



The good, the bad and the insignificant—assessing concept functions for conceptual engineering

Sigurd Jorem¹ 

Received: 18 February 2021 / Accepted: 17 January 2022 / Published online: 6 April 2022
© The Author(s) 2022

Abstract

Many theorists of conceptual engineering appeal to the functions, roles, purposes or aims of concepts to articulate how conceptual engineering ought to be done. The functional approach to conceptual engineering is well-motivated: It promises a good account of the limits of revision, and of what makes some concept good. In this paper, I raise a problem for the functional approach which concerns the existence of *harmful* and *methodologically insignificant* concept functions. I examine whether we can deal with these problematic functions by adopting a technical notion of function. I thus review the prospects for using the notions of a contextually stable function, of a designed function, of a proper function, and of a system function as our operative notion of function. None of them help us resolve the problem. On this basis, I argue that advocates of the functional approach should be committed to a comparatively weak claim, according to which functions must be assessed case-by-case, and that we are best served by employing an unsophisticated notion of function, according to which the function of a concept just is something the concept is used for.

Keywords Conceptual engineering · Conceptual revision · Function · Philosophical methodology

1 Introduction

To engage in conceptual engineering is to design and implement concepts that will, in the good case, serve us well in some capacity. Commonly, the aim is not just to add new concepts to our conceptual repertoire, but to replace or revise the concepts

✉ Sigurd Jorem
sigurd.jorem@gmail.com

¹ Department of Philosophy, Classics, History of Art and Ideas, University of Oslo, Oslo, Norway

we already possess. A conceptual engineer may e.g. seek to revise our concepts of gender, with the goal of promoting social justice (Haslanger, 2000); propose that we replace our actual concept of truth with concepts of truth that avoid the alethic paradoxes (Scharp, 2013); recommend concepts of causation that promise to improve our explanations of events and states of affairs (Woodward, 2003); or propose that we use some concept of knowledge that better serves our need for having a concept of knowledge (Craig, 1990; Hannon, 2019). These projects take a normative approach to our conceptual repertoire, to be contrasted with the descriptive approach characteristic of traditional conceptual analysis. While a conceptual analysis should accurately capture a concept in use, exercises in conceptual engineering are not so constrained. However, lifting the requirement to capture a concept in use does not mean that anything goes. Exercises in conceptual engineering are rather subject to different success conditions. The relevant success conditions should befit the aim of improving our conceptual repertoire. As a rough approximation, an exercise in conceptual engineering is successful only if

(*Goodness*) the engineered concept is good by standards that apply to the concept in question; e.g. it is useful or benign,

and

(*Limits*) it does not involve revising a concept beyond limits, viz. so that we merely “change the subject” or undermine legitimate purposes for having the concept that is being revised.

These conditions are not fleshed out, and they are not intended to be. In this paper, I examine a view of how we should develop them. According to what I call *the functional approach* to conceptual engineering, we should understand success in conceptual engineering in terms of the function(s) that a concept performs.¹ More precisely, the functional approach understands the limits of revision in terms of the function(s) currently performed by the concept to be revised, and what makes some concept good in terms of which function(s) the engineered concept is to perform.

The approach has many advocates. Recently Weinberg (2006), Brigandt (2010), Plunkett and Sundell (2013), Plunkett (2015), Prinzing (2018), Nado (2021a, 2021b), Simion and Kelp (2020), and Thomasson (2020a) have all advocated a functional approach to conceptual engineering. Some of them speak of the ‘role’, ‘purpose’ or ‘aim’ a concept serves, but we may interpret all of them as latching on to roughly the same idea; one that is described well with the word ‘function’. Looking further back in time, we may identify a functional approach to conceptual engineering in Strawson’s (1963) objection to Carnap—that using his method of explication to solve philosophical problems changes the subject rather than solving the problem—and

¹ *Goodness* and *Limits* are not supposed to be jointly sufficient for success. Arguably, unless the engineered concept is implemented somehow—e.g. becomes the semantic content of a word in a language, or used in some other way—the project is not yet fully successful. [redacted]. Meanwhile, I hold the working assumption that *Goodness* and *Limits* are necessary conditions on success in conceptual engineering. I set aside issues pertaining to implementation.

in the latter's response (Carnap, 1963). As Nado (2021b) observes, both Strawson's objection and Carnap's response are articulated in functional terms:

The kinds of concept we employ are not independent from the kinds of purpose for which we employ them; even though some concepts can fulfil more than one kind of purpose. (Strawson, 1963, 506)

The explicatum is intended to take the place of the explicandum, and that means, of course, that it is to be used for the same purpose as the explicandum. (Carnap, 1963, 936)

In addition to philosophers thinking about conceptual engineering, several philosophers engaged in the practice of conceptual engineering have appealed to function or purpose, including Craig (1990) and Hannon (2019) on knowledge, Haslanger (2000) on gender and race, and Woodward (2003) on causation.

However, it is not immediately clear what it is for concepts to perform functions. Proponents of the functional approach need to clarify their 'function'-talk. In doing so, they face several plausible demands. First, there need to exist functions in the sense being proposed. Second, the functions so construed need to help us understand the success conditions on conceptual engineering (i.e. *Goodness* and *Limits*). Third, the candidate functions should be functions we *ought* to preserve when we revise concepts. Reviewing extant proposals, I argue that there is probably no sense of 'function' that can meet these demands, so as to give us a fairly complete and non-question-begging account of *Goodness* and *Limits*. Whereas Cappelen (2018) argues on similar grounds that the functional approach fails, I aim to show that this conclusion is too hasty. Firstly, there are clear, non-empty notions of a concept function. If nothing else, we may understand a concept function simply as what we happen to use a concept for (on occasion).² However, not any given function performed by a concept should constrain how we may revise the concept. Not all functions are good and methodologically significant. Some are harmful and some are insignificant. But despite the existence of harmful and insignificant functions, there is comparatively weak claim the proponent of function can make, which still has methodological bite: For any given concept there are functions this concept performs that help determine what a successful revision must be like. What we cannot tell in advance, I argue, but must determine case-by-case, are exactly which functions we need to do conceptual engineering in the service of.

