



Truth in philosophy: a conceptual engineering approach

Jennifer Nado¹

Received: 12 May 2023 / Accepted: 21 February 2024
© The Author(s) 2024

Abstract

The focus of this paper will be to examine the implications that a “practical” approach to conceptual engineering might have for the “traditional” conception of philosophy as uncovering truths about phenomena of philosophical interest. In doing so, I will be building on the ideas of a figure that many take to be the first major philosopher to write on conceptual engineering: Rudolf Carnap. Though the current wave of interest in conceptual engineering goes back less than a decade, many conceptual engineers have found precedent for their views in Carnap’s characterization of what he called “explication.” Interestingly, however, not nearly as much attention has been paid to another Carnapian thesis which seems to me to have deep relevance to methodological questions about conceptual engineering. I have in mind here the distinction between internal and external questions proposed in “Empiricism, Semantics and Ontology” (1950b) and the accompanying claim that external questions lack cognitive content and can be cogently approached only as a matter of pragmatic decision-making. This is the aspect of Carnap’s views upon which I propose to build. I’ll first make the case that there is a suggestive similarity between Carnap’s claim that external questions are matters of pragmatic choice and the practical conceptual engineer’s claim that engineering success should be characterized in terms of suitability to a function or purpose. After that, I’ll look at potential worries about the Carnapian distinction—including, e.g., the concern that it relies on the analytic/synthetic distinction. Finally, I will propose a somewhat modified and expanded version of the distinction and examine its consequences for the thesis that philosophy aims at discovering truths about phenomena of philosophical interest.

Keywords Conceptual engineering · Rudolf Carnap · Truth

The following will likely strike most as a not-particularly-controversial claim: philosophers are primarily in the business of discovering truths about the natures of entities of philosophical interest. One such truth, for example, might be that

✉ Jennifer Nado
jennifernado@gmail.com

¹ Department of Philosophy, University of Hong Kong, Pokfulam, Hong Kong

knowledge is incompatible with being Gettiered. Another might be that killing is *ceteris paribus* more morally problematic than letting die. Yet another might be that water is necessarily identical to H₂O. Of course, whether these cases involve genuine truths is up for debate, but hopefully they serve well enough as examples to elucidate the general idea.¹ The claim that philosophers aim to discover such truths fits nicely with what we might call the “traditional” picture of how (analytic) philosophy operates. On that view, the philosopher aims at a conceptual analysis, ideally in terms of necessary and sufficient conditions, of the aforementioned entities of interest.² That is, she aims to produce something of the form “x is a case of F iff...”. Such analyses are, of course, presented as being true; hence the ubiquitous use of counterexamples as tests of accuracy. If a counterexample to a proposed analysis can be found, the analysis is (at least typically) declared false—and it’s back to the drawing table we go.

I’m not a big fan of this characterization of philosophy. Instead, I’m more inclined towards a recent (?), non-traditional characterization of the aims of philosophy. On this approach, the primary aim of philosophy is not *conceptual analysis*, but *conceptual engineering*. Thus, rather than generating an analysis of (e.g.,) knowledge, the goal is to *engineer improvements* to our concept of knowledge. There are multitudes of unresolved questions regarding exactly how conceptual engineering works, whether it’s even possible, and the conditions under which it should be considered successful.³ I’ll touch on some of these in passing, but my main goal in this paper will be to examine a particular family of views on conceptual engineering and the consequences such views might have for the idea that philosophers primarily aim at discovering truths about the natures of philosophical phenomena.

The family of conceptual engineering views I’ll be focused on is characterized by an emphasis on the functions, purposes, or roles that our concepts perform. Typically, views in this family hold that success in conceptual engineering is in some way or another related to improving a concept’s ability to serve its function or purpose, or to play its role. Just what these functions or roles might be is a matter of extensive debate, but a plausible example might be that our concept KNOWLEDGE has as one of its roles the identification of reliable informants (see, e.g., Craig, 1990; Hannon, 2019)—to categorize someone as a knower is in part to recognize or recommend them as a suitable source of testimony. Another example might be that the concept

¹ For much of the twentieth century, the idea that philosophers primarily aim to discover truths about entities of philosophical interest would likely have been construed as claiming that philosophers are interested in truths about the meanings of certain linguistic terms. On a more recently popular way of understanding this idea, however, it would be more accurate to claim that philosophers are concerned with truths about the natures of philosophically interesting phenomena “in the world”—that is, that they are interested in knowledge rather than “knowledge” (e.g.). The criticisms I give in this paper challenge the “discovery of truths” idea under both interpretations, though I’ll largely frame things in terms of the latter interpretation for expository ease.

² An influential defense of conceptual analysis and its role in philosophy can be found in Jackson (1998). Recently, it has become popular to cast doubt on the “conceptual” aspect of this characterization, on grounds that it reflects a problematic psychologization of philosophy—for further on this, see, e.g., Williamson 2007.

³ See, e.g., Cappelen & Plunkett (2020) and Isaac et al. (2022) for an overview of the territory.

CONSCIOUSNESS has as one of its functions picking out entities that deserve moral concern. Conceptual engineering approaches that focus on purposes or roles of this sort might be termed *practical* approaches,⁴ in the sense that they are focused on what a concept does for us, rather than merely on a concept's semantic properties.⁵

The focus of this paper, then, will be to examine the implications that a practical approach to conceptual engineering might have for the “traditional” conception of philosophy described above. In doing so, I'll be building on the ideas of a figure that many take to be the first major philosopher to write on conceptual engineering: Rudolf Carnap. Though the current wave of interest in conceptual engineering goes back less than a decade, many conceptual engineers have found precedent for their views in Carnap's characterization of what he called “explication”—hence the parenthetical question mark I appended to “recent” a few paragraphs back. As such, conceptual engineers have paid quite a bit of attention to the account of explication that Carnap sketched in his *Logical Foundations of Probability* (1950a). There, Carnap describes explication as the process of “transforming a given more or less inexact concept into an exact one” (Carnap, 1950a, 1950b, 3) and suggests that a successful explication should fulfil four desiderata—similar to the original concept, exactness, fruitfulness, and simplicity.⁶

Interestingly, however, not nearly as much attention has been paid to another Carnapian thesis which seems to me to have deep relevance to methodological questions about conceptual engineering. I have in mind here the distinction between internal and external questions proposed in “Empiricism, Semantics and Ontology” (1950b) (henceforth, ESO) and the accompanying claim that external questions lack cognitive content and can be cogently approached only as a matter of pragmatic decision-making. One proponent of practical conceptual engineering who does make use of the internal/external distinction is Amie Thomasson, who takes the distinction as a jumping-off point for her work on ontology (see Thomasson, 2015, 2016). The particular take on the distinction that I will later advocate is, in fact, much indebted to Thomasson. But although the internal/external distinction is fundamental to her overall philosophical approach, Thomasson has not devoted much explicit discussion to its specific application to conceptual engineering. In what follows, I hope to add to Thomasson's contribution in two ways: first, by examining how the relevance of the distinction can be extended beyond questions of ontology, and second, by using this broader application of the distinction to advocate a move away from a metaphysical focus on truth. Both of these extensions can, I hope, help illuminate the specific applicability of the distinction to conceptual engineering methodology. I must emphasize, though, that I do not intend to attribute any of these broader conclusions to Thomasson, nor to claim that they follow from her specific views. My aim is to provide what I hold to be a plausible (though by no means inevitable) extension of the *general* practical conceptual engineering approach.

⁴ See, e.g., Haslanger (2000, 2020), Simion & Kelp (2020), Thomasson (2020), and Nado (2021).

