response to Carnap's suggested method of explication. Since explication's recent resurgence, albeit now often labeled 'conceptual engineering', philosophers active in the debate have revoiced and responded to Strawson's worry (Cappelen, 2018, 2020; Haslanger, 2000, 2021; Nado, 2021; Pinder, 2020; Prinzing, 2018; Sawyer, 2018, 2020; Thomasson, 2020). I will discuss some of their proposals below. In this section, I want to highlight the specific background of Strawson's worry, which in my view shows that his own worry is rather discontinuous with how it has been characterized in the more recent literature.

The Carnap-Strawson debate is a clash between two different philosophical programs. Strawson objects to Carnap's method of explication not because there are some intrinsic difficulties or problems with explication per se, but because there is no, or little, room for explications in how Strawson understands philosophy in general. I understand his worry as taking roughly the following form[3]:

1. Typical philosophical problems are problems about the concepts of non-scientific discourse.
2. Explications do nothing to illuminate concepts of non-scientific discourse.
3. Therefore: Explications are unfit to solve typical philosophical problems.

The premise that is crucial for the soundness of this argument is (1). Strawson's endorsement of this premise can only be understood against the background of his picture of philosophy. Strawson is very explicit about this picture in later works. For instance, in his *Analysis and Metaphysics. An Introduction to Philosophy* (1992), he describes philosophy as analogous to the study of grammar:

[3] See Pinder (2020) for a more detailed account of Strawson's objection, which I take to be compatible with my characterization in the main text.

> just as the grammarian […] labours to produce a systematic account of the structure of rules which we effortlessly observe in speaking grammatically, so the philosopher labours to produce a systematic account of the general conceptual structure of which our daily practice shows us to have a tacit and unconscious mastery (p. 7).

On this picture, philosophy is a rather conservative discipline. Instead of providing us with new conceptual tools to make philosophical or scientific progress, its goal is to help us understand and account for the status quo. Against this background, it is not difficult to see why he places such little value on explications. Explications, in his view, do little to illuminate the status quo, and are instead attempts at changing it. In this sense, they are irrelevant to philosophy—assuming, again, that the goal of philosophy is to illuminate our actual concepts. So, explications might be called for in cases where there is absolutely no chance of giving a consistent account of our actual conceptual structure; but these cases should be restricted to the bare minimum.

Carnap, just like many of today's conceptual engineers, pursues a different philosophical program. In his view, concepts are not the primary targets of philosophical inquiry, but rather a means to an end. This explains his characterization of concepts as tools, which is also endorsed by many of today's conceptual engineers (Carnap, 1971/1950, p. 938f.). In Carnap's view, it is the philosopher's job to discover conceptual failures, to invent new and better concepts, and to replace our everyday concepts with more specialized ones that are more conducive to the specific purposes at hand. I am inclined to take Carnap's side in this grand metaphilosophical debate, but this is not the place to get into it. What matters for present purposes is simply to notice that the Carnap-Strawson debate is primarily one about what philosophy is for, instead of being about any specific procedural problem with explications. Furthermore, as we have seen, both Strawson's worry and his larger disagreement with Carnap can be stated without any reference to topics or subject matters.

The recent resurgence of the challenge from topic discontinuity, now applied to conceptual re-engineering instead of Carnapian explication, appears to be a version of Strawson's worry, but closer scrutiny shows that its motivation is rather different. I will call it the *neo-Strawsonian worry*. Consider first Sally Haslanger's version of the worry, which she considers in the context of her ameliorative analyses of race and gender concepts:

> In asking what race is, or what gender is, our initial questions are expressed in everyday vocabularies of race and gender, so how can we meaningfully answer these questions without owing obedience to the everyday concepts? Revisionary projects are in danger of providing answers to questions that weren't being asked. Given the difficulty of determining what "our" concept is, it isn't entirely clear when a project crosses over from being explicative to revisionary, *or when it is no longer even revisionary but simply changes the subject*. (Haslanger, 2000, p. 34; my italics)

Admittedly, "[r]evisionary projects are in danger of providing answers to questions that weren't being asked" sounds a lot like another expression of Strawson's worry.

However, the last bit of the quoted passage makes clear that she assumes there to be a boundary between (at least some) revisionary projects and changes of topic. Worrying about how to determine whether a project crosses the border from being revisionary to being a change of topic presupposes that it is possible for a revisionary project *not* to cross this border. The crucial question she is addressing here is how to locate this border, and how to make sure not to cross it. This is not what Strawson was worrying about. His concern was rather that revisionary projects are irrelevant to the goals of philosophy *as such*.

Cappelen's recent reformulation of the worry exhibits a similar shift in focus. While motivating what he calls the 'Strawsonian challenge', he writes:

> One reason the cluster of issues tied up with the Strawsonian challenge are so important is that they make us focus on the limits of revision: anyone who has thought seriously about conceptual engineering worries (or should worry) about *the limits of it*. *How much revision is too much*? What does 'too much revision' mean? Strawson worried about it, Carnap worried about it, Railton and Haslanger worry about it. Clark and Chalmers (and all other engineers of specific concepts) should be worrying about it. (Cappelen, 2018, p. 105f.; my italics)

Again, the gist of this worry is not that revisionary projects might be of no use to philosophy. Instead, the worry is that there is a sense in which revision can 'go too far', and that conceptual engineers should be careful not to take it too far. Just like in the case of Haslanger, this presupposes that there is a border between legitimate revisionary projects and those that change the topic.[4]

Now, why is this divergence between Strawson's original worry and the neo-Strawsonian worry so important? This is because, although both versions of the worry speak about topics (or subjects), the two worries put very different constraints on them. When Strawson said that explications amount to changes of subject, what he had in mind was that revised versions of a concept c do not contribute to the goal of illuminating our everyday concept of c. When the goal is to account for a given concept c, then introducing a revised concept c* is a change of topic because c is not identical to c*. The neo-Strawsonian worry requires a different notion of topic. For neo-Strawsonians, it is possible for two (slightly) different concepts c and c* to share a topic. The crucial question for them is how to design a more fruitful (or less defective) concept c* that shares its topic with c. This commits them to the claim that topics are not individuated by a one-to-one correspondence with particular concepts, but are, as Cappelen says, "more coarse-grained" than concepts or word meanings.