In Sect. 2, I show that the functional approach offers an appealing construal of *Goodness* and *Limits*. In Sect. 3, I raise a problem for the functional approach that I call *the selection problem*. According to the selection problem, there are functions that we should not engineer in the service of. Therefore, we have to make a selection

² The idea that we *use* concepts may provoke questions about how I conceive of concepts; perhaps driven by a worry that there is no good understanding of what concepts are, such that they are the sort of things we can use or engineer (Cappelen, 2018, 141, 200; Isaac, 2020; Koch, 2020). I assume, minimally, that concepts are the sort of things words can mean. Perhaps the set of concepts is broader than the set of linguistic meanings, but it will do no damage to my main line of argument if 'concept' is understood as the meaning of a substantive linguistic expression, and 'using a concept' as using some expression with a given meaning (or given speaker-meaning).

among functions, but it is not clear how one might specify the range of functions conceptual engineers need to pay heed to. In Sect. 4, I review whether extant ways of understanding function can deal with the selection problem without falling prey to other problems. In Sect. 5, I argue that the advocate of function should accept that the selection problem cannot be dealt with in a general manner. I elucidate the consequences this has for the functional approach.

2 Why appeal to function?

We use concepts to carve up the world, to explain and predict events, to evaluate actions and states of affairs, to prescribe thought and action, etc. The guiding thought behind the functional approach, as I understand it, is that what we *use* a given concept for matters for how we should assess revisions of that concept. On the side of *Goodness*, what a concept is to be used for matters for what kind of standards we should use to evaluate that concept. On the side of *Limits*, what we currently use a concept for matters for how we should circumscribe the limits of revision of that concept. In the following, I motivate and clarify the functional approach with respect to both of these claims; first *Goodness*, then *Limits*.

On a functional understanding of *Goodness*, what it is for a concept to be good is for it to perform its functions well. The functions a concept performs thereby (co-) determine the standards we should apply when evaluating the concept. As Thomasson writes,

no detailed evaluation can be made without an assessment of the functions that are to be served. Once a purpose (or multiple purposes) is/are identified, we can go on to use that in engaging in conceptual engineering—determining what sorts of rules or constraints that would best (or better) enable it to fulfill its function(s), going forward. (Thomasson, 2020a, 449)

We may use a comparative example to motivate this view of what it is for a concept to be good. It is sometimes claimed that concepts ought to “carve nature at its joints” (e.g. Sider, 2011). *Joint-carving* is a plausible candidate for a conceptual virtue or norm, i.e. as a standard of evaluation.³ Thus, take a concept that plays an important role in natural science, such as the concept of an electron. Presumably, there is a joint in nature that our concept of an electron carves out. This is, I assume, part of what makes it a good concept. Even so, it is dubious to evaluate all concepts according to whether they carve out a joint in nature. For instance, our concept of a table is a useful concept, despite failing badly to carve out a joint in nature (tables can be composed of any number of different materials, have any number of different shapes and sizes etc.). The functional approach can explain why this failure to carve out a joint in

³ I set aside the question of whether we should understand joint-carving in a deflationary or metaphysically heavyweight way (Thomasson, 2020a). On a heavyweight reading, facts about the joints of nature are deep facts about the metaphysical structure of the world. On a deflationary reading, talk of “the joints of nature” is something like a shorthand for talk about what it turns out to be useful to identify in the process of forming scientific explanations and predictions.

nature does not entail that our concept of a table is defective: The concept does not perform any functions that would justify us in evaluating it according to this standard. By contrast, many scientific concepts perform explanatory and predictive functions that do justify us in evaluating them so.

There are several other conceptual virtues that are important in science, but less important in non-scientific discourse. Consider being a quantitative concept vs. being a categorical concept, or being exact vs. being vague. Exact, quantitative concepts drive scientific progress. In everyday discourse, by contrast, we get along fine with categorical and vague concepts, such as our concept of being bald or our concept of friendship. We have no acute need to make these concepts more exact, or to replace them with quantitative alternatives. However, it seems arbitrary to set apart conceptual virtues that apply to scientific discourse from conceptual virtues that apply to non-scientific discourse unless there is *something* about scientific discourse that warrants our discrimination. Absent better explanations, we are justified in thinking of the relevant *something* as what we use scientific concepts for, i.e. generalization, explanation, prediction, etc. In turn, it makes sense to classify what we *use* scientific concepts *for* as functions performed by scientific concepts. Thereby, we find reason to endorse the following.

(*Functional Goodness*) The functions that a concept C is to perform help determine the standards of goodness for C.

What, then, of *Limits*? The need to address limits on conceptual revision came squarely in view with Strawson's objection to Carnap: Unless we preserve something from the concept to be replaced or revised, we run the risk of changing the subject rather than improving what we had. This bears upon a different success condition from *Goodness*, because a concept can be good on independent merits and yet fail to be acceptable as a revised state of, or replacement to, a given concept in use. To illustrate, our concept of belief is independently useful, but we would achieve nothing if we revised our concept of *desire* so as to attain the same content as our concept of belief. The functional approach to conceptual engineering can explain this by appealing to the fact that our concept of desire performs a different function from our concept of belief.

As with *Goodness*, the functional approach gains abductive support from its construal of *Limits*: It offers a good explanation of why we reject revisionary proposals that, intuitively, involve too much change and accept—or at least hesitate to reject—proposals that do not seem to involve too much change. To see this, compare two revisionary proposals. First, consider a (silly) proposal to revise our concept of knowledge so as to apply to all states of belief; not just those that are true and justified. Next, consider a proposal to revise our concept of knowledge so as to apply to all and only those states of belief that are true and justified (cf. Weatherston, 2003). Let 'KNOWLEDGE_{ETB}' name the concept that would result from the latter revision. Philosophers since Gettier (1963) have been painfully aware that our actual concept of knowledge is non-identical to KNOWLEDGE_{ETB}, but the existence of Gettier cases merely shows the non-identity of the two concepts, not that KNOWLEDGE_{ETB} would be too drastic of a departure from our current concept of knowledge. For that matter,

Gettier cases are few and far between. Extending our concept of knowledge to *all* states of belief is problematic in a way or to a degree that the KNOWLEDGETB revision is not. The functional approach to conceptual engineering offers a promising explanation for this. If we suppose—as is plausible—that we use our current concept of knowledge to extend epistemic praise, we can explain our diverging judgments towards the two proposals. If we were to revise our concept of knowledge so as to apply to all states of belief, we would undermine the function of epistemic praise. If, on the contrary, we revised our concept of knowledge so as to apply to all and only states of belief that are true and justified, we could still use the concept for this purpose.

Zooming out, it is plausible to think that the theoretical and practical value concepts have for us reside in the use we have for them. Unless there is *use* for a given concept—to carry out some set of tasks or perform some kind of function—there would be no reason to target that concept, as opposed to any other concept, as the object of revision. But if the value of a concept resides in the functions it performs for us, then it is plausible that the capacity to perform those functions is what we should preserve when we engage in revision. We are therefore justified in drawing the limits of revision accordingly:

(Functional Limits) A revision of a concept C is acceptable only if it preserves the capacity to perform functions already performed by C.