⁵ Though, of course, some engineers suggest a concept's functions/roles and its semantic features may be linked—see, for instance, Prinzing (2018).

⁶ For further discussion on Carnap's notion of explication and the desiderata he proposes, see, e.g., Hanna (1968), Maher (2007), and Novaes & Reck (2017).

Let me also emphasize that I'm not here primarily interested in Carnap exegesis (in part because I'm not even remotely a Carnap scholar). Instead, I want to ask whether something *in the spirit* of Carnap's distinction between internal and external questions might be utilized to characterize the methods and goals of (practical) conceptual engineering—for instance, to clarify the distinction between engineering and analysis. In the following sections, I'll first make the case that there is a suggestive similarity between Carnap's claim that external questions are matters of pragmatic choice and the practical conceptual engineer's claim that engineering success should be characterized in terms of suitability to a function or purpose. After that, I'll look at potential worries about the Carnapian distinction—including, e.g., the concern that it relies on the analytic/synthetic distinction. Finally, I'll propose a somewhat modified and expanded version of the distinction, largely following Thomasson (2015, 2016), and examine its consequences for the thesis that philosophy aims at discovering truths about the natures of phenomena of philosophical interest. I will also explore how this suggested move away from truth would impact existing metaphilosophical debates, such as those prompted by work in experimental philosophy.

1 Carnap on internal and external questions

Though it's not my primary aim, we'll need to start off with some amateur Carnap exegesis—both to get the broad outlines of the view I will be building on and to get a sense of why one might object to it. In ESO, Carnap claims that questions about existence come in two broad varieties—"internal" questions and "external" questions. As a quick first pass, internal existence questions are those asked from "within" a linguistic framework, and external existence questions are those that attempt to assess the "existence or reality of the framework itself" (Carnap, 1950b, 21–22). It would take a lot more Carnap scholarship than I am capable of to pin down exactly what exactly a linguistic framework is, but as a rough approximation, it is something like a language or language fragment containing the predicates, names, quantifiers, variables, etc. needed to talk about and quantify over the relevant (purported) entities, along with rules for the use of those predicates, names, quantifiers, variables, etc.⁷ Internal questions are, then, questions asked once one has already adopted the relevant framework. As for external questions, by the phrase "existence or reality of the framework itself," Carnap presumably really meant that external questions concern the reality of the *system of entities* the framework concerns. Thus, for instance, an external question about the "mathematics framework" would ask whether mathematical entities exist—and thus, whether the linguistic framework of mathematics succeeds in latching on to something "real."

To get a more careful sense of what Carnap has in mind, it will be helpful to examine the two primary examples Carnap hands us in ESO. The first is the example of the "thing-language"—that is, the linguistic framework that provides rules for

⁷ For discussion of an interpretation of Carnap along these lines, see Eklund (2013).

the use of terms describing or referring to ordinary, observable objects. According to Carnap, internal existence questions about things would include, for instance, “is there a white piece of paper on my desk?” or “are unicorns real?” Internal questions of this sort are, for Carnap, straightforwardly answerable through empirical and logical means. The question “is there a white piece of paper on my desk?,” for instance, can be answered simply by looking at my desk. Affirmative answers to more general internal existence questions, such as “do pieces of paper exist?” or “do physical objects exist?,” follow trivially from affirmative answers to questions like the above.

An external question, by contrast, is intended to question the reality of the framework (that is, the entities) as a whole. In the case of the thing-language, the intention would be to question whether physical things are “*really* real” rather than “just” “real according to the rules of the thing-language framework.” For Carnap, this is the sort of question that philosophers aim to pose when they ask “do physical objects exist?” in the context of a debate about realism regarding physical objects. It is what we would traditionally look upon as a “metaphysical” question. Contrast this with the “internal” interpretation of “do physical objects exist?” given above—Carnap takes the answer to this internal question to be a trivial consequence of the correct answers to other internal questions, and thus not plausibly the focus of protracted metaphysical debate.

As we’ve just seen, when “do physical objects exist?” is asked as an internal question, Carnap holds that the answer is trivial (yes!). But when “do physical objects exist?” is asked as an external question, Carnap claims that the question is malformed. He writes: “to be real in the scientific sense means to be an element of the framework; hence this concept cannot be meaningfully applied to the framework itself” (Carnap, 1950a, 1950b, 22). The only sense to be made of existence, in other words, requires adopting a linguistic framework. Carnap holds that external existence questions, asked “prior to” the adoption of the relevant framework, are meaningless—and as such, not truth-evaluable. As he puts the point: “[philosophers] have not succeeded in giving to the external question and to the possible answers any cognitive content. Unless and until they supply a clear cognitive interpretation, we are justified in our suspicion that their question is a pseudo-question” (Carnap, 1950b, 25).

Though external existence questions of this “theoretical” sort are (for Carnap) meaningless, Carnap holds that we *can* meaningfully consider another type of “external” question: “a practical question, a matter of a practical decision concerning the structure of our language” (Carnap, 1950a, 1950b, 23). That is, we can meaningfully debate the decision to adopt the thing-language, by reference to the practical advantages or disadvantages of such a decision. Our decision about whether to adopt a particular linguistic framework will depend on “the purposes for which the language is intended to be used” (Carnap, 1950b, 23). We might argue for adoption of the thing-language, for instance, on grounds that “[t]he thing language in the customary form works indeed with a high degree of efficiency for most purposes of everyday life” (Carnap, 1950b, 24).

At this point, the connection with the practical approach to conceptual engineering should be reasonably clear. Practical conceptual engineers frame conceptual engineering as a process of inventing or revising concepts with an aim to fit

functions or purposes; a successful case of engineering is one that produces a concept that effectively serves the intended purpose(s). For instance, they might aim to evaluate whether or not the concepts embedded in the “thing-language” really *do* work with a “high degree of efficiency for most purposes of everyday life.” Practical conceptual engineers hold, in other words, that philosophers ought to be in the business of pursuing something very much like Carnap’s “practical” external questions—only at the level of individual concepts/terms, rather than entire linguistic frameworks.⁸

I want to explore this connection, and to pursue answers to a few relevant questions in the vicinity. For instance: if the practical conceptual engineer adopts Carnap’s distinction between internal and external questions, should she also accept Carnap’s accompanying claim that “theoretical” external questions are meaningless? Alternatively, is this view just a remnant of verificationism that the practical conceptual engineer should avoid? If she should avoid it, then how should she do so while still upholding the spirit of the internal/external distinction? Finally, if philosophers should be doing conceptual engineering, would an identification of engineering with the attempt to answer Carnapian “practical” external questions imply that philosophy is mostly about pragmatic decision-making, rather than the discovery of truths?

2 Internal/external questions, verificationism, and the analytic/synthetic distinction

The most obvious reason to resist following Carnap too closely on these matters relates to, as we might say, his philosophical baggage. The internal/external distinction has been broadly criticized, most famously by Quine (1951), whose arguments play off of the distinction’s purported reliance on the analytic/synthetic distinction and a verificationist principle of meaning.⁹ However, while Carnap certainly held such commitments, I’ll argue that it’s not at all clear that the distinction itself requires them. Without such commitments, arguments such as Quine’s plausibly do not undermine the distinction.¹⁰

At least *prima facie*, Quine seems to argue that the internal/external distinction in some sense depends on or reduces to the analytic/synthetic distinction. He writes as follows:

⁸ This is not to claim that conceptual engineering is *inherently* limited to targeting individual concepts rather than entire conceptual or linguistic frameworks. On the contrary, I think a more holistic perspective on conceptual engineering might turn out to be quite fruitful. However, the overall tendency of current conceptual engineering literature is to focus on individual concepts.

