

Conceptual Engineering: For What Matters

SEBASTIAN KÖHLER 

Frankfurt School of Finance & Management, Germany
s.koehler@fs.de

HERMAN VELUWENKAMP 

University of Groningen, The Netherlands
h.m.veluwenkamp@rug.nl

Conceptual engineering is the enterprise of evaluating and improving our representational devices. But how should we conduct this enterprise? One increasingly popular answer to this question proposes that conceptual engineering should proceed in terms of the functions of our representational devices. In this paper, we argue that the best way of understanding this suggestion is in terms of normative functions, where normative functions of concepts are, roughly, things that they allow us to do that matter normatively (for example, things in virtue of which we have normative reasons to have these concepts). Not only does this introduce a novel view about functions to the literature. This proposal also fits more naturally than the alternatives with conceptual engineering as a normative enterprise, and it allows functions to play all of the explanatory roles assigned to them in the conceptual engineering literature. Our discussion of the explanatory advantages of normative functions also advances the understanding of the ways in which concepts can be authoritative, what this means for conceptual engineering, and highlights the importance of political philosophy for thinking about conceptual engineering in practice. Furthermore, the paper explicates what kind of role considerations about function can and should actually play in conceptual engineering.

1. Introduction

Conceptual engineering is the enterprise of evaluating and improving our representational devices. But how *should* we conduct this enterprise? One increasingly popular answer to this question proposes to take a *functional turn* (for example, [Queloz 2022](#), [Thomasson 2020](#), [Simion and Kelp 2020](#), [Eklund 2014](#), [Jorem 2022](#), [Riggs 2021](#)). According to the functional turn, conceptual engineering should proceed in terms of our representational devices' *functions*: 'to engineer a concept well, we must attend to its function' ([Thomasson 2020](#), pp. 440-1). How, though, should this suggestion be understood? This paper argues that conceptual engineering should focus on what we call *normative functions*,

which are, roughly, things that concepts allow us to do that matter *normatively* (for example, things in virtue of which we have normative reasons to have these concepts). This is an attractive account of conceptual functions that has not been articulated and championed in the debate. We show that this proposal fits more naturally than the alternatives with conceptual engineering as a normative enterprise, and that on this interpretation functions can play all of the explanatory roles assigned to them in the literature. Our discussion of the explanatory advantages of normative functions also substantially advances the understanding of the so far little-explored issue (raised by [Queloz 2022](#)) regarding the ways in which concepts can be authoritative and what this means for conceptual engineering. It also highlights the importance of political philosophy for thinking about conceptual engineering in practice. Furthermore, the paper explicates what role considerations about function can and should actually play in conceptual engineering.

The paper proceeds as follows: §2 introduces, briefly, the functional turn, what motivates it, and the most prominent views within it. We then turn to our account and why it offers the best version of the functional turn. The argumentation proceeds in two steps. As a first step, §3 introduces the account and offers initial motivation. We motivate the account by arguing that considerations based on the so-called ‘authority problem’, that [Matthieu Queloz \(2022\)](#) has recently advanced to motivate *concern-relative* functions, speak more naturally for normative functions. As a second step, §4 spells out the account’s explanatory benefits in more detail. To substantiate the claim that normative functions provide the best view on functions for conceptual engineering we show that normative functions can do the relevant explanatory work conceptual engineers have assigned to functions.

2. The functional turn

Conceptual engineering is the enterprise of evaluating and improving our repertoire of representational devices, where *representational devices* are words and concepts (by which we mean the constituents of thoughts). Conceptual engineering does not just take our repertoire of representational devices as given, but aims at *amelioration* and *innovation*: it critically assesses our existing representational devices to improve them (or eliminate them if they are in some central way defective and beyond repair), and considers which representational devices we could introduce to improve our intellectual lives.

Conceptual engineering raises tricky philosophical questions. One central question is this: how *should* we do conceptual engineering? Specifically, what should the *aim* of engineering a representational device be? The functional turn offers a (partial) answer to this question. Broadly speaking, its proponents hold that the *functions* of concepts determine how we should engineer them (Simion and Kelp 2020, Haslanger 2020, Thomasson 2020, 2021, Queloz 2022). This is taken to entail that we should not just focus on our representational devices' *representational* dimensions. Instead, we should take into account that these devices allow us to do all kinds of things, as they can have a variety of *functions*.

Focusing on functions is *prima facie* appealing, because conceptual engineers aim at ameliorating and improving our representational devices and engineering for functions seems to offer an attractive view of what this means (see, for example, Nado 2021). In addition to this, one original motivation for the functional turn was the idea that it would deal with the so-called 'Strawsonian challenge' (see Strawson 1963, Cappelen 2018, pp. 98-134). This is the challenge of showing that conceptual engineering does not just amount to changing the subject, but offers a better way to talk about the issues and questions we were concerned with. Here, the idea is that revising a representational device avoids changing the subject insofar as it preserves continuity of that device's function. Other benefits of the functional turn have also been highlighted. For example, functions are supposed to explain what we do when we engage in metalinguistic negotiation (Plunkett and Sundell 2013).

The functional turn itself, though, raises issues, the most fundamental being what Matthieu Queloz (2022, p. 1256) calls the 'function specification problem': if you suggest that conceptual engineering should focus on the functions of representational devices, you need to give an account of what *kinds* of functions we should focus on. This must be an account on which it is plausible that conceptual engineering should focus on such functions and which can deliver the promised theoretical goods.

One prominent way of responding to this problem is to appeal to *orthodox* understandings of 'function' found elsewhere in philosophy. For example, some (for example, Simion and Kelp 2020, Thomasson 2020) argue that we should understand 'functions' as *etiological* functions, where something's etiological function (for example, Millikan 1984, Wright 1973) is the effect that explains why it (or things of its kind) exists or persists. It has also been suggested (for example, Thomasson

2020, 2022) that ‘functions’ should be understood as *system* functions, where something’s system function (for example, Cummins 1975) is its contribution to the capacities of a system of which it is a part.

In recent work, Matthieu Queloz (2022) has argued that we should deviate from these orthodox understandings and opt for ‘concern-relative functions’, which he defines as follows:

A concept *X* has the [concern-relative function] of type *C* of producing effect *E* iff (1) users of *X* have among their present concerns—their needs, interests, desires, projects, aims, and aspirations—a concern of type *C*; (2) under propitious circumstances, applications of *X* produce *E*; (3) *E* stands in an instrumental relation to the concern of type *C*, which is to say that under propitious circumstances, producing *E* contributes to the satisfaction of a concern of type *C*. (p. 1261)

Concerns here can be anything from moral to self-interested concerns, but they have to be concerns that people have prior to engineering (p. 1253). However, Queloz only wants to include concerns of users that they themselves would endorse on reflection (p. 1260).

We agree that function-based conceptual engineering should not understand functions in terms of orthodox approaches. However, Queloz’s account is not the correct answer to the function specification problem either. Instead, the best way to understand ‘functions’ is as *normative* functions: a novel understanding of functions that has so far remained unnoticed in the conceptual engineering literature (and philosophy more generally). The remainder of this paper substantiates this thesis. The next section starts by pointing out that the considerations Queloz uses to support his account actually point in a different direction and then introduces our account.