As will become clear in the next section, there are problems with *Functional Limits*, in particular having to do with how to understand ‘functions’ and exactly which ones we ought to preserve. For now, let us observe that the claim is well-motivated.

Notably, *Functional Limits* makes no mention of “subjects” or “topics” (c.f. Capelen, 2018, chaps. 9–10); nor does it require that a concept preserves numerical identity through a revision process (cf. Prinzing, 2018). For reasons spelled out by Nado (2021a; MS), Koch (2021) and others, I think these ways of drawing the limits of revision are dead ends. Unless we specify an individuation convention for subjects or concepts in a way that piggybacks on facts or properties we care about for independent reasons, there is no good reason why conceptual engineers should not sometimes change the subject, or sometimes change *which* concept we use (as opposed to changing the concept as such). I have therefore specified *Functional Limits* without requiring sameness of topic or concept.

We seem to have good reason to accept *Functional Goodness* and *Functional Limits*. As these two commitments comprise the functional approach to conceptual engineering, we have *prima facie* reason to endorse the functional approach.⁴ However, our motivation has rested on an unsophisticated understanding of what it is for a concept to perform a function. In the next section, we attend to a problem that arises on this unsophisticated understanding.

⁴ Note that one could endorse a weaker version of the functional approach than one comprised by *Functional Goodness* and *Functional Limits*. One could e.g. think of function-preservation as a point in favor of a revision, to be weighted against other good-making features, rather than as a necessary condition. Pending arguments for taking a weaker approach, I find it more interesting to explore the prospects of the comparatively strong version encapsulated by *Functional Goodness* and *Functional Limits*.

3 The selection problem

On the understanding of function we used to motivate the functional approach, the function of a concept is what it is used for. In general, we use concepts for all sorts of tasks: To categorize, explain, predict, prescribe, evaluate, command etc. We may identify more specific tasks when we look at concrete examples. For example, Craig hypothesizes that we use our concept of knowledge to “flag approved sources of information” (Craig, 1990, 11). Woodward suggests that we use our concept of causation to identify “relationships that are potentially exploitable for purposes of manipulation and control” (Woodward, 2003, 25). These seem like plausible ascriptions of function. Moreover, it seems reasonable to evaluate revisionary proposals according to how well the engineered concept (of knowledge or causation) can perform these functions.

Importantly, however, for all we have assumed about concept functions so far, we have no reason to think that concepts perform exactly one function each. We use our concept of knowledge for a number of different tasks. Some of these tasks seem less well fit to determine what our concept of knowledge should be like: We use our concept of knowledge to discredit interlocutors (e.g. “And what do you know about that?”), to position ourselves socially (“I don’t *believe* that there is anthropogenic climate change, I *know* there is.”), to acknowledge mistakes (“I should have known better.”) and so on (cf. Cappelen, 2018, 182 on uses of ‘women’). Now, it may sound awkward to think that our concept of knowledge *has the function* of, e.g., discrediting interlocutors just because we sometimes use the concept for this purpose. I address this concern in the next section, where we explore more demanding notions of function. For now, we need only observe that it sounds dubious to ask that revisions of our concept of knowledge should aid us in discrediting interlocutors.

Recall the two claims that comprise the functional approach as I have construed it here:

(*Functional Goodness*) The functions that a concept C is to perform help determine the standards of goodness for C.

(*Functional Limits*) A revision of a concept C is acceptable only if it preserves the capacity to perform functions already performed by C.

These schematic statements make indefinite reference to the functions performed by a concept. They do not tell us exactly *which* functions of C that are supposed to determine the goodness of C or its limits of revision. And then, since any given concept performs a number of different functions, the functional approach lapses into obscurity if it does not clarify whether *all* the functions of a concept are supposed to count for *Goodness* and *Limits*, only some, or only one. In the event that only one function or only some functions count, we should want to know which one(s). As noted, one could supply a different, more demanding notion of function to replace the unsophisticated notion of function we have used so far. Then we could discount would-be functions such as *the use of our concept of knowledge to discredit interlocutors* as not being functions in the operative sense of ‘function’. But for now, I ask for patience with my equation of functions with uses as I spell out the present problem.

In the absence of a more sophisticated notion of function, we could still ask whether there is really a problem here. Is it really so bad to hold that *all* the functions of a concept are to determine its goodness and limits of revision? The answer is “yes”. As I show in the following, there are both *harmful* functions and functions that are otherwise *insignificant* (for the purposes of *Goodness* and *Limits*). The fact that our concepts perform harmful and insignificant functions means that we have to make a selection among the functions a given concept performs. Not all of them count toward *Goodness* and *Limits*. I call this *the selection problem*.

Some functions are methodologically insignificant. Take our concept of planet-hood as an example. Not too long ago, when Pluto was still considered a planet, our concept of planethood applied to nine objects in our solar system. We used the concept to think and speak about a property exhibited by nine objects in our solar system. Therefore, on our unsophisticated understanding of function, our concept of planethood had the function of categorizing nine objects in our solar system. In the wake of the International Astronomical Union’s decision to redefine ‘planet’ in 2006, our concept changed.⁵ According to the new definition, “A planet is a celestial body that [...] has cleared the neighborhood around its orbit.”⁶ On this criterion, Pluto no longer counts as a planet. While the revision preserved several functions performed by our pre-2006 concept of planethood (such as the function to categorize bodies of mass orbiting a star), it failed to preserve the function to categorize nine objects in our solar system. And yet the revision is acceptable. It follows that we would be wrong to require the preservation of the function considered. By consequence, we cannot require that revisions preserve all functions performed by the concept to be revised.

A similar example is due to the change from a pre-Linnaean concept of fish to the Linnaean concept of fish. The former concept functioned to classify whales as fish, but with Linnaeus’ taxonomy we transitioned to a concept of fish that does not apply to whales and that, by consequence, does not function to categorize whales. Yet we would consider this an acceptable revision. It turns out that some of the tasks we use concepts for are less important, methodologically speaking, than others. To classify whales as of the same kind of animal as cod and trout is not an important task. Neither is the task of denoting a property exhibited by nine objects in our solar system. I project that we can find many more examples of insignificant functions. In virtually any case of conceptual change we can probably specify something the old concept was used for that the new concept is no longer used for.