⁹ For historical context on the exchange between Carnap and Quine over ESO, see Isaacson (2004), especially pages 235–238.

¹⁰ Despite Quine’s arguments in his (1951), the pragmatic focus of Carnap’s perspective in ESO was certainly amenable to Quine’s overall philosophical perspective. As he writes in “Carnap and Logical Truth”: “I grant that one’s hypothesis as to what there is... is at bottom just as arbitrary or pragmatic a matter as one’s adoption of a new brand of set theory or even a new system of bookkeeping” (Quine 1960, 374).

No more than the distinction between analytic and synthetic is needed in support of Carnap's doctrine that the statements commonly thought of as ontological, viz. statements such as 'There are physical objects,' 'There are classes,' 'There are numbers,' are analytic or contradictory given the language.... The contrast which he wants between those ontological statements and empirical existence statements such as 'There are black swans' is clinched by the distinction of analytic and synthetic (Quine, 1951, 71).

Of course, as one might expect, this implies for Quine that the distinction Carnap presses is ill-founded: "if there is no proper distinction between analytic and synthetic, then no basis at all remains for the contrast which Carnap urges between ontological statements and empirical statements of existence" (Quine, 1951, 71). Quine concludes that statements of existence form a continuum rather than falling into two discrete kinds.

Unfortunately, Quine's argument for this alleged dependence on the analytic/synthetic distinction goes rather quickly, and the details are somewhat difficult to piece together. The problem cannot simply be that the internal/external distinction is *equivalent* to the analytic/synthetic distinction. Carnap explicitly indicates that the answers to internal questions might be *either* analytic or synthetic. Once a linguistic framework is introduced,

internal questions and possible answers to them are formulated with the help of the new forms of expressions. The answers may be found either by purely logical methods or by empirical methods, depending upon whether the framework is a logical or a factual one (Carnap, 1950b, 22).

In other words, at least as I understand Carnap here, he holds that a statement like "there are numbers" (construed internally) would be analytic, whereas a statement like "there are tables" (again, construed internally) would be synthetic.¹¹

Quine's full argument against the internal/external distinction in fact proceeds through a further step. Quine first attempts to reduce Carnap's internal/external distinction to the distinction between what he calls "category" questions and "subclass" questions. The former are "questions of the form 'Are there so-and-so's?' where the so-and-so's purport to exhaust the range of a particular style of bound variables" (Quine, 1951, 68–69). Thus, as I understand it, category questions are, e.g., questions such as "are there numbers?," asked in a context where the relevant domain is numbers, or "are there physical objects?," asked in a context where the relevant

¹¹ Some authors, such as Haack (1976) and Bird (1995), suspect that Quine's argument may have simply reflected a confusion over Carnap's views. Rather unfortunately, and no doubt reflective of the obscurity of Quine's argument, these authors hold precisely opposite views over what exactly the confusion amounts to. Haack claims that according to Quine, Carnap holds that the answers to all internal questions are analytic (or contradictory) and that the answers to all external questions are synthetic (Haack, 1976, 468). Bird, on the other hand, claims that: "[Quine's] suggestion is, with one minor proviso, that the contrast between the internal and the external just is that between the synthetic and the analytic. Internal existence claims are synthetic, and external existence claims are analytic" (Bird 1995, 53). As we have seen, neither of these interpretations squares with what Carnap explicitly says, and I suspect Quine did not get Carnap quite *that* wrong.

domain is physical objects. Subclass questions are, by contrast, questions where the “so-and-so’s” do not purport to exhaust the range of the bound variables. Thus, subclass questions include questions such as “are there *prime* numbers?,” asked in a context where the relevant domain is numbers.

However, Quine does note that these two pairs of categories (internal/external and category/subclass) do not wholly overlap in extension, writing that “the external questions are the category questions conceived as propounded before the adoption of a given language,” and that “the internal questions comprise the subclass questions and, in addition, the category questions when these are construed as treated within an adopted language as questions having trivially analytic or contradictory answers” (Quine, 1951, 69). In other words, Quine grants that category questions may be interpreted either internally or externally.

Quine’s real point, though, is that whether a question is a category question or a subclass question is *trivial*, depending merely on how we understand the ranges of the relevant variables. He writes that “[e]ven the question whether there are classes, or whether there are physical objects, becomes a subclass question if our language uses a single style of variables to range over both sorts of entities.” Thus, whether or not “there are physical objects” is fundamentally a different sort of statement from “there are unicorns” simply comes down to “the rather trivial consideration of whether we use one style of variables or two for physical objects and classes” (Quine, 1951, 69).¹²

It’s at this point that Quine suggests that Carnap in fact doesn’t really need anything beyond the analytic/synthetic distinction in order to maintain the claim that “statements commonly thought of as ontological... are analytic or contradictory given the language” and that they are “proper matters of contention only in the form of linguistic proposals” (Quine, 1951, 71). Thus, it seems that Quine’s argument is not so much that the internal/external distinction reduces to or depends on the analytic/synthetic, but instead that the internal/external distinction is trivial and the analytic/synthetic distinction is really what is doing the work in Carnap’s methodological picture—that is, in setting up a clear contrast between ontological and empirical questions. Interestingly, Quine goes on to say “there is in these terms no contrast between analytic statements of an ontological kind and other analytic statements of existence such as ‘There are prime numbers above a hundred’; but I don’t see why he [Carnap] should care about this” (Quine, 1951, 71). I take it that the idea here is that, while not all analytic statements will turn out to be ontological, all ontological statements will turn out to be analytic (or contradictory).¹³ On Quine’s picture of Carnap, then, the internal/external distinction does not exactly align with either the category/subclass distinction or the analytic/synthetic distinction. But Quine nonetheless thinks the internal/external distinction is unneeded for Carnap’s purposes.

¹² What’s more, Quine adds, Carnap’s view seems to presuppose something like a Russellian theory of types; in modern set theory, there simply is no fundamental distinction between styles of variables, and thus all questions ultimately become subclass questions.

¹³ As I will clarify below, a more accurate statement of the picture here is that all ontological statements when construed as answers to internal questions are analytic/contradictory.

Ultimately, the main suggestions that I would offer about this exchange between Quine and Carnap are the following. Quine believes that Carnap would identify questions “commonly thought of as ontological” with “category” questions, and that Carnap holds that *when such questions are interpreted internally*, their answers are either analytic or contradictory. As we’ve seen, this doesn’t mesh with Carnap’s own statements about claims like “there are physical objects”—such claims would, for Carnap, be synthetic (and empirical) when interpreted internally. By contrast, Quine seems to recognize, correctly, that Carnap would hold that category questions *when interpreted externally* are “properly to be construed as questions of the desirability of a given language form” (Quine, 1951, 69), and thus not plausibly analytic at all.

Oddly, though, Quine seems in these pages to suggest that Carnap’s primary aim is to support the claim that “the statements commonly thought of as ontological... are analytic or contradictory given the language” (Quine, 1951, 71). That is to say, he seems to think that Carnap’s primary focus is on the status of the answers to the internal questions. On the contrary, I take Carnap’s primary aim to be to claim that the statements commonly thought of as ontological are *meaningless*. In other words, Carnap is focused on the status of the answers to the external questions. While it is true, I take it, that Carnap holds that category questions *can* be interpreted internally, such questions are not the ones that entice the metaphysician:

From these [internal] questions we must distinguish the external question of the reality of the thing world itself. In contrast to the former questions, this question is raised neither by the man in the street nor by scientists, but only by philosophers. Realists give an affirmative answer, subjective idealists a negative one (Carnap, 1950a, 1950b, 22).