3. Toward normative functions for conceptual engineering

Queloz’s argument for the claim that conceptual engineering should prefer concern-relative functions over orthodox functions starts with the observation that conceptual engineers face the following challenge:

The authority problem: ‘why should we grant a novel or even a merely revised concept the power to shape and guide our thought and conduct? The problem is that of explaining, to the people who are urged to adopt an engineered concept, why they have reason to adopt this concept and structure their affairs in terms of it.’ (p. 1247)

Queloz argues that the authority problem supports the functional turn. This is so because we cannot deal with this challenge, if we hold that the reason to adopt a concept is solely due to its theoretical virtues compared with its alternatives, for example, its being more precise or consistent. For example, suppose our current concept PERSON (we follow the convention of denoting concepts using small capital letters) was less consistent than an alternative we could engineer, but that the latter licenses infanticide. Here the greater *theoretical* virtues of the latter clearly give us little reason to accept it over our current concept. Queloz's diagnosis is the following: 'The display of theoretical virtues such as consistency and precision is not the only thing we want from our concepts, and this makes it an open question whether vague, indeterminate, open-textured, or contested concepts might not sometimes serve us best' (p. 1252). If this is true, the authority problem *motivates* taking the functional turn.

However, whether a functional approach can deal with the authority problem depends on how it answers the function specification problem. And Queloz argues that accounts appealing to etiological or system functions cannot adequately answer it. This is unsurprising: the factors that explain why a concept(s) use) persists in a population of thinkers, or its role within a system, do not inherently provide us with reasons to adopt that concept. However, concern-relative functions seem able to deal with the authority problem: if a concept serves a concern-relative function for us, we at least *prima facie* have reason to use it, because it contributes to the satisfaction of one of our concerns. This is why we should conduct conceptual engineering in terms of concern-relative functions. Or so Queloz argues.

We think, though, that this argument does not motivate conducting conceptual engineering in terms of concern-relative functions. First, note that there is an ambiguity in Queloz's formulation of the authority problem. It has at least two readings. On one reading—what we will call the 'reflective endorsement' reading—it is the challenge of offering reasons to those who are urged to adopt a concept that they can reflectively endorse *on their own terms*. Here, we are looking for concepts that are authoritative in the sense that those who adopt them can accept that they have reasons to use them *from their own perspective*. On the other reading—what we will call the 'good reasons' reading—the authority problem is the challenge of offering those who are urged to adopt a concept good normative reasons to use it. Here, we are looking for concepts that are authoritative in the sense that there *are* good normative reasons for the relevant people to adopt them—irrespective of whether they can

accept that this is so from their own perspective. So, there is a difference here as to whether we demand that we must be able to get the relevant people to *understand* upon reflection that they have such reasons.

These readings can come apart, clearly, because we cannot guarantee that everyone can get to a point where they can see what reasons they have. Now, Queloz's concern-relative functions deal with the authority problem on the reflective endorsement reading. While this might be helpful for some theoretical and practical purposes, it does not motivate the idea that we should conduct conceptual engineering in terms of concern-based functions. Achieving authority in this sense does not matter in itself for how we should engineer our concepts. What should matter for conceptual engineering is whether the concept does something for us due to which we have good normative reasons to have the concept. *This* is what we should be engineering *for*. It should not matter in itself whether everyone who uses the concept is able to accept that they have such reasons on their own terms. After all, people's normative perspectives might be blind to the relevant normative reasons, for example, because their normative perspective is perverse, corrupt, or otherwise harmful or bad. And this blindness might not be something that would disappear on reflection. But neither of these facts should be relevant for what concepts they should use.

To make this point clear, consider an example: suppose there exists a community of speakers who use the concept MARRIAGE in such a way that it excludes the possibility of married same-sex couples. Now imagine that a proposal emerges within this community to engineer MARRIAGE to include same-sex marriage, and that there is decisive normative reason to do so. However, assume further that a subset of this community has concerns that would not be served at all, and might even be undermined, by this revision. In this case, we could not meet the authority problem on the reflective endorsement reading with regards to that group: we could not give them *any* reasons to adopt the revised concept on the basis of concerns they can reflectively endorse. But it seems hard to see why the community should not nevertheless engineer the concept in the suggested way.

Of course, we might assume that in real life such people always *have* concerns we can appeal to—such as a desire to live in a peaceful, stable community with others—so that we can give them reasons they can accept from their own perspective to use concepts that there are good normative reasons to use. But this response is beside the point: what the example highlights is that it does not matter in itself whether we can give everyone reasons for adopting a concept that they can endorse on the

basis of their reflectively endorsed concerns. What matters is whether there are good normative reasons *full stop* to adopt it.

Naturally, in the situation we are describing, the normative situation could be more complex: there could be strong normative reasons for the resistant group to use a different concept or relevant rights that the revision would violate. Nevertheless, while such considerations can be relevant for conceptual engineering (as we discuss below), they reinforce our central claim: it's the *normative* considerations that matter, *not* the ability to accommodate reflective concerns of all groups. If we stipulate such considerations *away* (for example, because the group's reflective concerns are *bad*), the reflective endorsement reading of the authority problem loses its relevance.

Naturally, when there is a group that we cannot give *any* reasons to adopt the revised concept on the basis of concerns they can reflectively endorse, it might be hard to *get* that group to use that concept. After all, if they do not regard themselves as having any reason to adopt the concept, they might never do so. However, this does not motivate caring about the reflective endorsement reading of the authority problem, because then it would be a version of the implementation challenge that makes additional robust assumptions about how that challenge must be met. The implementation challenge is the challenge of explaining how successful conceptual engineering can be achieved, and whether it is even possible (for example, [Cappelen 2018](#), [Koch 2020](#), [Deutsch 2020](#), [Nado 2020](#), [Jorem 2021](#)). Suppose we could meet that challenge in a way that does not require all members of the relevant community to have a reason that they can reflectively endorse for adopting the concept in question. So, suppose that it is possible to get people to adopt a concept, even if they do not have the relevant sort of reason. This is not implausible: the way conceptual change actually works suggests that people can have concepts in their repertoire that they have no reason to use in the light of their own reflective concerns. If this is true, though, the worry that it might be hard to get our resisting group to use the concept is no worry at all, but the real authority problem also seems left open. After all, the question *why* thinkers ought to grant the concept authority in their thought is not yet addressed. So, the worry here can only be a worry if we make specific assumptions about how the implementation challenge has to be met—assumptions we see little reason to grant at this point.

Of course, more can be said about the reflective endorsement reading (and we go a little deeper into this issue in §4.1). Let's assume

for the moment, though, that we are correct and that it is the authority problem's good reasons reading that matters for how we should conduct conceptual engineering. This suggests a different account of functions.