The explanatory role of topics according to the neo-Strawsonian worry is, as the authors quoted above indicate, to determine the 'limits of revision', or when revision 'goes too far'. But things can go too far in many ways, and from the outset it is neither clear that these ways will coincide, nor that they can be accounted for by a single theory of topics. As I see it, talk of the limits of revision or when it goes too

[4] See Prinzing (2018) for a similar argument.

far can be interpreted in three ways, each of them corresponding to one of the three questions outlined in the last section:

> *The normative reading*: Topics demarcate the area of acceptable revision. An account of topics is needed to answer the question: When does re-engineering stop being justified?
>
> *The metaphysical reading*: Topics demarcate the difference between replacement and revision. An account of topics is needed to answer the question: Where is the border between revision and replacement?
>
> *The terminological reading*: Topics demarcate the difference between legitimate and illegitimate uses of a word. An account of topics is needed to answer the question: If we revise or replace a given concept, then when should we introduce new terminology to express the re-engineered concept rather than using the old term(s)?

In the next section, I will argue that the metaphysical reading is not compelling as a theory of concepts and normatively inconsequential for conceptual engineers. In Sect. 5, will take a closer look at how the two leading accounts of topics fare with respect to the normative and the terminological readings. It will emerge that neither of them draws the right distinctions in either domain. Since it is unlikely that the problems faced by these accounts can be overcome, I will conclude that the best way forward is therefore to stop theorizing about topics and address the remaining key questions more directly.

4 The Alleged Role of Topics in Concept Identity

According to the metaphysical reading, we need topics to demarcate the difference between revision and replacement. Behind this reading is the idea that not all cases of conceptual re-engineering (and meaning change in general) are of the same kind: some preserve a stable topic or subject matter whereas others do not. Take the examples of 'meat' and 'fish' (Sawyer, 2018). Both expressions have changed their meanings in the course of the last couple of centuries—'meat' from referring to basically all types of solid food up to the thirteenth century to referring only to food from animal flesh, and 'fish' from including whales in the seventeenth century to excluding them now. Despite this similarity, there is an important difference between the two cases. If one of us met with someone from the thirteenth century and discussed whether potatoes are meat, we would simply talk past each other. Both of us would be right according to the meaning that 'meat' had for each of us. There would be no way of establishing that either of these meanings was superior to the other. However, if we instead discussed whether whales are fish, we would not simply be talking past each other. Although we would use 'fish' with different meanings, there is a sense in which we would be right and they would be wrong. After all, acknowledging that whales are not fish was an important scientific discovery. Proponents of the metaphysical reading claim that the difference between the two cases is a matter of topic

continuity: whereas the new meaning of 'fish' better captures the topic it had all along, 'meat' has simply changed from one topic to another.

There are different options of developing this view further. One option is to hold that concepts are individuated coarsely, such that they can survive certain changes in content, as long as these changes are topic-preserving (e.g. Prinzing, 2018; Richard, 2019; Sainsbury & Tye, 2013). Another option, recently developed by Sawyer (2018, 2020), is to divorce linguistic meanings from concepts, holding that words can change their meanings while still expressing the same concept that ties them to a topic. It is impossible to thoroughly reject all possible incarnations of this view here. Instead, I will argue for two more modest conclusions: first, that conceptual engineers can account for the intuitive difference between what happened to 'meat' and 'fish' without invoking any substantial notion of topics, and hence that nothing compels them to worry about topics; and second, that the difference between topic-preserving and topic-changing instances of conceptual re-engineering, even if it existed, would be normatively toothless, contrary to what the neo-Strawsonian challenge claims.

First off, note that there is a standard treatment of the 'fish' case that looks rather different from the above description. On standard externalist accounts, 'fish', but not 'meat', is a natural kind term that is anchored in the natural kind fish and whose intension is thus insensitive to, or at least not entirely determined by, our views about fish and our practice of using 'fish'. On this view, 'fish' has not changed its meaning *at all*; rather, we have come to discover what its actual extension has been all along and have adjusted our use of 'fish' accordingly. We can, of course, express this continuity by saying that 'fish' preserved, whereas 'meat' abandoned, its original topic. But if all we mean by this is that 'fish' has changed its extension and 'meat' has not, then we use 'topics' synonymously to 'extensions'—a technical term we already have at our disposal and whose utility no sensible philosopher would question. The upshot is that on standard externalist accounts, we neither need a substantial notion of topics (one that is not identifiable with extensions or references), nor a general distinction between revision and replacement to account for the difference between cases like 'fish' on the one hand, and 'meat' on the other.