The metaphysician asks what Carnap would call “theoretical” external questions, and Carnap holds such questions to be ill-formed. Thus, any purported answers to said questions would be neither analytic nor synthetic, but meaningless. Ultimately, Carnap’s suggestion is that we ought to cease asking such questions and instead ask such “practical” external questions as “would it be pragmatically useful to adopt the number-language?” In other words, rather than asking whether there are *really* numbers, we ought to *instead* ask the question of whether the number-language would be profitable to adopt, given our interests or purposes.¹⁴ Carnap is not recommending that metaphysicians settle their disputes via analysis; he is recommending that they *drop* their current disputes and instead focus on pragmatic concerns.

In any case, even if I’ve gotten my Carnap exegesis entirely wrong (which is well within the realm of possibility), it is certainly not the *conceptual engineer’s* primary aim to argue that the answers to ontological questions are analytic. The connection I’ve been exploring between the practical account of conceptual engineering and Carnap’s distinction between internal and external questions focuses instead on Carnap’s claims about the pragmatic nature of the choice between language

¹⁴ See Eklund (2013) for a good attempt to pin down exactly how to interpret Carnap here.

frameworks.¹⁵ Moreover, from the perspective of a practical conceptual engineer, the distinction Quine dwells on between category and subclass questions seems rather beside the point. As I'll argue in the next section, one could ask a practical external "subclass" question just as easily as one could ask a practical external "category" question. In a sense, Quine is quite right that the distinction is trivial; but for this exact reason, it's wholly orthogonal to the pragmatic concerns at issue.

With these points in mind, it seems to me that a conceptual engineer can at least make use of something *much like* Carnap's internal/external distinction without being committed to an analytic/synthetic distinction. My approach here will be much in line with that of Thomasson (2015, 2016), though there may be differences in detail, especially when I attempt to extend the distinction beyond ontological questions (for which, see below). Here is what I propose. External questions, understood theoretically, are meaningless; external questions, understood practically, should be approached by inquiring into the pragmatic advantages of the relevant conceptual framework. The conceptual engineer can follow Carnap on this. Now, Carnap held that the answers to some internal questions are analytic, while others are synthetic—but there is no need to follow Carnap here. I'm convinced by Boghossian's (1996) claim that analyticity, understood in the metaphysical sense of "truth in virtue of meaning alone," is plausibly incoherent: "How could the *mere* fact that S means that p make it the case that S is true? Doesn't it also have to be the case that p?" (Boghossian, 1996, 364). That is to say, *all* sentences are true *both* in virtue of what they mean *and* in virtue of the way the world is. This of course equally applies to *all* existence claims, internally understood.

As for the "epistemic" understanding of analyticity—the idea that a sentence is analytic if merely understanding it suffices to justify belief in its truth—I see no particular need to take a stance either way. Perhaps the notion of epistemic analyticity is legitimate, and the answers to some internal questions are epistemically analytic. But certainly not all are—recall that answers to internal questions include such clearly non-analytic claims as "There is a white piece of paper on my desk." The notion of analyticity, if it is legitimate in any form, does not in any way line up with the internal/external distinction—external (practical) questions and at least a great deal of internal questions are uncontroversially empirical, and whether or not there remains some residue of analytic internal questions seems to be neither here nor there.¹⁶

Carnap's commitment to analyticity, then, does not seem to be required for the purposes to which an engineer might put his internal/external distinction (or at least, a similar distinction in that spirit). What about Carnap's verificationism? Thus far, we have followed Carnap in claiming that external questions, understood theoretically, are meaningless. A commitment to verificationism might be inferred by interpreting Carnap as claiming that external questions are meaningless precisely

¹⁵ Arne Naess (1968, p. 64) similarly emphasizes the pragmatic emphasis of Carnap's overall picture, even going so far as to suggest that philosophical questions are pragmatic questions of language-choice, rather than internal questions or external "theoretical" questions. As will be seen below, I agree—the real goal of philosophy, I'll argue, is to make practical choices between representational systems.

¹⁶ Cf. Thomasson's take on the issue of analyticity in her (2015).

because, external to a linguistic framework, there are no verification conditions for sentences involving terms like “physical object.”

However, this verificationist commitment is stronger than what the practical conceptual engineer needs here. Internal questions are those asked “within a framework,” and their answers depend on the “rules” of that framework. But we don’t need to adopt exactly Carnap’s notion of a linguistic framework, whatever precisely that might be. We can (re-)interpret the idea of asking a question “within a framework” simply as follows. Take the internal question “are there physical things?” This question has an affirmative answer iff “physical thing” has a non-empty extension. The relevant “rules” are simply whatever linguistic facts determine the extension of “physical thing.” *Mutatis mutandis* for other categories we might ask internal questions about: the internal question “are there numbers?” has an affirmative answer iff “number” has a non-empty extension, and so forth.

To demand that “are there physical things?” be interpreted in an external (theoretical) sense, on this picture, would be something like demanding that we answer “are there spliggits?” without first introducing the word “spliggit” (and an accompanying meaning/extension for it) into the language. Again, this read on external questions is more or less the line that Thomasson (2015, 2016) takes—though she is careful to call it an “appropriation” of Carnap, rather than an interpretation. I’ll adopt the same caution. Here is Thomasson’s take on the position I have in mind:

In raising an existence question, we must use a term (‘number,’ ‘property,’ ‘proposition’...) to ask ‘are there numbers/properties/propositions?’ But if we are using those terms according to the rules of use by which they come to be introduced to the language, then those rules enable us to resolve the questions straightforwardly (through analytic or empirical means)... the question is an internal question. So, if the external question is not supposed to be so straightforwardly answerable (so it is not an internal question), then it must be aiming to use the terms in question without their being governed by the standard rules of use. But if they attempt to use the terms while severing them from these rules of use, they make the terms meaningless, and the questions pseudo-questions (Thomasson, 2016, 127).¹⁷

If this is right (and I think it is), then there is a perfectly reasonable, non-verificationist and analyticity-neutral sense to be given to the claim that the only sensible existence questions are ones that are asked “within a framework.” That is, the only reasonable (theoretical) existence questions are ones asked *given* a linguistic framework or *using* a linguistic framework, since the answer to an existence question depends on the meanings of the terms that are used in the question (as well as facts about the world, of course). In other words, you can’t evaluate the truth of “there are spliggits” unless you specify and hold fixed a meaning for “spliggit.” And once you’ve held fixed what “spliggit” means, you are “within” a linguistic framework.

¹⁷ Thomasson’s full view makes use of the use/mention distinction to clarify the difference between internal questions and (cogent) external ones. She credits Huw Price (2009) with initially proposing this connection. I will occasionally make use of the use/mention distinction below, though I have no strong opinions on the degree to which the distinction maps on to the internal/external distinction.

3 Conceptual engineering and external questions

We've concluded that the only answerable (theoretical) existence questions are internal. But of course, there's still some intuitive force to the idea that it's legitimate to ask whether (e.g.,) physical objects "*really*" exist, rather than asking whether they "exist-given-the-rules-of-thing-talk." Or, to state things perhaps more neutrally, there's intuitively a need for some way of inquiring into the overall *legitimacy* of "thing-talk."¹⁸ This is for more or less the reason that motivates conceptual engineering: what if our language/concepts are defective? What if, for example, some of our terms pick out "gruesome" or otherwise distasteful categories? On the interpretation that I've adopted, the answer to "are there grue things?," construed as an internal question, is a fairly trivial "yes." But don't we need the resources to express the idea that grue things aren't really as "legitimate," *qua* entities, as green things, numbers, or pieces of paper?