3.1 Normative functions

Our view is that conceptual engineering should be conducted in terms of *normative* functions. Normative functions of concepts are just what they are *for* in the sense that a concept's normative function is the effect of its use in virtue of which we ought (in the robustly, authoritatively normative sense) to have it in our conceptual repertoire. While there are different ways of formulating the view (for example, in terms of values, what we ought to do, and so on), we will focus on the following (for illustrative purposes):

Normative Function: A concept *X* has the normative function to produce effect *E* if and only if (1) in a relevant range of circumstances *C*, applications of *X* produce *E* and (2) users of *X* have *normative reason* to deploy *X* in thought and language *because X* produces *E* in *C*.

Some clarifications are in order. First, when we refer to 'normative reasons' we mean considerations that count in favour of actions, reactions, or attitudes (Scanlon 1998), in an authoritatively normative sense (McPherson 2018). On our view, this includes moral reasons tied to ethical considerations and values, epistemic reasons related to knowledge and understanding, prudential reasons connected to personal well-being, and considerations from all other authoritative normative domains (for example, the aesthetic, but not etiquette). So, on our account, normative functions can be grounded in *epistemic* reasons (for example, we have a normative reason to have a concept because it enhances our understanding), *moral* reasons (for example, we have reason to have it because it leads to *fairer* outcomes when deployed in relevant circumstances), and so on.

Second, the reason that is relevant for a concept's normative function is not a reason to use it on any particular occasion. Instead, it is a reason to *have* that concept in one's repertoire, that is, to be a person that conceptualises the world in terms of it. So, the idea is that a concept has a normative function to produce an effect if, when people use that concept, it produces that effect in a relevant range of circumstances such that this provides a reason to have that concept in one's conceptual tool kit.

Third, when we say that ‘in a relevant range of circumstances’ using the concept produces the effect, we just mean that in the concept’s normal use the effect will occur sufficiently often for us to have normative reason to adopt it.

Fourth, our main disagreement with the proponent of concern-based functions is in terms of what we should engineer for: effects that give us normative reasons or effects that speak to our concerns. Interestingly, with additional commitments, we can bring the views closer together. This is so on some internalist metaethical views about normative reasons (for example, [Schroeder 2008](#), [Williams 1980](#)), on which a person has normative reason to ϕ only if ϕ -ing is in some instrumental way connected to her motivational set. This is also so on unexplored weaker internalist views (not committed to externalism generally) on which a person has normative reason to *adopt a concept* only if doing so connects instrumentally to her motivational set. If any such internalist view is correct, our view and the concern-based view would single out the same features of concepts as their functions. The reason for this is that we would have normative reasons to use concepts only if this contributes to the satisfaction of a speaker’s concern. We would then still disagree on what it is that determines what we should engineer for, but would arrive at similar answers as to what the functions are.

In fact, however, we think that internalism about reasons for concept use is false. This is motivated by examples in which people with bad or perverse incentives figure, which are instances of what [Finlay and Schroeder \(2017\)](#) call the ‘Central Problem’ of internalism. We therefore think that we should engineer on the basis of normative reasons, *and* that externalism about normative reasons for concept use is correct. This also highlights that our disagreement with Queloz could, possibly, be grounded in a disagreement about internalism or externalism about reasons—though a focus on concern-based functions *is* perfectly compatible with externalism and could be motivated by arguments independently of internalism (for example, in the way we suggest in §4.1).

These last remarks highlight that we should add some final clarifications concerning the view about reasons we assume. We take it that the details of the best view still need to be worked out. While we cannot do so here, we will make explicit some assumptions we make, flagging optional ones, and highlighting where more work is required. For the most part, we will adopt a somewhat orthodox conception of (external) normative reasons (see, for example, [Parfit 2011](#), [Raz 2011](#)). First, we assume that we have normative reasons to adopt concepts because of facts about them, such as what their effects are when used, and that

values (which themselves are not (always) grounded in people's concerns) explain *which* facts are reason-giving. While our account should be compatible with other views about reasons (for example, reasons first views—see, for example, Scanlon 1998), we use this view because it is straightforward.

Second, we assume, as indicated above, that there are reasons to adopt concepts across all authoritatively normative domains, each grounded in distinct values. So, for example, if a concept's use enhances our knowledge, we've got reason to adopt it due to the (epistemic) value of knowledge. If using a concept leads to reduced discrimination against individuals, then the (moral) value of fairness or justice provides a reason to adopt it. This pattern applies across various domains, although there will likely be constraints of what counts as reasons of the *right kind* for adopting a concept (see Simion 2018 for discussion). Nevertheless, we assume that reasons from all authoritatively normative domains can meet these criteria.

Third, we assume that sets of reasons for adopting one version of a concept over another can be *weighed* (even across normative domains) and that there can be decisive reasons to adopt a concept when the weight of the set of reasons for adopting it are weightier than that of the sets of reasons for adopting any alternative. These assumptions about reasons are somewhat controversial (for example, Raz 1986, Chang 1997) and certainly need to be modified and qualified (see Cullity 2018). However, more complicated pictures should be compatible with our view, as long as principled deliberation about what concepts to have is, generally, possible.

Fourth, we assume that reasons for concept use can be *discovered*, both in principle and as a general rule. How? Through the standard methods of normative inquiry, the very methods conceptual engineers already employ to evaluate whether concepts can be improved along relevant normative dimensions. Naturally, if scepticism about normativity is warranted, this assumption will be false. However, we see no reason to grant such scepticism.

These assumptions are relatively common for reasons for *action*. What makes (some of) them *unorthodox* here is that we take them to hold for *adopting a (version of a) concept*. However, we take our picture to fit common assumptions in conceptual engineering, especially amongst those who take the functional turn. After all, conceptual engineers assume that we can *deliberate* about what concepts to adopt and that we *evaluate* our concepts on the basis of different dimensions—moral, political, epistemic—during that deliberation. Moreover, they

assume that the conclusions of these deliberations can influence what concepts we use. These assumptions fit naturally with the picture we've sketched and it seems that any plausible picture of those assumptions will have something like the shape of what we assume. Note that we do not assume that *adopting a concept* is an intentional action. We just assume that it is an attitude that is responsive to reasons of the relevant kind.

Undoubtedly, more could be said on all of this, which we leave for further debate. One important lesson to take away, though, is that it is worthwhile to think about these metaethical debates specifically in the context of conceptual engineering. For example, it will pay to think more explicitly in what ways our conceptual repertoire is responsive to reasons, as that will help us understand conceptual engineering more generally. Similarly, it would be valuable to investigate further whether there are special reasons to endorse internalism about concept use in particular. While our inclination is to reject this, we acknowledge the possibility for philosophical progress in this area.

With these clarifications in place, let us turn to the role of normative functions in conceptual engineering. Normative functions are things that a concept allows us to do that are in fact important in a robust normative sense (not just important *to us*). To explain how such functions figure in conceptual engineering, let us first clarify what we take to be the intended role of functions on any approach that takes the functional turn. Conceptual engineering, like any form of engineering or design, is engineering *for* something. This idea itself is a bit vague, though. For example, suppose we are designing a knife. There are many different things we might be designing *for*: how comfortable the grip is, how safely it can be used by children, or how nice it looks. However, there is one consideration that *guides* and *structures* our design effort in a specific way: this is the consideration of what the knife's point is in the first place, what it is *for* in this fundamental sense. Is it to cut things? Cut specific sorts of things, such as vegetables or bread? Is it to be used as a stabbing weapon? And so on. We take it that 'functions' in the context of conceptual engineering are supposed to play the equivalent role for concepts: a concept's function is that consideration that guides and structures our engineering effort in this way. In our picture, it is the concept's *normative* function that is to play that role. To be more precise: it is a concept's *most important* normative function that should play this role. A concept's most important normative function is, thereby, that normative function where the reasons we are provided with are strongest compared to all alternative normative functions.