One might of course grant this explanation but question whether it extends to other terms than paradigm instances of natural kind terms. What about the change in meaning of terms like 'marriage' or 'rape'? One strategy, endorsed by Haslanger (2006) and others, is to simply extend the externalist picture of meaning and reference that Kripke and Putnam introduced for natural kind terms to terms that stand for social kinds. On this view, terms like 'marriage' or 'rape' are externally anchored to pre-existing social kinds in much the same way that 'fish' is anchored to fish. Another strategy, recently defended by Plunkett and Sundell (2013), is to classify seemingly object-level disputes, e.g. over whether same-sex couples can get married, as (potentially covert) *metalinguistic negotiations*—disputes about how words should be used in light of relevant moral or other normative considerations. This strategy, too, allows for the possibility of genuine disagreement despite differences in meaning. The upshot is that conceptual engineers can accommodate the intuitive data points without invoking topics and a theoretical distinction between conceptual revision and replacement.

Second, even if there *were* a theoretically fruitful distinction between conceptual revision and replacement that is best drawn in terms of topics, this difference would not have the significance that it is commonly thought to have. As we have seen in the last two sections, both the Strawsonian and the Neo-Strawsonian challenge are commonly thought to carry *normative weight*. Strawson's concern with explications is that they are "utterly irrelevant" when it comes to answering philosophical questions. Cappelen argues that "all…engineers…should be worrying about" crossing the limits of revision. But if the role of topics were restricted to marking the difference between conceptual revision and replacement, then this would be misguided, because whether a given conceptual re-engineering proposal preserves or replaces the original concept *has no normative implications whatsoever.*

Let me explain this by use of the following thought experiment. In a close by possible world, two philosophers named Andy and David propose to re-engineer the concept of belief such that it includes information stored on easily accessible external devices.[5] Andy and David are unusually successful: they convince a great majority of their community to follow their proposal. But now the talented philosopher Susan comes along and proposes an ingenious account of concepts that equally convinces everybody. According to Susan's account, Andy and David's suggested concept changes the topic and therefore does not count as a revision of the original belief concept, but as a replacement thereof. Would their arguments to use the new concept instead of the old one lose their force in the face of this discovery? I submit that Susan's new theory of topics and concepts does not undermine Andy and David's reasons for re-engineering the concept of belief in the slightest. If their proposal was well supported in the beginning, then it remains so after learning that it replaces rather than revises the original concept. As far as Andy and David are concerned, what matters is that people use their proposed concept. Whether this amounts to replacing or revising the original concept is entirely irrelevant.

The above shows that the distinction between conceptual revision and replacement, even if it were real, would be *normatively toothless*: it would have no bearing whatsoever on whether a given re-engineering proposal should be endorsed or not, and it therefore ought not to cause sleepless nights to would-be conceptual engineers who propagate the use of certain concepts.

5 Accounts of Topics

If I am correct, then conceptual engineers ought not to worry about topics as markers of concept identity. In what follows, I will consider how the two most influential accounts of topics fare with respect to the two remaining candidates, i.e. the normative and the terminological reading. The two accounts I will look at are the

[5] See Clark and Chalmers (1998) for their actual and more sophisticated proposal.

samesaying account (Cappelen, 2018, 2020) and the *functionalist account* (Haslanger, 2000, 2021; Nado, 2021; Prinzing, 2018; Thomasson, 2020).[6]

5.1 The Samesaying Account of Topics

Cappelen (2018, 2020) argues, against the neo-Strawsonian challenge, that conceptual engineering without topic-change is possible. This argument ties topics to the phenomenon of samesaying and disquotational reports. The idea is, roughly, that a re-engineered term preserves the topic of the original term as long as the two can be used to say the same thing, i.e., as long as someone can use the re-engineered term to report what someone else using the non-engineered term has said. Suppose, for instance, that somebody re-engineered the meaning of the word 'marriage'. How can we decide whether this revision is topic-preserving? According to Cappelen, to answer this question we have to check whether someone could use the re-engineered term 'marriage' to report, disquotationally, what someone else who used the term 'marriage' with the old meaning has said. Based on this construal of topics, Cappelen runs the following argument for the possibility of topic-preserving conceptual engineering:

1. Conceptual engineering happens when people intentionally change a term's intension/extension.
2. Topic preservation goes hand in hand with 'samesaying' (disquotational reports).
3. Samesaying is possible despite differences in intensions/extensions.
4. Therefore: There can be topic preservation despite differences in intensions/extensions.
5. Therefore: Conceptual engineering is compatible with topic preservation.

(1) is part of Cappellen's austerity framework, according to which conceptual engineering does not involve concepts, but merely the intensions and extensions of words. (2) expresses his construal of topics. Cappelen offers two sorts of considerations in support of (3). First, the apparent legitimacy and success of disquotational reports involving context-sensitive expressions ('tall', 'smart', etc.) across different contexts. Here the idea is that terms like 'tall' and 'smart' have different intensions/extensions in different contexts, but can still be used in one context to report what someone said in another context. If this is true, then it shows that not every difference in intensions/extensions stands in the way of samesaying. Second, Cappelen appeals to Dorr and Hawthorne's (2014) thesis of semantic plasticity to argue that the same holds across different times as well. People today can use the word 'salad', for example, to report what people said when they used 'salad' in the eighteenth

[6] I should say that not all of the authors cited think of themselves as providing accounts of topics. In the first instance, what these authors do is to engage with the neo-Strawsonian worry. However, since their accounts aim to address the cluster of issues I identified as the explanatory role of topics in the last section, this poses no major challenge to the methodology I employ in this paper. Thanks to Sigurd Jorem for pointing this out to me.

century, even though the term underwent changes in intension/extension between then and now.[7] If this argument is sound, then it establishes that topic-preserving conceptual engineering is at least possible.[8]

For Cappelen, topics are semantic entities that are more coarse-grained than intensions and extensions. This is the reason why two slightly different intensions/extensions can still share a topic. Furthermore, topics go hand in hand with samesaying. The question I am interested in now is this. Which explanatory role can this notion of a topic play in theories of conceptual engineering, and in particular, how can this notion contribute to answering (one of) the three questions about conceptual engineering I have identified in Sects. 2 and 3?