Of course, we can note that "grue" does not pick out a natural kind; but even leaving aside the philosophical difficulties in characterizing natural kinds, it's certainly not the case that all and only natural kind expressions are "legitimate." Conceptual engineers have frequently been concerned with evaluating and critiquing all sorts of social categories—"marriage" being perhaps one of the most familiar examples.¹⁹ A conceptual engineer will want to claim that a version of the marriage concept that includes same-sex partnerships is superior, in some sense, to a version of the marriage concept that recognizes only male–female pairings. But, of course, neither plausibly picks out a natural kind. If we were instead to ask which best picks out a genuine *social* kind, the heteronormative version of the concept arguably has a better claim—most societies, historically, have recognized only male–female relationships as marriages. But obviously this does not plausibly show that the heteronormative concept is the concept we *ought* to have.

Here, then, is where a practical conceptual engineer might appeal to something in the spirit of Carnap's position in defense of a practical approach to engineering success. She might claim that there is a need to evaluate the legitimacy of our concepts "from the outside", and that the proper way to do so is by approaching such questions pragmatically, largely in the way that Carnap suggests approaching "external" existence questions. Thus, for instance, we would ask whether it is useful or beneficial to adopt a more inclusive marriage concept rather than a heteronormative one. However, we must take some care here. As we've seen, Carnap's concerns in ESO are largely ontological, and he frames his distinction in terms of linguistic frameworks rather than individual concepts. Thus, to really capture the typical sorts of projects that the conceptual engineer pursues, we need to make a few adjustments to Carnap's picture.

¹⁸ I don't intend here anything substantive or technical by the term "legitimate"—it serves as a placeholder for some more substantive normative notion that we can use to evaluate our concepts or classifications. For a practical conceptual engineer, a concept will be legitimate if it is well-suited to its purpose, but we need not presuppose such a view at this stage. All we need at this point is an intuitive sense that *some* notion of conceptual evaluation is needed.

¹⁹ See, for instance, Ball (2020), Ch. 2 of Cappelen (2018).

First, the internal/external distinction needs to be modified such that it applies to individual terms/concepts, rather than linguistic frameworks. I don't think there's any insurmountable barrier to doing this, though doing it rigorously might take some tinkering. The basic thought would be that our practical external questions, rather than asking about the pragmatic utility of adopting a framework, would ask about the pragmatic utility of adopting or altering a single concept/term within a conceptual framework or language.

Second, the internal/external distinction would need to be extended such that it applies more broadly, rather than merely to questions about ontology. Conceptual engineers, in recommending one concept over another, are not necessarily always in the business of making ontological claims. The sorts of questions one asks during the engineering process, moreover, are not always (or even frequently) questions about what exists. But as I'll now argue, the questions conceptual engineers ask *are* plausibly "external," in a sense that I imagine Carnap would be quite happy to accept.

Consider a standard question that a conceptual analyst might ask about knowledge—something like "does knowledge require being non-Gettiered?" Such a question, though not in any obvious sense about ontology, can be viewed as a straightforwardly internal question. According to the interpretation we settled on earlier, the internal *ontological* question "does knowledge exist?" has an affirmative answer iff "does 'knowledge' have a non-empty extension?" does. Let me now suggest that the *non-ontological* question "does knowledge require being non-Gettiered?," considered as an "internal" question, has an affirmative answer iff "is it necessarily true that, if a mental state/relation/etc. is in the extension of 'knowledge,' then it is non-Gettiered?" does. This "internal" reading, I would argue, is the only cogent interpretation of such a question.

This is really just to re-apply what we said earlier about internal existence questions: the answers to such questions hinge on the meanings of the terms involved (as well as the facts about the world). They are not asked "from the outside." In the same way, standard questions asked by conceptual analysts hinge on the meanings of the terms involved. We need not claim that such questions are trivial—after all, Carnap didn't claim that all internal questions are trivial, just that certain entailments are (e.g., "there exists a piece of paper" to "there exists a physical thing"). We also don't need to claim that this shows conceptual analysis to be solely a matter of discovering truths about language—that is, that it concerns truths about "knowledge" rather than truths about knowledge. We can instead say that, given the truism that "knowledge" refers to knowledge, these in practice amount to the same thing. By discovering that that knowledge necessarily requires being non-Gettiered, one also discovers that it is true that, necessarily, if an entity is in the extension of "knowledge," then it is non-Gettiered.

The real moral here is that settling this "internal" question (from a position of semantic ascent or otherwise) does not settle the "external" question of whether we should retain the "knowledge" concept, or revise it. Consequently, insofar as conceptual analysis questions are of this general internal type, they do not settle the sort of "external" questions that Carnap, and conceptual engineers, are interested in. What a conceptual engineer wants to know is whether the fact that knowledge

requires being non-Gettiered is a fact worth bothering about, or whether it is a fact on par with the fact that grueness requires being green if observed before time t . The claim “knowledge requires being non-Gettiered” and the claim “grueness requires being green if observed before time t ” are both true; but is one more “legitimate”?

Similarly, attempting to argue for one of the two marriage concepts discussed earlier using “internal” claims would seem to simply miss the point. Right-wing politicians have long been fond of backing their opposition to marriage equality by baldly asserting things like “marriage is between one man and one woman.” Whether this claim (understood internally) is true or not depends (in part) on the meaning of “marriage.” But the conceptual engineer is not interested in what “marriage” *does* mean—she is interested in what it *should* mean, or alternately whether there is some *other* term with a different meaning that we ought to be employing instead. Perhaps the true metasemantic theory will entail that “marriage” really does apply to only heterosexual pairings, and thus that the internal question “is marriage solely between a man and a woman?” should be answered in the affirmative. This will be of little consequence from a conceptual engineering perspective; it simply further demonstrates the need for amelioration.²⁰

So then, what exactly would be the relevant *external* question here? That is to say—what is the “external” analogue to “does knowledge require being non-Gettiered?,” or “is marriage solely between a man and a woman”? For internal questions, the meaning of the concepts/terms involved is already “fixed.” By contrast, an external question is meant to be one where a key concept/term is itself up for debate—that is, its meaning is *not* fixed. Attempting to frame a straightforward “theoretical” external analogue to the internal knowledge question just mentioned, then, would be to ask something like: “Does knowledge *really* require being non-Gettiered?,” where this is intended to be answered, somehow, “from the outside”—prior to whatever facts fix the meaning of “knowledge.” As we saw earlier, such a question is arguably meaningless, for the fairly trivial and non-verificationist reason that it by hypothesis contains a meaningless component.

Such a question is meaningless, but we *can* make use of semantic ascent to ask a “nearby” question: “should we, in contexts where we currently use ‘knowledge,’ adopt/retain a term/concept whose extension necessarily excludes Gettier cases?” By doing so, we can ask what Carnap would characterize as a “practical” external question. In fact, I’m tempted to claim that if this sort of extension of Carnap’s

²⁰ The view I’m advocating here, essentially, is that metasemantic facts do not constrain the conceptual engineer in the way that many authors have claimed. Many proponents of conceptual engineering offer accounts which are explicitly interwoven with their preferred metasemantic theory (e.g., Cappelen 2018; Prinzing 2018; Sawyer 2020). Moreover, several authors have raised worries about whether or not conceptual engineers can plausibly hope to alter the intensions or extensions of our terms (see, e.g., Cappelen 2018, Koch 2021, Pinder 2021 for this “implementation problem”). But our response to practical external questions need not involve modifying the meanings of terms we currently have; we can adopt new terminology, with the relevant semantic rules fixed by stipulation. The only necessary metasemantic claim, then, would be that meanings can in some cases be fixed via stipulation (which many externalists would be happy to accept). In fact I suspect even this degree of metasemantic constraint may be unneeded—I have more fully explored the possibility of a non-metasemantic view of conceptual engineering in Nado (2021).