While we should not be obliged to preserve insignificant functions, we actively want to avoid preserving harmful functions. Consider a concept of rape that does not apply to intra-marital sexual transgressions. Let us call it ‘the extra-marital concept

⁵ Is this a genuine example of conceptual change, as opposed to a change in belief? Haslanger (2012, 398) sometimes describes her project as revealing what we meant all along by terms like ‘woman’ and ‘black’. Externalist views, especially temporal externalist views, may lead to such verdicts. See (Ball, 2020) for a defense applied to conceptual engineering. Following Ball, one might think that our concept of planethood disappplied to Pluto all along. However, even if we accept this, we may still evaluate the example as a hypothetical case of conceptual change. Our interest is with what it takes for a revision to be acceptable, and this interest is equally well served by considering hypothetical cases.

⁶ C.f. Resolution B5—Definition of a Planet in the Solar system: https://www.iau.org/static/resolutions/Resolution_GA26-5-6.pdf (accessed October 2020).

of rape'. Unfortunately, we do not have to go far back in time to find evidence that the extra-marital concept of rape was our actual concept of rape. Edward Schiappa quotes Edwin Silberstang on US legal practice from only a few decades ago: "A husband cannot be guilty of an actual rape, or an assault with intent to rape his wife even if he has, or attempts to have, sexual intercourse with her forcibly and against her will." (Silberstang, 1972, 775; quoted in Schiappa, 2003, 54). In the legal context, use of an extra-marital concept of rape has helped protect wrong-doers from liability to punishment. Indeed, the extra-marital rape *performed the function* of protecting wrong-doers from liability to punishment. The extra-marital concept of rape thereby performed a harmful function. What I take to be our current concept of rape is marriage neutral. It does not perform the function of protecting wrong-doers from liability to punishment. Despite not performing this function, it remains an acceptable replacement to the extra-marital concept of rape. Indeed, the very point of broadening our concept of rape would be to have it cease performing the function considered.

The extra-marital concept of rape is not a unique case. Concepts of race may function to consolidate social injustice in virtue of underwriting dubious explanations of socio-economic differences. A 'both sides' concept of objective journalism may function to distort our ability to identify and disseminate facts. The former is a morally harmful function and the latter is an epistemically harmful function. We do not want concepts to perform these functions. We should not require their preservation or judge concepts to be good according to how well they perform them.

Conceptual revisions cannot be normatively constrained by insignificant functions or harmful functions. We need a way to discount them, preferably in a principled and non-question-begging manner. If we achieve this, we will have solved the selection problem.

4 Can we solve the selection problem by refining our notion of function?

It is fair to object to some of the function ascriptions I made in order to raise the selection problem. I have flouted the principle that "All functions are uses; but not all uses are functions." (Prinzinger, 2018, 868) To be sure, many of the function ascriptions I have made turn on what I have called an unsophisticated notion of function and would not count as functions on more demanding accounts. In this section, I review strategies for solving the selection problem that are based on existing ways of operationalizing what it is for a concept to perform a function.

Some advocates of the functional approach hint at a specific understanding of function, but do not make it fully explicit. E.g. Haslanger (2000, 35) speaks of the "central functions" of a term, but does not spell out what it takes to be a central function. Brigandt (2010, 2011) speaks of "the epistemic goals" of scientific concepts, and can be interpreted as meaning the (epistemic) function *intended* to be served by those who deploy or have developed the concepts in question. Thomasson proposes explicit accounts of function—Millikan's (1984) account of proper function and Cummins' (1975) account of system function—but does not commit to using either to develop the functional approach (Thomasson, 2020a, 444–46). Haslanger

adopts Cummins' notion of system function to understand conceptual engineering, and ascribes functions to concepts based on what they contribute to the workings of the (social) systems in which they are used (Haslanger, 2020a, 2020b). Simion and Kelp (2020), meanwhile, propose that the conceptual engineer is to construct a concept with a designed function—some task the concept is designed to serve—with the goal that the designed function becomes that which explains why the concept sustains and proliferates.⁷ If this happens, the designed function becomes an etiological function of the concept, which is, in essence, the same as what Millikan calls a proper function. Nado, on the other hand, argues that “neutrality on the nature of function is here not only permissible, but appropriate.” (Nado, 2021a). Riggs (2021) advocates a similar position. Arguing that the notion of a concept's function is being asked to play too many different explanatory roles, Riggs proposes that we do not need a sophisticated notion. He concludes,

Talk of the function, point, or aim of a concept isn't a way of referring to a useful piece of the theoretical machinery, but instead is a way of directing our attention to what *matters* in a given inquiry, which will change depending on what is at issue. (Riggs, 2021)

I agree with Nado and Riggs that we do not need the more demanding notions of function. One reason for this is that the more demanding notions do not, as I show in this section, help us deal with the selection problem.⁸

4.1 Contextually stable functions

As a first stab at restricting the operative sense of 'function', one could narrow it down to *contextually stable functions* (cf. Cappelen, 2018, 182). That is, we could restrict which functions we count as methodologically significant to those functions that are performed across all its contexts of use. By doing so, we avoid several problematic function candidates. For instance, although we do use our concept of knowledge to discredit interlocutors and to acknowledge mistakes, it is not the case that every time we use our concept of knowledge it performs these functions. By

⁷ The sort of success condition Simion and Kelp thereby identify is different from what *Functional Goodness and Functional Limits* purport to capture. If a concept is engineered to serve a designed function, and its serving this function becomes that which explains its sustained use and proliferation, one might say that the concept has been *implemented* as intended, but it is not necessarily a *good* concept for this reason (cf. Section 4.3). Moreover, success at implementation has nothing to do with the limits of revision as such.

⁸ Thus, the present argument is different from the argument Riggs (2021) gives. Riggs considers three categories of explanatory roles that the notion of a concept's function is supposed to play. One of these is the one I focus on here, of accounting for what revisions of a concept should be like. He argues that no notion of function can play all the explanatory roles they are supposed to play, and further, that we therefore ought to deflate the notion of function to something contextually variable, such as what we care about in a particular context. The present argument concerns only the normative role of 'function'-talk about conceptual engineering. I show that sophisticated notions of function do not help us solve the selection problem, and this lends some support, as I explain later, to the conclusion Riggs draws, i.e. that we should deflate 'function' talk.

narrowing the operative notion of function down to contextually stable functions, then, the functional approach avoids misattributing significance to what are in fact insignificant functions.