As argued above, Cappelen's talk of 'the limits of revision' or 'revision going too far' can be interpreted via the normative, the metaphysical, or the terminological reading. These interpretations correspond to three possible explanatory roles for topics. We can therefore assess the samesaying account by considering whether samesaying can play either of these roles:

Samesaying-Normative: We shouldn't re-engineer beyond samesaying.

Samesaying-Metaphysical: Concepts cannot be revised beyond samesaying.

Samesaying-Terminological: Usage of the original term shouldn't go beyond samesaying.

As argued in the last section, the demand for a marker of concept identity posed by the metaphysical reading is disputable as a theory of concepts and normatively inconsequential for conceptual engineers. In what follows, I will therefore focus the discussion on *Samesaying-Normative* and *Samesaying-Terminological.* If *Samesaying-Normative* or *Samesaying-Terminological* is true, then the notion of topic delineated by the samesaying account does indeed play a significant explanatory role, because it helps us to determine at least one sense in which revision can go too far. If, on the other hand, they are not true, then the notion of topic delineated by the samesaying account does not play an explanatory role and, in the absence of some plausible alternative reading, may thus be discarded.

The main reason for not accepting *Samesaying-Normative* is that there can be good reasons to revise the meaning of an expression beyond the limits of samesaying. It is just not true that the limits of samesaying mark the boundaries of acceptable revision. To see this more clearly, let us consider Kate Manne's (2018) recent revisionary account of misogyny. In her own words, her aim is "to offer an ameliorative proposal about how we *ought* to understand misogyny, at least for many purposes" (p. 63). The account she ends up offering goes as follows:

[7] For concerns about Cappelen's evidence for P3, see Koch (2019).

[8] Despite arguing for the possibility of topic-preserving conceptual engineering, Cappelen claims that there are no clear conditions for topic preservation, as these conditions are themselves up for grabs. He calls this the *Contestation Theory of the Limits of Revision*. The details of this will not matter in what follows.

> MISOGYNY$_{new}$
>
> Constitutively speaking, misogyny in a social environment comprises the hostile social forces that
>
> (a) will tend to be faced by a (wider or narrower) class of girls and women because they are girls and women in that (more or less fully specified) social position; and
>
> (b) serve to police and enforce a patriarchal order, instantiated in relation to other intersecting systems of domination and disadvantage that apply to the relevant class of girls and women (e.g., various forms of racism, xenophobia, classism, ageism, transphobia, homophobia, ableism, and so on). (p. 63)

As she herself claims, MISOGYNY$_{new}$ deviates from our unameliorated concept of misogyny (the "naïve conception", p. 32), which might roughly be characterized as follows:

> MISOGYNY$_{old}$.
>
> [M]isogyny is primarily a property of individual agents (typically, although not necessarily, men) who are prone to feel hatred, hostility, or other similar emotions toward any and every woman, or at least women generally, simply because they are women. (p. 32)

MISOGYNY$_{old}$ and MISOGYNY$_{new}$ are rather different concepts. Roughly speaking, the former refers to a particular character vice, typically of men, whereas the latter refers to a social force that serves to police and enforce a patriarchal order. These two concepts clearly have different extensions. Furthermore, it is rather unlikely that two people who use the expression 'misogyny' with these two different meanings are saying the same thing. After all, one is talking about a character trait, whereas the other is talking about social forces that are not necessarily grounded in any character traits at all. Because of this, we cannot simply report Kate Manne's view by saying e.g. 'Kate Manne claims that misogyny is a social force rather than a character trait'. Assuming, with Manne, that 'misogyny' currently expresses MISOGYNY$_{old}$, this statement is false if MISOGYNY$_{old}$ applies to character traits. But Manne's actual view is compatible with MISOGYNY$_{old}$ applying to character traits. So to report Manne's view in the proposed way would be to misreport it. This shows us that MISOGYNY$_{old}$ and MISOGYNY$_{new}$ do not allow for samesaying. But does this mean that Manne's proposal to re-engineer MISOGYNY in this way is off the table? Not quite. Whether her re-engineering proposal is worthwhile depends on whether she is right that MISOGYNY$_{new}$ picks out "the phenomenon we need to be thinking about" (p. 62). If this is indeed the case, then it does not matter whether MISOGYNY$_{old}$ and MISOGYNY$_{new}$ can be used to say the same thing.[9]

At this point, you might object that my statement of *Samesaying-Normative*, and thus the target of my attack, is unduly lax in that it does not properly distinguish the two widely recognized variants of conceptual re-engineering: revision

[9] Pinder (2021, p. 159) makes essentially the same point.

and replacement.[10] I have treated Cappelen as being committed to the view that we should *neither* revise *nor* replace beyond the boundaries of samesaying. But perhaps he merely thinks that we should not *revise* concepts beyond samesaying, where this is compatible with the legitimacy of *replacing* them beyond those limits. However, this response fails for the reasons outlined when discussing the metaphysical reading in the last section. To repeat, the distinction between replacement and revision is normatively inconsequential for conceptual engineers. The above discussion has shown that *Samesaying-Normative* draws the wrong normative classifications: it unduly dismisses Manne's proposal as unjustified. For the distinction between replacement and revision to secure a way out for *Samesaying-Normative,* this distinction itself would have to be normatively significant, which, according to my arguments in the last section, it is not.