notion of a practical external question is sound, then the resulting “broad” version of the internal/external distinction might have been much more acceptable to Quine.²¹ Consider Quine’s closing remarks in “On Carnap on Ontology”: “Carnap maintains that ontological questions, and likewise questions of logical or mathematical principle, are questions not of fact but of choosing a convenient conceptual scheme or framework for science; and with this I agree only if the same be conceded for every scientific hypothesis” (Quine, 1951, 72). It’s beyond the scope of this paper to fully consider whether the evaluation of *every* hypothesis (scientific or otherwise) enjoins us to consider the pragmatic upshots of our choice of classificatory scheme. But certainly, not-obviously-philosophical statements like “Pluto is a planet” can and have been approached “externally” in the sense just outlined. I suspect there is something of a spectrum. We’re unlikely to feel compelled to question the pragmatic merits of the concept “cat” when determining whether the cat is on the mat—unless, say, the alleged cat in question is so atypical that we’re prompted to reexamine our biological taxa. In any case, my proposed extension of the internal/external distinction seems to come closer to what Quine had in mind—at the very least, a *great many* non-ontological hypotheses can be approached as involving choices between conceptual schemes.

Unfortunately, the ubiquitous pressures of space constraint prevent us from considering any further whether this move would have satisfied Quine’s qualms. So let’s instead return to the practical external question just outlined: “should we, in contexts where we currently use ‘knowledge,’ adopt/retain a term/concept whose extension necessarily excludes Gettier cases?” Answering this question, I would argue, does not require fully determining the nature of knowledge, or fully determining the meaning/extension of “knowledge”—in other words, it does not require being in possession of a true conceptual analysis. In particular, it does not require determining whether knowledge necessarily excludes Gettier cases and thus whether “knowledge” applies to Gettier cases. The question can be successfully answered even if we are seriously mistaken or ignorant about knowledge/ “knowledge” (imagine, for instance, that through some metasemantic quirk, “knowledge” somehow refers to all and only true beliefs—we can still debate whether we *ought* to use a term that excludes Gettierization).

The answer to our practical external question doesn’t hinge on the meaning of “knowledge,” at least not in the same way internal questions about knowledge do. But it might, of course, hinge on facts about the contexts in which “knowledge” is used. Substituting “belief” for “knowledge” would not be beneficial from a practical perspective—but to explain why, we need to appeal to usage facts about “knowledge,” such as (for example) the fact that people attribute knowledge in contexts where they wish to endorse another speaker’s value as a source of testimony. Nonetheless, such usage facts need not constitute a counterexample-proof conceptual *analysis*. The relevant usage facts, for a pragmatic conceptual engineer at least, are those that provide insight into the purposes for which a given concept is employed. These will presumably overlap with the usage facts that would be needed to provide

²¹ Thanks to an anonymous reviewer for comments which prompted me to consider this angle.

a correct analysis, but I suspect the overlap will be nowhere near complete. What's more, I would argue that the usage facts required for practical conceptual engineering are clearly empirical—they are better ascertained via linguistics or experimental philosophy than via the armchair. Now, perhaps the usage facts needed for successful conceptual analysis are similarly empirical—but if so, that would cause its own problems for the traditional picture of philosophical methodology.

Thus, an explanation of why we should not substitute “belief” for “knowledge,” though it in some sense hinges on the meaning (or more precisely, the usage) of “knowledge,” would still not be “internal” in the sense we have adopted—we can easily construct our explanation such that it merely mentions, rather than uses, the term “knowledge.” In fact, we can likely avoid even *mentioning* “knowledge”—we can, for example, simply ask what sort of concept would best serve the purpose of expressing endorsement of a source of testimony. The question of whether we ought to use a term that excludes Gettier cases (in such-and-so contexts), then, can and should be answered “from the outside.”

The internal/external distinction, I've claimed, can be applied perfectly well not only to existence questions but also to questions of other sorts with which conceptual engineers are concerned. What's more, this demonstrates exactly why Quine's category/subclass distinction is wholly irrelevant to the spirit of the internal/external distinction, at least as we've appropriated that distinction for conceptual engineering purposes. Quine seems to hold that Carnap takes “general” existence questions to be “category” questions—that is, questions like “are there numbers?” rather than “are there prime numbers?,” or “are there physical objects?,” rather than “are there blue physical objects?” According to this Quinean interpretation of Carnap, “are there blue physical objects?” would, by contrast, be a subclass question. Quine further seems to hold that these category questions are the only ones for which Carnap recognizes an “external” interpretation.

It's not clear to me whether Quine would claim that, for Carnap, “is there knowledge?” is a category question (though I suspect he would not). But I do think it's clear that “does knowledge require being non-Gettiered?” is certainly not a category question in anything like the sense Quine has in mind. Nevertheless, as we've just discussed, there is a plausible (practical) external equivalent to “does knowledge require being non-Gettiered?” This is the question of whether we should adopt a term that excludes Gettier cases. Now, perhaps Quine would not count “does knowledge require being non-Gettiered?” as a subclass question either, since subclass questions seem to be limited to questions about existence. But certainly, the existence question “is there Gettiered knowledge?” would be classified by Quine as a subclass question, in the same sense that “are there prime numbers?” or “are there black swans?” is. And similarly, there is a plausible external equivalent to “is there Gettiered knowledge?” This would be something like “should we, in contexts where we currently use ‘knowledge,’ adopt a term/concept ‘knowledge*’ such that ‘there is Gettiered knowledge*’ turns out true?”²² If all this is correct, then it is perfectly

²² This last should be read as allowing for the possibility that “knowledge*” is identical in meaning to “knowledge,” thus allowing for the possibility that our current concept should be retained as-is.

possible to ask (practical) external subclass questions, as well as (practical) external category questions. And, assuming that the category/subclass distinction is meant only to apply to existence questions, it is also possible to ask (practical) external questions that fall outside of that distinction's scope. Thus, the "trivial" distinction Quine accuses Carnap of relying on is simply orthogonal to the internal/external distinction—at least in our adapted version thereof.

Let's now generalize from the knowledge case. I suggest that the sort of questions that fans of a practical approach to conceptual engineering should pursue are roughly of the following form: "Should we, in contexts *x*, *y*, *z*..., adopt a term/concept with features *F*, *G*, *H*...?"²³ The "features" here will typically be related to the term's meaning, e.g., necessarily excluding Gettierization, applying to both homosexual and heterosexual partnerships, and so on. But conceivably, said features could also be non-semantic: for instance, features like being easy to learn or being associated with positive connotations. Moreover, practical engineers should hold that these external questions ought to be approached via consideration of the purposes the term/concept in question would be intended to serve. For "knowledge," this could include, e.g., the purpose of picking out reliable testifiers, acting as the norm of assertion, and so forth. Determining the relevant purposes is, of course, not a trivial issue—but this is a general problem for practical conceptual engineers, and one that must be left for another day.

4 Truth in philosophy

Let's finally return to the notion of truth. What does all of this imply about the role of truth in philosophy? Should we join Carnap in suggesting that the conceptual engineer's practical external questions are "not of a cognitive nature" (Carnap, 1950b, 23)? If so, then does it follow that these questions are not truth-evaluable and thus that the practical conceptual engineer is not ultimately interested in discovering truths?