Does every concept have a normative function? On our account, a concept has a normative function insofar as we have normative reasons to have it in our conceptual repertoire, due to some effect it has when deployed in relevant circumstances. So, concepts would lack normative functions if they are so flawed or useless that there are *no* normative reasons to have them in one's repertoire, or if there is no such thing as a normative reason to have a concept in one's repertoire in the first place. The former is certainly possible, but such concepts would only be relevant for conceptual engineering to the extent that it is good to identify and either eliminate or avoid adding them to our conceptual repertoire. We take it that the latter is ruled out by the aforementioned assumption underlying conceptual engineering that we can influence what concepts we use. If we can change what concepts we use, we can also have reasons to opt for one change over another.

Does every concept have a *most important* normative function? We assume that reasons can be weighed. However, there could still be cases where no reason is stronger than all others: multiple reasons could be equally strong, for example, or the issue might be indeterminate. In those cases, there won't be one most important normative function. What does our view say for those cases? In such cases, engineers should still engineer *for* all equally important functions. However, what this means in concrete cases is tricky, because it depends on the details: it could be best to have multiple, but different versions of the same concept or to engineer for a version that serves multiple functions equally well. In what follows, we ignore this complication and speak only as if concepts have one most important function, but this caveat should be kept in mind.

So, we suggest that conceptual engineering's aim should be to engineer for a concept's most important function, ideally in a way that serves that function best (in the sense that no other version of the concept would serve that function better). Note that what a concept's most important normative function is can depend on the context and can differ across time or groups of people. As in the case of the knife, it may be that for different contexts, different ways of weighing the considerations lead to different optimal solutions. For example, a cutlery drawer may contain a very sharp knife for filleting fish, while a less sharp knife is better suited for use by children. Similarly, it may be good to have different versions of a concept at our disposal for use in different contexts or for different groups to have different concepts. This means that the normative function account can explain why concepts can have a wide range of functions (we come back to this in §4.1).

Note that it is also compatible with our view that conceptual engineering has to take other things into account, besides a concept's most important normative function: less important normative functions, for example. Moreover, there might be things we have normative reasons to avoid; for example, we do not want a racist concept of personhood. However, this is compatible with normative functions' role. Our suggestion is not that we should *only* engineer for what would best serve a concept's most important normative function. Rather, the most important function should be the (suitably constrained) *focus* of engineering, what the concept is *for* that should guide and structure our engineering efforts. Note also that our account is compatible with many different ways in which conceptual engineering can proceed: we can, for example, start by thinking about what important things a concept could do for us and then invent a concept that does that job. Or we could consider what our current concepts can do and then revise them accordingly.

So, this is our suggestion: 'functions' in conceptual engineering should be normative functions. This suggestion seems highly compelling, particularly because conceptual engineers are typically focused on *evaluation* and *improvement*. Specifically, conceptual engineers are normally not just interested in improving our concepts' theoretical virtues by, for example, making our concepts consistent, removing vagueness, and so on, but in improving our conceptual tool kit along all kinds of normative dimensions (see, for example, the examples in Cappelen 2018, pp. 9-27 or Marques 2020): moral, prudential, or more robustly epistemic. It also resonates especially well with those who endorse the functional turn in conceptual engineering. By focusing on the *important* things concepts allow us to do, that is, their normative function, we align our engineering efforts with what is actually valuable. It makes less sense to focus on past effects or roles within larger systems (that is, etiological and system functions), as these are often not normatively significant. In fact, even engineering for our concerns seems like a bad place to take as fixed when engineering our concepts, given that our concerns might be bad. When we engineer, our deliberations should not be *inward*: we should not ask 'Why do we *care* about adopting the concept?' Instead, our deliberations should be *outward*: we should ask 'What *reasons* do we have to adopt the concept?'—unless the former question is asked as a proxy for the latter.

Normative functions are, consequently, a more natural candidate for filling the role that functions are supposed to play in conceptual engineering compared to other kinds of functions. Of course, other kinds can still be relevant for conceptual engineering. First, concern-relative

functions can be relevant as proxies or because our concerns themselves can sometimes give us normative reasons. Second, etiological or system functions *can* be what we should engineer for. For example, we might have good reasons to promote what explains why a concept's use persists. And in such cases, it might well then be that a concept's etiological function is (a part of) its normative function. The scenarios where this is the case might explain the initial plausibility of such views. However, by looking at cases where orthodox and normative functions come apart, we can see that normative functions do the real explanatory work.

Furthermore, there is one way in which it *matters*, at least for etiological functions, that engineered concepts acquire them. It is plausible to require that engineered concepts be *stable*, in the sense that their use will be taken up and persist in the relevant population. But stable concepts acquire etiological functions: there will be causal effects that explain their persistence. So, conceptual engineers should aim to create concepts capable of acquiring etiological functions. However, it would be wrong to infer from this that etiological functions should, then, be those that *guide* and *structure* conceptual engineering efforts in the way we have explicated above (as opposed to the claim defended by, for example, [Simion and Kelp \(2020, p. 990\)](#), who suggest that conceptual engineering's aim should be that what a concept was designed for becomes its etiological function). This role should still be played by the most important thing a concept allows us to do, that is, its normative function, irrespective of whether this will eventually be the concept's etiological function. While it would be nice if a concept's normative function *also* explained why people employ it, there is no guarantee this will happen and little reason to aim for it. After all, many different things can come to causally explain why people (continue to) use a concept (for helpful examples, see [Nimtz's \(2021\)](#) work on solving the implementation challenge) and for different people this might be different things. Conceptual engineers must take *some* considerations of implementation into account at the design stage (for example, whether the concept is even useable and how difficult it is to use), but their *guiding* concern should be the concept's most important normative function.

4. Benefits of conceptual engineering for normative functions

At the end of the last section we offered initial reasons for finding conceptual engineering in terms of normative functions attractive. In this section we argue that normative functions can do the explanatory work

assigned to functions. Over the course of the paper we have encountered four explanatory tasks for conceptual engineers that functions are supposed to be able to shoulder. First, to offer an attractive account of what conceptual engineering should aim for, one that identifies features to engineer our concepts for that would lead to genuine improvements (for example, [Nado 2021](#)). Second, to allow us to deal with the authority problem, the challenge of explaining why people have reasons to use an engineered concept ([Queloz 2022](#)). Third, to help us with the Strawsonian challenge, the challenge of explaining how a proposed revision does not just amount to changing the subject, but provides an improved way to talk about the issues that concern us (for example, [Thomasson 2020](#)). Fourth, to explain what we are doing when we engage in metalinguistic negotiations (for example, [Plunkett and Sundell 2013](#)).