What I've argued in the case of 'misogyny' generalizes to other re-engineering projects as well. The underlying issue, I take it, is the following one. The area of acceptable re-engineering delineated by *Samesaying-Normative* is extremely narrow. Even granting that samesaying is sometimes possible despite differences in intensions, any significant re-engineering project will cross this boundary. But it is simply not plausible to regard so many significant re-engineering projects as non-starters. There seems to be no good reason to restrict the conceptual engineer's choice of replacement concepts (or ways of revising an existing concept) so radically.

Does samesaying settle the *terminological question* about when one is licensed to retain the old word through the process of re-engineering? Considerations about samesaying should certainly play a role in our terminological choices. After all, if samesaying breaks down, then confusion, miscommunication, and verbal disputes may ensue. But such considerations are not strong enough to warrant *Samesaying-Terminological.* Even though it is *prima facie* bad if conceptual re-engineering leads to confusion, miscommunication and verbal disputes, there are other things that factor into the equation and that can sometimes override these considerations. I argued earlier that retaining the lexical effects of a particular term is sometimes important, and it is not difficult to imagine that this can be more important than avoiding minor kinds of confusion, miscommunication and verbal disputes. A further consideration that might override samesaying has to do with what we might call 'focus shifting'. When we communicate transparently that we use an old term with a new meaning, we invite people to shift their attention from the term's earlier meaning to its new one. We tell people that *this,* rather than *that,* is the phenomenon we should concern ourselves with. This aspect of re-engineering proposals is lost if we limit our terminological choices in the way suggested by the samesaying account.

Consider again the example of 'misogyny', and suppose that samesaying does not work for $\textsc{Misogyny}_{\text{old}}$ and $\textsc{Misogyny}_{\text{new}}$. If *Samesaying-Terminological* were true, this alone would tell us that Manne (and others) should not use 'misogyny' to express $\textsc{Misogyny}_{\text{new}}$, but introduce a new term instead. But this seems false. Whether or not it would be better to introduce a new term rather than using the old one depends on the overall consequences. It could well be the case that retaining

[10] Thanks to an anonymous reviewer of this journal for prompting me to consider this objection.

the word 'misogyny' has better consequences overall than introducing a new one. Perhaps calling the phenomenon Manne is interested in 'misogyny' would get many more people interested in engaging with her view, which would eventually have positive consequences for society. Perhaps retaining this word is the easiest way to latch on to the negative emotional responses that Manne aims to elicit. Perhaps this is the most successful way of communicating that what we currently take to be explained by $\textsc{Misogyny}_{old}$ is actually explained by $\textsc{Misogyny}_{new}$. And so on. To be clear, I am not arguing that any of these possible consequences do in fact occur. My point is merely that these and potentially other kinds of consideration ought to guide our terminological choices. Although samesaying might play a role therein, it is surely not the only relevant consideration.

5.2 The Functional Account of Topics

Functionalists about topics tie topic-preservation to the function, purpose, or point of a concept.[11] Roughly, the idea is that revisions are topic-preserving as long as they retain the concept's original function, point, or purpose. This view is currently very popular and has been defended in different versions. The central distinction between these views is whether functions are taken to individuate concepts or not. Following Prinzing (2018), I will call the former the 'principled view' and the latter the 'unprincipled view'.

According to the principled view, a topic is preserved through conceptual revision iff the revision preserves the identity of the concept. So, the principled view is basically a view of concepts. On this view, concepts have essential and accidental features. If a revisionary project saves the essential ones, then it counts as identity-preserving, and thus as preserving the original topic. The essential features of concepts are taken to be their functions, i.e., what they are for. According to the unprincipled view, functions do not account for concept identity. Instead, conceptual re-engineering is taken to be topic-preserving just in case it retains the original function of the concept, regardless of whether this amounts to preserving the identity of the concept or not.[12]

Again, as with the samesaying account, let us see which explanatory role conceptual functions can play for us in the context of determining the 'limits of revision'. I have argued above that there are three possible ways to understand what is meant by this phrase. Applied to the functionalist account of topics, this gives us the following three claims:

[11] Or at least they think that functions play the role that I earlier identified topics as playing (see footnote 6).

[12] There are of course different accounts of conceptual functions which advocates of either of these views can choose from. The literature on functions distinguishes between three different paradigms: design functions, system functions, and etiological functions (Houkes and Vermaas, 2010). There is currently no agreement among functionalists about which of these fits the conceptual engineer's purposes best: Prinzing (2018) appeals to design functions; Thomasson (2020) leans more toward an etiological view; Haslanger (forthcoming) speaks of system functions; and Simion and Kelp (2019) defend the view that conceptual engineering is successful insofar as design functions turn into etiological functions.

Functionalism-Normative: We shouldn't re-engineer beyond the original function(s).

Functionalism-Metaphysical: Concepts cannot be revised beyond their original functions.

Functionalism-Terminological: Usage of the original term shouldn't be extended to concepts with different functions.