Following Cappelen (2018), however, it is not obvious that there are any good candidates for contextually stable functions. Are there specifiable tasks that we always carry when we use a given concept? If there are not, the wanted functions do not exist, and this strategy fails. If, on the other hand, there are contextually stable functions, the strategy does not thereby succeed. Cappelen considers what he calls a trivial version of the functional approach: “According to this view the only *universal*, i.e., stable, function of a concept ‘C’ is to denote Cs. The function of the concept ‘tiger’ is to denote tigers. The function of the concept ‘salad’ is to denote salads, the function of the concept ‘woman’ is to denote women, and so on.” (Cappelen, 2018, 182).

Although these denotational functions might count as contextually stable functions,⁹ they will not help us understand *Goodness* and *Limits*. Requiring the preservation of a denotational function would only beg the question of what it takes for a revised concept ‘C*’ to denote Cs. It is not helpful to be told that revisions of our concept of knowledge should still denote *knowledge*. On one interpretation of that demand, the revised concept should be qualitatively identical to our current concept of knowledge. But this interpretation of the demand is unacceptable, as it would preclude revision across the board. If the request is rather that the revised concept should still denote some phenomenon or state worthy of the label ‘knowledge’ we are not much wiser, since if we knew the range of phenomena worthy of the label ‘knowledge’, we would already know the limits of revision of our concept of knowledge. Thus, denotational functions do not help us understand *Limits*.

The kinds of functions *practitioners* in conceptual engineering appeal to are not perfectly stable. Consider the functions Craig and Woodward have appealed to. Craig engineers a concept of knowledge to serve the purpose of *flagging approved sources of information* (Craig, 1990, 11). This is not a contextually stable function. When we speak about knowledge in general terms—e.g. “Knowledge has great value.”—we are not using our concept of knowledge to flag an approved source of information. Woodward, meanwhile, engineers a concept of causation to serve the purpose of *identifying relationships that are potentially exploitable for purposes of manipulation and control* (Woodward, 2003, 25). Again, this is not a perfectly stable function. When we use our concept of causation to query the hypothesis that impact from an asteroid caused the mass extinction of dinosaurs, we are arguably not trying to identify a relationship that we could potentially exploit (cf. Woodward, 2003, 11). Any given concept can be used in an indefinite range of contexts. Chances are, there are some contexts in which a given function candidate is not performed. If there is some function that is performed across all contexts of use—as witness highly general functions such as the function of *parsing objects into categories*—chances are it will not set appropriate limits of revision or standards of goodness.

⁹ Note, meanwhile, that there are good reasons to resist the idea that every concept or piece of language serves a *descriptive* or *representational* function. See e.g. Brandom (1994), Price (2011) and Thomasson (2020b).

I believe both Craig and Woodward have identified important functions served by our concepts of knowledge and causation. That is to say, if a revision of either concept resulted in a concept that could not perform the respective function, this would give us reason against endorsing that revision. However, the importance of the respective functions does not entail that every kind of use of our concepts of knowledge and causation, silly and serious, are in their service.¹⁰ Recall that we are asking not only for the existence contextually stable functions, but for them to be such that we get a better grasp of *Goodness* and *Limits*, and for the set of contextually stable functions to be such that we ought to preserve them when we engage in revision. Therefore, we had better look to other notions of function to solve the selection problem.

4.2 Designed functions

Most concepts are occasionally employed for questionable purposes. As noted, our concept of knowledge could be employed for discrediting an interlocutor. Meanwhile, a fork could be used as a murder weapon, but we would not say that the function of a fork is to kill people. Forks are created as tools for eating. That is, a fork could function as a murder weapon on a particular occasion of use, but that does not mean that *the function of the fork* is to kill people. As Wright observes, “The function of a telephone is effecting rapid, convenient communication, but there are many other things telephones do: take up space on my desk, disturb me at night, absorb and reflect light, and so forth.” (Wright, 1973, 141). In these examples, *the function of X* may be understood as something like the purpose for which X was created. Following this line, we could amend the operative notion of function to *the purpose(s) for which a given concept were created*. Let us call these functions ‘designed functions’. Brigandt appears to apply this notion to account for the rationality of semantic change in science (Brigandt, 2010; 2011).¹¹ Prinzing flirts with a notion of designed function when he writes that “the function of a concept is what it was designed for” and illustrates with the functions that can openers and hammers were designed for (Prinzing, 2018, 869).¹² Although they do not use it to spell out *Goodness* and *Limits*, Simion and Kelp (2020) also use the notion of a designed function to define a way in

¹⁰ The respective functions could be important in an additional sense: They could help explain the proliferation of our concepts of knowledge and causation. Playing this role is again compatible with the function not being served on every occasion of use. What Millikan calls “the stabilizing and standardizing direct proper function of a language device” is not a contextually invariant function, but still explains proliferation and sustained use of the language device (Millikan, 1984, 31–32).

¹¹ The following comment from Brigandt suggests that he may have something less deliberate/intentional in mind than what the notion of a designed function indicates: “First, while scientists may be aware of what epistemic goal they pursue with the use of a certain term (especially if other scientists use the term for somewhat different purposes), I do not require that the epistemic goal is a belief explicitly held by these scientists. Instead, the epistemic goal is constituted implicitly by how a scientific community uses a term.” (Brigandt, 2011, 183).

¹² Prinzing qualifies: “‘Design’, of course, need not mean ‘conscious design’. Darwin’s great lesson, after all, was that the presence of a design does not entail the existence of a designer.” (Prinzing, 2018, 869). This suggests a notion of function that applies indiscriminately to products of evolution and conscious design. Millikan’s notion of proper function is one such notion; I discuss the strategy of using it in the next subsection.

which conceptual engineers may succeed in implementing the concept they prescribe (see f.n. 7).

Intuitively, narrowing down the operative notion of function to designed functions would preclude the function of protecting wrong-doers from liability to punishment from counting towards *Goodness* and *Limits*: Our concept of rape was not constructed for this purpose. It would intuitively preclude the function to denote a property exhibited by nine objects in our solar system, since we would have had our concept of planethood even if Pluto never existed. These are desirable consequences. However, there are major problems facing this strategy. There are reasons to dispute, first, the existence or prevalence of designed functions of concepts, second, our epistemic access to them, and third, their adequacy for generating norms on conceptual engineering.

The notion of a designed function is an intentional notion. It denotes the intentions that some individual(s) had when they designed a device. However, it is dubious that many concepts, let alone all, have been designed intentionally. Consider our concepts of knowledge and causation: Who designed them and what were their intentions? It is plausible to think that these concepts arose out of interactions between environmentally situated individuals, guided by practical needs for survival, coordination, welfare, etc., but without the metaconceptual awareness of a concept *designer*. Some concepts might have designed functions, but we want our account of success in conceptual engineering to cover more than a narrow range.