It is noteworthy that although different proponents of the functional turn allude to different explanatory advantages, no one has argued that a single understanding of ‘function’ can do all of this explanatory work. In fact, [Jared Riggs \(2021\)](#) has argued that there cannot be a notion of function that can do this. We take it that the last section made clear that normative functions accommodate the first feature assigned to functions. In what follows, we take the remaining three explanatory tasks in turn. What our arguments will show is that Riggs is mistaken: *normative functions can* do all the relevant explanatory work.

4.1 Normative functions and the authority problem

As we spelled out above, the authority problem has at least two versions: first, the challenge of offering reasons to those who are urged to adopt a concept, which they can reflectively endorse on their own terms—the ‘reflective endorsement’ reading—and second the challenge of being able to offer them good normative reasons to adopt the concept—the ‘good reasons’ reading. Above we’ve suggested that the problem only matters in terms of the good reasons reading. In this section we go into more detail. First, we investigate the reflective endorsement reading more closely: what motivates it and what implications this has for normative functions. Second, we explain how and why the normative functions view deals with the authority problem on the good reasons reading. Third, we uncover yet another reading of the authority problem, discussion of which will clarify more systematically what work functions can do in determining how we should engineer our concepts.

What reason is there to focus on the reflective endorsement reading as the core version of the problem for conceptual engineering? Here is what we take to be a *prima facie* attractive line, based on considerations

in virtue of which Queloz himself (2022, pp. 1264-1272; see also his other work, for example, Queloz and Bieber 2021) seems to think that it is an important problem for conceptual engineers: it is an undeniable fact that people disagree on normative matters. Indeed, even reasonable people will disagree about all kinds of normative questions—what Rawls (1993) calls the ‘fact of reasonable pluralism.’ However, one might think that given such disagreement—and given that we need to concede that we are just as fallible on normative matters as others—we owe people a justification they can accept on their own terms, when asking them to adopt certain concepts. If we are unable to do so, we can only make people adopt new concepts in, for example, coercive or manipulative ways that are incompatible with a ‘humanistic commitment to the sovereignty and autonomy of human beings’ (Queloz 2022, p. 1269). Think about the example we used above, of the group that is resistant to a reform of MARRIAGE. We might urge that there is *this* important sense in which the reformed concept lacks authority over that group. If we do not offer them something they can accept as a reason from their own perspective, it seems that the only ways to make them adopt the reformed concept would violate their ‘sovereignty and autonomy.’ Thus, we might have to conclude that *we* (those with concerns served by reform) should introduce the reform for *us*, but that *they* should use a different concept.

Following this, it might be argued that it is because of this commitment to respecting people’s diverse normative perspectives that we should take the reflective endorsement reading seriously as articulating a requirement on how we should conduct conceptual engineering. And it speaks specifically for Queloz’s view that it satisfies it. Of course, due to reasonable pluralism, there is also pluralism about conceptual authority, if authority is understood along these lines, in the sense that different concepts will be authoritative to different groups in the sense at issue. Again, Queloz’s view naturally accommodates this, which we might think is also an advantage. After all, it is quite natural to think that whatever sense of authority matters for conceptual engineering, there won’t just be one set of concepts that is authoritative for all human beings (Queloz 2022, p. 1270). Different groups are situated in different circumstances and, hence, different groups can have different concepts without a failure of rationality.

We think that this line of response highlights important things for conceptual engineers to think about, issues that so far have not gained sufficient attention in the debate. Note that the response is based on a distinctive first-order *normative* commitment with respect to how

conceptual engineers *should in practice* respond to the fact that different things matter to different people. This is the commitment that in such circumstances, conceptual reform must introduce only concepts for which everyone asked to adopt them can accept that they have reasons to do so from their own normative perspective. This normative commitment seems in tension with our view, but fits naturally with Queloz's. However, we think that the *normative* lessons to draw from the response's starting observations about authority are different, that they do not support Queloz's view, and that the normative functions view has plausible implications in relation to *these* kinds of worries about authority.

First, the commitment seems plausible, because it seems familiar from political philosophy: it is a version of the demand of being able to *justify*—in the face of normative disagreement—a particular measure that has significant implications for people's life to those who will be subject to it (Vallier 2022, §2.3). However, it is crucial to note that one important way of meeting this demand is by focusing on *institutional mechanisms* (or *procedures*) for resolving normative disagreements. For example, the demand for justification at the heart of political philosophy does *not* mean that a particular tax policy must be justified to everyone subject to it on grounds they can accept from their own normative perspective—except in the sense that the mechanism that generates this policy can be justified. If we follow this line of reasoning, normative disagreement should at most motivate conceptual engineers to think about what mechanisms for generating and implementing particular conceptual engineering proposals would be justifiable, not whether any particular proposals can be justified to all those asked to use the concept without recurrence to such a mechanism.

This reveals another ambiguity in the authority problem, namely, whether it is sufficient to meet it by appealing to such *procedural* reasons. We take it that the authority problem's reflective endorsement reading, *especially if it is motivated along the lines suggested here*, is more plausible if it *allows* such reasons. However, what it then seems to motivate is at most that the practical implementation of conceptual reform has to happen against an institutional background that can be justified to those affected by it. But this does not support the view that what we should engineer *for* is something that provides reasons that people can reflectively endorse. In fact, plausibly, within the relevant institutional settings, what conceptual engineers should engineer *for* is what makes the concept important in the first place. So, the lesson to take away is not that we need a different account of functions from the one we suggested,

but that more work needs to be done to think about justifiable institutional arrangements for translating conceptual engineering into practice. While a complete account needs to address this, it is hard to see why plausible accounts on this matter should be incompatible with the normative functions account.

Second, on one of the most prominent and plausible interpretations of this demand for justification in political philosophy (even for the relevant institutional mechanisms), it concerns only disagreements between *reasonable* people, and so at most requires that we offer *reasonable* people reasons they can accept on their own terms (Rawls 1993). This is a weaker commitment than the one underlying the authority problem's reflective endorsement reading and provides no reason to take that reading seriously. After all, what could be justified to reasonable people need not be something everyone can be brought to reflectively endorse from their own normative perspective. For example, consider T.M. Scanlon's (1998) view of morality, which requires us to consider what could be justified to reasonable people, but where such justification depends on people's concern-*independent* reasons. In fact, on some accounts of what makes someone 'reasonable', being asked to give someone reasonable something they can accept as a reason from their own normative perspective just *is* offering them good normative reasons. Of course, in this sense, normative functions can meet the authority problem.

Third, there are also distinctive *ethical* reasons to care about authority in the specific sense that Queloz's view offers: in the face of pluralism, it is ethically preferable when we can offer people reasons to accept a conceptual reform they can reflectively endorse *over and above* the legitimacy of some relevant disagreement-resolving mechanism. For example, your family might have an agreed-upon method for resolving conflicts with regards to what to do on the weekend. Still, when a decision is made it seems ethically preferable to give everyone involved a reason they can endorse to accept the resolution that is independent of the fact that it was arrived at with an agreed-upon method. The same can hold for communities and decisions about what concepts to adopt. Thus, it *can* make sense to take into account whether and to what extent people can be given such reasons to adopt a concept. However, surely there is a limit to the ethical desirability of giving such reasons. Furthermore, that desirability is not sufficient to generate a *requirement* to take such reasons into account—the requirement that underlies the normative commitment supporting the reflective endorsement reading. Specifically, it is hard to see how such desirability would support the

view that this is the thing we should be engineering *for*, the concept's function.