Functionalism-Metaphysical is an instantiation of the metaphysical reading. But as argued in the last section, the metaphysical reading is disputable as a theory of concepts and normatively inconsequential for conceptual engineers, which is why conceptual engineers have little reason to accept *Functionalism-Metaphysical*. I will therefore focus the discussion on *Functionalism-Normative* and *Functionalism-Terminological.*

We can distinguish between a strong and a weak version of *Functionalism-Normative.*[13] According to the strong version, the original functions of a concept must be the *only* functions of the successor concept. According to the weak version, the original functions of a concept must merely be preserved by the successor concept. The weak version is thus compatible with the addition of new functions, whereas the strong version is not. In my view, the strong version of *Functionalism-Normative* is clearly false, because, no matter how exactly you construe the function of a concept, there will almost certainly be (actual or possible) cases of legitimate re-engineering proposals in which new functions are added to a concept (or a replacement thereof). Think of Haslanger's project of re-engineering gender concepts. As she herself claims, the goal of an ameliorative analysis such as her own is to identify how "we might usefully revise what we mean [by certain terms] for certain theoretical and political purposes" (Haslanger, 2000, p. 34). In Haslanger's case, the relevant theoretical and political purposes are feminist purposes; in particular, the purposes to identify and remediate gender-based oppression. Now, is it among the central functions of our ordinary gender concepts to identify and remediate gender-based oppression? Haslanger (2000) voices some optimism, but does not argue this point in detail. Personally, I do not think that our ordinary gender concepts even come close to having these functions, but I admit that this is ultimately an empirical question that cannot be settled on a priori grounds (as Haslanger would surely agree). Suppose it turned out that, as a matter of empirical fact, our gender concepts do not have these functions. If *Functionalism-Normative* in its strong version were true, then this would immediately show that her re-engineering proposal comes out unjustified. But this seems wrong. For even if our current gender concepts do not serve these functions, it is still an open question whether they *should* serve these functions, and ought to be revised in ways that allow them to do so efficiently. In

[13] I am thankful to an anonymous reviewer of this journal for suggesting this distinction.

fact, we can easily imagine a Haslanger-style ameliorative project being justified precisely on these grounds.[14]

What about *Functionalism-Normative* in its weaker version? Is it sometimes legitimate *not* to preserve certain functions of a concept through the process of re-engineering it? Almost certainly. The simple reason is that there is no guarantee that all the functions that our concepts are endowed with are good functions. Where they are not, it is perfectly legitimate—arguably even required—to make sure that a re-engineered concept does not have them. Suppose, for example, that what functions a concept has is a matter of what made people use this concept in the past, as an etiological view of functions might suggest. We can easily imagine that people have used concepts for bad reasons—because they helped them to sustain or even reinforce sexist and racist oppression, or simply because people failed to see that these concepts stood in the way of scientific progress. In such cases, we ought not to preserve the relevant functions. Basically the same holds for proponents of a design view about functions. Concepts can be designed for all sorts of purposes, not all of them good or legitimate. It would be bizarre to think, as per *Functionalism-Normative*, that we ought not to correct our mistakes of the past.

Notice, once again, that alluding to the distinction between revision and replacement will not be of help to proponents of *Functionalism-Normative* (in neither of its versions). The issue is that *Functionalism-Normative* draws the wrong normative classifications. For the distinction between revision and replacement to be of help here, this distinction would have to make a normative difference. But as argued in the last section and repeated in the discussion of *Samesaying-Normative*, this is not the case. Even if the distinction between replacement and revision were a real one (for which I have provided some doubts in the last section), whether a given re-engineering proposal counted as a revision or a replacement would not affect the reasons that justify it.

The general lesson is the following. It is often a legitimate goal of conceptual engineers to design concepts with an eye on rather specific theoretical or practical purposes—be they feminist, scientific, or whatever else. But doing this sometimes requires adding new functions or abandoning old ones. There is no good reason to think that this practice is a non-starter. It therefore seems that, pace *Functionalism-Normative*, there are justified instances of conceptual re-engineering in which new functions are added to a concept or old ones are abandoned.[15]

[14] This point is also recognized by Nado (2021, p. 1519): "Has Haslanger changed the subject? Probably. But she has done so in a principled fashion, rejecting functions she takes to need rejecting while retaining uses that still hold value."

[15] In fact, this point is sometimes recognized even by proponents of functionalism. Prinzing (2018) writes that "Sometimes a change in subject is precisely what we need. We shouldn't be too afraid to concede discontinuity" (footnote 14). Thomasson (2020) takes it to be a "crucial step…to undertake work that is more explicitly in conceptual ethics: determining what functions (if any) these concepts should serve, and are to serve going forward, given the goals and purposes we have" (pp. 448–449). Nado (2021) holds that "replacing radically defective concepts with ones that 'change the subject'—should be a permissible strategy for an engineer" (p. 1514).

Now let us consider *Functionalism-Terminological.* Here the case is similar to that of *Samesaying-Terminological.* Terminological choices depend on many different factors: success in communication, lexical effects, communicating shifts of interest, the intention to disrupt, etc. Even if functions should play a role in how we ought to make these choices, this role will certainly not be strong enough to warrant *Functionalism-Terminological*.

To see this, consider again Haslanger's proposal for re-engineering our gender concepts, and suppose that her proposal does not preserve all and only their original function(s).[16] If *Functionalism-Terminological* were true, this alone would tell us that she should not use the words 'woman' and 'man' to express her newly designed concepts. But this seems false. Gender concepts are so contested and ubiquitous in our everyday communication that their lexical effects can hardly be overestimated. If these effects are conducive to Haslanger's engineering project, then retaining the words 'man' and 'may well be justified. To be clear, my point is not that this is in fact the case. My point is merely that considerations about a concept's function will not (completely) settle the important terminological question at stake in revisionary projects such as Haslanger's.