Even if a concept of interest were designed intentionally, we would not thereby be in a position to divine which purpose the concept was designed to serve. For instance, our concepts of knowledge and causation have a long history, and even assuming that they were intentionally designed, we do not know who developed them, much less what their intentions were. This would leave us unable to assess revisions of our concepts of knowledge and causation, since we would not have access to the proclaimed determinants of their *Goodness* and *Limits*. Finally, concepts may be designed with malicious intent, to perform what are in fact harmful functions. If a concept is designed to serve a harmful function, revisions of that concept should not be functionally continuous. The notion of a designed function thus fails several of our desiderata.

4.3 Proper functions

In answering an objection from Cappelen (2018) about relying on an intentional notion, Thomasson (2020a, 444) appeals to Millikan's (1984) notion of a proper function. The notion of a proper function does not require that there be a conscious purpose behind the candidate function. Thus, we can use it to make sense of our ascriptions of function to biological items, e.g. hearts, kidneys and instinctive behaviors. Millikan herself uses the notion of a proper function to develop a general theory about language and content. It is only natural to consider whether the functional approach to conceptual engineering can be articulated with her notion of a proper

function. Although the ensuing argument focuses on Millikan's account, it also applies to other etiological accounts of function, such as Wright's (1973).¹³

On Millikan's account, proper functions crop up in virtue of copying effects in social and natural environments. A copying effect ("reproduction") occurs when there is a causal relation between two items, such that one item attains one or more features of the other. This could be the effect that occurs when the genes of a parent are copied in the offspring, the mass production of a commercial item, or some individual's copying another individual's linguistic behavior. Such copying effects give rise to what Millikan calls "reproductively established families": Sets of items that share features in virtue of a copying effect. Proper functions are defined on this basis. Roughly, if a copied feature contributes to the survival or proliferation of members of the reproductively established family, then what the copied feature does for members of the family will count as a proper function for those members.¹⁴ On Millikan's view, then, we may view our (token) concepts as members of reproductively established families. Over the course of history, certain (type) concepts have enjoyed proliferation and sustained use, e.g. our concepts of knowledge and causation. Other concepts have been discarded, e.g. the concept of phlogiston. Whether a concept "survives" or not depends on what it does for us, or what it lets us do. The proper function of a concept C, then, is whatever we (and our ancestors) were able to do with C that explains how it enjoys sustained use.

Using the notion of a proper function as our operative notion lets us discount several abhorrent function candidates from misguiding our efforts in conceptual engineering. Consider the pre-2006 concept of planethood. The fact that it applied to Pluto meant that we could use this concept to think and speak about a property exhibited by nine objects in our solar system. However, our being able to do this cannot explain why we had that particular concept of planethood: The pre-2006 concept of planethood was developed before Pluto had been discovered and, presumably, the discovery of Pluto did not add any evolutionary advantage to the concept.¹⁵ Therefore, the function to denote a property exhibited by nine objects in our solar system cannot be a proper function of our pre-2006 concept of planethood. The present way of operationalizing *Functional Limits* thus avoids the undesirable consequence that the new concept of planethood was unacceptable as a replacement to the old concept. I assume that we can write off many problematic function ascriptions in a similar fashion. There is a general reason for thinking this: Selection effects tend to favor concepts that perform valuable functions. If a particular function is practically worthless, it is less likely that appealing to that function can help explain why we have a concept that can perform it. In this way, a function's counting as a proper function will tend to coincide with it being a function worth serving. By consequence, proper functions tend to be the functions we *ought* to do conceptual engineering in the service of.

¹³ There are differences between Millikan's and Wright's accounts of function that are significant for other purposes, cf. (Millikan, 1993 ch. 1).

¹⁴ Millikan's account is more complex than what I can do justice to here. See (Millikan, 1984, 18–49).

¹⁵ I assume that the discovery of Pluto was an ordinary empirical discovery and did not constitute conceptual change from the concept of planethood we operated with before the discovery.

There is one minor problem and one major problem with the current strategy. The minor problem has to do with our epistemic access to the proper functions of our concepts. If we take seriously the idea that we need to identify a proper function to ascertain whether an instance of conceptual revision is acceptable, conceptual engineers will have to engage in some hard empirical work. Sure enough, if we identify some valuable task *F* that a concept *C* enables us to do, we have the beginnings of an explanation for why we have *C*, and this is *prima facie* evidence in favor of believing that *F* is a proper function of *C*. But the belief is highly defeasible. There are any number of reasons why a concept could proliferate or survive. The reason why we have a concept of insanity could be, echoing Foucault (1973), that it aids us in confining socially undesirable people, rather than identifying a psychological condition that is detrimental to the deviant individual as such. It is not clear how we should go about to assess the quality of these competing explanations. Moreover, if our route to identifying a proper function is to identify some valuable task *F* that a concept enables us to do, then we should start to wonder why we have to take the detour of determining whether *F* is a proper function. If we already know that *F* is a valuable task enabled by the target concept, that is arguably all we need to know to be justified in evaluating revisions of the concept according to whether or how well they let us carry out *F*.

This brings us to the major problem with articulating *Functional Goodness* and *Functional Limits* in terms of proper function: The fact that a concept has proliferated and survived thanks to a function *F* does not entail that *F* is worth (pre-)serving. Recall the extra-marital concept of rape. It is in the interest of actual and potential offenders not to be liable to punishment for sexual violence against their spouses. This can explain why we had a concept of rape that did not apply to acts within the bonds of marriage. Arguably, the extra-marital rape had the *proper* function of protecting wrong-doers from liability to punishment. But might does not make right, and an explanation is no justification. It would be wrong to require that this function be preserved by revisions to the extra-marital concept of rape.

Millikan avows that “[t]he task of the theory of proper functions is to define this sense of ‘designed to’ or ‘supposed to’ in *naturalist*, *nonnormative*, and nonmysterious terms.” (Millikan, 1984, 17, emphasis added). Millikan’s goal is to develop a theory of language and content that, among other things, can aid *explanations* for why we have the concepts that we in fact find ourselves with. Meanwhile, our goal is to define *normative* limits on conceptual revision and what it is for a concept to be *good*. For this purpose, the fact that a function can at once be harmful *and* explain why we have the target concept is deeply problematic.