Carnap indeed construes answering practical external questions as being "a matter of decision" and holds that the questions themselves are not "theoretical" (Carnap, 1950b, 23). But although the decision itself is "not of a cognitive nature," it isn't wholly divorced from considerations of truth. This is because, as Carnap observes, it will "usually be influenced by theoretical knowledge" (about, e.g., the effectiveness of the framework for a given purpose). However, "these

²³ A reviewer suggests that, while my use of the internal/external distinction does not depend on the analytic/synthetic distinction, it may depend on a distinction between questions of meaning and questions of fact. But the "external" questions I've posed here need not be asked in terms of meaning. We might ask, for instance: "how should we categorize Gettier cases? Should we include them in a grouping with standard cases of justified true belief, or separate them?" Or, alternately, we might ask something like "Which phenomenon is more worth talking about—knowledge or knowledge*?" (where the latter is some alternate epistemic phenomenon like justified true belief, etc.). It's true one would have to stipulate a meaning for "knowledge*" before asking this latter question, but once that has been done, the question itself is not about meanings. Similarly, on the view I've presented, the questions "concerning facts" can equally well be stated in terms of meaning ("does inclusion in the extension of 'knowledge' require being non-Gettiered?").

questions cannot be identified with the question of realism. They are not yes–no questions but questions of degree... it would be wrong to describe this situation by saying: ‘The fact of the efficiency of the thing language is confirming evidence for the reality of the thing world’; we should rather say instead: ‘This fact makes it advisable to accept thing language’” (Carnap, 1950b, 24). That is to say, there are truths we ought to be concerned with—truths about the practical utility of adopting a bit of language—but these truths do not carry ontological consequences regarding the referents of the expressions under discussion.

Let’s apply these claims to the picture I have defended of practical conceptual engineering as asking “pragmatic” external questions about our concepts. In the opening paragraph of this paper, I claimed that most philosophers would find it uncontroversial that philosophers aim at discovering truths like “knowledge is incompatible with being Gettiered.” But typically, such claims are put forward as answers to what I’ve argued are “internal” questions—questions about, say, the nature and/or essential features of knowledge. A practical conceptual engineer, however, focuses her efforts not on internal questions, but on external ones. She is only concerned with whether knowledge excludes Gettier cases if it turns out that knowledge is a category worth bothering about—if it is a category whose use would be well-suited to our purposes. It’s this latter question—whether knowledge is well-suited to our purposes—that is the primary target.

But as we’ve just seen, the philosopher qua conceptual engineer *will* inevitably be concerned with certain truths—just not the ones we expected, and perhaps not as the ultimate goal of philosophical activity. We won’t be primarily concerned with ascertaining the truth of claims like “knowledge requires being non-Gettiered.” Instead, we will be concerned with truths which will help us determine the effectiveness of our concepts for various purposes—that is, with truths that fill the role Carnap gives to “theoretical knowledge” in answering external practical questions. Thus, the most pertinent truths for a practical conceptual engineer will be ones like “within the contexts in which we currently use ‘knowledge,’ adopting a concept/term/category which requires being non-Gettiered would be conducive to purposes *x*, *y*, *z*.” Discovering the truth of such claims will not be the endpoint of philosophical activity, but a means to decision-making; they will guide us in determining which concepts to adopt. And plausibly, assessing the truth of such claims will not always be a task for the philosopher—in some cases, said assessment will largely hinge on empirical facts that must be investigated by (e.g.,) the physical or social sciences. Of course, the practical conceptual engineer will accept that investigations of the natures of philosophical phenomena will still be *one* of philosophy’s goals—but only *after* those phenomena are selected from amongst the plenitude of things that one might investigate. Knowledge, for instance, may or may not be one of the phenomena that ultimately demands philosophical scrutiny; perhaps, some custom-tailored epistemic concept would help us “latch on” to a better target. We *may* end up interested in whether knowledge is incompatible with being Gettiered; but only if knowledge proves to be the best category for the various purposes for which it is invoked. Suppose it turns out, for instance, that knowledge* better predicts and explains human action, better guides practical reasoning, and better facilitates identification of reliable sources

of testimony. Under such circumstances, the investigation of the essential features of knowledge would start to look rather less pressing. From the perspective of a practical conceptual engineer, I would argue, most of the “internal” debates that philosophers engage in should instead be approached externally. Rather than ask whether free will is compatible with determinism, we should ask whether it would best suit our purposes if “free will” were used in such a way that “free will is compatible with determinism” comes out true. Rather than ask whether justification is internal, we should ask whether an internalist or externalist version of “justification” would be more useful. It’s only after those external questions are settled that we will know which internal questions need asking.

What’s more, depending on how the phenomena of (practical) philosophical interest are pinned down, there may be little residual “internal” philosophical investigation to be done. In many cases, a carefully stipulated definition might fully obviate the need for anything like conceptual analysis—informative necessary and sufficient conditions for membership in the referent category might already be laid bare by the stipulation. If the conceptual engineer has already argued that an externalist notion of justification—let us call it justification-E—better serves the purposes to which “justification” is put, then there is obviously no longer any need for the conceptual analyst to determine whether or not justification-E is externalist. That’s not to say that there will be *no* philosophical work *at all* to be done after the engineer has had her say. But the remaining philosophical questions will not likely be of the form “is such-and-so hypothetical case a case of F?” or “is possessing feature x necessary for being an F?”.

I’ve advocated a shift from viewing philosophy as primarily a matter of *discovery* to viewing it as primarily a matter of *decision* (theoretically informed decision, of course). If this picture of philosophy is on the right track, then beyond being interesting in of itself, I would argue that the shift has important implications in other areas of metaphilosophical debate. One in particular that I’ll mention is this: it complicates debates over the reliability of intuition as a source of evidence for philosophical theorizing. Many (though by no means most) proponents of experimental philosophy have suggested that their findings call into question intuition’s reliability (for just a few examples, see, e.g., Alexander & Weinberg, 2007; Nadelhoffer & Feltz, 2008; Tobia et al., 2013; Machery, 2017). Experimental evidence arguably suggests that judgments in response to thought experiment prompts tend to vary as a function of irrelevant factors like cultural background and order of presentation. Sensitivity to such irrelevant factors is, (some) experimentalists claim, indication of unreliability. Others, however, have argued that intuition is actually more stable than the studies suggest, or that it is nonetheless on the whole reliable even if it sometimes errs (see, e.g., Wright, 2010; Nagel, 2012; Knobe, 2019).

On the picture I’m proposing, intuition’s reliability on standard thought-experiment judgments like “is (case described in vignette) a case of knowledge” is of tangential methodological import at best. Such questions are internal questions—and I’ve argued that the philosopher’s primary goal should not be to seek true answers to such questions. Even were it to turn out that our intuitions on such cases were highly reliable, it would not mitigate the possibility that the truths thus attained are couched in sub-standard concepts. If “knowledge” picks out a category that is not

well-suited to the various purposes to which it is put, having reliable intuitions about knowledge will be about as epistemically valuable as having reliable intuitions about grue. The real concern should be whether or not our intuitions (and/or our words or concepts) track a category that is worth tracking. If this is right, then argumentation on both sides may need to be reframed—reliability is simply not the right metric for assessing intuition’s epistemic worth.

Ultimately, rejecting a standard reliability-focused examination of intuition’s epistemic credentials may be more of a boon to “negative” experimentalists than to “traditionalists.” Sensitivity to irrelevant factors quite plausibly still indicates that intuition is not an ideal guide to concept choice; carefully arguing for this modified conclusion, however, might require a few extra steps. Exploring the needed changes further would, unfortunately, take us outside the scope of this paper. It will be less straightforward to re-frame the defenders’ arguments. Demonstrating the robustness or overall reliability of our intuitions, after all, doesn’t seem to be particularly direct evidence for the pragmatic utility of the underlying concepts. Other defenses of intuition face similar problems. Nearly all such defenses focus on supporting the claim that intuition justifies belief; but even if intuition justifies me in believing that (it is true that) Gettier cases are not cases of knowledge, this does not obviously address the external question of whether knowledge is a pragmatically useful concept to have. For a fan of the traditional intuition-based view of philosophical method, a new approach entirely might be called for.