Let us take stock. The main explanatory role of topics within theories of conceptual engineering is to determine the limits of revision. I have argued that there are three ways of conceiving of such limits—a normative, a metaphysical and a terminological one. The normative and the terminological readings, but not the metaphysical one, give us a notion of topics that is salient in the context of conceptual engineering. In this section, I have argued that neither of the two most popular approaches to topics—samesaying and functionalism—determine the limits of revision in either of these senses.

The following common patterns have emerged. Conceptual re-engineers can pursue many different goals: to find a concept that serves a given function better, to dispose of the function of a given concept, to design concepts for very specific purposes, etc. This makes it highly implausible that samesaying, functionalism, or indeed any other fixed criterion will be decisive with respect to the normative question. The same is true of the terminological question. This question demands highly case-sensitive answers that draw on a large variety of factors that must be carefully balanced. This makes it difficult to see how samesaying, functionalism, or any other fixed criterion of topic-continuity will adequately address this question.

6 What is Left of Topics, Functions and Samesaying

If what I have argued so far is correct, then there is a strong pull toward eliminating topics, or subject matters, from our theorizing about conceptual engineering. This is because the envisaged role of topics, namely to determine the limits of revision in a sense relevant for conceptual engineering, is not well defined to begin with—it is

[16] Here, I will leave it as an exercise to the reader to adjust this argument for a potential weaker reading of *Functionalism-Terminological* that is compatible with the addition of new functions.

often unclear whether such limits are construed in terms of normativity, the individuation of concepts, or terminological choice—and none of the suggested renderings is both independently plausible and well explained by either of the extant accounts of topics in the literature. Given the complexity and case-sensitivity of what justifies conceptual re-engineering proposals and the relevant terminological choices, it is also unlikely that any other account of topics will do better than the ones discussed here.

This raises two important further questions. First, is there anything left of the neo-Strawsonian worry, according to which good cases of conceptual engineering preserve something of the original concept? Second, are there other explanatory roles for samesaying and conceptual functions beyond delineating topics?

To answer the first question, consider how Quine describes the method of explication:

> We fix on the particular functions of the unclear expression that make it worth troubling about, and then devise a substitute, clear and couched in terms to our liking, that fills those functions. Beyond those conditions of partial agreement, *dictated by our interests and purposes*, any traits of the explicans come under the head of "don't-cares" . (Quine, 1960, p. 258–9; my italics)

Some writers take this passage as evidence that Quine held a functional view about topic preservation (Nado, 2021). I want to suggest a different reading: Quine effectively says that through the process of explication, we need to preserve only that which is *dictated by our interests and purposes*. The interests and purposes we have vary from context to context. Sometimes, we might want to preserve a concept's central function. Sometimes, we might instead want to preserve one of its deviant uses. Sometimes, we might be interested in extensional overlap between the old and the new concept. And at yet other times, we might want to preserve little more than the lexical effects of the term expressing the old concept in a particular language. These issues are highly sensitive to the context of a particular re-engineering proposal. And although providing a good justification for re-engineering a particular concept for a particular purpose is far from trivial, I do not think that there is much to say about this from a purely general perspective. This has both a positive and a negative upshot. Conceptual engineers can expect little guidance in how to make conceptual choices from meta-theories of conceptual engineering. On the other hand, their choices will also not be restricted by meta-theories (despite what members of the pro-topics-camp might suggest).

Second, what remains of samesaying and conceptual functions? First of all, notice that I do not claim that samesaying or conceptual functions are irrelevant in the context of conceptual engineering. I have only argued that neither of them sets the limits of acceptable revision, which is what advocates of samesaying and functionalism typically want them to do. As mentioned earlier, it is possible, and indeed quite plausible, that considerations of samesaying do matter for how we ought to make terminological decisions. It is often a good idea to avoid confusion, miscommunication and verbal disputes. If Cappelen is right, then not every instance of retaining the old term through the process of conceptual replacement or revision has these negative consequences. Depending on which other factors enter the equation in a given

instance, arguing that a particular re-engineering proposal falls within the boundaries of samesaying may sometimes tip the scale toward retaining the old term.

It seems that functions, too, have a legitimate role to play in conceptual re-engineering projects. Nado (2021), for instance, thinks of functions as playing two separate roles in conceptual engineering projects. The first is to give us the desired continuity between the old and the new concept, and the second is to account for what makes conceptual engineering successful, either at the level of conceptual design, or at the level of implementation, as Simion and Kelp (2019) have recently argued. Nothing of what I've said goes against this latter reason for endorsing conceptual functions. Indeed, it might as a matter of fact be the case that we sometimes *do* want to preserve the original concept's function. Sometimes there is no problem with a given concept's function as such, and the problem is rather that the concept does not serve its function particularly well. I do not mean to object to any of this. My key claim is just that we do not *always* want to preserve a given concept's function, which is why we cannot use functions to demarcate the limits of acceptable revision.

Acknowledgements Earlier versions of this material were presented at the conceptual engineering and inconsistent concepts workshop in Bochum (10/2019), at Thomas Grundmann's epistemology research colloquium in Cologne (11/2019), at ECAP.10 in Utrecht (8/2020), at Christian Nimtz' research colloquium in Bielefeld (1/2021), and at Katja Stoppenbrink's conceptual engineering seminar in Münster (2/2021). I would like to thank the audiences of these events – especially Sarah Sawyer and Cyrill Mamin – for their helpful feedback. I would also like to thank Joachim Horvath, Sigurd Jorem, Eve Kitsik, Jennifer Nado and Mark Pinder, who commented on earlier versions of this paper. My work on this project was supported by a grant of Deutsche Forschungsgemeinschaft for the Emmy Noether Independent Junior Research Group 'Experimental Philosophy and the Method of Cases: Theoretical Foundations, Responses, and Alternatives' (EXTRA), project number 391304769.