4.4 System functions

As noted earlier, Thomasson (2020a) suggests that we might use Cummins’ notion of system function to flesh out the functional approach to conceptual engineering. Haslanger (2020a, 2020b) endorses this idea. Cummins’ main motivation for developing an account of function was to wrest the notion free from etiological and teleological presupposition (Cummins, 1975). On Cummins’ view, an ascription of function to an item does not (and should not) entail that the functional item *exists* or is present in a system *because* it performs the target function. Instead, Cummins thinks

of function more austere, in terms of the effects an item has on the capacity of a system. His example to illustrate the account is an assembly-line production (Cummins, 1975, 760). The production line constitutes a system with the capacity to produce a commodity. To explain this capacity of the system, we may appeal to the capacities of machines or workers along the line. On Cummins' view, the various tasks performed by components of the system count as functions. Specifically, the capacity of the components count as functions relative to an explanation of the higher-level capacity of the system itself (Cummins, 1975, 758–65).

Although system functions are different in kind from proper functions, the proposal to use them as the operative notion in *Functional Goodness* and *Functional Limits* runs into the same kind of problems. First, there is an epistemic problem. What exactly is the system, the components of which are concepts, and what performance of the system are we trying to explain by appeal to the workings of those components? Haslanger writes that “social, legal, and religious systems lay claim to the concept of MARRIAGE. But the concept has different functions relative to those systems.” (Haslanger, 2020b, 253). However, Haslanger does not tell us what the functions might be in this case, and it is not obvious how we may specify them relative to the systems Haslanger mentions. Recall that what the concept of marriage does for us (socially, legally, religiously) only counts as a system function relative to an explanation of the system's capacity, and it is not clear which capacities those are in this case or how the concept contributes to them. What is the *capacity* of the system which the concept is supposed to contribute to, and what kind of capacity of the concept makes a contribution to the system's capacity? A production line has an easily definable capacity, but this is not true of the system candidates that concepts are components of. This leaves us in a poor position to say what the function of a concept is, because the function of a concept is only a system function relative to an explanation of its contribution to the capacity of the system of which it is part.

The bigger problem will come as no surprise: Operationalizing *Functional Goodness* and *Functional Limits* with the notion of a system function does not go one step toward discounting harmful and insignificant functions. Whatever we consider to be the system of which a concept of interest is a component, it cannot be something perfectly benign or optimally useful. To illustrate, legal and religious systems both have a capacity for oppression. Insofar as a concept consistently contributes to the system's capacity for oppression, the concept is performing a harmful system function. Indeed, the details of Cummins' account are not important on this score, because the notion of a system function does not even begin to select for the functions we ought to engineer in the service of. As the notion of a proper function, the notion of a system function was designed for explanatory purposes—for use in scientific explanation, specifically—not for normative purposes such as determining what revisions of a concept must be like. For our purposes, the notion does little more than add unhelpful technical detail to our account.

5 Consequences for the functional approach

We do not have a notion of function that solves the selection problem for us. None of the notions surveyed in the previous section latch on to all and only the functions we want our concepts to serve. The technical and semi-technical notions we have reviewed may have independent merit, even as applied to concepts, but they do not add much if anything to our effort to solve the normative issues of present concern: of determining what makes a concept good and of determining the limits of revision.

In saying that the limits of revision are defined by functional continuity, or in saying that how a concept should be evaluated is in part determined by its function, it may sound like we are solving normative, methodological problems regarding conceptual engineering by identifying some non-normative property of the target concept, viz. its function. However, the solution is clearly incomplete. First, we have no reason to think that our concepts perform exactly one function each. Any given concept is used for a plurality of different tasks, and is bound up with our speech, thought and action in innumerable ways. On an unsophisticated way of thinking about concept function, the functions of a concept just are the tasks and purposes for which we use a concept. On my understanding, there is nothing problematically unclear about this way of understanding function, but the tasks and purposes for which we use concepts are many and varied. Some are good, some are bad and some are insignificant. This gives rise to the selection problem: The unsophisticated notion of function applies to many functions we do not need, and often better not, engineer in the service of.

Since advocates of the functional approach have suggested more demanding notions of function, it would be foolish not to consider what these notions can do for us in our attempt to spell out *Goodness* and *Limits*. However, a review of the notions on offer makes it clear that whatever notion of function we opt for, we need to make further evaluation of the target functions. It does not suffice to certify that they are functions in this or that technical sense of the word ‘function’.

For the specific task of defining normative constraints on conceptual engineering—as opposed to, say, the task of explaining why we have the concepts we currently do—we do not have any real reason to favor the technical notions of function over an unsophisticated notion. Therefore, the present investigation offers some support to the deflated view of function advocated by Riggs (2021) and Nado (2021a). This view may be summed up by the suggestion to “think of a concept’s functions as being merely extrinsic, relational properties like ‘is used by x for y’.” (Nado, 2021a, 1522) The notion of what we use something for is unsophisticated yet clear. The fact that it applies to harmful and methodologically insignificant function candidates is less problematic than what it might seem, because no matter how technical we make the operative notion of function, we still have to assess whether the target functions really are worth performing and preserving. Indeed, if we judge by the functions practitioners in conceptual engineering have actually appealed to, there is nothing to suggest that they are using a more fancy notion. Woodward suggests that we use our concept of causation to identify “relationships that are potentially exploitable for purposes of manipulation and control.” (Woodward, 2003, 25), and designs concepts of causation that lets us do this. Something similar may be said for Craig, who

assumes that we possess our concept of knowledge in order to “flag approved sources of information” (Craig, 1990, 11).

According to the argument presented, there is no general, non-question-begging way to specify the range of functions we should do conceptual engineering in the service of. Crucially, this does not spell defeat for the functional approach. As argued in Sect. 2, there are good reasons to understand *Goodness* and *Limits* in terms of what we use a concept for. In addition, there clearly exist concept functions in several senses of the word ‘function’. Many of them are worth performing, just not all. The advocate of function therefore has warrant to claim that there exist concept functions, and that some of them play a normatively constraining role for exercises in conceptual engineering. These are comparatively weak claims, but they still have methodological significance. Once we have identified some benign and worthwhile function of a concept, we may evaluate revisions of the concept according to whether and how well the revised concept performs the function in question. In view of the selection problem and the absence of a neat solution to it, conceptual engineers just have to assess functions on a case-by-case basis. Conceptual engineering was never supposed to be easy.

Acknowledgements I thank Delia Belleri, Herman Cappelen, Mirela Fuš, Manuel Gustavo Isaac, Steffen Koch, Mark Pinder, David Plunkett, Joey Pollock, Amie Thomasson, an anonymous reviewer, the audience at a workshop at Ruhr University Bochum and the participants to an online workshop with Concept-Lab (University of Oslo) for valuable comments and discussions.