As noted above, in actual practice, there will always be many different reasons that speak for or against a conceptual engineering proposal. Why does the concept matter? How many people will understand and be responsive to it? What is the concept's uptake and how far can we deviate from engineering it for the thing in virtue of which it matters to increase our uptake? And so on. But when conceptual engineers ask after a concept's function or practical point, they are not after all of these reasons. What they are interested in is what the concept can do for us that really matters. But when we pose *this* question it seems irrelevant whether we can give everyone reasons to adopt the concept that they can reflectively endorse.

Fortunately, the normative functions view has the resources to accommodate all of this: in cases where we have strong (for example, ethical) reasons to offer people relevant reasons they can endorse, the view can allow that those are relevant for engineering. However, the view also accommodates the fact that even in those cases, these considerations do not determine what we should engineer for in the sense of being the guiding consideration provided by the concept's function. In fact, the view naturally provides guidance for how to deliberate on these matters. We need to weigh how important the normative function is and how much the ability of a concept to perform that function would be constrained by taking into account other relevant reasons.

So, once we unpack them further, the considerations that might be taken to support the relevance of the reflective endorsement reading in fact fail to do so. Once we think about the motivating considerations in the different light we have suggested, they are no longer in tension with the normative functions view. Of course, more could and should be said about these issues and maybe there are other ways to defend the reflective endorsement reading. We leave this for further debate.

Let us finish this discussion by talking, briefly, about pluralism about conceptual authority. As should be clear, our account—as opposed to Queloz's—does not assign any special importance to pluralism about conceptual authority in the sense that reflective people differ in what concepts they can regard themselves as having reasons to use—pluralism about authority in the *reflective endorsement* sense. That different people *regard* different concepts as authoritative matters, on our account, only in the ways just discussed. However, we do not think that this is problematic, because the plausible observation about pluralism

we flagged above does not licence a focus on that reading and can be accommodated by the normative functions account.

The plausible observation was this: it *is* natural to think that there won't just be *one* set of concepts that is authoritative for *all* humans. It must be acknowledged that different people can have reasons to have different concepts and, hence, can adopt different concepts without rationally failing. The normative functions account can accommodate this. After all, what it is that a concept allows people to do that gives them (most) reason to use it can plausibly differ between groups. For example, in a society marked by racial inequalities, the most important thing race concepts allow people to do might be to help them eliminate such inequalities. But, in a society without such inequalities, there might be nothing important such concepts do. So, in the first society a specific concept of race would be authoritative for people, but no such concept would be authoritative for people in the second society. Furthermore, even for groups for which the most important normative function is the same, it can be true that they have reasons to use different versions of that concept. For example, in a post-AI society, people might need a concept of responsibility that avoids so-called 'responsibility-gaps' (Matthias 2004). But a different concept might be better suited in a pre-AI society. So, the normative functions account is compatible with and straightforwardly accommodates the observation that different people can have reasons to use different concepts in virtue of what the concept does for them and as such can *legitimately* have different views about what concepts to use. But this is in virtue of normatively relevant differences in circumstances, not because of differing concerns.

Let us now turn to the authority problem on the good reasons reading. The discussion so far should make clear that etiological and system functions, as well as concern-relative functions, perform poorly in relation to this challenge. None of these functions guarantee that if a concept possesses that function, people have good normative reasons to use it. In fact, what explains why a concept's use emerges or persists, what role it plays in some larger system or what concerns it serves, might give us very good reasons *not* to use the concept. Thus, none of these functions are plausible candidates for what we should be engineering *for* and none of these functions can plausibly meet the challenge. Normative functions, on the other hand, can deal with the authority problem on the good reasons reading: normative functions *are* things in virtue of which we have normative reason to have the concept in our conceptual repertoire. In fact, on our suggestion the function that should be engineered for has a very strong sort of authority in this sense, as it is the

thing in virtue of which we have most reason to have the concept. Of course, some people might not be positioned to recognize that there is such a reason. As we've discussed above, this can raise a host of issues that are separate from the authority problem on the good reasons reading—issues about which proponents of the normative functions account have plausible things to say, even if their full exploration will have to be left for future work.

We should flag, though, that the authority problem's formulation allows for an additional ambiguity, which leads to another reading normative functions *cannot* address. On one reading, the authority problem is concerned with *pro tanto* reasons. Queloz often talks as if he is concerned with giving people *a* reason to adopt a revised concept. However, this seems a bit weak. Consider again the more consistent concept PERSON that licenses infanticide. That the concept is more consistent *is a pro tanto* reason to adopt it. This reason is nevertheless defeated by the much stronger reason not to use a concept that licenses infanticide. This might suggest a different reading of the authority problem on which it requires the reasons to be *decisive* or *sufficient* (we'll focus on *sufficient* reasons, but everything we say should hold for decisive reasons too).

To see how normative functions fare in relation to this version of the challenge, we must distinguish two questions, which will help clarify some issues and reveal yet another dimension of the authority problem. First, what reason do we have to adopt a concept in the first place? Second, what *version* of the concept should we adopt? Take the example from above: on the one hand, there is the question whether to adopt the concept PERSON (or to phrase it differently, whether to adopt *any* version of the concept PERSON). Then, though, there is the question of what version of the concept to adopt: should we stick to our inconsistent version or opt for a more consistent one? Normative functions are relevant to both questions, but in different ways.

Normative functions straightforwardly relate to the first question: they give us the reasons why we should opt for or retain the concept in the first place. But they do not necessarily give us *sufficient* reasons to do so. This depends on the strength of the reasons that pertain to the concept's most important function, but also on relevant defeaters. Suppose, for example, that a concept has good effects, but is also in some significant way objectionable. Stipulate furthermore that no revision could get rid of the way it is objectionable. Then we might not have sufficient reason to adopt or retain the concept.

Normative functions' relation to the second question is more complicated. Normative functions give us reasons to opt for one version of a

concept over another, insofar as that version better fulfils the concept's normative function. But normative functions do not *necessarily* give us *sufficient* reasons to opt for one version of the concept. This emphasises something we have touched on above: we should not *just* engineer for the most important normative function. There are often other reasons we have to consider. Let us return to the analogy of designing a knife. Suppose the most important function of the knife we aim to design is to cut food. One of the design choices we face is how sharp the blade should be. If the knife's normative function was our only consideration, we should probably make it as sharp as possible to allow it to perform its most important function best. There are, however, other relevant considerations: for example, how safe the knife would be for its users.