Either way, the general moral is that a practical conceptual engineer should urge that we re-frame how we approach methodological debates over intuition. The question should not be whether our intuitions reveal truths about, e.g., knowledge, but whether a systematization of our intuitive classification judgments would generate categories that are pragmatically useful. My suspicion is that such categories might be *reasonably* well-suited to our purposes, but certainly not optimally so. Again, a shift of focus may be warranted—rather than asking whether our intuitions are true or false, we should be asking whether the classificatory verdicts they generate make our concepts better or worse.

Summing up, then, my suggestion in this paper is that the practical conceptual engineer ought to adopt a Carnapian stance on the primary aim of philosophy. On such a view, the primary aim of philosophy it is to *make decisions* about which concepts to adopt, rather than to *discover truths* about, e.g., what knowledge is, or whether it exists—much less the elusive “truths” about what knowledge “really” is, or whether it “really” exists. Like Carnap, I would suggest that we not view the practical utility of, e.g., the “knowledge” concept as evidence that knowledge “really” exists (in the external sense), but rather merely as a consideration in favor of the adoption of “knowledge.” The nature and existence of knowledge may be worth debating, but only if “knowledge” is the concept we ought to be using—and it’s this latter question that will take the most philosophical work. Truths about the nature of knowledge are simply of secondary interest to the question of whether knowledge is the right category to be investigating in the first place.

Data availability No data was generated or analyzed as part of this project.

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

- Alexander, J., & Weinberg, J. M. (2007). Analytic epistemology and experimental philosophy. *Philosophy Compass*, 2(1), 56–80.
- Ball, D. (2020). Revisionary analysis without meaning change (or, could women be analytically oppressed?). In A. Burgess, D. Plunkett, & H. Cappelen (Eds.), *Conceptual Engineering and Conceptual Ethics* (pp. 35–58). Oxford University Press.
- Bird, G. H. (1995). Carnap and Quine: Internal and external questions. *Erkenntnis*, 42(1), 41–64.
- Boghossian, P. (1996). Analyticity Reconsidered. *Noûs*, 30(3), 360–391.
- Cappelen, H., & Plunkett, D. (2020). A guided tour of conceptual engineering and conceptual ethics. In A. Burgess, H. Cappelen, & D. Plunkett (Eds.), *Conceptual Ethics and Conceptual Engineering* (pp. 1–26). Oxford University Press.
- Cappelen, H. (2018). *Fixing language: An essay on conceptual engineering*. Oxford University Press.
- Carnap, R. (1950a). *Logical foundations of probability*. University of Chicago Press.
- Carnap, R. (1950b). Empiricism, semantics, and ontology. *Revue Internationale de Philosophie*, 4(11), 20–40.
- Craig, E. (1990). *Knowledge and the state of nature: An essay in conceptual synthesis*. Clarendon Press.
- Eklund, M. (2013). Carnap's metaontology. *Noûs*, 47(2), 229–249.
- Haack, S. (1976). Some preliminaries to ontology. *Journal of Philosophical Logic*, 5(4), 457–474.
- Hanna, J. F. (1968). An explication of "explication." *Philosophy of Science*, 35(1), 28–44.
- Hannon, M. (2019). *What's the point of knowledge? A function-first epistemology*. Oxford University Press.
- Haslanger, S. (2000). Gender and race: (What) are they? (What) do we want them to be? *Noûs*, 34(1), 31–55.
- Haslanger, S. (2020). How not to change the subject. In T. Marques & A. Wikforss (Eds.), *Shifting Concepts: The philosophy and psychology of conceptual variation* (pp. 235–259). Oxford University Press.
- Isaac, M., Koch, S., & Nefdt, R. (2022). Conceptual engineering: A road map to practice. *Philosophy Compass*, 17, e12879.
- Isaacson, D. (2004). Quine and logical positivism. In R. Gibson (Ed.), *The Cambridge Companion to Quine* (pp. 214–269). Cambridge University Press.
- Jackson, F. (1998). *From metaphysics to ethics*. Oxford University Press.
- Knobe, J. (2019). Philosophical intuitions are surprisingly robust across demographic differences. *Epistemology & Philosophy of Science*, 56(2), 29–36.
- Koch, S. (2021). The externalist challenge to conceptual engineering. *Synthese*, 198, 327–348.
- Machery, E. (2017). *Philosophy within its proper bounds*. Oxford University Press.
- Maher, P. (2007). Explication defended. *Studia Logica*, 86(2), 331–341.
- Nadelhoffer, T., & Feltz, A. (2008). The actor–observer bias and moral intuitions: Adding fuel to Sinnott-Armstrong's fire. *Neuroethics*, 1(2), 133–144.
- Nado, J. (2020). Taking control: Conceptual engineering without (much) metasemantics. *Inquiry*. <https://doi.org/10.1080/0020174X.2020.1850342>
- Nado, J. (2021). Conceptual engineering, truth, and efficacy. *Synthese*, 198, 1507–1527.
- Naess, A. (1968). *Four modern philosophers: Carnap, Wittgenstein, Heidegger*. University of Chicago Press.

- Nagel, J. (2012). Intuitions and experiments: A defense of the case method in epistemology. *Philosophy and Phenomenological Research*, 85(3), 495–527.
- Novaes, C. D., & Reck, E. (2017). Carnapian explication, formalisms as cognitive tools, and the paradox of adequate formalization. *Synthese*, 194(1), 195–215.
- Pinder, M. (2021). Conceptual engineering, metasemantic externalism and speaker-meaning. *Mind*, 130, 141–163.
- Price, H. (2009). Metaphysics after Carnap: The ghost who walks? In D. Chalmers, R. Wasserman, & D. Manley (Eds.), *Metametaphysics: New essays on the foundations of ontology* (pp. 320–346). Oxford University Press.
- Prinzing, M. (2018). The revisionist's rubric: Conceptual engineering and the discontinuity objection. *Inquiry*, 61(8), 854–880.
- Quine, W. V. O. (1951). On Carnap's views on ontology. *Philosophical Studies*, 2(5), 65–72.
- Quine, W. V. O. (1960). Carnap and logical truth. *Synthese*, 12(4), 350–374.
- Sawyer, S. (2020). Truth and objectivity in conceptual engineering. *Inquiry*, 63, 1001–1022.
- Simion, M., & Kelp, C. (2020). Conceptual innovation, function first. *Noûs*, 54(4), 985–1002.
- Thomasson, A. L. (2015). *Ontology made easy*. Oxford University Press.
- Thomasson, A. L. (2016). Carnap and the prospects for easy ontology. In S. Blatti & S. Lapointe (Eds.), *Ontology After Carnap* (pp. 122–145). Oxford University Press.
- Thomasson, A. L. (2020). A pragmatic method for conceptual ethics. In H. Cappelen, D. Plunkett, & A. Burgess (Eds.), *Conceptual Ethics and Conceptual Engineering* (pp. 435–458). Oxford University Press.
- Tobia, K., Buckwalter, W., & Stich, S. (2013). Moral intuitions: Are philosophers experts? *Philosophical Psychology*, 26(5), 629–638.
- Williamson, T. (2007). *The philosophy of philosophy*. Blackwell.
- Wright, J. C. (2010). On intuitional stability: The clear, the strong, and the paradigmatic. *Cognition*, 115(3), 491–503.