Author’s Contribution N/A

Funding No funding

Open access funding provided by University of Oslo (incl Oslo University Hospital)

Availability of Data and Material N/A

Declarations

Conflicts of Interest/Competing Interests N/A

Code availability N/A

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ball, D. (2020). “Revisionary Analysis without Meaning Change (Or, Could Women Be Analytically Oppressed?).” In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 35–58. Oxford: Oxford University Press
- Brandom, R. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, Mass: Harvard University Press
- Brigandt, I. (2010). “The Epistemic Goal of a Concept: Accounting for the Rationality of Semantic Change and Variation.” *Synthese*, 177(1), 19–40. <https://doi.org/10.1007/s11229-009-9623-8>
- Brigandt, I. (2011). “Natural Kinds and Concepts: A Pragmatist and Methodologically Naturalistic Account.” In *Pragmatism, Science and Naturalism*, edited by Jonathan Knowles and Henrik Rydenfelt, 171–96. Frankfurt am Main: Peter Lang Publishing
- Cappelen, H. (2018). *Fixing Language: An Essay on Conceptual Engineering*. Oxford: University Press
- Carnap, R. (1963). “Replies and Systematic Expositions.” In *The Philosophy of Rudolf Carnap*, edited by Paul Arthur Schilpp, 859–1013. La Salle, Ill: Open Court
- Craig, E. (1990). *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon
- Cummins, R. (1975). “Functional Analysis.” *The Journal of Philosophy*, 72(20), 741–765. <https://doi.org/10.2307/2024640>
- Foucault, M. (1973). *Madness and Civilization: A History of Insanity in the Age of Reason. Folie et Dérailson: Histoire de La Folie à l'âge Classique*. New York: Vintage
- Gettier, E. L. (1963). “Is Justified True Belief Knowledge?” *Analysis*, 23(6), 121–123
- Hannon, M. (2019). *What's the Point of Knowledge? A Function-First Epistemology*. Oxford: Oxford University Press
- Haslanger, S. (2000). “Gender and Race: (What) Are They? (What) Do We Want Them To Be?” *Noûs*, 34(1), 31–55. <https://doi.org/10.1111/0029-4624.00201>
- Haslanger, S. (2012). *Resisting Reality: Social Construction and Social Critique*. Oxford: Oxford University Press. <https://www.oxfordscholarship.com/view/https://doi.org/10.1093/acprof:oso/9780199892631.001.0001/acprof-9780199892631>
- Haslanger, S. (2020a). “Going On, Not in the Same Way.” In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 230–60. Oxford: Oxford University Press
- Haslanger, S. (2020b). “How Not to Change the Subject.” In *Shifting Concepts—The Philosophy and Psychology of Conceptual Variability*, edited by Teresa Marques and Asa Wikforss, 235–59. <https://doi.org/10.1093/oso/9780198803331.003.0013>
- Isaac, M. G. (2020). “How to Conceptually Engineer Conceptual Engineering?” *Inquiry*, 1–24. <https://doi.org/10.1080/0020174X.2020.1719881>
- Koch, S. (2020). “Engineering What? On Concepts in Conceptual Engineering.” *Synthese*. <https://doi.org/10.1007/s11229-020-02868-w>
- Koch, S. (2021). “Why Conceptual Engineers Should Not Worry About Topics.” *Erkenntnis*. <https://doi.org/10.1007/s10670-021-00446-1>
- Millikan, R. G. (1984). *Language, Thought, and Other Biological Categories*. Cambridge, Mass: MIT Press
- Millikan, R. G. (1993). *White Queen Psychology and Other Essays for Alice*. Cambridge, Mass: MIT Press
- Nado, J. (2021a). “Conceptual Engineering, Truth, and Efficacy.” *Synthese*, 198, 1507–1527. <https://doi.org/10.1007/s11229-019-02096-x>
- Nado, J. (2021b). “Conceptual Engineering via Experimental Philosophy.” *Inquiry* 64(1–2):76–96. <https://doi.org/10.1080/0020174X.2019.1667870>
- Nado, J. (MS). “Strawson’s Challenge and the Problem of Multiple Successors.” *Unpublished Manuscript*
- Plunkett, D. (2015). “Which Concepts Should We Use?: Metalinguistic Negotiations and The Methodology of Philosophy.” *Inquiry*, 58(7–8), 828–874. <https://doi.org/10.1080/0020174X.2015.1080184>
- Plunkett, D., and Tim Sundell (2013). “Disagreement and the Semantics of Normative and Evaluative Terms.” *Philosopher’s Imprint*, 13(23), 1–37
- Price, H. (2011). *Naturalism Without Mirrors*. Oxford: Oxford University Press
- Prinzling, 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>
- Riggs, J. (2021). “Deflating the Functional Turn in Conceptual Engineering.” *Synthese*. <https://doi.org/10.1007/s11229-021-03302-5>

- Scharp, K. (2013). *Replacing Truth*. Oxford: Oxford University Press
- Schiappa, E. (2003). *Defining Reality: Definitions and the Politics of Meaning*. Carbondale, IL: Southern Illinois University Press
- Sider, T. (2011). *Writing the Book of the World*. Oxford: Clarendon Press
- Silberstang, E. (1972). "Rape." *American Jurisprudence* (2nd ed., 65 vol., pp. 757–838). Rochester, NY: Lawyers Cooperative
- Simion, M., and Kelp, C. (2020). "Conceptual Innovation, Function First." *Noûs*, 54(4), 985–1002. <https://doi.org/10.1111/nous.12302>
- Strawson, P. F. (1963). "Carnap's Views on Conceptual Systems versus Natural Languages in Analytic Philosophy." In *The Philosophy of Rudolf Carnap*, edited by Paul Arthur Schilpp, 503–18. La Salle, Ill: Open Court
- Thomasson, A. L. (2020a). "A Pragmatic Method for Normative Conceptual Work." In *Conceptual Engineering and Conceptual Ethics*, edited by Alexis Burgess, Herman Cappelen, and David Plunkett, 435–58. Oxford: Oxford University Press
- Thomasson, A. L. (2020b). *Norms and Necessity*. New York: New York: Oxford University Press. <https://doi.org/10.1093/oso/9780190098193.001.0001>
- Weatherston, B. (2003). "What Good Are Counterexamples?" *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 115(1), 1–31
- Weinberg, J. M. (2006). "What's Epistemology For? The Case for Neopragmatism in Normative Metaepistemology." In *Epistemological Futures*, edited by Stephen Hetherington, 26–47. Oxford: Clarendon Press
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press
- Wright, L. (1973). Functions. *The Philosophical Review*, 82(2), 139–168. <https://doi.org/10.2307/2183766>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.