In addition, there might be normative reasons not to engineer for a specific thing. Let us stick with the knife example. The knife's most important normative function is what it is *for*, in the sense that this should guide and structure our engineering efforts. However, there might be moral side-constraints that restrict the designs that fulfil the knife's normative function. The knife should, for example, not be made out of toxic materials. Not being toxic is *not* part of a knife's function. However, we still have normative reasons not to engineer toxic knives. When designing a knife we should make sure that it performs its most important function sufficiently well, but also take such other normative considerations into account. Whether we have sufficient reason to design our knife in a specific way is determined by all these factors.

The same applies to the (re-)engineering of concepts. Of course, we first need to determine what the concept's most important normative function is, which gives us reasons that play the relevant guiding and structuring role in our engineering effort. In addition, though, concepts might have normative functions that are weaker than the most important ones. And there might be normative reasons for having a version of a concept not based on normative functions. For example, [Nicholas Smyth \(2022\)](#) has argued this for identity-constituting concepts: for example, gender and race concepts. A concept is identity-constituting if it is part of a person's practical identity. According to Smyth, we have normative reasons to refrain from revising identity-constituting concepts, irrespective of the consequences that possessing such concepts has.

So, to determine if there is sufficient normative reason to have (a version of) a concept, we need to consider many different considerations. An appeal to normative functions alone will rarely address this. However, notice that this is going to be true on *any* account of functions.

Our discussion importantly clarifies what the role of functions *can* be for a conceptual engineer. Functions can play the relevant structuring and guiding role in conceptual engineering in telling us what the concept is *for*. But functions cannot settle all questions. This should be true, no matter how we understand functions. However, normative functions give us *very good, yet defeasible* reasons to adopt a concept and they give us *very good, yet defeasible* reasons to opt for one specific version of a concept. We take it that on a charitable reading of what the most significant and plausible version of the authority problem can require, this is sufficient.

4.2 Normative functions and the Strawsonian challenge

The Strawsonian challenge reoccurs in the conceptual engineering literature. The central worry is that if conceptual engineering involves changing our concepts' extensions and intentions, then a proposed revision will not put us in a better position to answer the questions we were concerned with in the first place. Instead, we would be changing the subject. The functionalist response to this worry is that continuity of subject consists in continuity of function (Sundell 2020, Thomasson 2020). It argues that concepts have a specific function. We revise a concept because we deem the revised version more suitable to perform the function the original version performed. Functions are supposed to explain when two concepts are 'on the same subject', that is, versions of the *same* concept. However, whether functions are indeed able to explain this depends on our answer to the function specification problem. Riggs (2021) has argued, for example, that this explanation fails if we take functions to be etiological or system functions (pp. 11572-3).

The main problem Riggs identified is that two concepts that are intuitively on the same subject can have very different etiological or system functions. Consider this for etiological functions: suppose that we found out that there is a community that uses a concept which is psychologically identical to our concept MARRIAGE. However, the effects that explain the persistence of this concept in the community are very different from ours. The etiological function account implies that the other community uses a concept that differs greatly from our concept. But this would be difficult to accept, as in this case it seems plausible that the concepts *are* the same. A similar argument can be made for system functions (Riggs 2021, p. 11573). We agree with Riggs that the orthodox accounts of function cannot explain the sameness of subject after all. However, we believe, unlike Riggs, that this argument does not generalise to all kinds of functions. First, it is important to be clear what

it would take to address the Strawsonian challenge. We take it that the use of phrases like ‘change the subject’ does not effectively delineate the real issue underlying the challenge. The challenge isn’t just about maintaining semantic continuity; it’s about preserving the concept’s practical, theoretical, or normative significance. The challenge arises most clearly when we revise a concept to avoid certain problems. For example, we might revise our concept FREE WILL to side-step worries about the compatibility of free will with determinism. This revision would be futile, though, if the reasons that we cared about free will in the first place were lost in the revision, such that the revision does not really address why people were worried about the potential incompatibility. Specifically: there is something *important* about free will and the revision should preserve *that* to really address the worries we had in the first place. More generally, what a revision should do is to preserve why the concept matters.

There are different ways in which we could cash this out. One suggests that we must still be talking about the same thing (in which case the Strawsonian challenge cannot be met) or that we at least must preserve the general subject matter. Proponents of the functional turn propose, instead, that preserving the *function* of FREE WILL is crucial. This perspective reveals the shortcomings of orthodox accounts of function in addressing the Strawsonian worry. After all, the concept’s causal history or systemic role is irrelevant unless that history or system itself is important. So, appealing to these functions is insufficient to address the Strawsonian worry: we are not interested in just *any* sort of continuity, but in a continuity in what matters. We can now also easily appreciate Riggs’ point in more general terms: there is little reason to assume that orthodox functions and the reasons why our concepts matter in the first place go hand in hand. Hence, sameness of orthodox function doesn’t necessarily yield sameness of relevant topic.

This also shows quite clearly, though, that normative functions are ideally suited to address the Strawsonian challenge: when we have continuity of normative function, we literally preserve what is most important about the concept. If we have this, then we have the most plausible and attractive sense in which we can have ‘continuity of topic’. In fact, notice that it should really be irrelevant if we are talking about something different after the revision, if we have preserved what was important about the concept. To go back to our earlier example, even if our revision of FREE WILL changes the intension and extension of the concept, this should not really matter if the revision preserves why the concept was important in the first place (if that is compatible with a

change in intension or extension, something that needs to be argued by those aiming to revise).

It should also be clear that this response is immune to Riggs' criticism of orthodox accounts. While there might be ways of thinking about 'continuity of topic' that can come apart from continuity in normative function, we take our discussion here to highlight that the important sort of continuity cannot, because normative functions are *designed* to preserve why our concepts matter in the first place. The upshot of this discussion is therefore that the functionalist can meet the Strawsonian challenge if we interpret functions as normative functions.

4.3 Normative functions and metalinguistic negotiation

The final aspect supposed to be explained by functions is metalinguistic negotiation. Speakers engage in metalinguistic negotiation when they disagree about how they should use a concept in a specific context. A frequent example in the conceptual engineering literature is that of a couple disagreeing about whether specific behaviour constitutes cheating, for example, having an emotional affair (Riggs 2021) or having sex with a humanoid sex doll (Jorem and Löhr 2022). When we interpret this kind of disagreement as metalinguistic negotiation, we take the disagreement not to be about the extension or intension of our current concept of cheating. It's about whether we *should* use 'cheating' in a way that includes the behaviour in question.

Plunkett and Sundell (2013) argue that disagreements about concept use actually centre on which version of the concept should play the functional role of the concept in question. This presupposes that people have some idea of what the functional role of 'cheating' is, and disagree whether the problematic behaviour is included in the concept that should play this functional role (Plunkett & Sundell 2013, p. 21). Riggs (2021, pp. 11570-2) argues that none of the proposed understandings of 'functional role' in the literature is suitable for explaining metalinguistic negotiation. His main argument comprises two premises:

- 1 people usually have no idea what the etiological or system function of a concept is, and
- 2 the metalinguistic disagreement cannot be about something on which people 'do not have views' (p. 11570).