Funding Open Access funding enabled and organized by Projekt DEAL.

References

Belleri, D. (2021). Downplaying the change of subject objection to conceptual engineering. *Inquiry*. https://doi.org/10.1080/0020174X.2021.1908161

Brun, G. (2016). Explication as a method of conceptual re-engineering. *Erkenntnis, 81*, 1211–1241. https://doi.org/10.1007/s10670-015-9791-5

Burgess, A., & Plunkett, D. (2013a). Conceptual ethics I. *Philosophy Compass, 8*, 1091–1101.

Burgess, A., & Plunkett, D. (2013b). Conceptual ethics II. *Philosophy Compass, 8*, 1102–1110.

Cappelen, H. (2018). *Fixing language. An essay on the foundations of conceptual engineering*. Oxford University Press.

Cappelen, H. (2020). Conceptual engineering. The master argument. In H. Cappelen, D. Plunkett, & A. Burgess (Eds.), *Conceptual engineering and conceptual ethics* (pp. 132–151). Oxford University Press.
Carnap, R. (1971/1950). *Logical foundations of probability* (2nd ed.). Chicago Ill.: Univ. of Chicago Press.
Chalmers, D. (2011). Verbal disputes. *Philosophical Review, 120*(4), 515–566.
Cummins, R. (1975). Functional analysis. *Journal of Philosophy*, *72*, 741–765.
Dorr, C., & Hawthorne, J. (2014). Semantic plasticity and speech reports. *Philosophical Review, 123*(3), 281–338.
Haslanger, S. A. (2000). Gender and race: (What) are they? (what) do we want them to be? *Nous, 34*, 31–55. https://doi.org/10.1111/0029-4624.00201
Haslanger, S. A. (2006). What good are our intuitions? *Aristotelian Society Supplementary, 80*, 89–118. https://doi.org/10.1111/j.1467-8349.2006.00139.x
Haslanger, S. A. (2021). How not to change the subject. In T. Marques & A. Wikforss (Eds.), *Shifting concepts* (pp. 235–256). Oxford University Press.
Houkes, W., & Vermaas, P. E. (2010). *Technical functions. On the use and design of artefacts*. Dordrecht: Springer.
Koch, S. (2019). Cappelen, H. 2018. Fixing Language. An Essay in Conceptual Engineering. Oxford University Press. 224 pp. ISBN: 978-0-198-81471-9. *History of Philosophy & Logical Analysis*, *22*(1), 248-256.
Knoll, V. (2020). Verbal disputes and topic continuity. *Inquiry*. https://doi.org/10.1080/0020174X.2021.1908161
Manne, K. (2018). *Down girl. The logic of misogyny*. Oxford University Press.
Millikan, R. G. (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT Press.
Millikan, R. G. (1993). *White queen psychology and other essays for Alice*. MIT Press.
Nado, J. (2021). Conceptual engineering, truth, and efficacy. *Synthese, 198*, 1507–1527. https://doi.org/10.1007/s11229-019-02096-x
Neander, K. (1991). The teleological notion of "function". *Australasian Journal of Philosophy, 69*, 454–468. Reprinted in Allen et al. (1998, pp. 313–333).
Pinder, M. (2020). On Strawson's critique of explication as a method in philosophy. *Synthese, 197*(3), 955–981. https://doi.org/10.1007/s11229-017-1614-6
Pinder, M. (2021). Conceptual engineering, metasemantic externalism and speaker-meaning. *Mind, 130*(517), 141–163. https://doi.org/10.1093/mind/fzz069
Plunkett, D., & Sundell, T. (2013). Disagreement and the semantics of normative and evaluative terms. *Philosophers' Imprint, 13*(23), 1–37.
Prinzing, M. (2018). The revisionist's rubric: Conceptual engineering and the discontinuity objection. *Inquiry, 61*(8), 854–880. https://doi.org/10.1080/0020174X.2017.1385522
Quine, W. V. O. (1960). *Word and object*. MIT press.
Richard, M. (2019). *Meaning as species*. Oxford University Press.
Sainsbury, R. M., & Tye, M. (2013). *Seven puzzles of thought and how to solve them*. Oxford University Press.
Sawyer, S. (2018). The importance of concepts. *Proceedings of the Aristotelian Society, 118*(2), 127–147.
Sawyer, S. (2020). Thought and talk. In H. Cappelen, D. Plunkett, & A. Burgess (Eds.), *Conceptual engineering and conceptual ethics* (pp. 379–395). Oxford University Press.
Simion, M., & Kelp, C. (2019). Conceptual engineering, function-first. *Nous, 54*(4), 985–1002.
Sterken, R. (2020). Linguistic intervention and transformative communicative disruptions. In H. Cappelen, D. Plunkett, & A. Burgess (Eds.), *Conceptual engineering and conceptual ethics* (pp. 417–434). Oxford University Press.
Strawson, P. F. (1991/1963). Carnap's view on constructed systems versus natural languages in analytic philosophy: The two methods. In P. A. Schilpp (Ed.), *The philosophy of Rudolf Carnap* (pp. 502–519). La Salle/Ill.: Open Court.
Strawson, P. F. (1992). *Analysis and metaphysics: An introduction to philosophy*. Oxford University Press.
Thomasson, A. L. (2020). A pragmatic method for normative conceptual work. In H. Cappelen, D. Plunkett, & A. Burgess (Eds.), *Conceptual engineering and conceptual ethics* (pp. 435–458). Oxford University Press.