We find both premises plausible and agree with Riggs that this is a problem for existing proposals to explain metalinguistic negotiation. However, *normative* functions are well-suited to explain this

phenomenon. People usually have ideas about what it is a concept allows us to do, that is, which normatively significant effects having the concept has. For example, it is plausible that the normative function of judging that ‘x is cheating on y’ is something like that of signalling that x and y have a special, romantic relation and that x has acted in a way that undermines this relation. Speakers are aware that something like this is the *normative* function of having the concept of cheating, in the sense that this is an important thing it allows us to do. Moreover, it is plausible that when speakers engage in metalinguistic negotiation they disagree about which version of the concept should perform this function. In the current case, the disagreement is probably a substantive disagreement about whether specific behaviour undermines the special, romantic relation between individuals. Hence, if we understand functions as normative functions, then it is plausible to interpret the disagreement in meta-linguistic negotiation as a disagreement about functions, despite Riggs’ worries. Given that this means that normative functions can *also* perform the relevant work in meta-linguistic negotiation, we conclude that normative functions can, in fact, do the explanatory work that are assigned to conceptual functions in the literature.¹

References

- Cappelen, Herman 2018, *Fixing Language. An Essay on Conceptual Engineering* (Oxford University Press)
- Chang, Ruth 1997, *Incommensurability, incomparability, and practical reason* (Harvard University Press)
- Cullity, Garrett 2018, ‘Weighing Reasons’, in Daniel Star (ed.), *The Oxford Handbook of Reasons and Normativity*, 423–442 (Oxford University Press)
- Cummins, Robert 1975, ‘Functional Analysis’, in *Journal of Philosophy* 72: 741–65
- Deutsch, Max 2020, ‘Speaker’s Reference, Stipulation, and a Dilemma for Conceptual Engineers’, in *Philosophical Studies* 177: 3935–57

¹ We are very grateful to Hannah Altehenger, Graham Bex-Priestley, Susanne Burri, James Brown, William Gamester, Rachel Handley, Johannes Himmelreich, Jessica Isserow, Maximilian Kiener, Leo Menges, Sven Nyholm, Andreas Schmidt, Christine Tiefensee, two reviewers for *MIND* (one of which gave incredibly detailed feedback), as well as The Editors of *MIND*. Very special thanks goes to Peter Königs and Irina Schumski for the great discussions on functions in our reading group on conceptual engineering and technology. Both authors contributed equally to this paper.

- Eklund, Matti 2014, 'Replacing Truth?' in Brett Sherman and Alexis Burgess (eds.), *Metasemantics* 293–310 (Oxford University Press)
- Haslanger, Sally 2020, 'Going On, Not in the Same Same Way', in Alexis Burgess, David Plunkett, and Herman Cappelen (eds.), *Conceptual Engineering and Conceptual Ethics* 230–60 (Oxford University Press)
- Jorem, Sigurd 2021, 'Conceptual Engineering and the Implementation Problem', in *Inquiry*
- 2022, 'The Good, the Bad and the Insignificant - Assessing Concept Functions for Conceptual Engineering' in *Synthese* 200
- Jorem, Sigurd, and Guido Löhr 2022, 'Inferentialist Conceptual Engineering' in *Inquiry*
- Koch, Steffen 2020, 'There Is No Dilemma for Conceptual Engineering. Reply to Max Deutsch', in *Philosophical Studies* 178: 2279–91
- Marques, Teresa 2020, 'Amelioration vs perversion' in Teresa Marques and Asa Wikforss (eds.), *Shifting Concepts: The Philosophy and Psychology of Conceptual Variability* 260–284 (Oxford University Press)
- Matthias, Andreas 2004, 'The Responsibility Gap: Ascribing Responsibility for the Actions of Learning Automata' in *Ethics and Information Technology* 6: 175–83
- McPherson, Tristram 2018, 'Authoritatively Normative Concepts' in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 13 (Oxford University Press)
- Millikan, Ruth 1984, *Language, Thought, and Other Biological Categories* (MIT Press)
- Nado, Jennifer 2020, 'Taking Control: Conceptual Engineering Without (Much) Metasemantics' in *Inquiry*
- 2021, 'Conceptual Engineering, Truth, and Efficacy' in *Synthese* 198: 1507–27
- Nimtz, Christian 2021, 'Engineering Concepts by Engineering Social Norms: Solving the Implementation Challenge' in *Inquiry*
- Parfit, Derek 2011, *On What Matters* (Oxford University Press)
- Plunkett, David, and Tim Sundell 2013, 'Disagreement and the Semantics of Normative and Evaluative Terms' in *Philosophers' Imprint* 13 (23)
- Queloz, Matthieu 2022, 'Function-Based Conceptual Engineering and the Authority Problem' in *Mind* 131 (524): 1247–78
- Queloz, Matthieu, and Friedemann Bieber 2021, 'Conceptual Engineering and the Politics of Implementation' in *Pacific Philosophical Quarterly*
- Rawls, John 1993, *Political Liberalism* (Columbia University Press)

- Raz, Joseph 1986, *The Morality of Freedom* (Oxford University Press)
- Raz, Joseph 2011, *From Normativity to Responsibility* (Oxford University Press)
- Riggs, Jared 2021, 'Deflating the Functional Turn in Conceptual Engineering' in *Synthese* 199: 11555–86
- Scanlon, Thomas M. 1998, *What We Owe to Each Other* (Harvard University Press)
- Schroeder, Mark 2008, *Slaves of the Passions* (Oxford University Press)
- Finlay, Stephen, and Mark Schroeder 2017, 'Reasons for Action: Internal vs. External.' in Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy* (Fall 2017), <<https://plato.stanford.edu/archives/fall2017/entries/reasons-internal-external/>>
- Simion, Mona 2018, 'The 'Should' in Conceptual Engineering' in *Inquiry* 61: 914–928
- Simion, Mona, and Christoph Kelp 2020, 'Conceptual Innovation, Function First' in *Noûs* 54: 985–1002
- Smyth, Nicolas 2022, 'The Return of the Philosopher-King: Conceptual Engineering, Consequentialism and Social Authority' Paper Presented at the British Society for Ethical Theory Annual Conference 2022
- Strawson, Peter F. 1963, 'Carnap's Views on Conceptual Systems versus Natural Languages in Analytics Philosophy' in *The Philosophy of Rudolf Carnap* 503–18 (Open Court Publishing)
- Sundell, Tim 2020, 'Changing the Subject' in *Canadian Journal of Philosophy* 50: 580–93
- Thomasson, Amie 2020, 'A Pragmatic Method for Conceptual Ethics' in Herman Cappelen, David Plunkett, and Alexis Burgess (eds.), *Conceptual Engineering and Conceptual Ethics* (Oxford University Press)
- 2021, 'Conceptual Engineering: When Do We Need It? How Can We Do It?' in *Inquiry*
- 2022, 'How Should We Think about Linguistic Function?' in *Inquiry*
- Vallier, Kevin 2022, 'Public Justification', in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Winter 2022 Edition), <https://plato.stanford.edu/entries/justification-public/>
- Williams, Bernard 1980, 'Internal and External Reasons' in Ross Harrison (ed.), *Rational Action* (Cambridge: Cambridge University Press)
- Wright, Larry 1973, 'Functions' in *The Philosophical Review* 82: 139–68